30. The computer system of claim 29, wherein the web server module signals control information to the media playback component and a web browser component of the Internet enabled device, to cause the media playback component of the network enabled device to programmatically load and playback the media web resource while the web browser displays a web page of a web site corresponding to the media web resource being played back.

31. A computer system for playing back media from a network, the computer system comprising:

a network enabled platform comprising a media playback component;

a database comprising a plurality of addresses, each address locating a media resource on the network, each address accessing a media resource; and

a network server module that coupleable to the media playback component and the database, the network server being able to receive a request for media playback from the network enabled device, to select multiple addresses from the database, to signal the multiple addresses to the network enabled device, and to control a media playback component on the network enabled device to use the addresses to automatically access and play back the media resource associated with the addresses.

32. The computer system of claim 31, wherein the network server module resides on the network.

33. The computer system of claim 31, wherein at least a portion of the network server module resides on the network enabled device.

34. The computer system of claim 31, wherein the network server module comprises a sequence of code that is signaled to the network enabled device from another device on the network, the sequence of code being executable on the network enabled device.

35. The computer system of claim 31, wherein the network enabled device is an Internet enabled device comprising the media playback component and a

network browser component, the database comprises a plurality of links to media web resources, wherein the network server module programmatically controls the media playback component of the Internet enabled device to access and playback the media web resource of each link signaled to the Internet enabled device.

36. The computer system of claim 31, wherein the network enabled device further comprises a web browser component, and wherein the network server module programmatically controls the web browser component to display a web page associated with the media web resource being played back.

37. The computer system of claim 31, further comprising a play-list module coupleable to the network server module, the play-list module generating a play-list based on one or more criterias designated by a user of the network enabled device, the play-list comprising a plurality of addresses that are signaled to the network enabled device in response to the request for media playback.

38. The computer system of claim 31, wherein the network enabled device is a terminal device selected from a group of devices consisting of a personal computer, a handheld computer, a wireless Internet enabled device, and a handheld computer.

39. The computer system of claim 31, wherein the database comprises a plurality of links to Internet sites, wherein each link accesses a media web resource and contains no other internal link.

40. The computer system of claim 31, wherein all or a portion of the database resides on the network enabled device.

41. The computer system of claim 31, wherein all or a portion of the database is accessible to the network enabled device over the network.

42. A network enabled device configured to playback media from a network, the network enabled device being coupleable over to a database

comprising a plurality of addresses, each address locating a media resource on a network, the network enabled device comprising:

a user-interface to prompt for a search request;

a network interface to signal the request to a network server module that is communicatable with the database, and to receive one or more addresses in the database that match the search request; and

a media playback component being configured to be programmatically controllable by the network server module to automatically load the media resource associated with each address signaled to the network enabled device upon accessing the media resource.

43. The network enabled device of claim 42, wherein the user interface is controllable by the network server module after the network server module signals the response, the user-interface providing one or more prompts to enable an end user to control the media played back from the one or more media resources associated with the addresses signaled from the network server module.

44. The network enabled device of claim 42, wherein the user-interface prompts for a user to signal the network server module to immediately access a next address signaled to the network enabled device from the network server module.

45. The network enabled device of claim 44, wherein the user-interface prompts for the end user to signal the media playback component to repeat playing back a previously played media resource.

46. The network enabled device of claim 42, wherein the user interface prompts for a specific media resource from the media resources signaled from the network server module for playback, and to pause playback from the media resources signaled from the web server module.

47. The network enabled device of claim 42, wherein the network access component is a web browser residing on an Internet enabled device.

48. The network enabled device of claim 42, wherein all or a portion of the database is accessible to the network enabled device over the network.

49. The network enabled device of claim 42, wherein all or a portion of the database resides on the network enabled device.

50. The network enabled device of claim 42, wherein the network device is Internet enabled, wherein the network interface signals the request to a web server module that is communicatable with the database to receive one or more links to web sites containing media web resources, the one or more links containing no internal links, and wherein the Internet enabled device further comprises a web browser component, the web browser component being controllable by the web server module to automatically access and display a web page associated with the media web resource being played back by the media playback component.

51. The network enabled device of claim 50, wherein the network interface includes a communication port between the network enabled device and the network.

52. A network enabled device configured to play back media from a network, the network enabled device being coupleable over the network to a database comprising a plurality of addresses, each address locating a media resource on the network, the network enabled device comprising:

a user-interface including a plurality of user-interactive features, a first user-interactive feature in the plurality of user-interactive features prompting to receive a search request for media playback;

a network interface that signals the request to a network server module upon the first user-interactive feature receiving the search request for media playback, the network interface being communicatable with the database to receive one or more addresses in the database that match the search request; and

a media playback component being configured to be programmatically controllable by the network server module to automatically load the media resource associated with each address signaled to the network

enabled device upon accessing the media resource, a playback of the media playback component being controllable by one or more control user-interactive features in the plurality of user-interactive features.

53. The network enabled device of claim 52, wherein the first user-interactive feature in the plurality of user-interactive features is selected from a group consisting of selectable icons, text-entry field, and menu items.

54. The network enabled device of claim 53, wherein the control user-interactive features are selected from a group of selectable icons having functions corresponding to skipping to a next address selected to be signaled to the network enabled device, pausing the media resource being played back; repeating the media resource previously played back, and stopping playback from the media playback component.

55. The network enabled device of claim 53, wherein the user-interface includes a plurality of user-interactive features to select a characteristic of a media playback from the media playback component.

56. The network enabled device of claim 55, wherein the user-interface features include the plurality of user-interactive features to select a genre, an artist, or a title of a network media resource for playback .

57. The network enabled device of claim 56, wherein the search request is composed by selecting one or more of the user-interactive features of the user-interface.

58. A method for sharing media playback on a first terminal coupled to a network with other terminals coupled to the network, the method comprising:

signaling a link in a play-list to the first terminal, the play-list including one or more links to media files on the network;

receiving a signal from the first terminal to message a second terminal while the first terminal is playing back media from one of the media files;

signaling the second terminal a selectable link to enable the second terminal to receive one or more links in the play-list; and

signaling the one or more links in the play-list to the second terminal upon the selectable link being selected by the second terminal.

59. The method of claim 58, wherein signaling the second terminal a selectable link includes messaging the second terminal the link to the play-list.

60. The method of claim 58, wherein signaling the second terminal a selectable link includes signaling a message to the second terminal that causes the second terminal to access the play-list once the message is opened.

61. The method of claim 58, wherein signaling the one or more links in the play-list to the second terminal includes signaling each link in the play-list to the first terminal and to the second terminal concurrently.

62. The method of claim 58, further comprising receiving the signal from the first terminal to message the second terminal while the first terminal is playing back media located by a link in the play-list, and causing the media playback component of the second terminal to load the media file being played by the first terminal.

63. The method of claim 58, wherein signaling a link to the second terminal that locates the play-list module includes signaling the link to include arguments that locate the address in the play-list corresponding to the media file being played back on the first terminal.

64. A system for sharing media playback from a network between a plurality of terminals, the plurality of terminals including a first terminal and a second terminal, the system comprising:

a play-list component locatable on the network by a selectable link, the play-list component identifying a plurality of links to form a play-list, each link in the play-list locating a media file on the network;

a network server module signaling the plurality of links that form the play-list to the first terminal, the network server module being able

to receive a signal to transmit the selectable link to the second terminal to enable the second terminal to locate the play-list module.

65. The system of claim 64, wherein the network server module signals the plurality of links that form the play-list to the second terminal upon the selectable link being transmitted to the second terminal is selected.

66. The system of claim 64, wherein the messaging module transmits the selectable link as a content of a message.

67. The system of claim 64, wherein the play-list module transmits the selectable link with a corresponding message to the second terminal so that the selectable link is selected when the message opened.

68. The system of claim 65, wherein the play-list module identifies the plurality of links to form a plurality of play-lists, and wherein the messaging module transmits the selectable link with identifiers that identify the play-list being signaled to the first terminal.

69. The system of claim 68, wherein the messaging module transmits the selectable link with identifiers that identify the media file being played back on the first terminal.

70. The system of claim 69, wherein the network server module identifies the link in the play-list being played back on the first terminal and transmits the identified link with a remainder of the links in the play-list to the second terminal, so that the first terminal and the second terminal simultaneously playback the same media files while the links in the play-list are signaled to each of the first terminal and the second terminal.

71. A network enabled device comprising:

a media playback component configured to communicate with a network-side module to receive a first plurality of links, each of the first plurality of links locating a media file on a network;

a web browser component configured to receive a second plurality of links, each of the second plurality of links hosting a media file

located by one of the first plurality of links, the web browser component displaying the web site for each of the second plurality of links that corresponds to the media file being played back by the media playback component.

72. The network enabled device of claim 71, further comprising a user-interface including user-interactive features for controlling media output from media files located by the first plurality of links.

73. The network enabled device of claim 72, wherein the user-interface is displayed in a first window, and the web browser component displays the web site for each of the second plurality of links in a second window.

74. The network enabled device of claim 73, wherein the media playback component is configured to sequentially receive the first plurality of links from the network-side module, and the web browser displays the web sites hosting media files for each of the web sites in the sequence.

75. A system for playing back media from a network, the system comprising:

means for receiving a search criteria from a network enabled device;

means for accessing a database comprising a plurality of network addresses, the database associating each address with one or more classes of information, each address accessing a media resource;

means for selecting at least one address in the database using the search criteria;

means for signaling the selected address to the network enabled device; and

means for controlling the network enabled device so as to automatically access and play back the media resource of the selected address

76. A system for playing back media from a network, the system comprising:

means for receiving a request for media playback from a network enabled device;

means for accessing a database comprising a plurality of network addresses, each address accessing a media resource;

means for identifying at least two addresses from the database;

means signaling each identified address to the network enabled device; and

means for controlling the network enabled device to access and automatically play back the media resources of each of the signaled addresses.

77. A method to locate web resources on the Internet, the method comprising:

accessing a web site;

identifying a plurality of links using a web browser component, each of the plurality of links being selectable to open a corresponding web resource of a specified data type on the web site; and

making the plurality of links available to a plurality of Internet enabled devices that select one or more of the links.

78. The method of claim 77, further comprising arranging the plurality of links in a database that is accessible over the Internet to the plurality of Internet enabled devices.

79. The method of claim 78, further comprising identifying one or more classes of information about each of the plurality of links using a media playback component, and storing each of the plurality of links in association with the one or more classes of information.

80. The method of claim 79, further comprising making the one or more classes of information available to the plurality of Internet enabled devices.

81. The method of claim 77, wherein accessing a web site includes accessing a web site identified by a metacrawler.

82. The method of claim 77, wherein identifying a plurality of links includes identifying links that only access web resources including audio data types, video data types, or a combination of audio and video types.

83. The method of claim 77, wherein identifying a plurality of links using a web browser component includes reconfiguring a commercially available web browser.

84. The method of claim 77, wherein identifying a plurality of links using a web browser component includes programmatically controlling the web browser component through an application programmable interface of a web browser to access the web resource for each of the plurality of links.

85. The method of claim 84, wherein programmatically controlling the web browser through an application programmable interface includes programming the web browser to bypass a shell of the web browser.

86. The method of claim 85, wherein programmatically controlling the web browser through an application programmable interface includes programming the web browser to reduce functionality of the web browser in displaying and translating data structures on each web site.

87. The method of claim 84, wherein programmatically controlling the web browser component through an application programmable interface further comprises using the web browser component to parse a web page that provides access to each link in the plurality of links to identify the plurality of links programmatically.

88. The method of claim 77, further comprising verifying that each of the plurality of links are selectable to open the corresponding web resource of the specified data type upon being selected.

89. The method of claim 88, wherein verifying that each of the plurality of links are selectable includes verifying that the web resources of each available on the network.

90. The method of claim 89, wherein verifying that each of the plurality of links are selectable to open the corresponding web resource of the specified data type includes loading each of the plurality of media web resources into a media playback component.

91. A system for locating web resources on the Internet, the system comprising:

a web browser component;

a database;

a search module that controls the web browser component to access at least one web site, the search module controlling the web browser component to identify a plurality of links to media web resources at the web site, each of the plurality of links being selectable to open a media web resource, the search module storing the plurality of links in the database.

92. The system of claim 91, wherein the search module signals an application programmable interface of the web browser an address of the web site so as to cause the web browser component to load and parse the web site for the plurality of links to media web resources.

93. The system of claim 91, wherein the search module signals an application programmable interface of the web browser an address of the web site so as to cause the web browser component to load and parse the web site for the plurality of links to media web resources without displaying the media web resources.

94. The system of claim 92, wherein the search module signals the application programmable interface of the web browser the address of the web site so as to cause the web browser component to load and parse the web site for the plurality of links to media web resources of a video data type, an audio data type, or a combination of the audio and video data types.

95. The system of claim 94, wherein the search module causes the web browser component to load and parse the plurality of links to media web resources of data types including animation and/or still images

96. The system of claim 91, further comprising a metacrawler module to locate a plurality of web sites for the search module.

97. The system of claim 91, wherein the search module controls a media playback component to extract metadata about the media web resources of each of the plurality of links.

98. A method to locate web resources on the Internet, the method comprising:

accessing a database including a plurality of links, the plurality of links being selectable to open a corresponding web media resources;

programmatically identifying metadata information about the web media resource corresponding to each of the plurality of links; and

making the plurality of links accessible to a plurality of Internet enabled devices, the plurality of Internet enabled devices being able to select one or more of the links to open the corresponding media web resource.

99. The method of claim 98, wherein programmatically identifying metadata information about the web media resource corresponding to each of the plurality of links includes signaling each link to a media playback component so as to cause the media playback component to load the corresponding media resource.

100. The method of claim 99, wherein signaling each link to a media playback component includes loading streaming audio media into the media playback component for sound recordings.

101. The method of claim 99, wherein signaling each link to a media playback component includes loading streaming video media in the media playback component for video recordings, including video recording having sound.

102. The method of claim 100, wherein signaling each link to a media playback unit includes determining at least one of a duration of the web media resource, and a title of the web media resource.

103. The method of claim 100, wherein signaling each link to a media playback unit includes determining a broadcast quality of a web media resource.

104. The method of claim 100, wherein signaling each link to a media playback unit includes determining an image size of a web media resource having video or image.

105. The method of claim 100, wherein signaling each link to a media playback unit includes determining at least one of an artist, a genre, and a data type of the web media resource.

106. The method of claim 98, further comprising manually extracting the metadata information after programmatically identifying the metadata information.

107. The method of claim 106, further comprising storing the identified and extracted metadata information with each of the plurality of links in a database, and then making the database available to the plurality of Internet enabled devices.

108. A method to locate web resources on the Internet, the method comprising:

accessing a database including a plurality of links to media web resources;

programmatically verifying whether each of the plurality of links opens a corresponding web media resource; and

making each verified link in the plurality of links accessible to a plurality of Internet enabled devices that select one or more of the links to open the corresponding media web resource.

109. The method of claim 108, wherein programmatically verifying whether each of the plurality of links opens a corresponding web media resource includes loading each link into a media playback component.

110. The method of claim 109, programmatically verifying whether each of the plurality of links opens a corresponding web media resource includes verifying each of the plurality of links makes available a web media resource for playback on one of the Internet enabled devices.

111. The method of claim 108, wherein programmatically verifying whether each of the plurality of links opens a corresponding web media resource includes loading each link into a web browser component

112. The method of claim 108, further comprising programmatically identifying metadata information about the web media resource corresponding to each of the plurality of links using the media playback component.

113. The method of claim 112, wherein programmatically identifying metadata information includes loading each of the verified links into the media playback component and then receiving editorial input about the web media resource loaded onto the media playback component.

114. A system for locating web resources on the Internet, the system being accessible by a plurality of Internet enabled devices, the system comprising:

a first indexed data structure comprising a plurality of links;

a media playback component coupleable to the database, the media playback component loading each of the plurality of links to verify whether the link is selectable to open a media web resource;

a second indexed data structure stores each verified link in the plurality of links, the second indexed data structure being available to the plurality of Internet enabled devices.

115. The system of claim 114, wherein the media playback component identifies one or more classes of information about the web media resource upon loading the link corresponding to the web media resource.

116. The system of claim 114, further comprising a search module that signals each of the plurality of links in the first indexed data structure to the media playback component, and wherein for each verified link, the search module signals the verified link to the second indexed data structure.

117. The system of claim 112, further comprising an editor interface to receive an editor input, the editor input including metadata information that was not identified from the media playback component loading the link.

118. The system of claim 117, further comprising an editor interface to receive an editor input, the editor input including metadata information identified from the media playback component loading the link.

119. The system of claim 118, further comprising an editor interface to receive an editor input, the editor input including metadata information identified from the media playback component loading the link, wherein the search module signals the metadata information about that verified link to the second indexed data structure, so that each verified link is associated with the corresponding metadata received from the editor input.

120. A method for providing links for use in a media search engine, the method comprising:

identifying a plurality of internal links on a network site;

determining whether the network site makes a network resource of a specific data type accessible for a network enabled device;

extracting the internal links that are selectable to open the network resource of the specific data type;

storing the external link in a database; and

automatically signaling one or more of the links to a media playback component in response to receiving a search request from the network enabled device.

121. The method of claim 120, wherein extracting the internal links that are selectable to open the network resource of the specific data type includes using a web browser component.

122. The method of claim 120, wherein determining whether the network site makes a network resource of a specific data type accessible for a network enabled device includes using a media playback component.

123. The method of claim 120, including extracting the internal links that are selectable to open the network resource of the specific data type for playing back media.

124. A method for providing links for use in a media search engine, the method comprising

a) receiving from a first indexed data structure a first external link to a first network site;

b) initializing a second data structure to be empty;

c) determining if the first network site contains at least one internal link; and

d) storing the at least one internal link contained on the first network site that is not in the first indexed data structure and not in the second indexed data structure as another external link in the first indexed data structure;

e) identifying the internal links contained on the first network site that are selectable to open a network resource of a specific data type or types;

f) moving the first external link from the first indexed data structure to the second indexed data structure; and

g) repeating steps a) through f) until the first indexed data structure is empty.

125. The method of claim 124, wherein prior to receiving from a first indexed data structure a first link to a first site, adding one or more of the first plurality of links to the first indexed data structure using an editor interface.

126. The method of claim 125, wherein prior to receiving from a first indexed data structure a first link to a first site, adding one or more of the first plurality of links to the first indexed data structure using a metacrawler.

127. A computer system to search for links to streaming media playback on a network, the network being accessible to a network enabled device, the computer system comprising:

a metacrawler to locate one or more media sites containing streaming media;

a media search module coupled to the metacrawler be signaled the one or more media sites, the media search module being able to identify a plurality of media links located by the one or more media sites, each of the plurality of media links being selectable to open streaming media network resource;

a metadata extraction module that accesses each media link identified by the media search engine to extract metadata about the identified media link; and

a database comprising the plurality of media links identified by the media search engine, and the metadata extracted about each identified media link, the database enabling the network enabled device to access the plurality of media links.

128. The computer system of claim 127, wherein the metacrawler locates one or more media sites in media directories, each media directory comprising a plurality of media sites.

129. The computer system of claim 127, where the metadata extraction module comprises a media playback component that loads the streaming media network resource.

130. The computer system of claim 127, wherein the metadata extracted from each identified link is sorted in the database according to one or more classes of information.

131. The computer system of claim 130, wherein the metadata extracted is sorted into one or more classes of information selected from classes consisting of genre, sub-genre, artist, and title.

132. The computer system of claim 131, further comprising a verification module that verifies each media link identified by the media link module opens the corresponding streaming media network resource.

133. A computer system to search for links to streaming media playback on a network, the network being accessible to a network enabled device, the computer system comprising:

means for accessing a web site;

means for identifying a plurality of links using a web browser component, each of the plurality of links being selectable to open a corresponding web resource of a specified data type on the web site; and

making the plurality of links available to a plurality of Internet enabled devices that select one or more of the links.

134. A system for providing media from a network to a terminal, the terminal including a media playback component, the system comprising:

a first network site and a second network site, each network site locating one or more media network resources, each media network resource being locatable on the network by a corresponding address that accesses the media network resource; and

a network server module that communicates with the media playback component, the network server module identifying a first media network resource from the first network site and a second media network resource from the second network site, the network server module signaling the corresponding address of the first media network resource to the media playback component with control signals to cause the playback component to automatically load the first media network resource, the network server module automatically signaling the corresponding address of the second media network resource to the media playback component with control signals to cause the playback component to automatically load the second media network resource.

135. The system of claim 134, wherein the network server module signals the corresponding addresses of the first and second media network resource to the terminal with control signals to automatically playback the first media network resource and then the second media network resource on the terminal.

136. The system of claim 134, further comprising a database including the corresponding addresses of the one or more media network resources on each of the first and second network sites, wherein the network server module selects the first media network resource and the second media network resource from the database using a selection criteria.

137. The system of claim 134, wherein the network server module is coupleable to the user terminal through the network.

138. The system of claim 134, wherein the network server module is located on the user terminal.

139. The system of claim 136, wherein the database is coupleable to the user terminal through the network.

140. The system of claim 136, wherein the database is located on the user terminal.

141. The system of claim 136, wherein the database includes metadata about each of the media network resources, and wherein the network server module selects the first network media resource and the second network media resource using metadata associated with the corresponding address of the first media network resource and the second media network resource.

142. The system of claim 141, wherein the network server module identifies the first and second media network resources by matching a criteria signaled from the terminal to metadata information provided in the database.

143. The system of claim 134, wherein the network server module identifies the first media network resource and the second media network resource through a random selection.

144. The system of claim 134, wherein the network server module signals the first media network resource, the second media network resource, and other media network resources to the terminal in a predetermined order, the other media network resources being located on the first network site, the second network site, or an other network site.

145. The system of claim 144, wherein the network server module signals the terminal to access a next media network resource in the predetermined order automatically if any of the first media network resource, the second media network resource, or the other media network resources are unavailable.

146. The system of claim 134, wherein the first media network resource and the second media network resource are locatable over the Internet.

147. The system of claim 146, wherein the first media network resource includes an Internet addresses having a first domain, and the second media network resource includes an Internet address having a second domain.

148. The system of claim 134, wherein the first media network resource and the second media network resource include a combination of music and/or video.

149. The system of claim 134, wherein the corresponding addresses of the one or more media network resources on the first network site and on the second network site each contain no internal links.

150. A media playback system for the Internet, the system comprising:
an end terminal including a media playback component; and
a web server module coupleable to the end terminal through the Internet,
the web server module having access to one or more media web resources on a first web site, and to one or more media web resources on a second web site,
the web server module signaling a first link that locates a first media web resource on the first web site, and a second link

that locates a second media web resource on the second web site;

the web server module providing control signals to the end terminal to cause the end terminal to access and load the first media web resource and the second media web resource into the media playback component.

151. The system of claim 150, wherein the web server module provides control signals to the playback component of the end terminal to cause the end terminal to playback the first media web resource after loading the first media web resource, and to playback the second media web resource after loading the second media web resource.

152. The system of claim 151, wherein the end terminal includes a web browser, the web browser receiving the first link and the control signals from the web server module to cause the web server module to display a web page of the first web site concurrently while the first media web resource is being played back, the web browser receiving the second link and the control signals from the web server module to cause the web server module to display a web page of the second web site concurrently while the second web media resource is being played back.

153. A media playback system for the Internet, the system comprising:

a user terminal including a media playback component; and

a web server module coupleable to the user terminal through the Internet, the web server module having access to a plurality of web sites including a first web site and a second web site, each of the plurality of web sites including one or more media web resources, the web server module signaling a first link to a first media web resource on the first web site and a second link to a second media web resource on the second web site, the web server module providing control signals to the user terminal to cause the user terminal to access and load the first media web resource and the second media web resource; and

a database including links to each of the plurality of media web resources, the web server module identifying the first link and the second link from the database.

154. The system of claim 153, wherein each of the plurality of links in the database are selectable to open a corresponding media web resource in the plurality of media web resources.

155. The system of claim 154, wherein the database includes metadata about each of the plurality of media web resources, and wherein the web server module identifies the first link by using metadata associated with the first media web resource, and the web server module identifies the second link by using metadata associated with the second media web resource.

156. The system of claim 155, wherein the web server module identifies at least the first link by matching a criteria signaled from the terminal to metadata information in the database corresponding to the first media web resource.

157. The system of claim 155, wherein the web server module identifies at least the first link to the first media web resource and the second link to the second media web resource through a random selection.

158. The system of claim 153, wherein the web server module signals the terminal to access another media web resource automatically if an accessed media web resource is unavailable, the other media web resource being located on either the first web site, the second web site, or a third web site.

159. A media playback system for the Internet, the system comprising:
a terminal including a media playback component and a user-interface;
a web server module coupleable to the user terminal through the Internet, the web server module having access to a plurality of links, each of the plurality of links locating a media web resource, the plurality of links being accessible on a plurality of web sites, the web server module signaling the plurality of links to the user terminal in a designated order to cause the terminal to load the media web

resource located by each of the plurality of links into the media playback component; and

a database comprising the plurality of links, the database being accessible to signal the plurality of links to the web server module in the designated order;

wherein the user-interface signals one or more inputs from a user to the web server module, the one or more inputs directing the web server to alter the designated order in which the database signals the plurality of links to the web server module.

160. The system of claim 159, wherein the user-interface signals a first input from the user to skip a subsequent link in the designated order.

161. The system of claim 160, wherein the user-interface signals a second input from the user to cause the web server module to replace a first designated order with a second designated order.

162. The system of claim 161, wherein the first designated order includes a selection of links stored in the database belonging to a first category, and the second designated order includes a selection of links stored in the database belonging to a second category.

163. The system of claim 162, wherein the web server module selects at least some of the plurality of media web resources by matching a criteria signaled from the terminal to metadata information stored in the database, the metadata being sorted to correspond to one or more of the plurality links.

164. The system of claim 159, wherein the web server module randomly selects at least some of the plurality of links.

165. The system of claim 159, wherein the web server module signals another link in the plurality of links to the terminal to as to cause the terminal to access another media web resource automatically if an accessed media web resource is unavailable.

166. A system for providing media play-back on a network, the system comprising:

a terminal coupleable to the network;

a play-list module coupleable to the terminal, the play-list module storing a first play-list signaled from the terminal, the first play-list comprising a plurality of network addresses, including a first network address that locates a first media network resource on a first network site, and a second network address that locates a second media network resource on a second network site;

a network server module coupleable to the terminal and to the play-list module, the network server module signaling the first play-list to the terminal, the network server module controlling the terminal to cause the terminal to access the media network resource associated with each network address in the first play-list, and to automatically load each respective media network resources into the media playback component.

167. The system of claim 166, wherein the play-list module receives and stores a plurality of play-lists from a plurality of terminals, each play-list comprising a plurality of network addresses, including a first network address that locates a first media network resource on a first network site, and a second network address that locates a second media network resource on a second network site, wherein the network server module responds to a request from a second terminal for a select play-list in the plurality of play-lists by signaling the plurality of network address in the select play-list to the second terminal.

168. The system of claim 167, wherein the media network resource located by each network address in one of the plurality of play-lists is loaded into a playback component on the requesting terminal to playback media having a combination of music and/or video.

169. A method for providing media to a terminal coupled to a network, the terminal comprising a media playback component, the method comprising:

programmatically directing the terminal to access a first network site in the plurality of network sites;

causing the media playback component on the terminal to automatically load a first media web resource located at the first network site to playback a first media;

programmatically directing the terminal to access a second network site in the plurality of network sites; and

causing the media playback component on the terminal to automatically load a second media web resource located at the second network site to playback a second media.

170. The method of claim 169, further comprising:

programmatically directing the terminal to access one or more additional network sites in the plurality of network sites; and

causing the media playback component on the terminal to automatically load each media web resource located on each of the one or more additional network sites so as to sequentially playback media from each of the network sites accessed by the media playback component.

171. The method of claim 170, wherein programmatically directing the terminal to access a first network site in the plurality of network sites includes loading a first web address for the first network site into a web browser component of the terminal, and wherein programmatically directing the terminal to access a second network site in the plurality of network sites includes loading a second web address for the second network site into the web browser component of the terminal.

172. The method of claim 169, wherein programmatically directing the terminal to access a second network site in the plurality of network sites includes programmatically directing the terminal to access the second network site having an address that is external to the first address.

173. The method of claim 169, wherein programmatically directing the terminal to access the first and second network sites in the plurality of network

sites includes programmatically directing a web browser component and the media playback component from a web server module coupled to the terminal via the network.

174. A method for providing media to a terminal coupled to the Internet, the method comprising:

accessing a database comprising a plurality of links, each link opening a corresponding media web resource;

selecting a first link from the database, the first link being located on a first network site;

selecting a second link from the database, the second link being located in a second network site, the second network site being external to the first network site;

signaling the selected links to a media playback component on the terminal to sequentially access the media web resources associated with the selected links; and

automatically signaling the media playback component on the terminal to load each of the media web resource accessed from the selected links so as to playback a media corresponding to each media web resource.

175. The method of claim 174, wherein accessing a database comprising a plurality of links to media web resources includes accessing the database comprising metadata information about the plurality of links, and selecting the first and second links from the database includes selecting the one or more links using the metadata information.

176. The method of claim 174, further comprising detecting if the media playback component fails to load one of the media web resource in the selected media web resources, and then signaling the media playback component to use another selected link from the database to load another media web resource.

177. The method of claim 175, further comprising detecting if the media playback component fails to playback a media corresponding to one of the media web resource in the selected media web resources, and then signaling the

media playback component use another link to load another selected link from the database to load another media web resource.

178. A system for providing media to a terminal coupled to a network, the method comprising:

means for accessing a plurality of links, each link opening a corresponding media network resource;

selecting a first link from the database, the first link being located on a first network site;

selecting a second link from the database, the second link being located on a second network site;

means for sequentially accessing the media network resources associated with the selected links; and

means for automatically playing back media from the media network resources.

179. A rating system for rating media network resources on a network, the network being coupleable to a plurality of terminals, the rating system comprising:

a database comprising a plurality of addresses, each address locating a corresponding media network resource on the network;

a network server module coupleable to the plurality of terminals, the network server module being able to access the database to signal one or more addresses from the database to the plurality of terminals; and

a rating module coupleable to the plurality of terminals, the rating module being able to receive a rating input from each of the plurality of terminals, the rating module associating the rating input with a selected address in the database.

180. The rating system of claim 179, wherein the database includes a rating associated with each of the plurality of addresses

181. The rating system of claim 179, wherein the rating module accesses the database and reconfigures the rating associated with the selected address in response to receiving the rating input for the selected address.

182. The rating system of claim 179, wherein the database further includes one or more classes of information associated with each of the plurality of addresses.

183. The rating system of claim 179, further comprising a play-list module, the play-list module selecting a combination of the plurality of addresses to form a play-list, the network server module accessing the play-list to signal the combination of addresses to one or more of the terminals.

184. The rating system of claim 183, wherein the rating module receives the rating input from one or more of the terminals for each address that comprises the play-list signaled to that terminals.

185. The rating system of claim 183, wherein the rating module receives a play-list rating input for the play-list from one or more of the terminals receiving the play-list.

186. The system of claim 183, wherein the play-list module selects the combination of addresses using the rating associated with each address in the database.

187. The system of claim 183, wherein the play-list module selects the combination of addresses using the rating of each address in the database that share a same class of information.

188. The system of claim 187, wherein the same class of information is selected from a class group consisting of genre, sub-genre, title, and artist.

189. The system of claim 187, wherein the play-list module selects the combination of addresses using the rating of each address in the database that share two or more of the same classes of information, each same class of

information being selected from a group consisting of genre, sub-genre, title, and artist.

190. The rating system of claim 179, wherein the rating associated with each of the plurality of addresses is a ranking, including a maximum ranking and a minimum ranking.

191. The rating system of claim 190, wherein the rating module receives the ranking from one of the terminals, accesses an existing ranking from the database, and uses an averaging formula to reconfigure the ranking in the database.

192. The rating system of claim 191, wherein the averaging formula is a weighted average.

193. The rating system of claim 191, wherein the averaging formula is a moving average.

194. A rating system for rating media network resources on a network, the network being coupleable to a plurality of terminals, the rating system comprising:

a database comprising a plurality of addresses, each address locating a corresponding media network resource on the network, the database including one or more classes of information associated with each of the plurality of addresses;

a network server module coupleable to the plurality of terminals, the network server module communicating with each of the plurality of terminals to receive a search request, the network server module signaling the database to retrieve one or more addresses from the database in response to the search request, the retrieved addresses being associated with a class of information matching the search request; and

a rating module coupleable to the plurality of terminals, the rating module being able to receive a rating input from each of the plurality

of terminals, the rating module associating the rating input with a selected address in the database.

195. The rating system of claim 194, wherein the network module signals the database to retrieve one or more addresses from the database in response to the search request by matching a criteria in the search request with a parameter included with the one or more classes of information.

196. The rating system of claim 195, wherein the one or more classes of information stored in the database with the plurality of addresses is a rating.

197. The rating system of claim 194, wherein the search request also includes a rating standard that selects addresses in the database using a rating associated with each address.

198. A rating system for rating media network resources available over a network, the media network resources being located on the network by a plurality of terminals, the rating system comprising:

a database comprising a plurality of addresses, each address locating a corresponding media network resource on the network, the database including a rating associated with each of the plurality of addresses;

a network server module coupleable to each of the plurality of terminals, the network server module being able to access the database to signal one or more addresses from the database to the plurality of terminals;

a rating module coupleable to each of the plurality of terminals, the rating module being able to receive a rating input from one of the terminals for each of the plurality of addresses in the database, wherein in response to receiving the rating input from one of the plurality of terminals for a selected address in the database, the rating module accesses the database and reconfigures the rating associated with the selected address; and

a play-list module, the play-list module accessing the addresses to select one or more combinations of addresses, each combination of addresses forming a play-list, wherein the play-list module signals

the play-list to the network server module as addresses to be signaled to one or more of the plurality of terminals.

199. The rating system of claim 198, wherein the play-list module forms the play-list by using the ranking of the addresses in the database.

200. The rating system of claim 198, wherein the play-list module forms the play-list by arranging each address in the play-list in a designated order so that the addresses in the play-list are accessed by the one or more terminals in the designated order.

201. The rating system of claim 199, wherein the designated order is determined by a ranking of each of the addresses in the play-list relative to rankings of the other addresses in the play-list.

202. The rating system of claim 200, wherein the designated order is determined by arranging each of the addresses in the play-list according to a ranking from highest to lowest.

203. The rating system of claim 198, wherein the database further includes one or more classes of information associated with each of the plurality of addresses, and the play-list module forms a plurality of play-lists using the one or more classes of information.

204. The rating system of claim 198, wherein the play-list module forms a plurality of play-lists by arranging a combination of addresses in the plurality of addresses together in one play-list, wherein the combination of addresses share a same parameter of the one or more classes of information.

205. The rating system of claim 199, wherein the play-list module arranges the addresses in the play-list according to an order determined by a ranking of each of the addresses forming the play-list.

206. A method for ranking media sources on a network, the method comprising:

accessing a database comprising a plurality of addresses, each address locating a media resource on the network and each address being associated with a rating;

signaling a selected address from the database to a terminal coupled to the network;

receiving a rating input from the terminal after signaling the selected address to the terminal; and

adjusting the rating associated for the selected address in response to receiving the rating input.

207. The method of claim 206, wherein adjusting the rating associated for the selected address in response to receiving the rating input includes averaging the rating input with an existing rating for the selected address in the database.

208. The method of claim 207, wherein averaging the rating input with an existing rating includes implementing a weighted average between the rating input and the existing rating.

209. The method of claim 206, further comprising signaling one or more selected addresses to a plurality of terminals coupled to the network, and receiving a rating input from each of the plurality of the terminals after signaling the one or more selected addresses to the terminal.

210. The method of claim 209, further comprising adjusting the rating associated for each of the one or more selected addresses in response to receiving a rating input from each terminal by calculating a weighted average reflecting an aggregate of the plurality of terminals.

211. The method of claim 206, wherein each address locates a music source on the Internet.

212. The method of claim 206, wherein the rating input corresponds to a ranking.

213. A method for ranking media sources on a network, the method comprising:

accessing a database comprising a plurality of addresses, each address locating a media resource on the network and each address being associated with a rating;

selecting a combination of addresses to form a play-list;

signaling the play-list to a terminal coupled to the network;

receiving a ranking from the terminal after signaling the one or more of the addresses in the play-list to the terminal; and

adjusting the rating of each address signaled to the terminal from the play-list in response to receiving the ranking.

214. The method of claim 213, wherein selecting a combination of addresses for a play-list includes selecting multiple addresses from the database that share one or more classes of information.

215. The method of claim 213, wherein selecting a combination of addresses for a play-list includes arranging each address in the play-list in an order, the order being determined by the ranking of each address in the play-list.

216. A method for ranking media sources on a network, the method comprising:

accessing a database comprising a plurality of addresses, each address locating a media resource on the network and each address being associated with a rating;

signaling a selected address from the database to a terminal coupled to the network;

receiving a ranking from the terminal after signaling the selected address to the terminal; and

adjusting the rating associated for the selected address in response to receiving the ranking.

217. The method of claim 216, wherein each address locates a music source.

218. The method of claim 216, wherein each address locates a video or animation source.

219. A rating system for rating media network resources available over a network, the media network resources being located on the network by a plurality of terminals, the rating system comprising:

means for accessing a database comprising a plurality of addresses, each address locating a media resource on the network and each address being associated with a rating;

means for signaling a selected address from the database to a terminal coupled to the network;

means for receiving a ranking from the terminal after signaling the selected address to the terminal; and

means for adjusting the rating associated for the selected address in response to receiving the ranking.

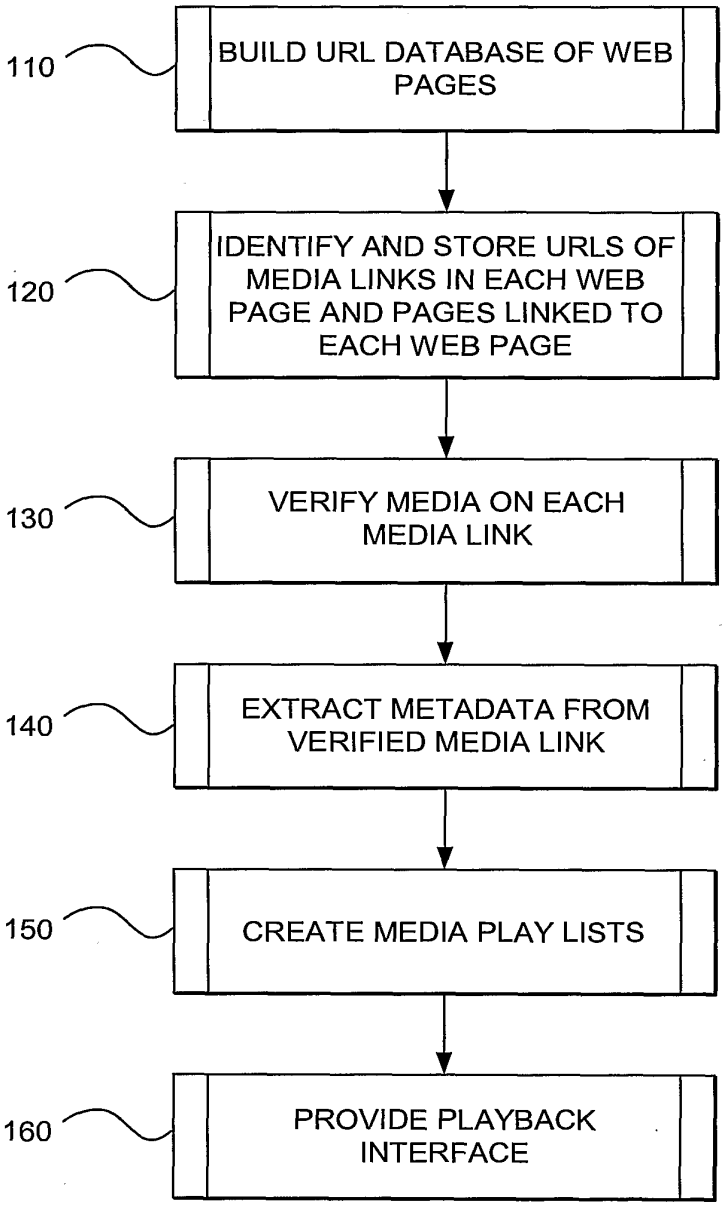

110
BUILD URL DATABASE OF WEB PAGES
120
IDENTIFY AND STORE URLS OF MEDIA LINKS IN EACH WEB PAGE AND PAGES LINKED TO EACH WEB PAGE
130
VERIFY MEDIA ON EACH MEDIA LINK
140
EXTRACT METADATA FROM VERIFIED MEDIA LINK
150
CREATE MEDIA PLAY LISTS
160
PROVIDE PLAYBACK INTERFACE

FIG. 1

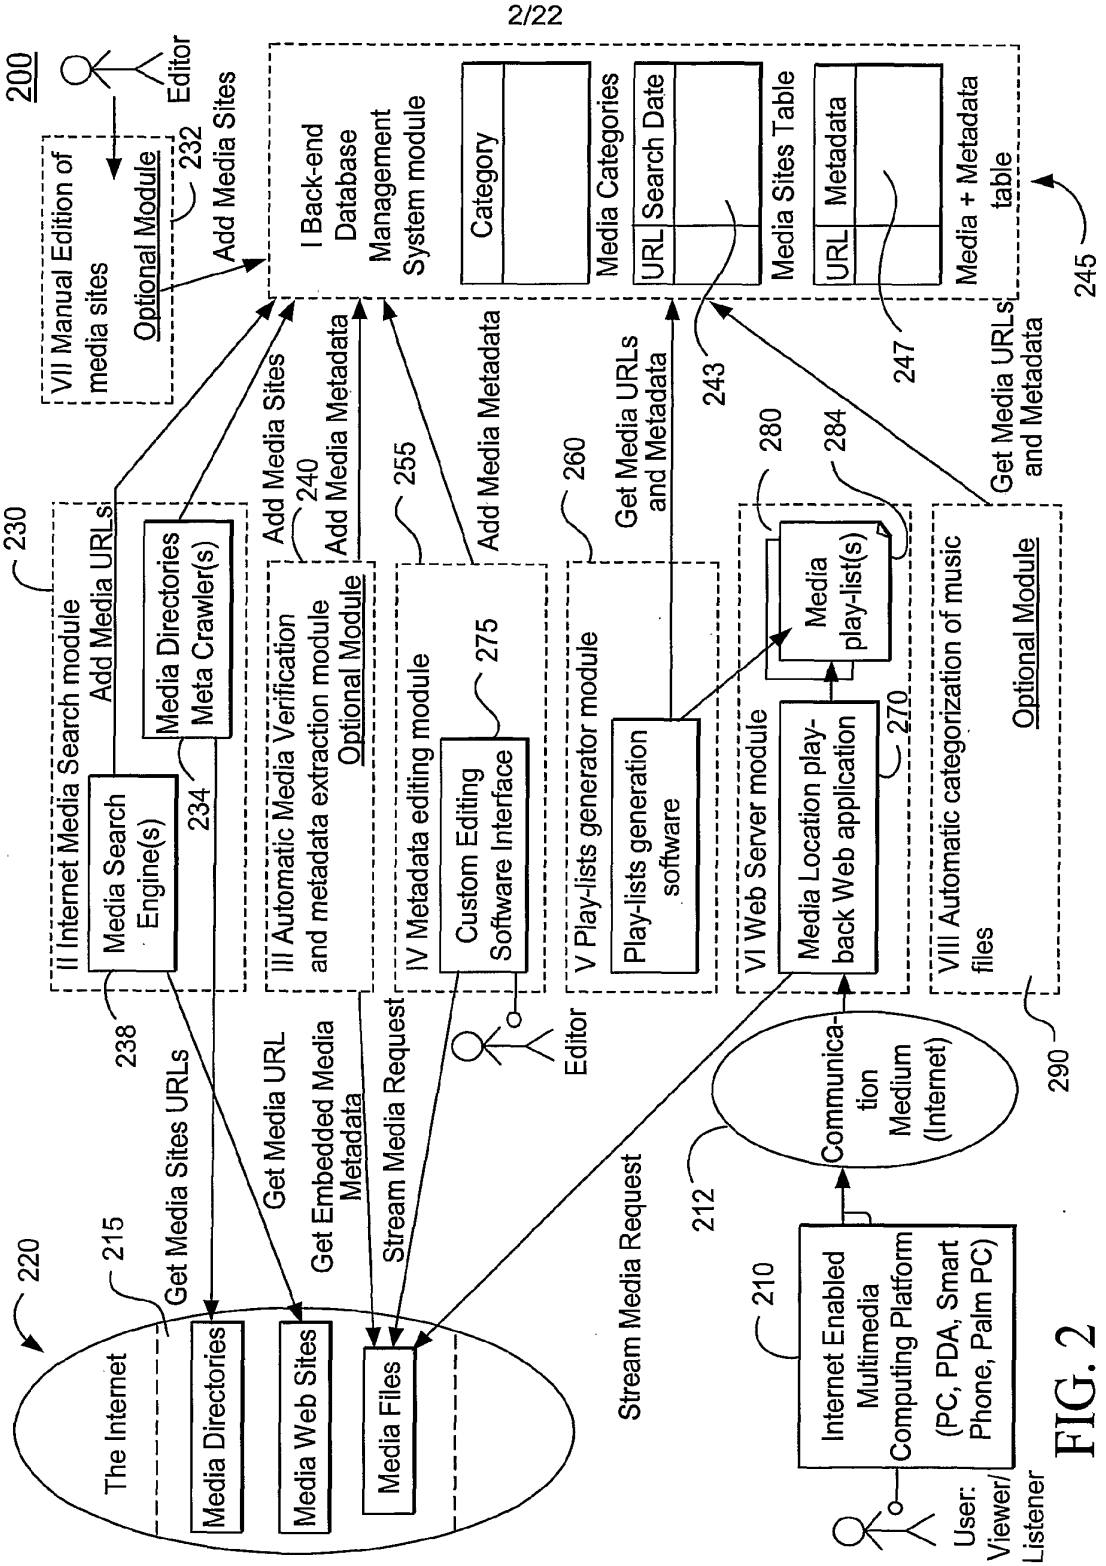

200
Editor
VII Manual Edition of media sites
Optional Module
232
Add Media Sites
I Back-end Database Management System module
Category
Media Categories
URL
Search Date
Media Sites Table
URL
Metadata
Media + Metadata table
245
247
243
230
II Internet Media Search module
Add Media URLs
Media Search Engine(s)
234
Media Directories Meta Crawler(s)
238
Add Media Sites
240
Add Media Metadata
III Automatic Media Verification and metadata extraction module
Optional Module
255
Add Media Metadata
IV Metadata editing module
275
Custom Editing Software Interface
Editor
260
V Play-lists generator module
Play-lists generation software
Get Media URLs and Metadata
280
284
VI Web Server module
Media play-list(s)
Media Location play-back Web application
270
VIII Automatic categorization of music files
Optional Module
Get Media URLs and Metadata
290
Get Media Sites URLs
Get Media URL
Get Embedded Media Metadata
Stream Media Request
Stream Media Request
212
Communica-tion Medium (Internet)
210
Internet Enabled Multimedia Computing Platform (PC, PDA, Smart Phone, Palm PC)
User: Viewer/ Listener
220
215
The Internet
Media Directories
Media Web Sites
Media Files

FIG. 2

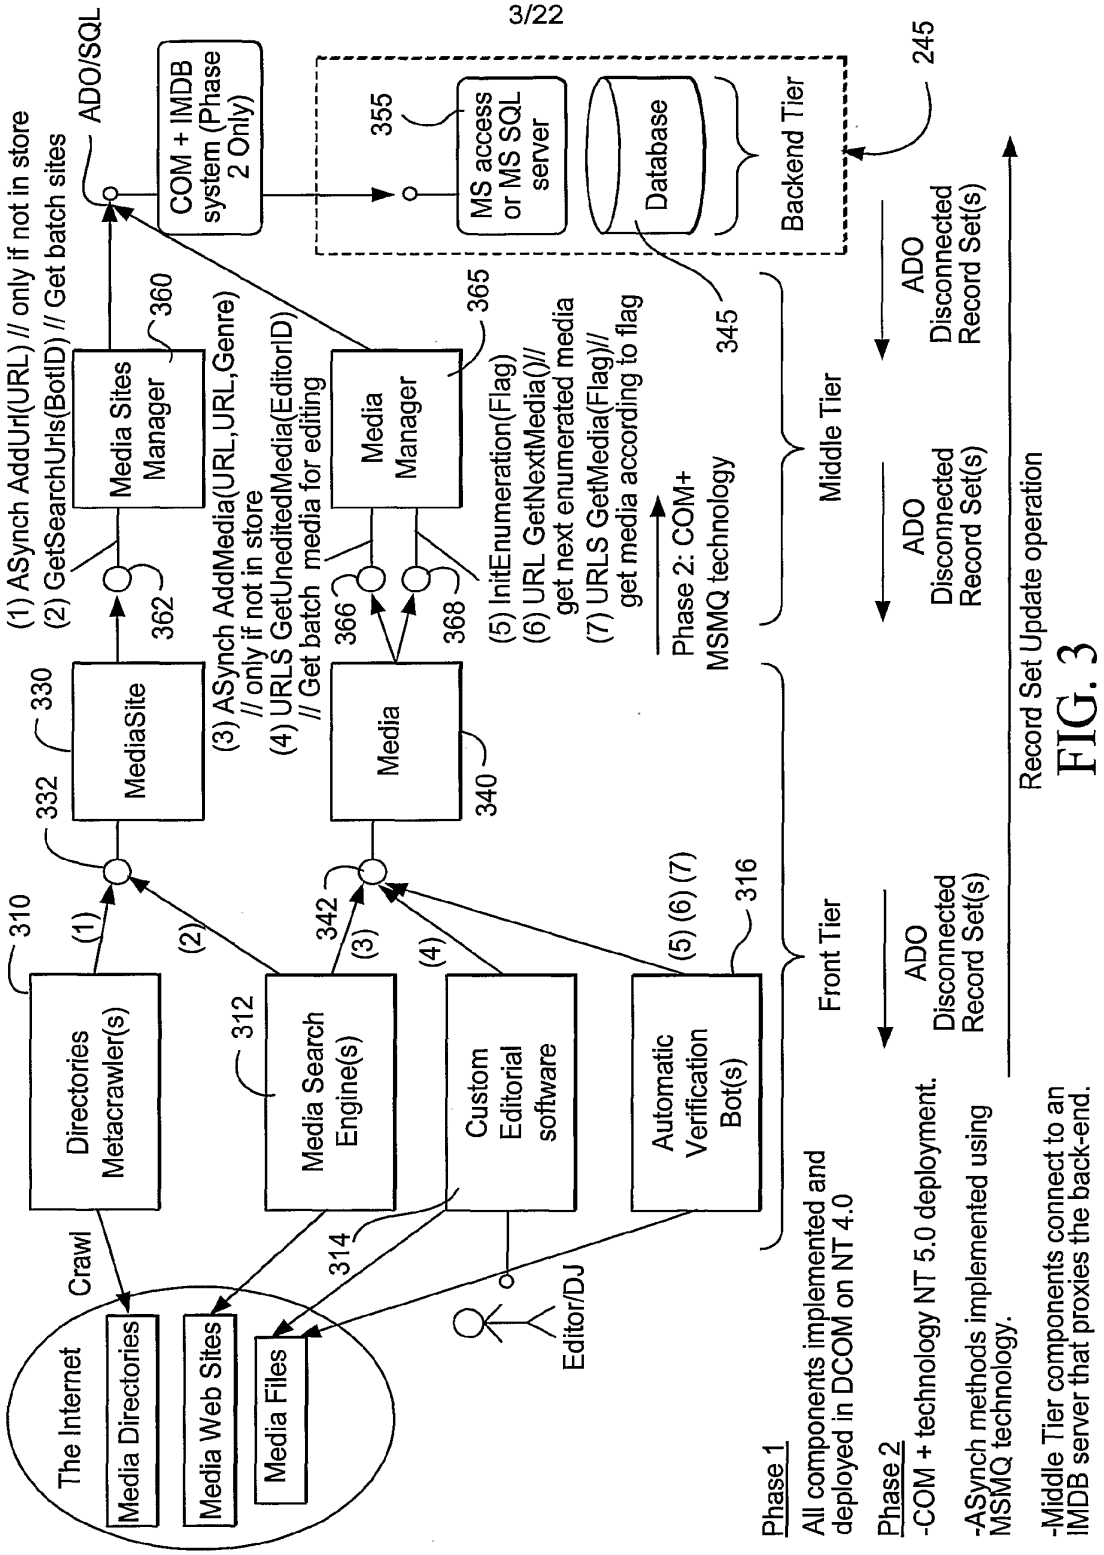

The Internet
Media Directories
Media Web Sites
Media Files
Crawl
Editor/DJ
Directories Metacrawler(s)
Media Search Engine(s)
Custom Editorial software
Automatic Verification Bot(s)
310
312
314
316
(1)
(2)
(3)
(4)
(5) (6) (7)
342
332
MediaSite
330
340
Media
362
366
368
Media Sites Manager
360
Media Manager
365
(1) ASynch AddUrl(URL) // only if not in store
(2) GetSearchUrls(BotID) // Get batch sites
(3) ASynch AddMedia(URL,URL,Genre) // only if not in store
(4) URLS GetUneditedMedia(EditorID) // Get batch media for editing
(5) InitEnumeration(Flag)
(6) URL GetNextMedia()// get next enumerated media
(7) URLS GetMedia(Flag)// get media according to flag
ADO/SQL
COM + IMDB system (Phase 2 Only)
355
MS access or MS SQL server
Database
345
Backend Tier
245
Phase 2: COM+ MSMQ technology
Middle Tier
ADO Disconnected Record Set(s)
ADO Disconnected Record Set(s)
ADO Disconnected Record Set(s)
Front Tier
Record Set Update operation
Phase 1
All components implemented and deployed in DCOM on NT 4.0
Phase 2
-COM + technology NT 5.0 deployment.
-ASynch methods implemented using MSMQ technology.
-Middle Tier components connect to an IMDB server that proxies the back-end.

FIG. 3

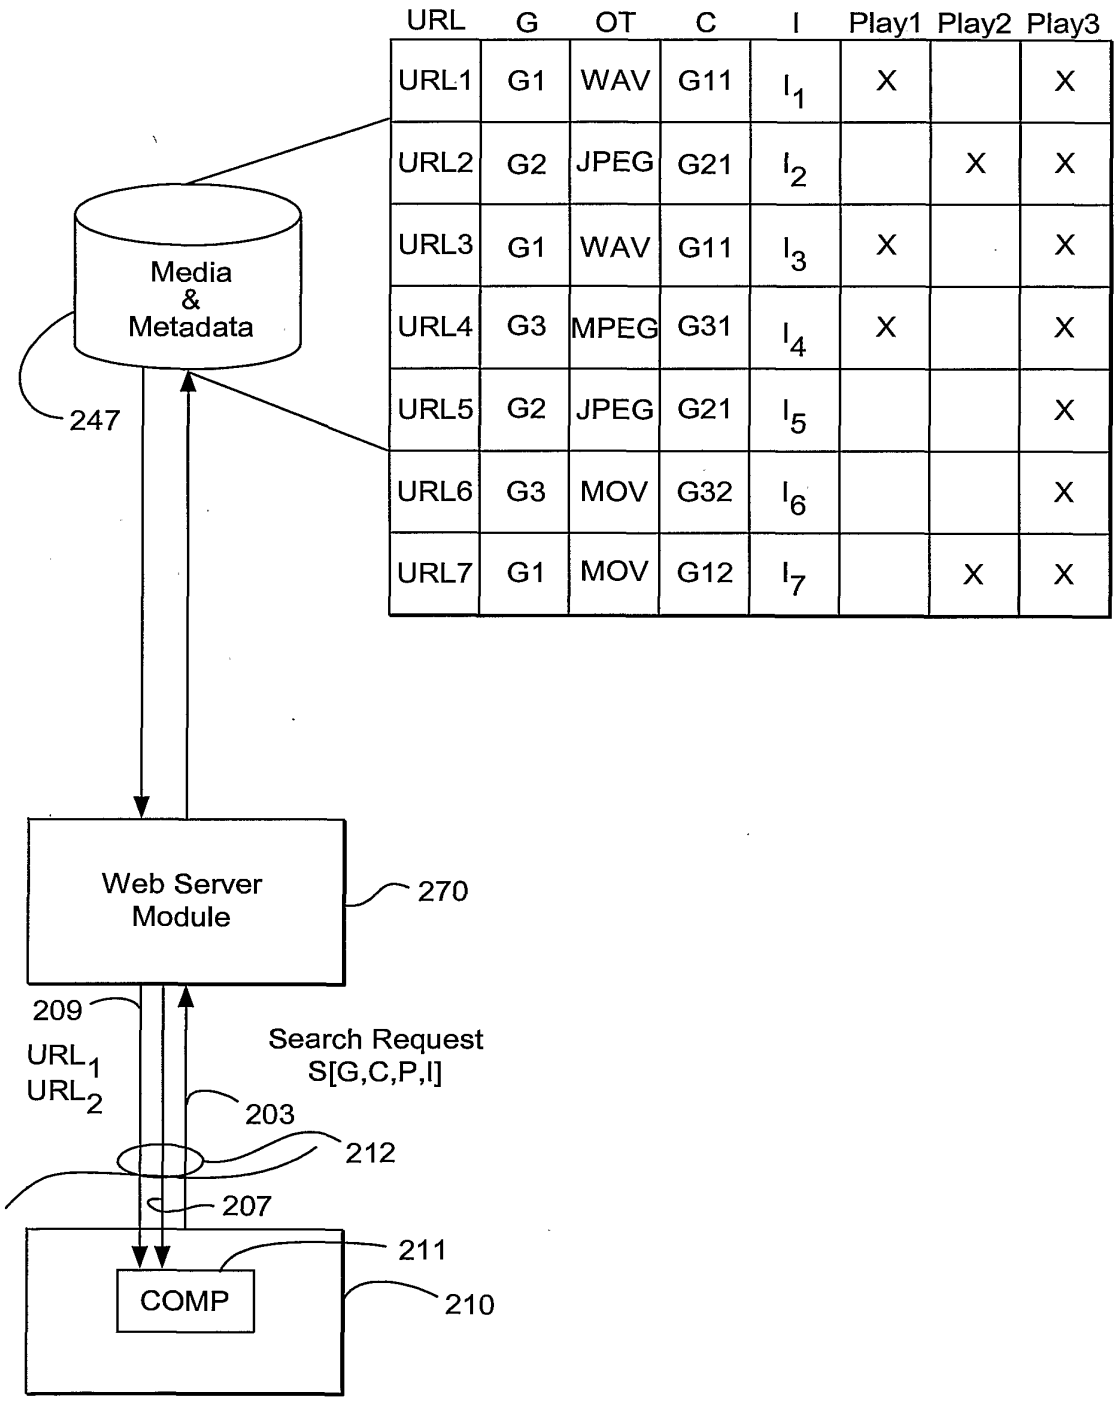

URL G OT C I Play1 Play2 Play3
URL1 G1 WAV G11 I_1 X X
URL2 G2 JPEG G21 I_2 X X
URL3 G1 WAV G11 I_3 X X
URL4 G3 MPEG G31 I_4 X X
URL5 G2 JPEG G21 I_5 X
URL6 G3 MOV G32 I_6 X
URL7 G1 MOV G12 I_7 X X
Media & Metadata
247
Web Server Module
270
209
URL_1
URL_2
Search Request S[G,C,P,I]
203
212
207
211
COMP
210

FIG. 4

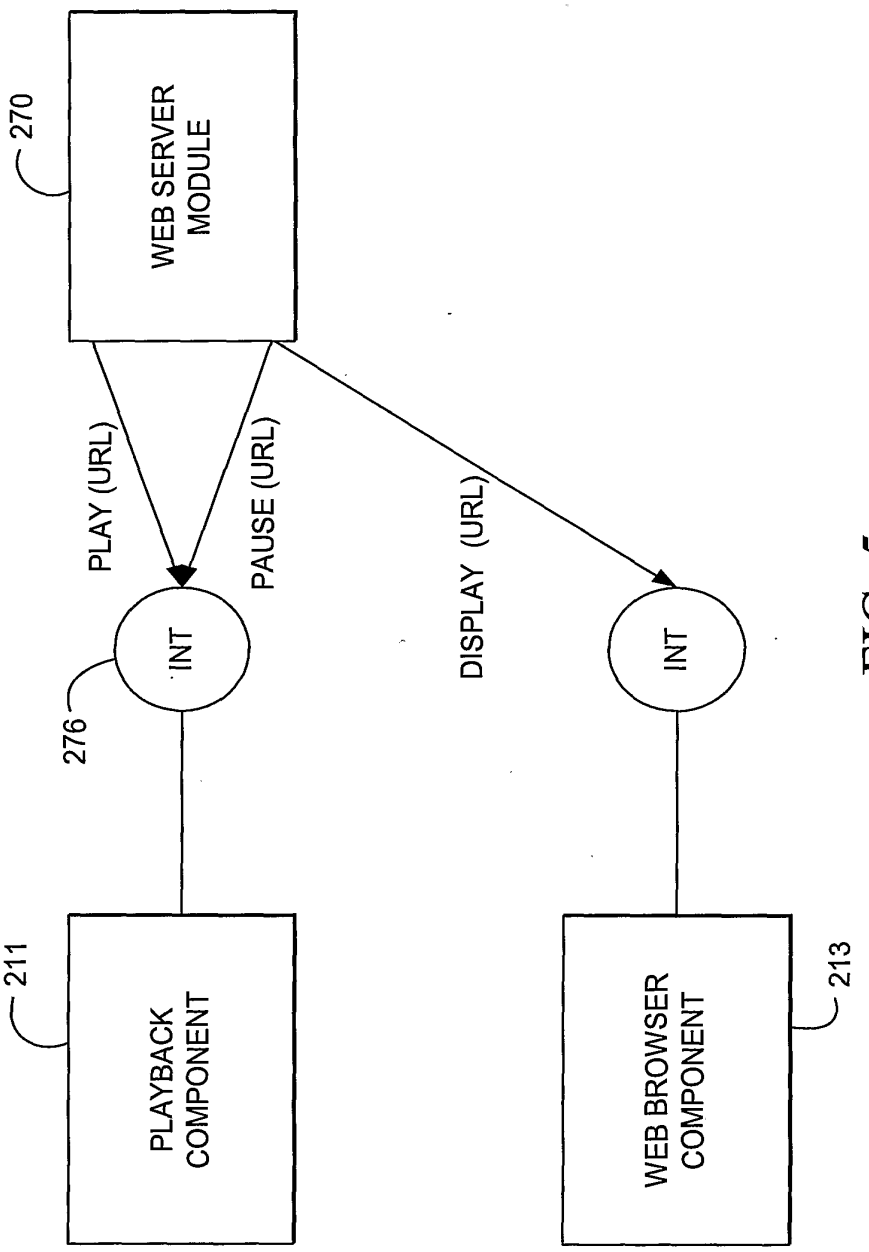

270
WEB SERVER MODULE
PLAY (URL)
PAUSE (URL)
DISPLAY (URL)
INT
276
INT
211
PLAYBACK COMPONENT
213
WEB BROWSER COMPONENT

FIG. 5

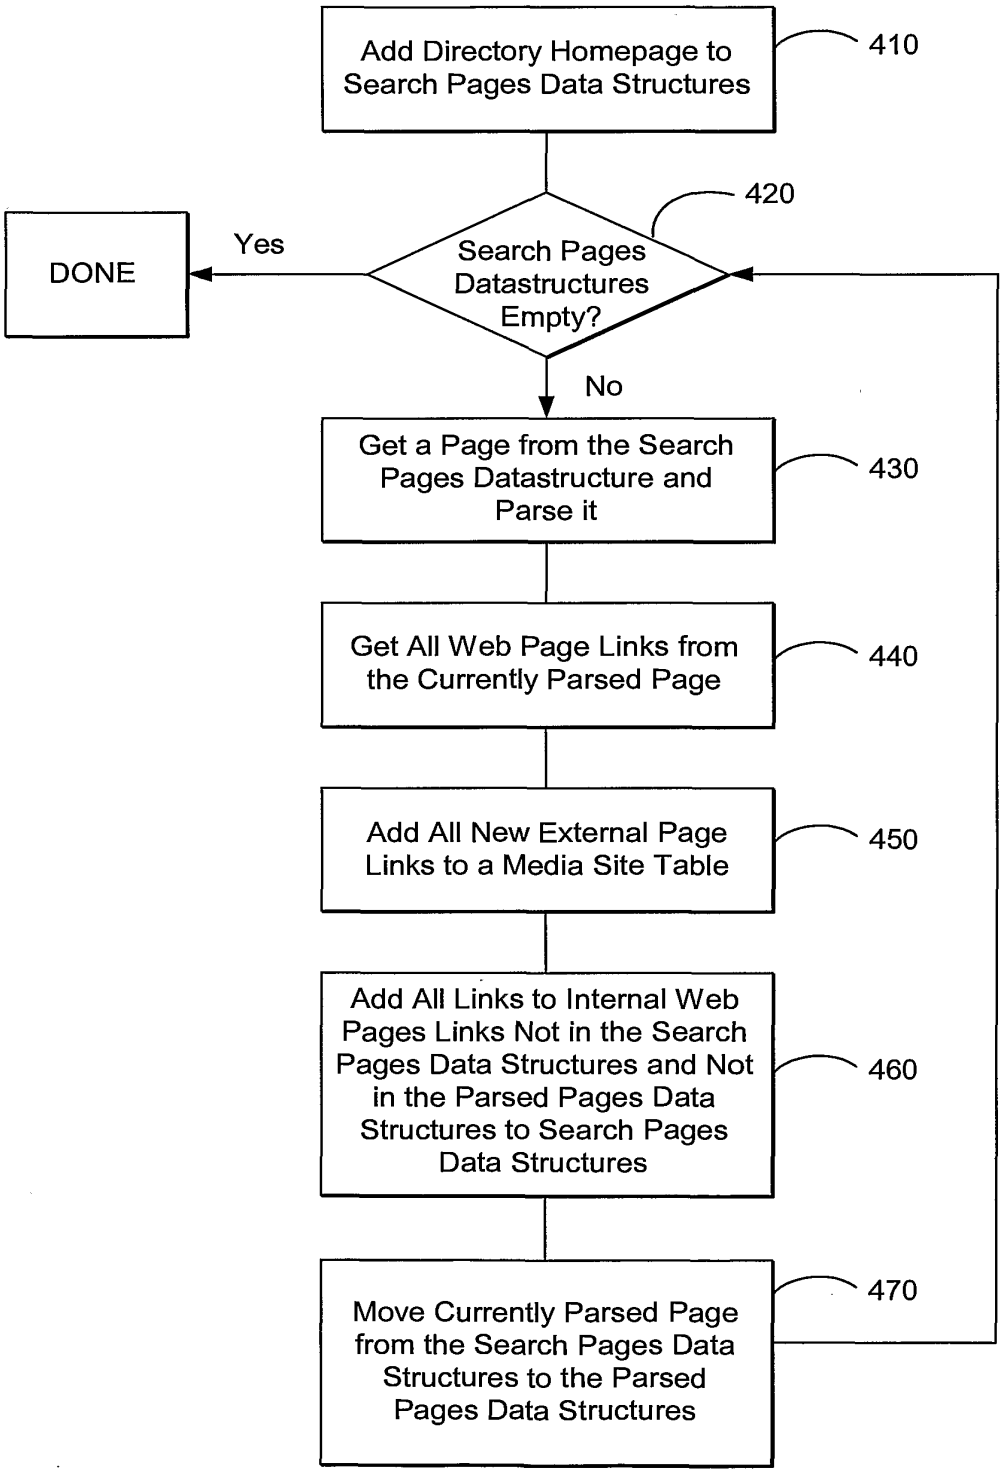

Add Directory Homepage to Search Pages Data Structures
410
Search Pages Datastructures Empty?
420
Yes
DONE
No
Get a Page from the Search Pages Datastructure and Parse it
430
Get All Web Page Links from the Currently Parsed Page
440
Add All New External Page Links to a Media Site Table
450
Add All Links to Internal Web Pages Links Not in the Search Pages Data Structures and Not in the Parsed Pages Data Structures to Search Pages Data Structures
460
Move Currently Parsed Page from the Search Pages Data Structures to the Parsed Pages Data Structures
470

FIG. 6

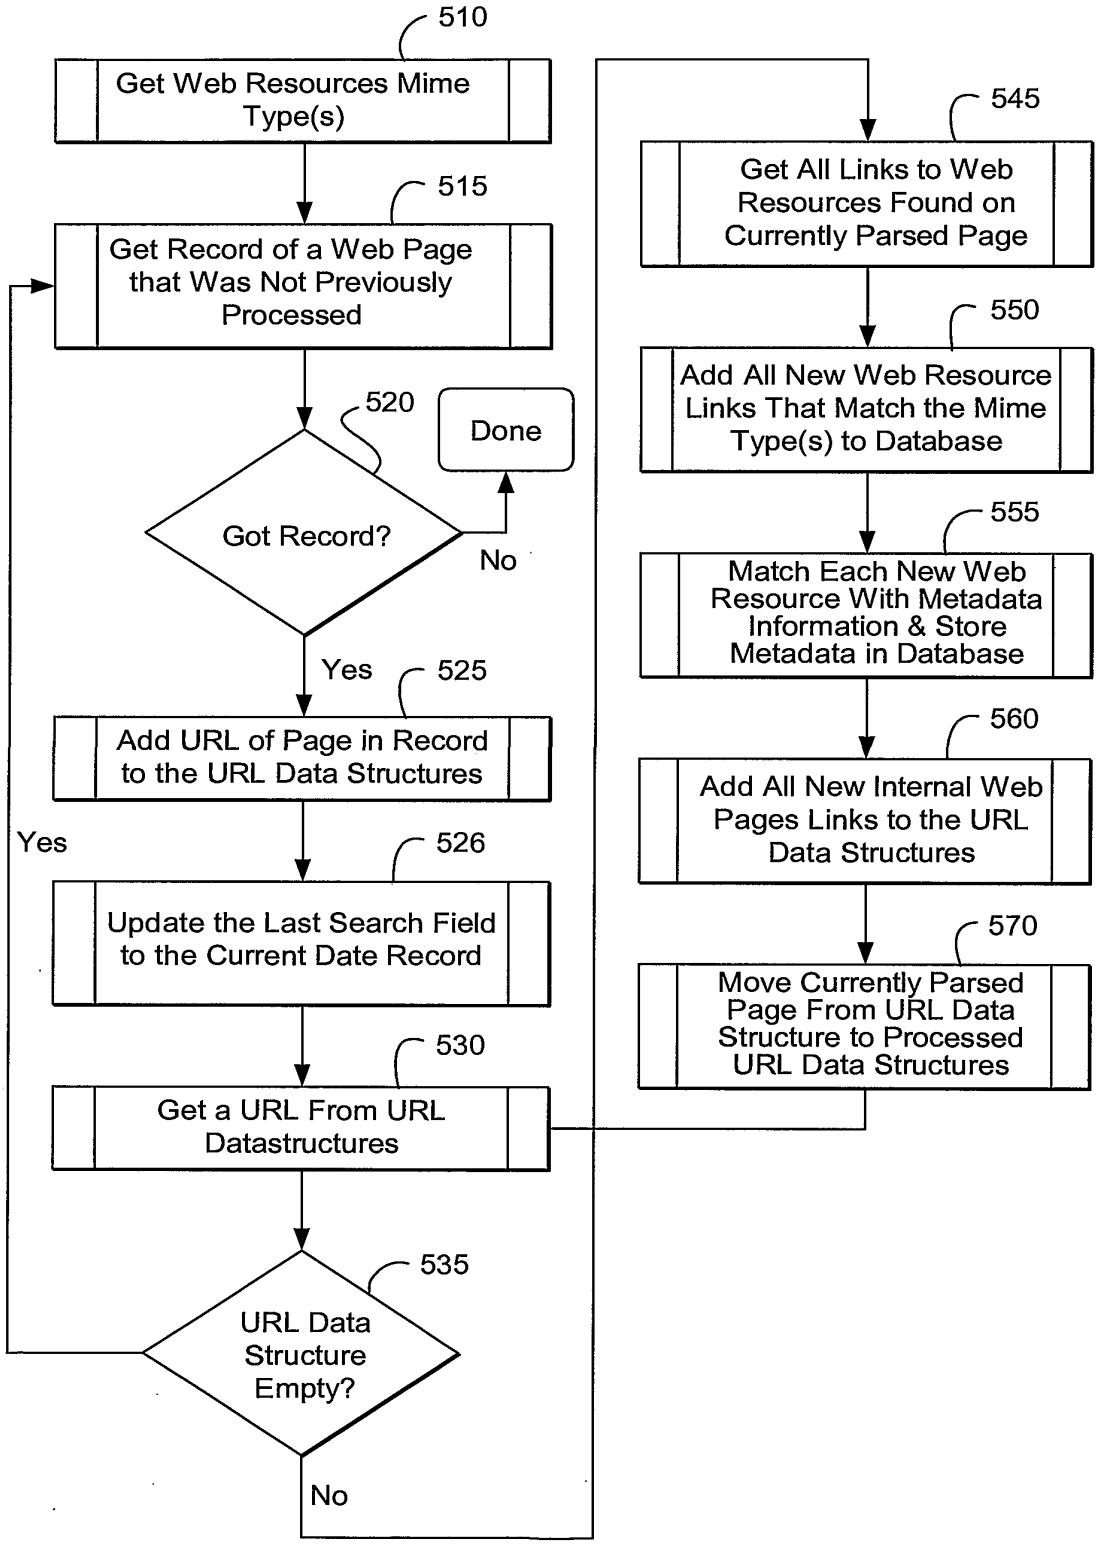


FIG. 7

SUBSTITUTE SHEET (RULE 26)

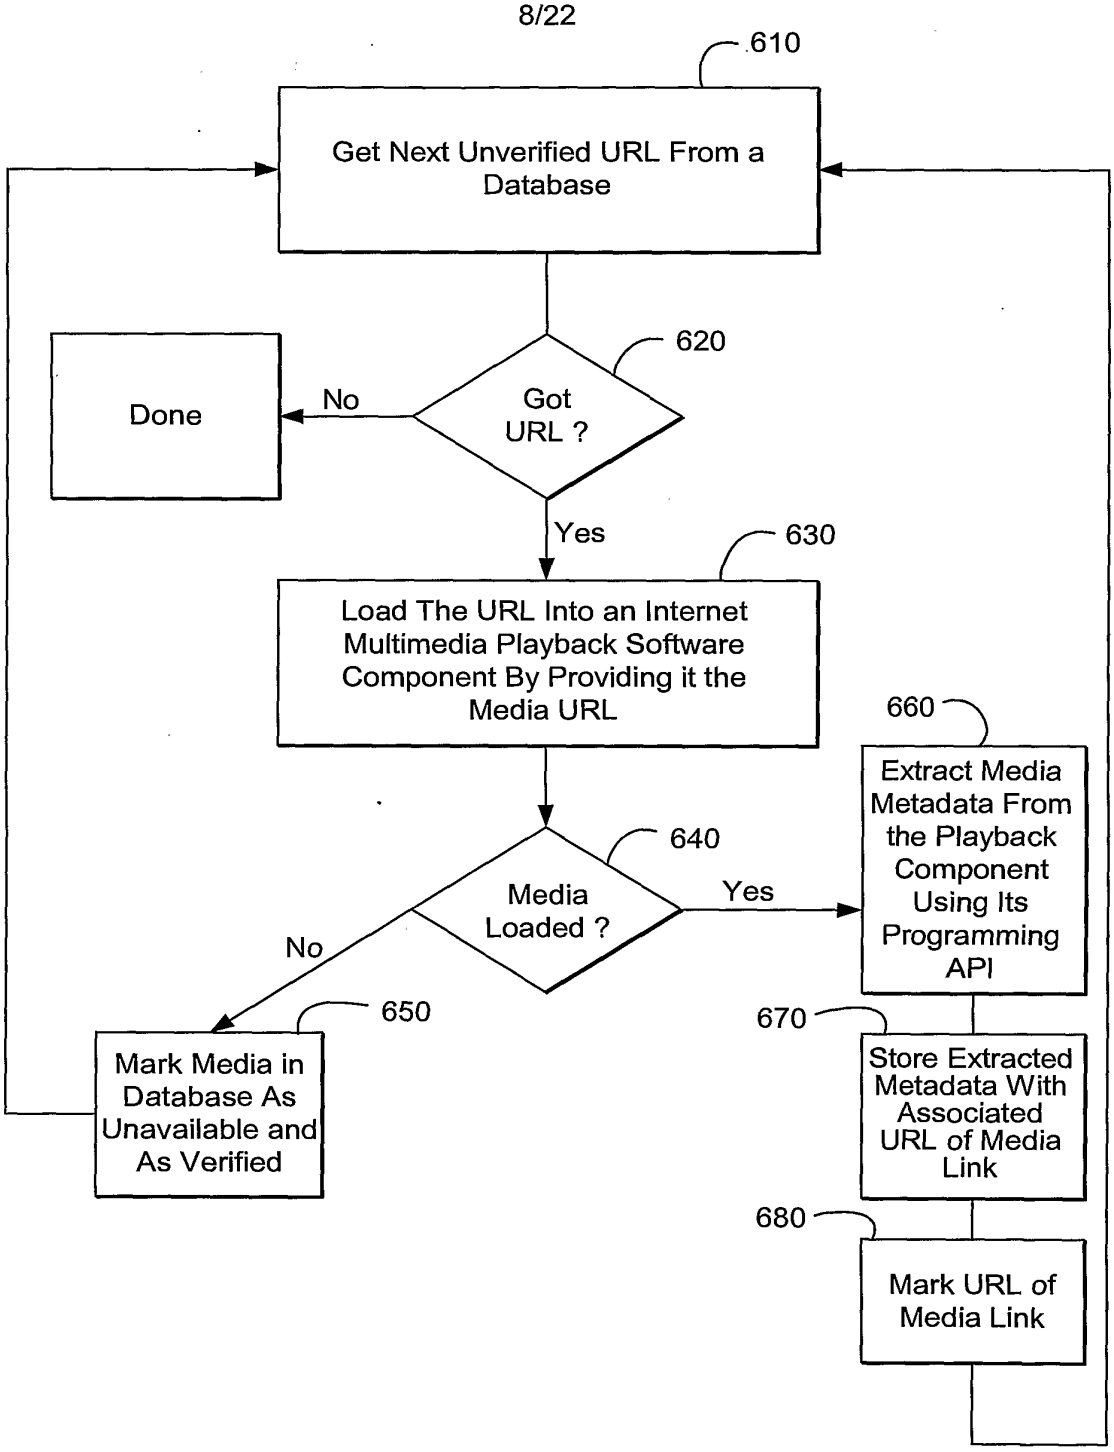

610
Get Next Unverified URL From a Database
620
Got URL ?
No
Done
Yes
630
Load The URL Into an Internet Multimedia Playback Software Component By Providing it the Media URL
640
Media Loaded ?
No
Yes
650
Mark Media in Database As Unavailable and As Verified
660
Extract Media Metadata From the Playback Component Using Its Programming API
670
Store Extracted Metadata With Associated URL of Media Link
680
Mark URL of Media Link

FIG. 8

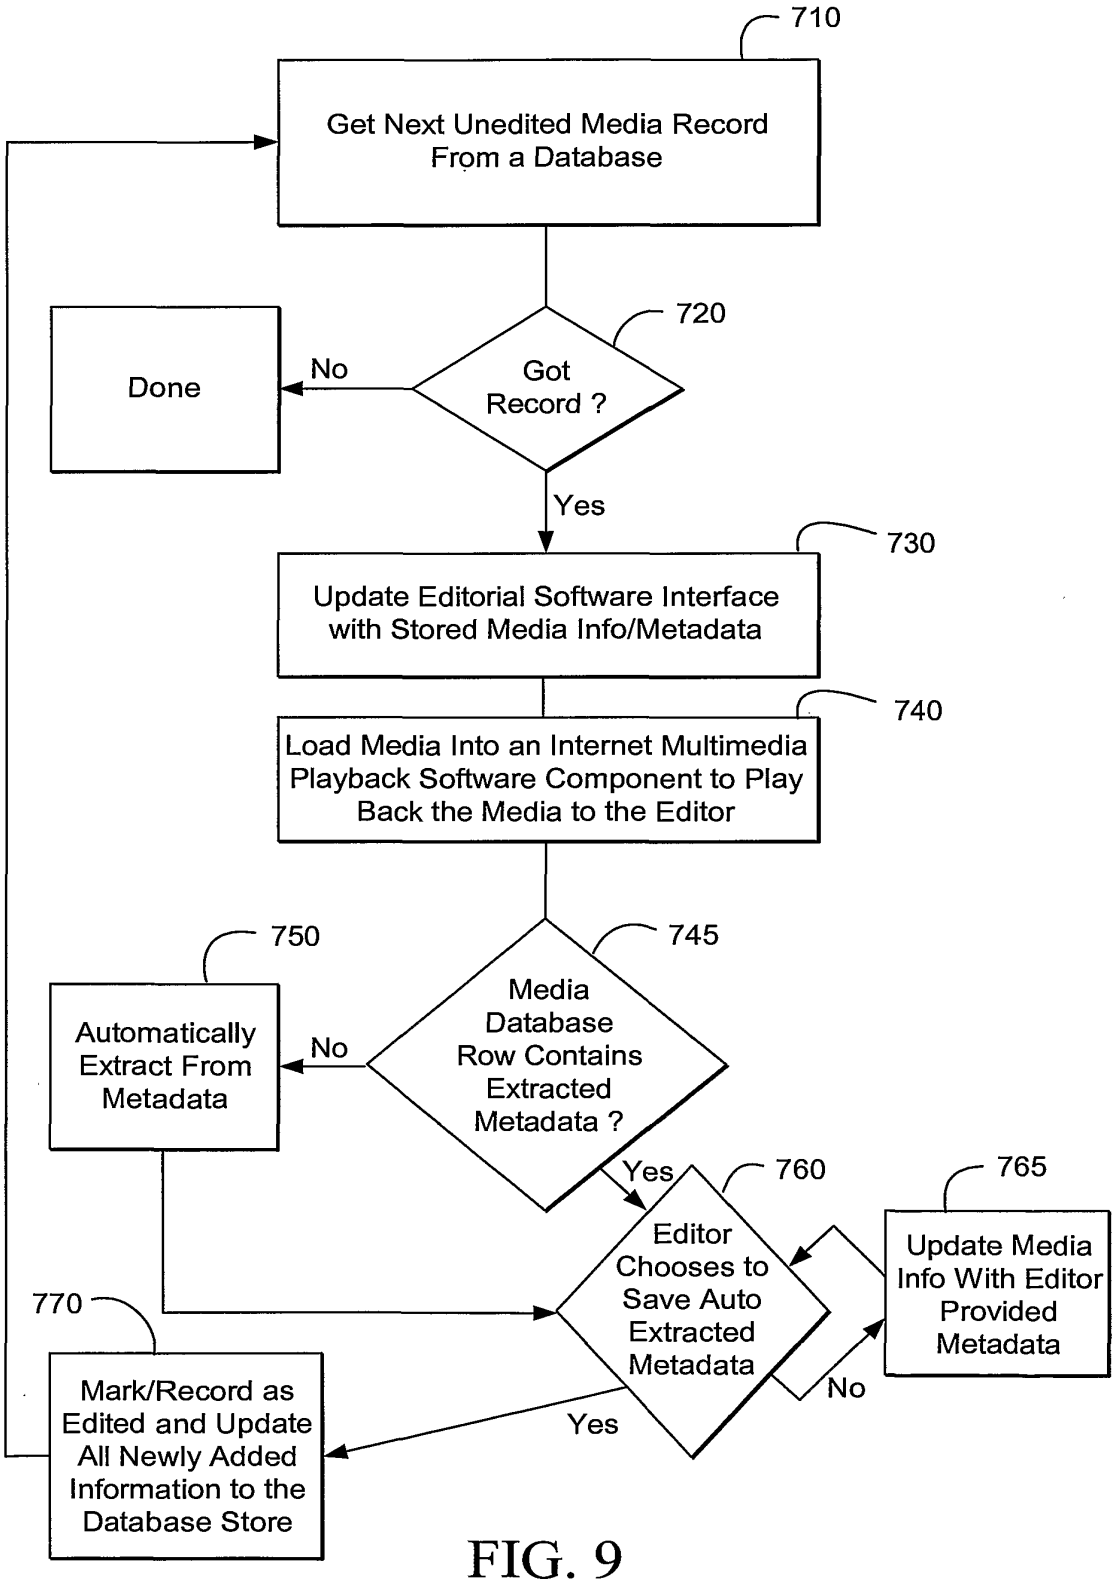

710
Get Next Unedited Media Record From a Database
720
Got Record ?
No
Done
Yes
730
Update Editorial Software Interface with Stored Media Info/Metadata
740
Load Media Into an Internet Multimedia Playback Software Component to Play Back the Media to the Editor
745
Media Database Row Contains Extracted Metadata ?
No
750
Automatically Extract From Metadata
Yes
760
Editor Chooses to Save Auto Extracted Metadata
No
765
Update Media Info With Editor Provided Metadata
Yes
770
Mark/Record as Edited and Update All Newly Added Information to the Database Store

FIG. 9

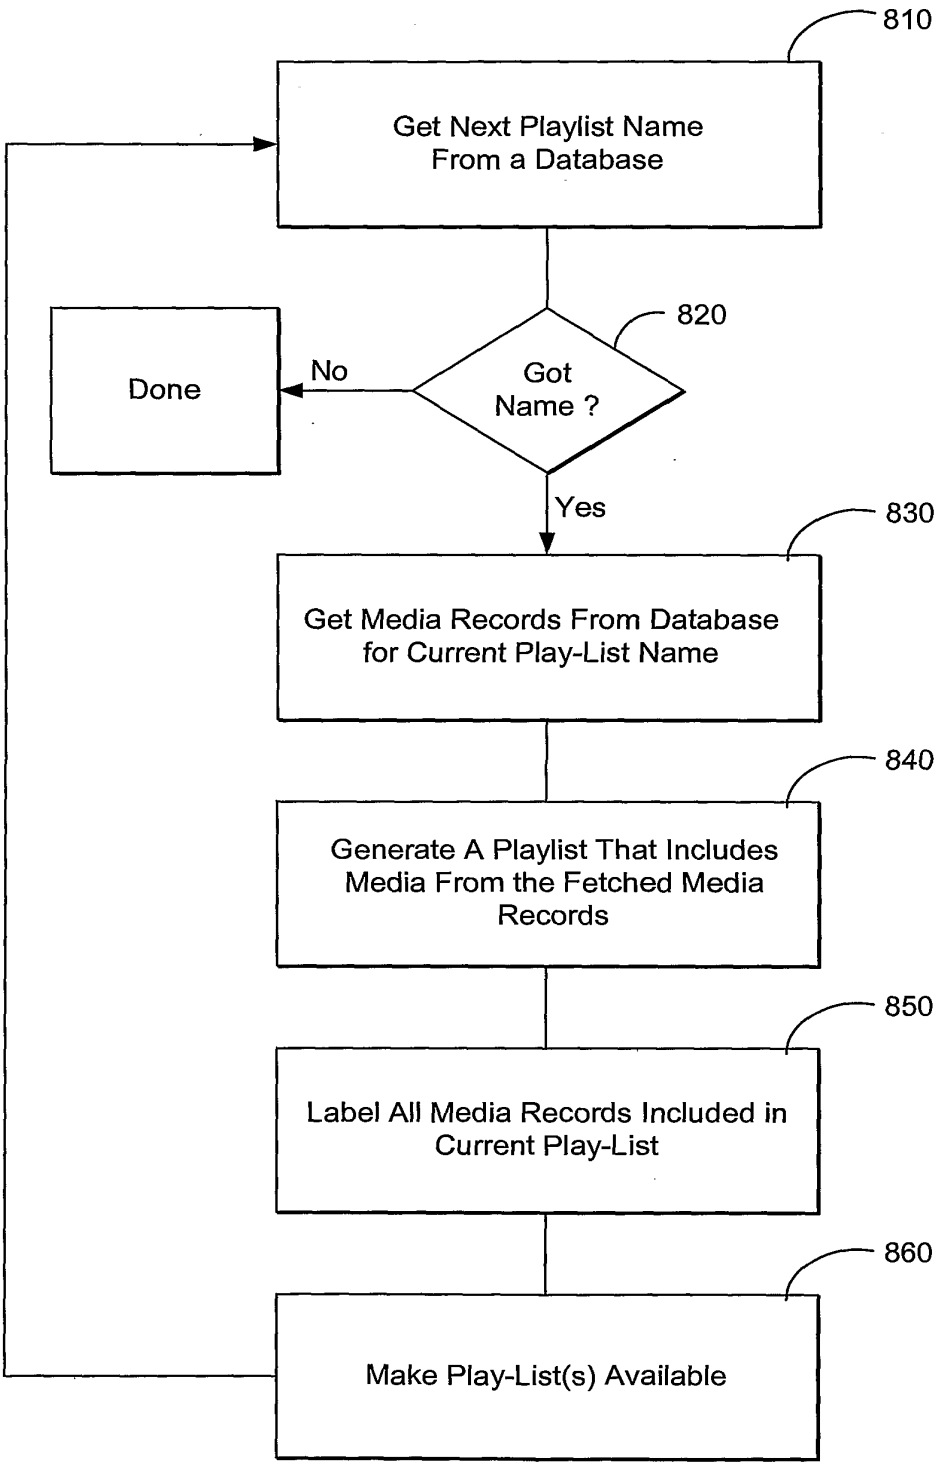

810
Get Next Playlist Name From a Database
820
Got Name ?
No
Done
Yes
830
Get Media Records From Database for Current Play-List Name
840
Generate A Playlist That Includes Media From the Fetched Media Records
850
Label All Media Records Included in Current Play-List
860
Make Play-List(s) Available

FIG. 10

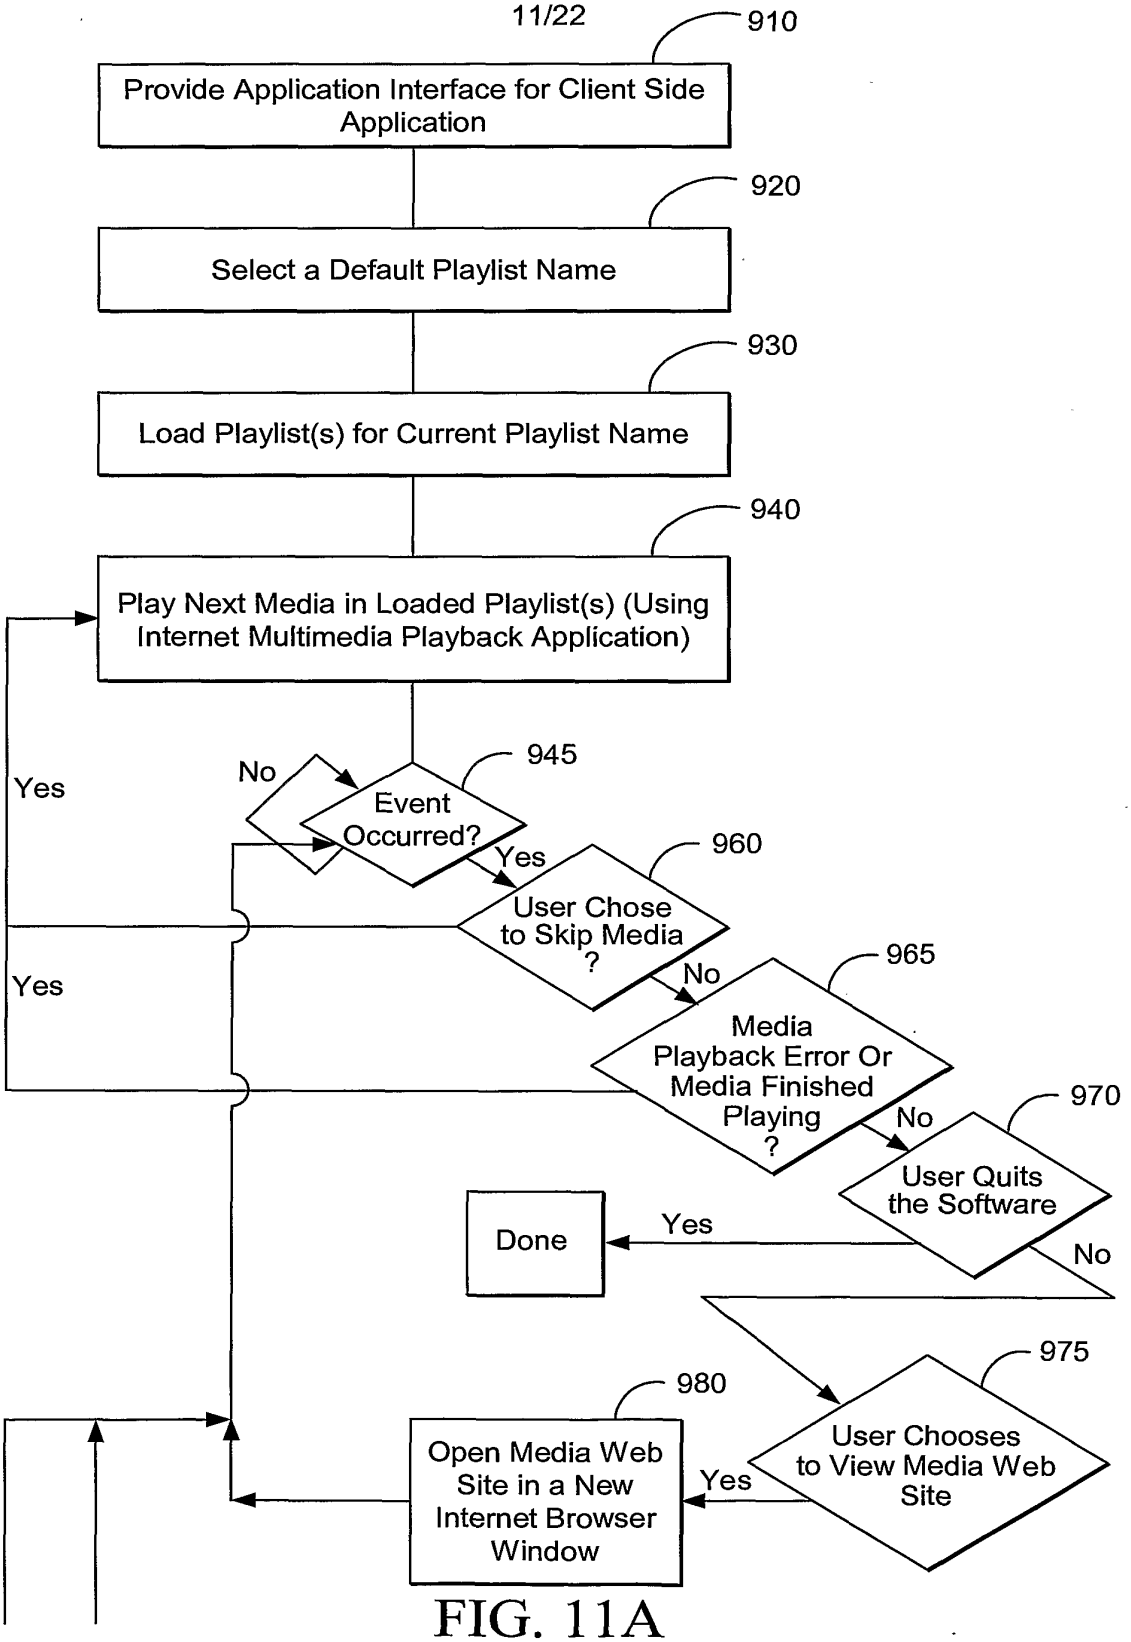


FIG. 11A

SUBSTITUTE SHEET (RULE 26)

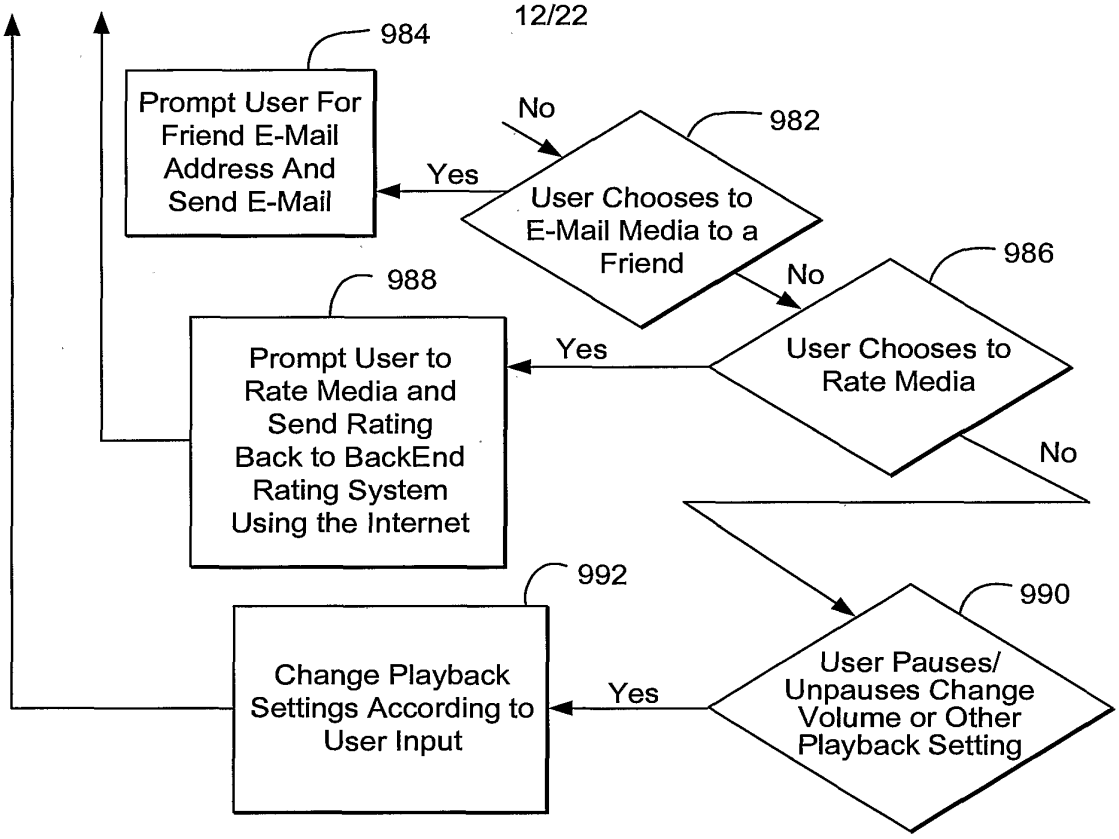
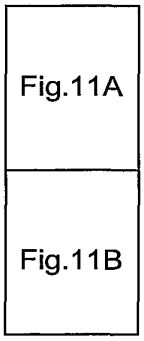

984
Prompt User For Friend E-Mail Address And Send E-Mail
No
Yes
982
User Chooses to E-Mail Media to a Friend
No
986
User Chooses to Rate Media
988
Prompt User to Rate Media and Send Rating Back to BackEnd Rating System Using the Internet
Yes
No
992
Change Playback Settings According to User Input
Yes
990
User Pauses/ Unpauses Change Volume or Other Playback Setting

FIG. 11B


Fig.11A
Fig.11B
Key

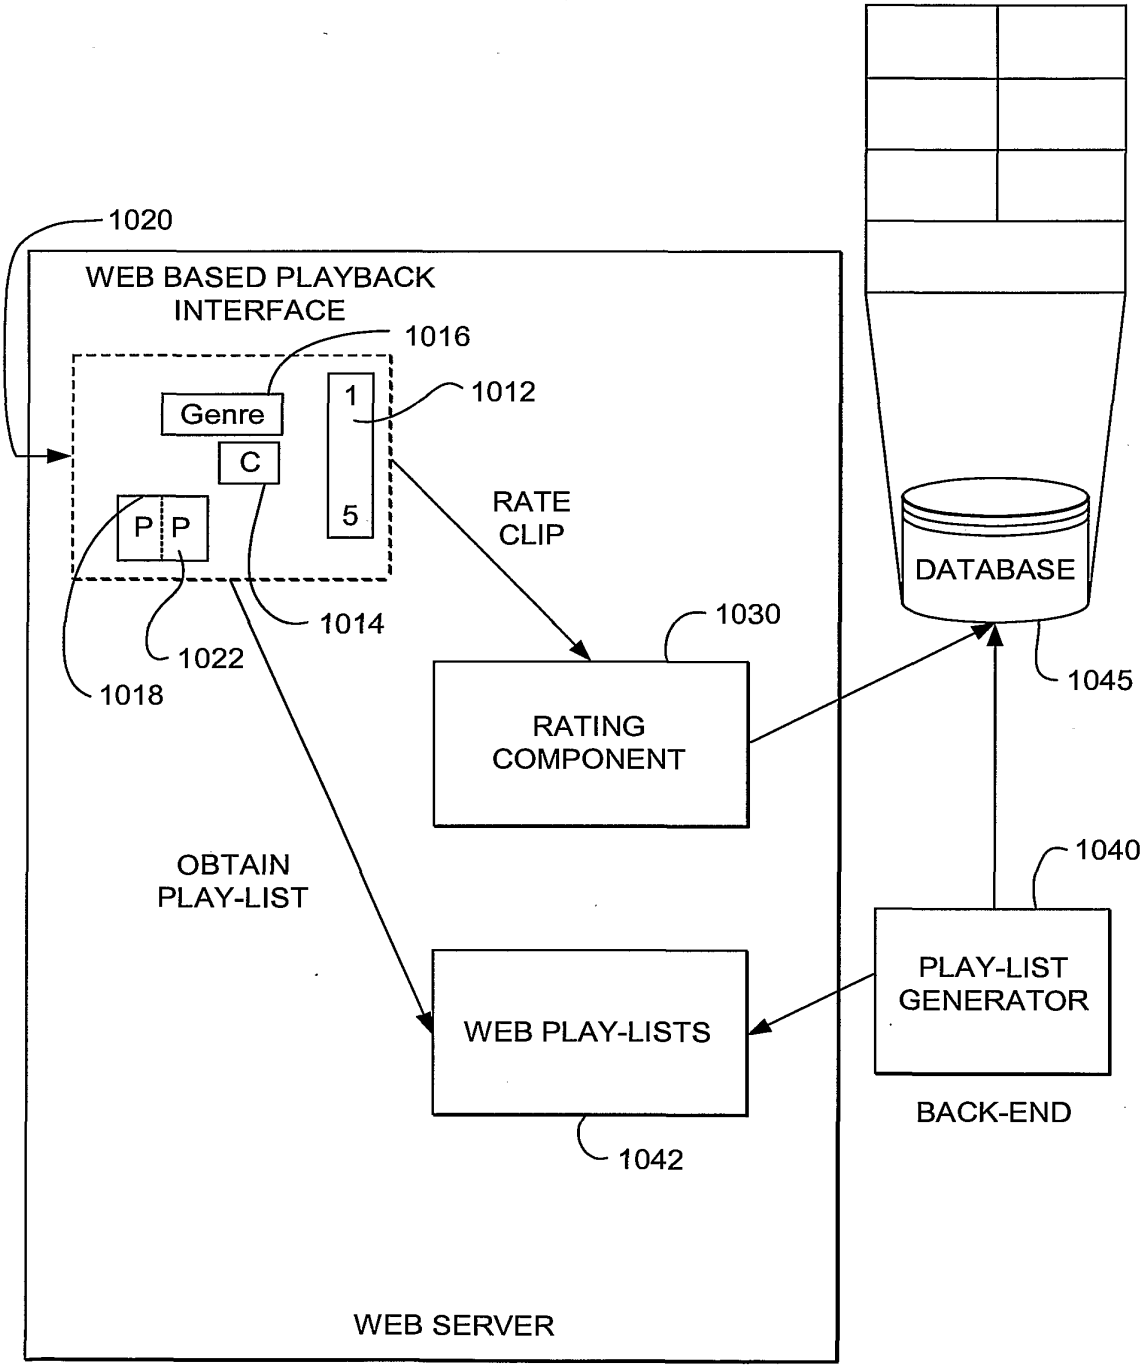

1020
WEB BASED PLAYBACK INTERFACE
1016
Genre
C
1
5
1012
P P
1014
1022
1018
RATE CLIP
1030
RATING COMPONENT
OBTAIN PLAY-LIST
WEB PLAY-LISTS
1042
WEB SERVER
DATABASE
1045
1040
PLAY-LIST GENERATOR
BACK-END

FIG. 12

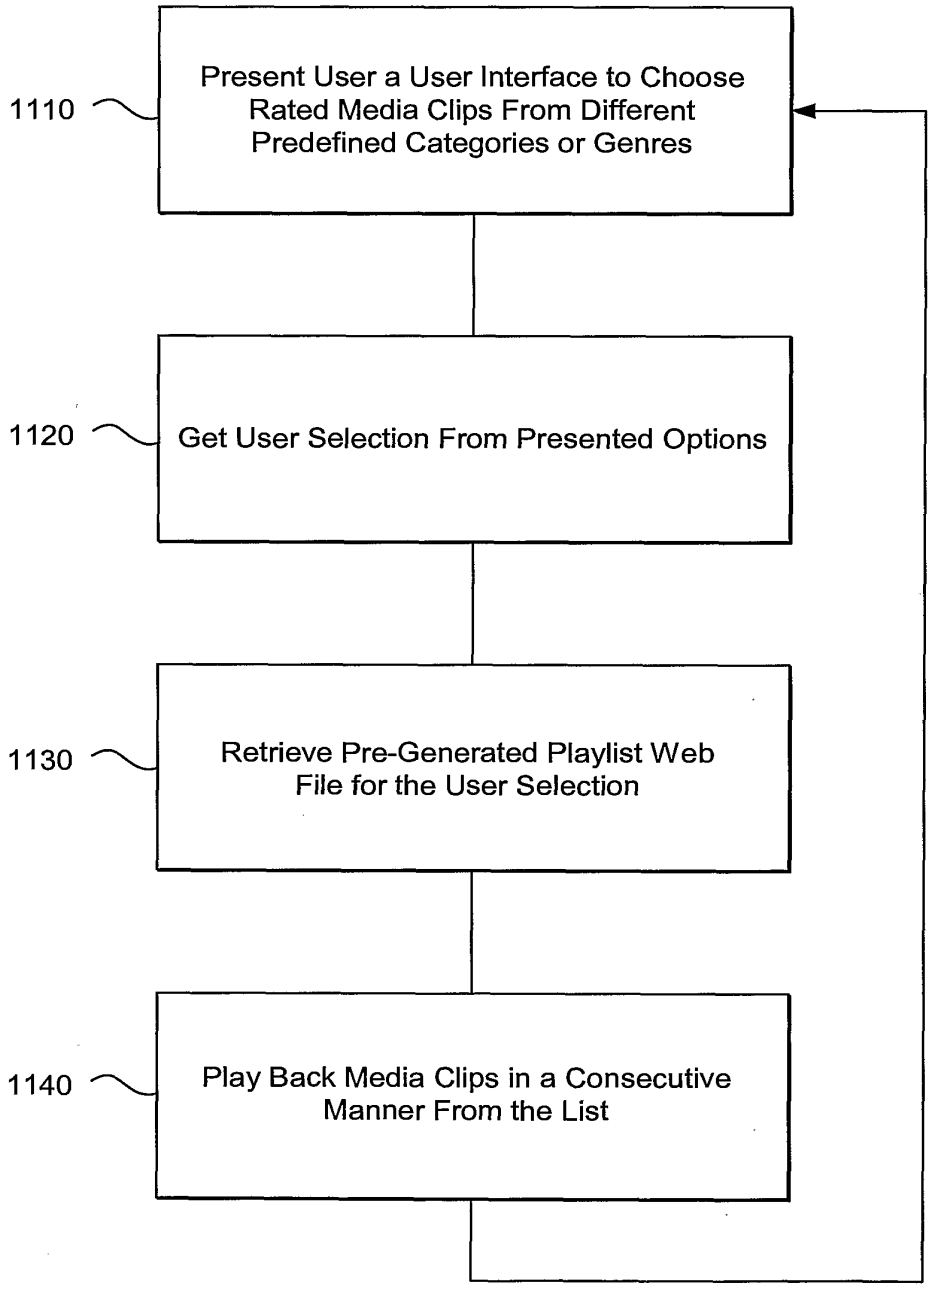

1110
Present User a User Interface to Choose Rated Media Clips From Different Predefined Categories or Genres
1120
Get User Selection From Presented Options
1130
Retrieve Pre-Generated Playlist Web File for the User Selection
1140
Play Back Media Clips in a Consecutive Manner From the List

FIG. 13

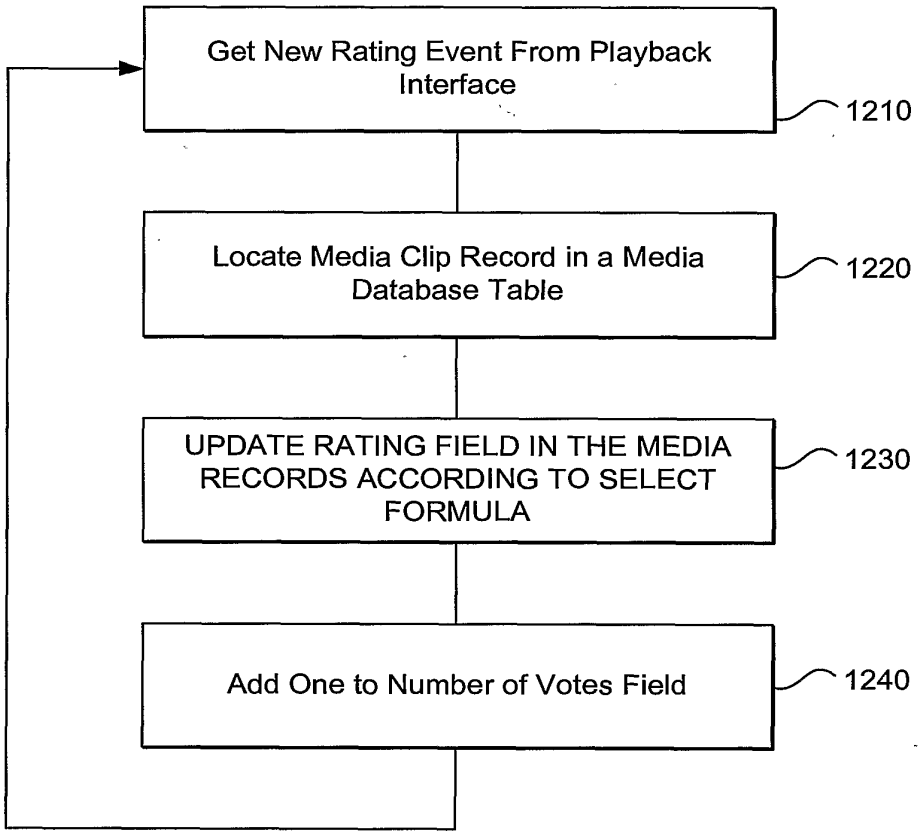


FIG. 14

MEDIA DATABASE TABLE

Media URL	Number of Votes	Rating	Additional Info
URL	O . . . N	O . . . Max Rating	

FIG. 15

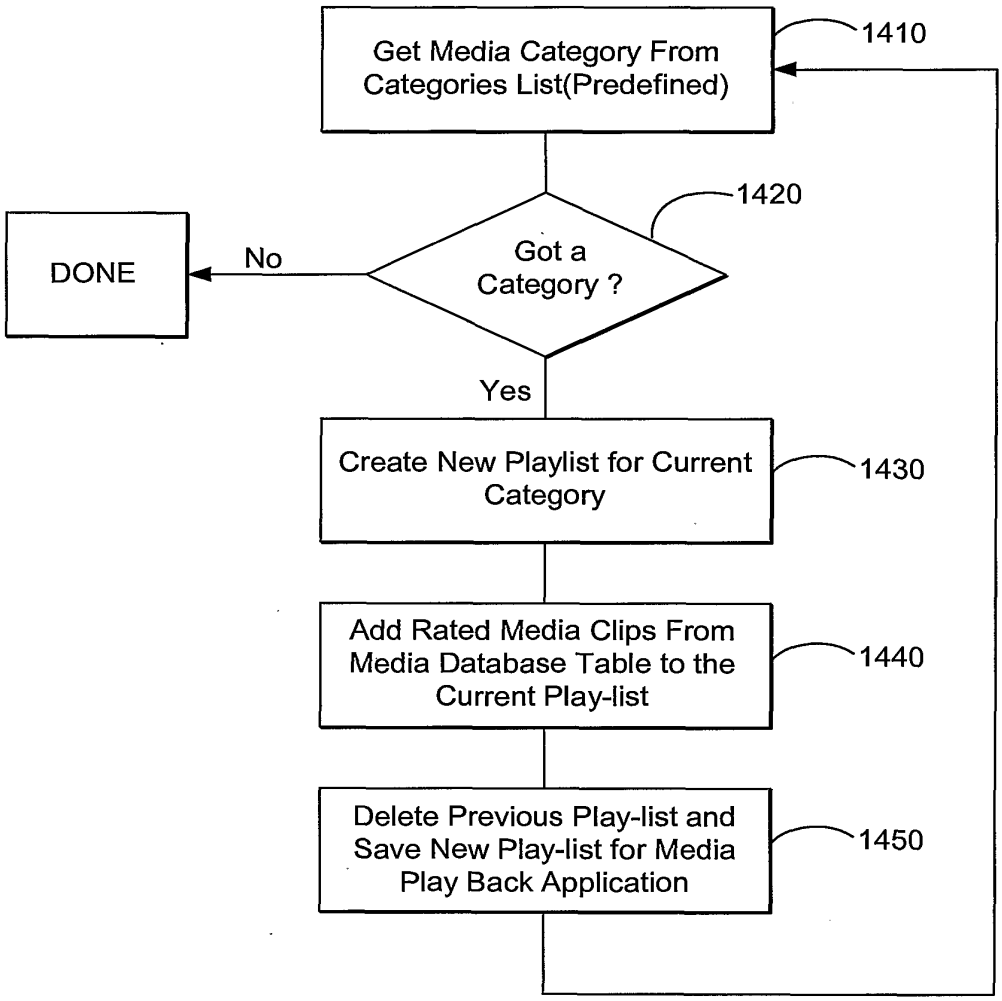

Get Media Category From Categories List(Predefined)
1410
1420
Got a Category ?
No
DONE
Yes
Create New Playlist for Current Category
1430
Add Rated Media Clips From Media Database Table to the Current Play-list
1440
Delete Previous Play-list and Save New Play-list for Media Play Back Application
1450

FIG. 16

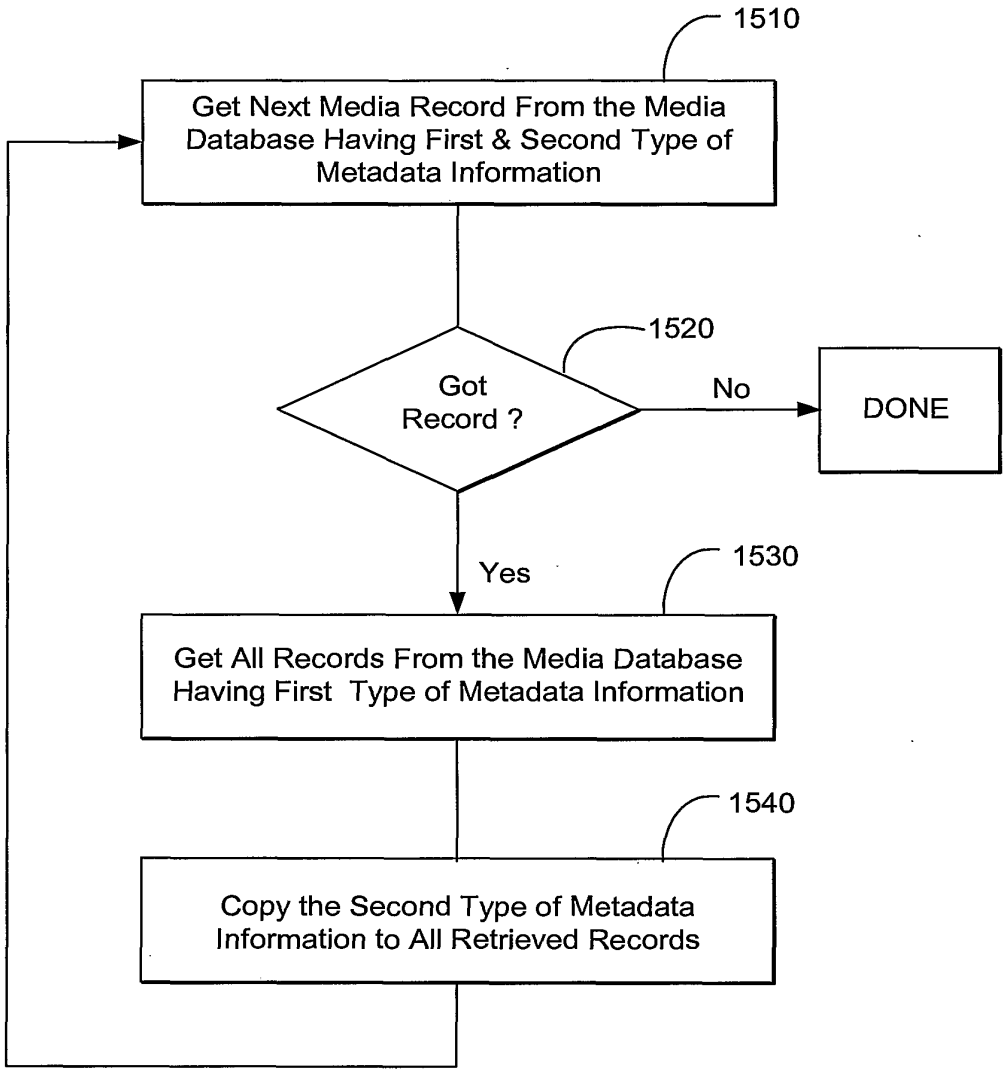

1510
Get Next Media Record From the Media Database Having First & Second Type of Metadata Information
1520
Got Record ?
No
DONE
Yes
1530
Get All Records From the Media Database Having First Type of Metadata Information
1540
Copy the Second Type of Metadata Information to All Retrieved Records

FIG. 17

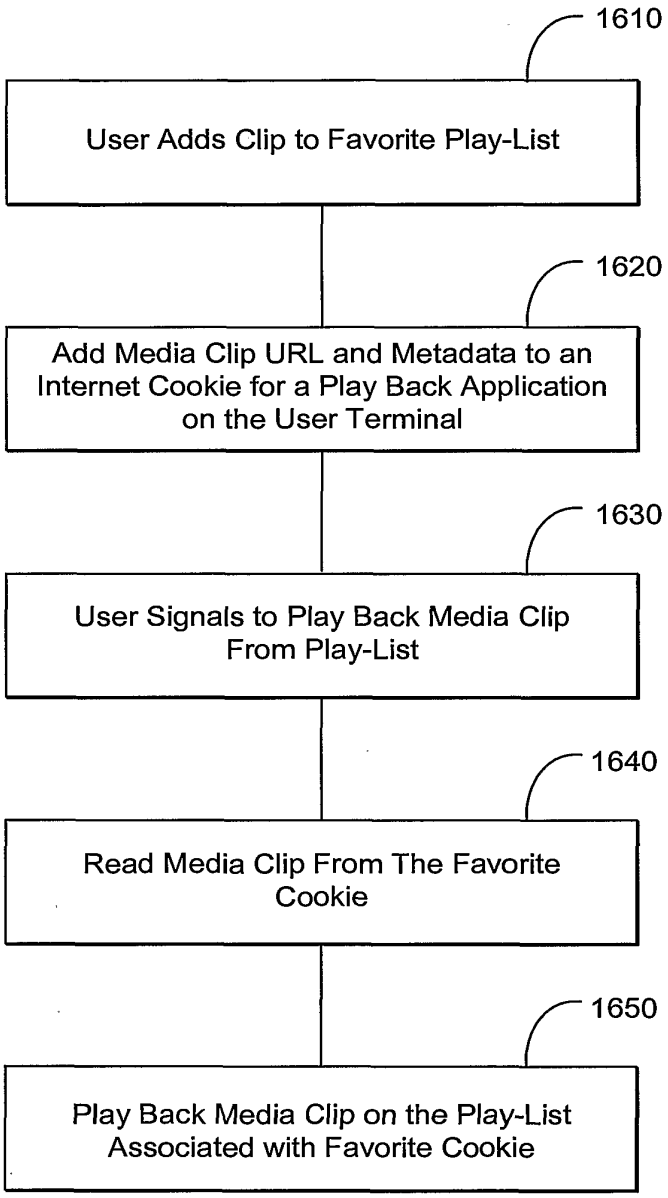

1610
User Adds Clip to Favorite Play-List
1620
Add Media Clip URL and Metadata to an Internet Cookie for a Play Back Application on the User Terminal
1630
User Signals to Play Back Media Clip From Play-List
1640
Read Media Clip From The Favorite Cookie
1650
Play Back Media Clip on the Play-List Associated with Favorite Cookie

FIG. 18

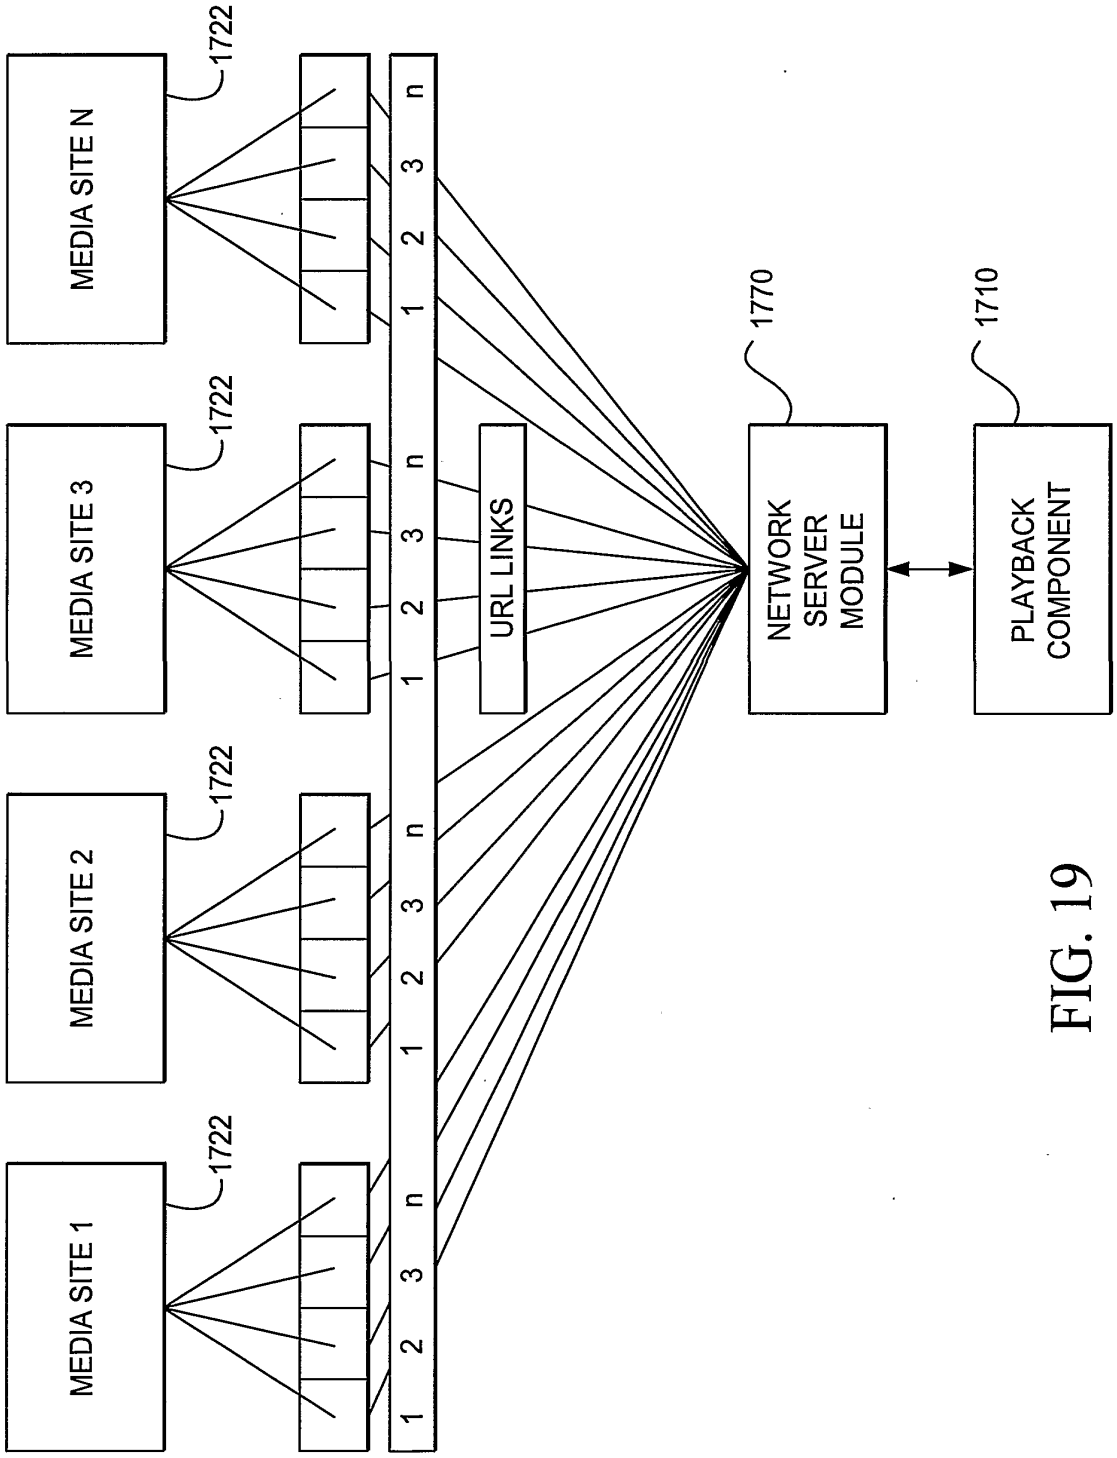

MEDIA SITE N
MEDIA SITE 3
MEDIA SITE 2
MEDIA SITE 1
1722
1722
1722
1722
1 2 3 n
1 2 3 n
1 2 3 n
1 2 3 n
URL LINKS
1770
NETWORK SERVER MODULE
1710
PLAYBACK COMPONENT

FIG. 19

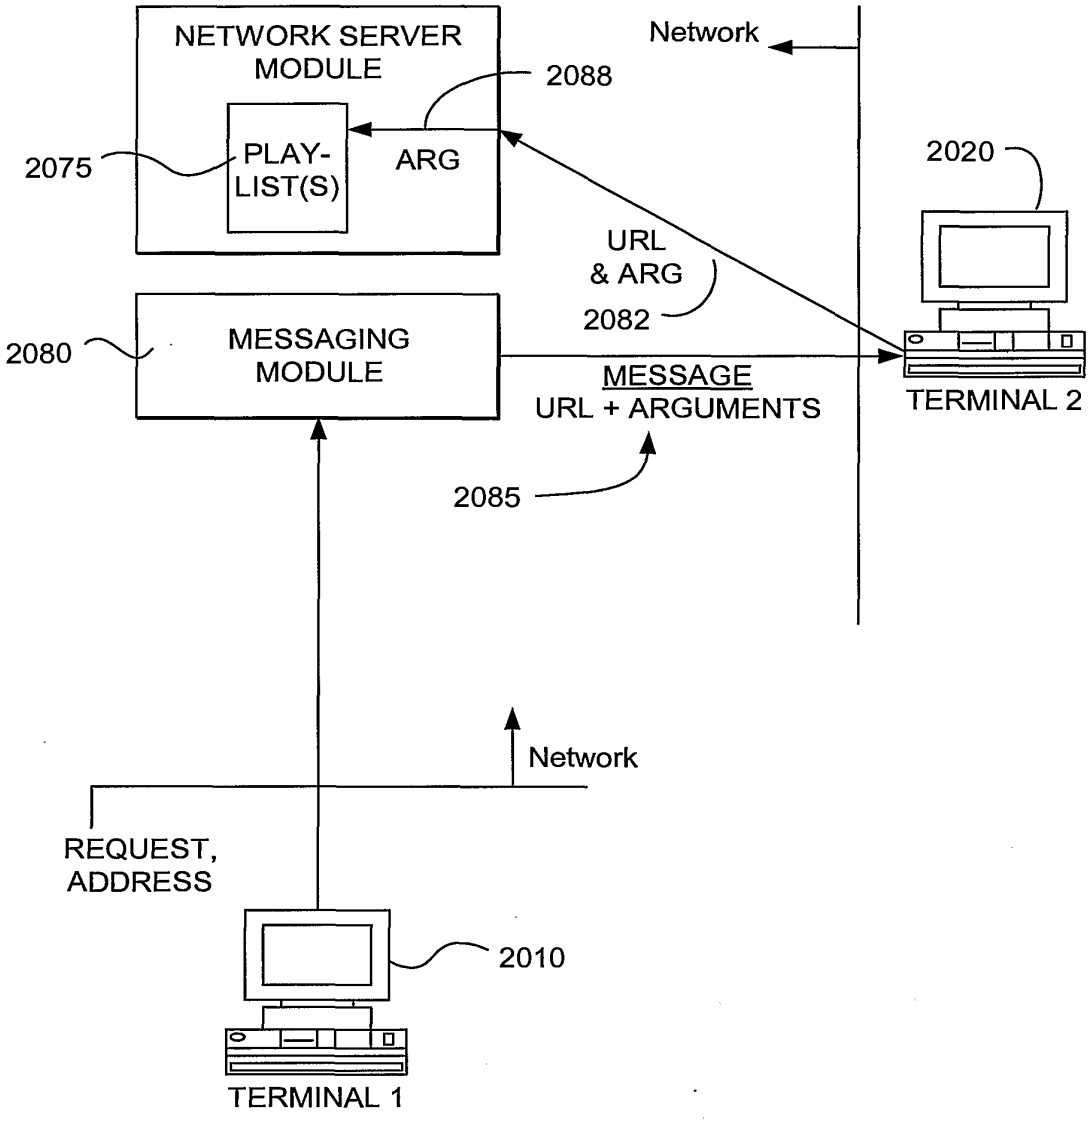

NETWORK SERVER MODULE
2088
Network
2075
PLAY-LIST(S)
ARG
2020
URL & ARG
2082
2080
MESSAGING MODULE
MESSAGE
URL + ARGUMENTS
TERMINAL 2
2085
Network
REQUEST, ADDRESS
2010
TERMINAL 1

FIG. 20

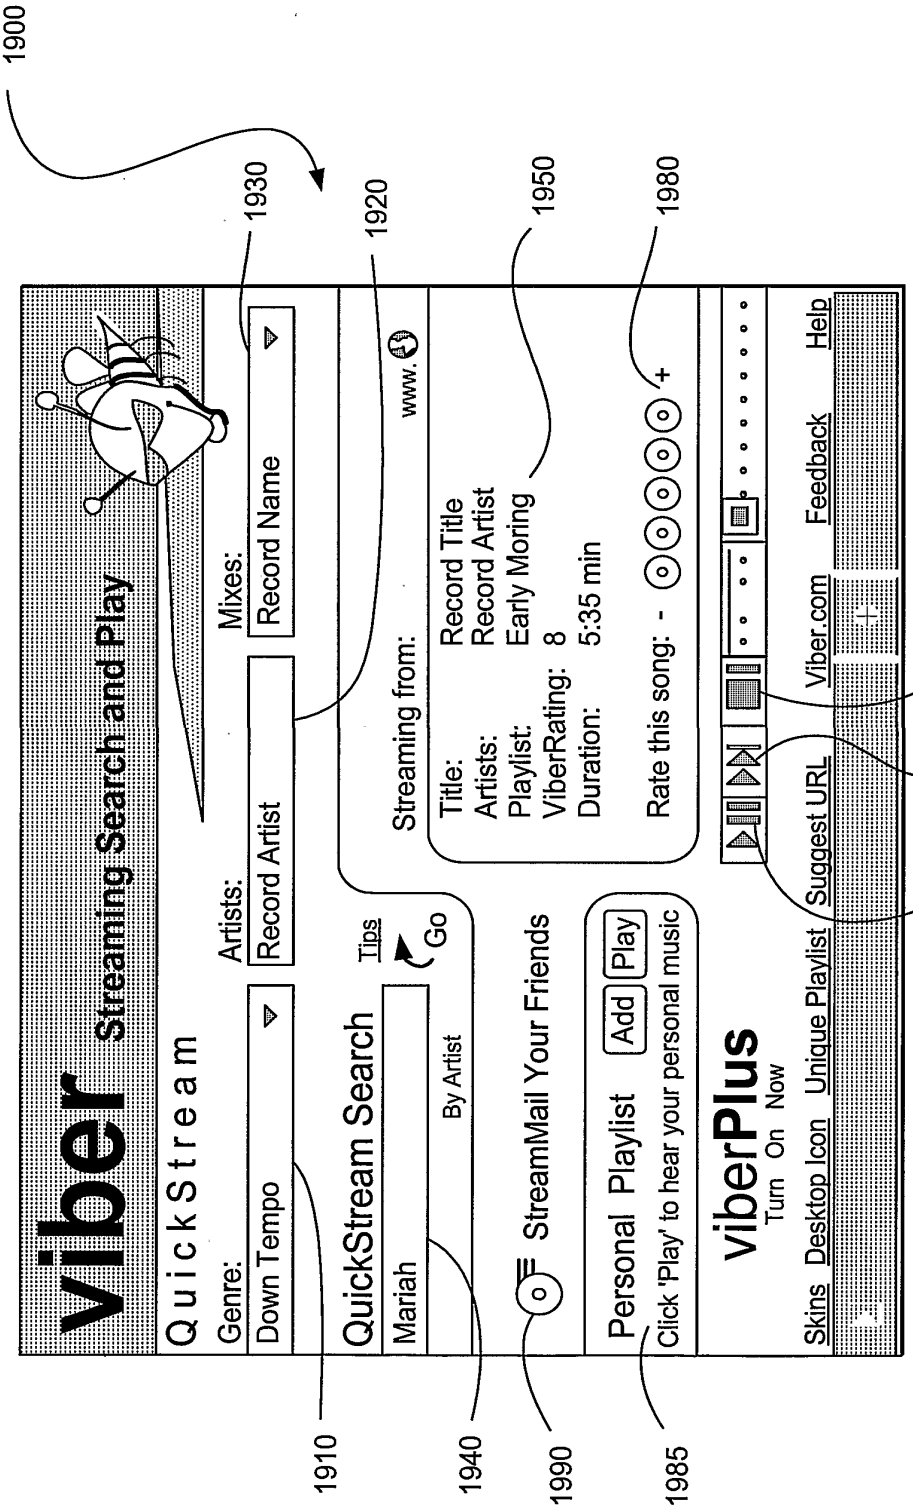

1900
viber Streaming Search and Play
QuickStream
Genre:
Down Tempo
Artists:
Record Artist
Mixes:
Record Name
Tips
Go
QuickStream Search
Mariah
By Artist
StreamMail Your Friends
Personal Playlist
Add
Play
Click 'Play' to hear your personal music
Streaming from:
www.
Title: Record Title
Artists: Record Artist
Playlist: Early Moring
ViberRating: 8
Duration: 5:35 min
Rate this song: -
+
viberPlus
Turn On Now
Skins
Desktop Icon
Unique Playlist
Suggest URL
Viber.com
Feedback
Help
1910
1920
1930
1940
1950
1972
1974
1976
1980
1985
1990

FIG. 21

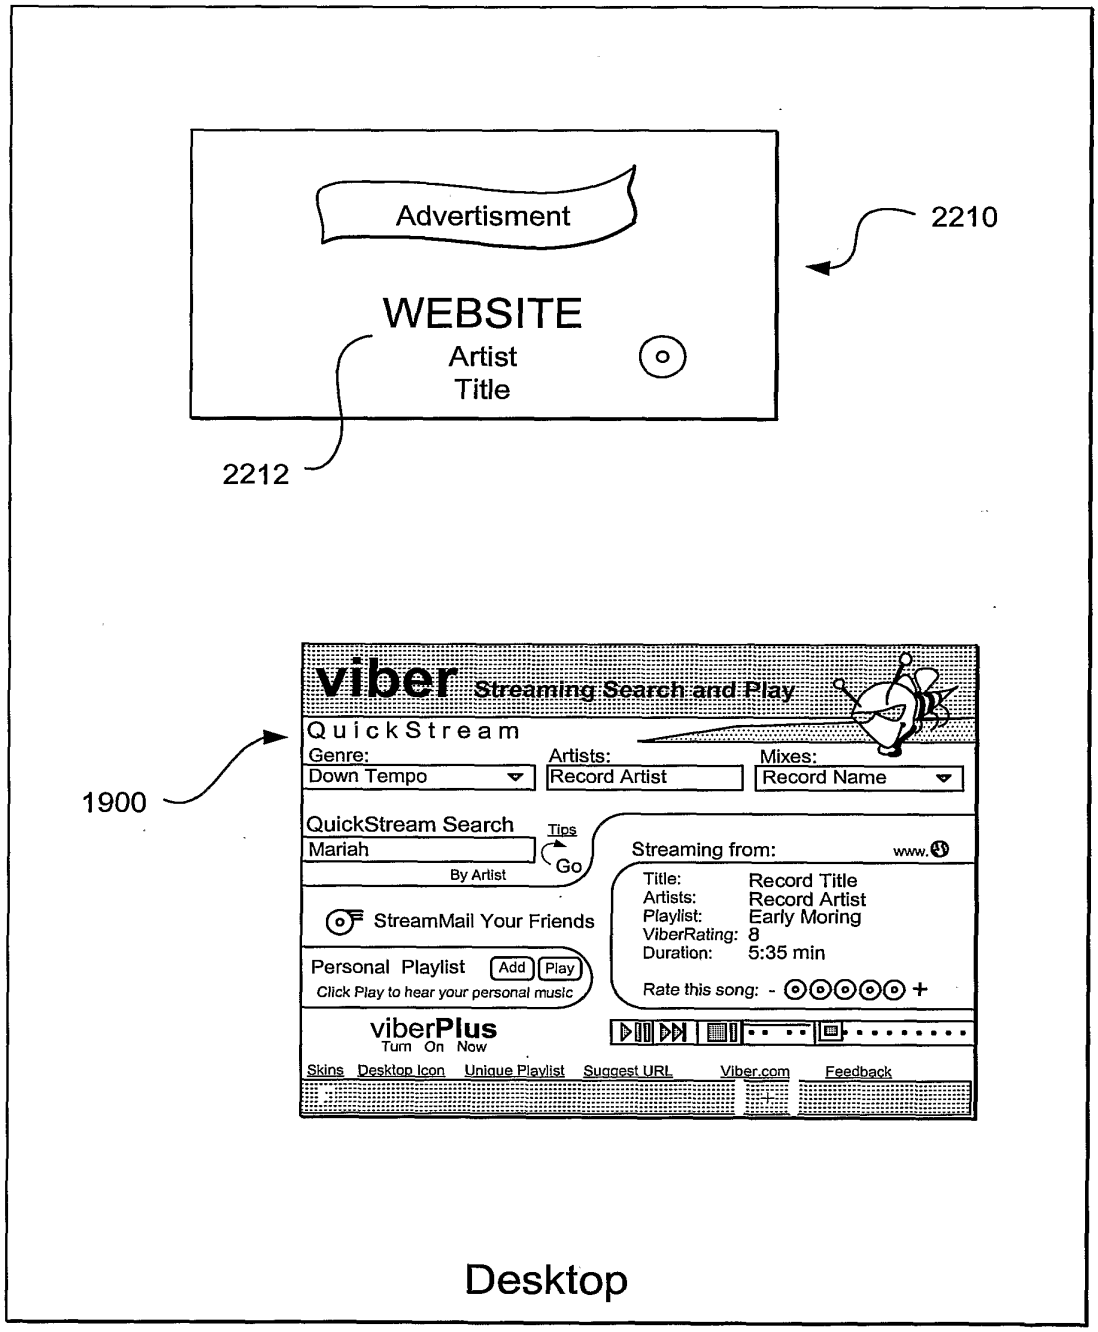

Advertisment
WEBSITE
Artist
Title
2210
2212
1900
viber Streaming Search and Play
QuickStream
Genre:
Down Tempo
Artists:
Record Artist
Mixes:
Record Name
QuickStream Search
Mariah
By Artist
Tips
Go
Streaming from:
www.
Title: Record Title
Artists: Record Artist
Playlist: Early Moring
ViberRating: 8
Duration: 5:35 min
Rate this song: - +
StreamMail Your Friends
Personal Playlist
Add
Play
Click Play to hear your personal music
viberPlus
Turn On Now
Skins
Desktop Icon
Unique Playlist
Suggest URL
Viber.com
Feedback
Desktop

FIG. 22

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	13864248
Filing Date	2013-04-17
First Named Inventor	MILLINGTON, NICHOLAS A. J.
Art Unit	2442
Examiner Name	SURVILLO, OLEG
Attorney Docket Number	SNS-71301US10 (13-0401)

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	5696896		1997-12-09	BADOVINATZ; Peter Richard et al.	
	2	5787249		1998-07-28	BADOVINATZ; Peter Richard et al.	
	3	6839752		2005-01-04	MILLER; Robert et al.	
	4	6907458		2005-06-14	TOMASSETTI; Stephen Robert et al.	
	5	7412499		2008-08-12	CHANG; Chun-Shi et al.	
	6	7743009		2010-06-22	HANGARTNER; Rick et al.	
	7	8020023		2011-09-13	MILLINGTON; Nicholas A.J. et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	MILLINGTON, NICHOLAS A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301US10 (13-0401)

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20010027498		2001-10-04	VAN; De Meulenhof Dennis et al.	
	2	20020184310		2002-12-05	TRAVERSAT; Bernard A. et al.	
	3	20030055892		2003-03-20	HUITEMA; Christian et al.	
	4	20030204273		2003-10-30	DINKER; Darpan et al.	
	5	20030204509		2003-10-30	DINKER; Darpan et al.	
	6	20030227478		2003-12-11	CHATFIELD; Keith M.	
	7	20040133689		2004-07-08	VASISHT; Samrat	
	8	20050216556		2005-09-29	MANION; Todd R. et al.	
	9	20070192156		2007-08-16	GAUGER; Derek K.	
	10	20090193345		2009-07-30	WENSLEY; Paul et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	MILLINGTON, NICHOLAS A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301US10 (13-0401)

	11	20130197682		2013-08-01	MILLINGTON; Nicholas A. J.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	2003037585	JP		2003-02-07	SONY CORP		☐
	2	2003101958	JP		2003-04-04	TOSHIBA CORP		☐
	3	02073851	WO		2002-09-19	PULSE-LINK, INC.		☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	"Advisory Action mailed 18 September 2008 for U.S. Application No. 10/816217 filed 01 April 2004, 8 pages".	☐
	2	"JAPANESE INTELLECTUAL PROPERTY OFFICE, "Decision of Rejection," issued in connection with Japanese patent application No.2012-178711, mailed on July 4, 2014, 3 pages".	☐
	3	"JAPANESE INTELLECTUAL PROPERTY OFFICE, "Office Action Summary," issued in connection with Japanese patent application No. 2012-178711, mailed on November 19, 2013, 5 pages".	☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	MILLINGTON, NICHOLAS A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301US10 (13-0401)

	4	"Non-Final Office Action mailed 19 April 2010 for U.S. Application No. 11/801468 filed 09 May 2007, 16 pages".	☐
	5	"Non-Final Office Action mailed 25 July 2014 for U.S. Application No. 14/184526 filed 19 February 2014".	☐
	6	"Non-Final Office Action mailed 29 November 2010 for U.S. Application No. 11/801468 filed 09 May 2007, 17 pages".	☐
	7	"Non-Final Office Action mailed 31 July 2014 for U.S. Application No. 13/533105 filed 26 June 2012".	☐
	8	"Notice of Allowance mailed 06 May 2011 for U.S. Application No. 11/801468 filed 09 May 2007, 10 pages".	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	MILLINGTON, NICHOLAS A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301US10 (13-0401)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Robert J. Irvine III/	Date (YYYY-MM-DD)	2014-08-11
Name/Print	Robert J. Irvine III	Registration Number	41865

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt	
EFS ID:	19851029
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301US10 (13-0401)
Receipt Date:	14-AUG-2014
Filing Date:	17-APR-2013
Time Stamp:	17:26:44
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Foreign Reference	JP2003037585.pdf	1724624 9e7ba737dab5d1356d5b52153dbe3a9d9cab25e0	no	25

Warnings:
Information:

2	Foreign Reference	JP2003101958.pdf	950403 7831b89a239a5dd9c19be87e7ea69120d63f97ab	no	12
Warnings:					
Information:					
3	Foreign Reference	WO02073851A1.pdf	916609 6920fe9531909d9f0d7af0933a26248436978166	no	26
Warnings:					
Information:					
4	Non Patent Literature	10816217AdvisoryAction09-18-08.pdf	312262 897746fab6918e88f71ff5ae0615cc3935b0c3fa	no	8
Warnings:					
Information:					
5	Non Patent Literature	JP2012-17811-DecisionofRejection7-8-14.pdf	82313 a679627963f695ead81556e982639ff783490ac1	no	3
Warnings:					
Information:					
6	Non Patent Literature	JP2012-17811OfficeActionSummary-11-19-13.pdf	145837 0747263e6ba99a33fdd3382a19ab12154cfe1cb1	no	5
Warnings:					
Information:					
7	Non Patent Literature	11801468-Non-FinalOA04-19-10.pdf	655403 1615233f75130acc4090e548f6575832611c595a	no	16
Warnings:					
Information:					
8	Non Patent Literature	14184526-Non-FinalOA-07-25-14.pdf	259390 246899dd54263c7304b05966e7032010d6795cf6	no	9
Warnings:					
Information:					
9	Non Patent Literature	11801468Non-FinalOA11-29-10.pdf	692235 7382d1289f9bcf6333083af286322028ec2bd558	no	17
Warnings:					
Information:					
10	Non Patent Literature	13533105-Non-FinalOA07-31-14.pdf	1034996 66ae4ac9d0c13b66770b7ae5a3a519233973e89d	no	31
Warnings:					
Information:					

11	Non Patent Literature	11801468NOA05-06-11.pdf	451323 f4346fdd7726aa5033684c6ed8e8e6844ff8c93e	no	10
Warnings:					
Information:					
12	Information Disclosure Statement (IDS) Form (SB08)	SNS-71301US10-SupplDS-AUG2014.pdf	707822 335ec0d9cfe0972fe7c0f4ab7dbc3abb128979eb	no	6
Warnings:					
Information:					
		Total Files Size (in bytes):	7933217		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Espacenet

Bibliographic data: JP2003037585 (A) — 2003-02-07

NETWORK SYSTEM, OUTPUT DEVICE, AND METHOD FOR SYNCHRONIZING NETWORK SYSTEM

Inventor(s): FUJINAMI YASUSHI ± (FUJINAMI YASUSHI)

Applicant(s): SONY CORP ± (SONY CORP)

Classification: - international: ***G06F1/14; G06F13/00; H04B1/20; H04L12/28; H04L12/413; H04L7/00; H04N7/62; H04S3/00;*** *H04J3/06; H04L12/56;* (IPC1-7): G06F1/14; G06F13/00; H04L7/00

- cooperative: **G06F1/14**; **H04B1/205**; **H04L12/2803**; **H04L12/2807**; **H04L12/2838**; **H04L47/28**; **H04N21/4305**; **H04S3/008**; H04J3/0638

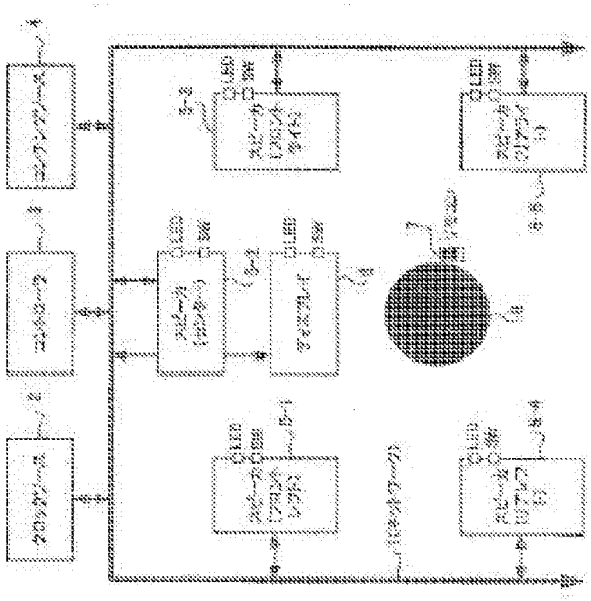
Application number: JP20010224983 20010725

Priority number(s): JP20010224983 20010725

Also published as: JP3591493 (B2) EP1320213 (A1) EP1320213 (A4) US2004010727 (A1) TW577231 (B) more

Abstract of JP2003037585 (A)

PROBLEM TO BE SOLVED: To provide a clock synchronized between reproducing devices con nected to the digital network. SOLUTION: A clock is transmitted to a sink device group like speakers 5 (5-1 to 5-5) connected to a network 1 and a display 6. And also this network system is equipped with a clock source 2 which transmits time information per prescribed time, a contents source 4 which provides the contents for the sink device group through the network 1, and a controller 3 which provides for the contents source 4 the delay time based on the network delay when the content data are transmitted to the sink device group and the decode delay in the sink device group when reproducing the contents.



(19)日本国特許庁（ＪＰ）　(12) **公開特許公報**（Ａ）　(11)特許出願公開番号
特開2003－37585
（P2003－37585A）
(43)公開日　平成15年２月７日(2003.2.7)

(51)Int.Cl.[7]	識別記号	ＦＩ		テーマコード（参考）
Ｈ０４Ｌ　7/00		Ｈ０４Ｌ　7/00	Ｚ	５Ｂ０８９
Ｇ０６Ｆ　1/14		Ｇ０６Ｆ　13/00	３５１Ｃ	５Ｋ０４７
13/00	３５１	1/04	３５１Ｂ	

審査請求　未請求　請求項の数13　ＯＬ　（全 23 頁）

(21)出願番号　特願2001－224983(P2001－224983)

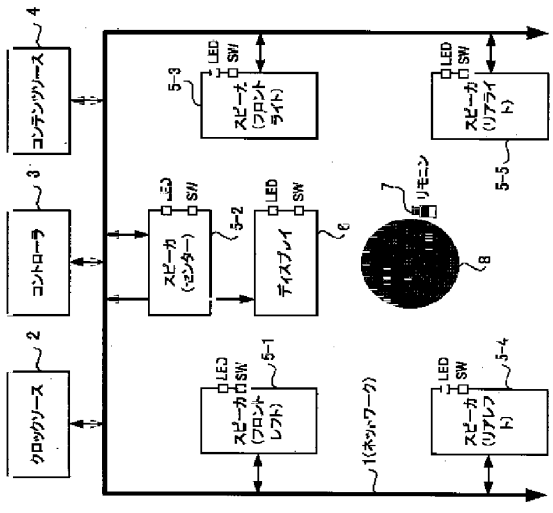
(22)出願日　平成13年７月25日(2001.7.25)

(71)出願人　000002185
ソニー株式会社
東京都品川区北品川６丁目７番35号
(72)発明者　藤波　靖
東京都品川区北品川６丁目７番35号　ソニー株式会社内
(74)代理人　100104880
弁理士　古部　次郎　（外１名）

Ｆターム(参考)　5B089 GA23 JB11 KA12 KB04 KE10
5K047 AA18 DD02 GG01 GG06

(54)【発明の名称】　ネットワークシステム、出力機器、およびネットワークシステムの同期方法

(57)【要約】
【課題】　ディジタルネットワークに接続された各再生機器の間において、クロックと位相の同期をとる。
【解決手段】　ネットワーク１に接続されるスピーカ５(５－１～５－５)やディスプレイ６等のシンク機器群に対してクロックを送信すると共に所定時間ごとに時刻情報を送信するクロックソース２と、このシンク機器群に対してネットワーク１を介してコンテンツを提供するコンテンツソース４と、コンテンツを再生するにあたり、コンテンツデータがシンク機器群に送られるときのネットワークディレイとシンク機器群におけるデコードディレイとに基づく遅延時間をコンテンツソース４に対して提供するコントローラ３とを備えた。



【特許請求の範囲】
【請求項１】　複数の機器がディジタルネットワークに接続されるネットワークシステムであって、
前記複数の機器に対してクロックを配信すると共に時刻情報を配信するクロック配信手段と、
前記クロック配信手段により配信された前記クロックおよび前記時刻情報に基づき、各々の機器においてクロックを調整するクロック調整手段と、
前記ディジタルネットワーク上をストリームが通信される際に生じるネットワークディレイを考慮して前記複数の機器に対してディレイ補正を施すディレイ補正手段とを備えたことを特徴とするネットワークシステム。
【請求項２】　前記ディレイ補正手段は、前記複数の機器が前記ストリームに関する変換に際して生じる変換ディレイを考慮してディレイ補正を施すことを特徴とする請求項１記載のネットワークシステム。
【請求項３】　前記ディレイ補正手段は、前記機器が置かれた位置によって生じるディレイ補正を含むことを特徴とする請求項１記載のネットワークシステム。
【請求項４】　ディジタルネットワークに接続されるシンク機器群に対してクロックを送信するクロックソースと、
前記ディジタルネットワークに接続される前記シンク機器群に対して当該ディジタルネットワークを介してコンテンツを提供するコンテンツソースと、
前記コンテンツを再生するにあたり、コンテンツデータが前記シンク機器群に送られるときのネットワークディレイと当該シンク機器群におけるデコードディレイとに基づく遅延時間を前記コンテンツソースに対して提供するコントローラとを備えたことを特徴とするネットワークシステム。
【請求項５】　前記クロックソースは、前記シンク機器群に対して所定時間ごとに時刻情報を送信することを特徴とする請求項４記載のネットワークシステム。
【請求項６】　前記コンテンツソースは、前記コントローラから提供された前記遅延時間に基づく遅延情報メッセージを作成し、コンテンツデータに伴って前記シンク機器群に対して配信することを特徴とする請求項４記載のネットワークシステム。
【請求項７】　ディジタルネットワークに接続され、当該ディジタルネットワークを介して提供されるコンテンツデータをデコードする出力機器であって、
前記ディジタルネットワークを介して受信した基準時刻信号に基づいてクロックを再生するクロック再生手段と、
前記ディジタルネットワークを介して受信した前記コンテンツデータに対して必要なディレイを施し、デコードして出力するストリーム再生手段とを備えたことを特徴とする出力機器。
【請求項８】　内部にて用いられるクロックを発振するクロック発振手段とを更に備え、
前記クロック再生手段は、受信された前記基準時刻信号と前記クロック発振手段からの出力の値とを比較して、当該クロック発振手段の発振周波数を調整することにより、前記クロックを再生することを特徴とする請求項７記載の出力機器。
【請求項９】　前記ストリーム再生手段は、前記コンテンツデータに伴って受信される当該コンテンツデータを出力すべき時刻情報から、自らのデコードディレイ時間を考慮してデコードを開始することを特徴とする請求項７記載の出力機器。
【請求項１０】　ディジタルネットワークに接続される複数の機器にて入力または出力の同期をとるためのネットワークシステムの同期方法であって、
前記複数の機器に対してクロックと共に時刻情報を配信し、
配信される前記クロックと前記時刻情報とに基づいて、共通な時計を前記複数の機器で動作させ、
ネットワークディレイおよび前記複数の機器における変換ディレイを考慮して前記時計の時刻を用いた入力タイミングまたは出力タイミングを当該複数の機器に使用し、
使用された前記入力タイミングまたは前記出力タイミングと各々の機器における変換ディレイとに基づいて、当該各々の機器における変換を開始することを特徴とするネットワークシステムの同期方法。
【請求項１１】　前記入力タイミングまたは前記出力タイミングは、前記ネットワークディレイの最大時間に、前記複数の機器における前記変換ディレイの最大時間を加えて決定されることを特徴とする請求項１０記載のネットワークシステムの同期方法。
【請求項１２】　ディジタルネットワークに接続される複数の機器にて出力の同期をとるためのネットワークシステムの同期方法であって、
前記ディジタルネットワークを介してクロックと共に時刻情報を受信し、
受信した前記クロックおよび前記時刻情報に基づいてクロックを調整し、
コンテンツデータと共に当該コンテンツデータの再生を開始する時刻を示す情報を受信し、
受信した前記情報に基づき、前記コンテンツデータをデコードする際に生じるディレイに基づいてデコードの開始タイミングを決定し、
決定した前記開始タイミングによりデコードを開始して前記コンテンツデータを再生することを特徴とするネットワークシステムの同期方法。
【請求項１３】　前記ディジタルネットワークを介してコンテンツソースとしてのネットワークアドレスの指定を受け、
前記ネットワークアドレスに対してなされるディレイの

指示に基づいて前記開始タイミングを決定することを特徴とする請求項１２記載のネットワークシステムの同期方法。

【発明の詳細な説明】

【０００１】

【発明の属する技術分野】本発明は、複数の機器を備えたネットワークシステム等に係り、特に、ネットワークに接続された複数の機器において同期をとるネットワークシステム等に関する。

【０００２】

【従来の技術】近年、通信ネットワークが企業、学校、家庭に浸透しつつ有り、このインフラを利用してオーディオビデオのコンテンツを配信しようという試みがなされている。　これをもう一歩進めて、家庭内のオーディオビデオの配線をネットワークで置き換えることが当然考えられる。しかしながら、新しい技術分野であるために、未だ充分な検討が行われていなかった。

【０００３】家庭内のオーディオビデオの配線を通信ネットワークで置き換えることは、例えばＤＶＤ(Digital Versatile Disk)、ディスプレイ、スピーカなどをそれぞれ通信ネットワークに接続することを意味する。例えば、ＤＶＤを再生すると、ディジタル通信ネットワークを介してディスプレイ、スピーカからそれぞれ動画と音声が出力されるということになる。

【０００４】

【発明が解決しようとする課題】しかしながら、現在の通信ネットワークでは、その上を流れるデータにおいて、配信されるタイミングが保証されないという問題点がある。即ち、複数の入出力機器間の入出力位相関係(時間的前後関係)にて、現在のディジタルネットワークでは、伝送遅延が固定ではないことから、十分に小さい位相誤差で出力されるという保証がない。また、ディジタル通信ネットワークを介した場合に、アナログ接続に比べて遅延の絶対値が大きくなる。かかる場合、上記の例では、ＤＶＤから発信されたコンテンツにおいて、ディスプレイや複数存在するスピーカに到着する時刻がまちまちになってしまう。これをそのまま再生すると、結果としてまちまちなタイミングで再生が行われることになる。

【０００５】これは、即ち、アナログケーブルで接続されていた場合に実現されていた出力信号の同期が達成されないということになる。ディジタルネットワーク化によるメリットは数多く存在するが、ディスプレイとスピーカ、あるいはスピーカ同士の出力位相、更にはマイクとカメラの入力位相を合わせることが出来なければ、家庭内におけるオーディオビデオの配線をディジタル通信ネットワークで置き換えることは困難となる。

【０００６】本発明は、以上のような技術的課題を解決するためになされたものであって、ディジタルネットワークに接続された各機器の間において、クロックを同期させると共に、入出力位相を一致させることを主たる目的とする。

【０００７】

【課題を解決するための手段】かかる目的のもと、本発明は、複数の機器がディジタルネットワークに接続されるネットワークシステムであって、複数の機器に対してクロックを配信すると共に時刻情報を配信するクロック配信手段と、配信されたクロックおよび時刻情報に基づき各々の機器においてクロックを調整するクロック調整手段と、ディジタルネットワーク上をストリームが通信される際に生じるネットワークディレイ、複数の機器がストリームに関する変換に際して生じる変換ディレイを考慮して複数の機器に対してディレイ補正を施すディレイ補正手段とを備えたことを特徴としている。

【０００８】また、本発明が適用されるネットワークシステムは、ディジタルネットワークに接続されるシンク機器群に対してクロックを送信するクロックソースと、このシンク機器群に対してディジタルネットワークを介してコンテンツを提供するコンテンツソースと、コンテンツを再生するにあたり、コンテンツデータがシンク機器群に送られるときのネットワークディレイとシンク機器群におけるデコードディレイとに基づく遅延時間をコンテンツソースに対して提供するコントローラとを備えたことを特徴としている。

【０００９】ここで、このクロックソースは、シンク機器群に対して所定時間ごとに時刻情報を送信することを特徴とすれば、位相調整の前提として、クロックを同期させることができる点で好ましい。また、このコンテンツソースは、コントローラから提供された遅延時間に基づく遅延情報メッセージを作成し、コンテンツデータに伴ってシンク機器群に対して配信することを特徴とすれば、各シンク機器群では、自らのデコードディレイを考慮してデコードを開始することで、シンク機器群の間における再生タイミングを合わせることが可能となる。

【００１０】他の観点から把握すると、本発明は、ディジタルネットワークに接続され、このディジタルネットワークを介して提供されるコンテンツデータをデコードする出力機器であって、ディジタルネットワークを介して受信した基準時刻信号に基づいてクロックを再生するクロック再生手段と、ディジタルネットワークを介して受信したコンテンツデータに対して必要なディレイを施し、デコードして出力するストリーム再生手段とを備えたことを特徴としている。

【００１１】ここで、内部にて用いられるクロックを発振するクロック発振手段とを更に備え、クロック再生手段は、受信された基準時刻信号とクロック発振手段からの出力の値とを比較して、クロック発振手段の発振周波数を調整することにより、クロックを再生することを特徴とすることができる。

【００１２】一方、本発明を方法のカテゴリから把える

と、本発明は、ディジタルネットワークに接続される複数の機器にて入力または出力の同期をとるためのネットワークシステムの同期方法であって、複数の機器に対してクロックと共に時刻情報を配信し、配信されるクロックと時刻情報とに基づいて、共通な時計を複数の機器で動作させ、ネットワークディレイおよび複数の機器におけるデコードディレイやエンコードディレイ等の変換ディレイを考慮して時計の時刻を用いた入力タイミングまたは出力タイミングを複数の機器に使用し、使用された入力タイミングまたは出力タイミングと各々の機器における変換ディレイとに基づいて、各々の機器における変換を開始することを特徴とすることができる。

【００１３】また、本発明は、ネットワークに接続されるシンク機器にて適用される発明として把握できる。即ち、本発明におけるネットワークシステムの同期方法は、ディジタルネットワークを介してクロックと共に時刻情報を受信し、受信したクロックおよび時刻情報に基づいてクロックを調整し、コンテンツデータと共にコンテンツデータの再生を開始する時刻を示す情報を受信し、受信した情報に基づき、コンテンツデータをデコードする際に生じるディレイに基づいてデコードの開始タイミングを決定し、決定した開始タイミングによりデコードを開始してコンテンツデータを再生することを特徴としている。

【００１４】

【発明の実施の形態】以下、添付図面に示す実施の形態に基づいて本発明を詳細に説明する。図１は、本実施の形態におけるネットワークシステムの全体像を示した構成図である。図１に含まれる各機器は、ネットワーク１により相互に接続されている。リスニングポジション８は、特定の機器ではないが、このネットワークシステムにおけるユーザの存在位置を模式的に示したものである。ユーザは、リスニングポジション８にて操作用のリモコン７を用いてネットワークシステムに対する指令を行うことができる。

【００１５】ネットワーク１には、システムにクロックを供給するクロックソース２、システム全体を制御するコントローラ３、コンテンツデータ(信号)を送り出すコンテンツソース４、ユーザに対して音声信号を供給する５つのスピーカ５(５－１～５－５)、ユーザに対し映像信号を供するディスプレイ６、の各入出力機器が接続されている。ここで、ディスプレイ６とスピーカ５とを、コンテンツソース４に対して(コンテンツ)シンク機器群と呼ぶことができる。また、各スピーカ５は、図１に示すような位置、即ち、フロントレフト(５－１)、センター(５－２)、フロントライト(５－３)、リアレフト(５－４)、リアライト(５－５)の各位置に設置されている。各スピーカ５およびディスプレイ６は、各々ＬＥＤとプッシュスイッチを備えており、ＬＥＤを用いたユーザへの働きかけとプッシュスイッチによるユーザからの反応を得ることができる。

【００１６】図２は、クロックソース２における第１の構成例を示した図である。この第１の構成例におけるクロックソース２は、クロック発振器１１、カウンタ１２、タイマ１３、ラッチ１４、およびネットワークインタフェース１５を備えており、このネットワークインタフェース１５を介してネットワーク１と接続されている。クロック発振器１１の出力はカウンタ１２とタイマ１３に各々入力される。カウンタ１２の出力(時刻信号)はラッチ１４に入力される。またラッチ１４の出力はネットワークインタフェース１５に入力される。更に、タイマ１３の出力(トリガ信号)は、ラッチ１４のラッチ信号端子とネットワークインタフェース１５の送信信号端子に入力されている。クロックソース２は、ネットワーク１によって相互に接続された機器に対し、このような構成を用いてクロックを送信している。

【００１７】図３は、クロックソース２における第２の構成例を示した図である。この第２の構成例におけるクロックソース２は、クロック発振器２１、カウンタ２２、タイマ２３、およびラッチ２４を備えている。またクロックソース２は、制御装置としてのＣＰＵ２５、ＲＯＭ２７、ＲＡＭ２８を備えており、ネットワークインタフェース２６を介してネットワーク１と接続されている。このＣＰＵ２５とＲＯＭ２７、ＲＡＭ２８は、システムバス２０により相互に接続されたマイクロコンピュータを構成している。また、ネットワークインタフェース２６はシステムバス２０に接続されている。クロック発振器２１の出力はカウンタ２２とタイマ２３に入力される。カウンタ２２の出力(時刻信号)はラッチ２４に入力される。ラッチ２４の出力はシステムバス２０に接続されている。またタイマ２３の出力(トリガ信号)はラッチ２４のラッチ信号端子とＣＰＵ２５の割り込み端子に入力される。クロックソース２は、ネットワーク１によって相互に接続された機器に対し、このような構成を用いてクロックを送信している。

【００１８】図４は、図１に示したコントローラ３の構成を説明するための図である。本実施の形態におけるコントローラ３は、ＣＰＵ３１、ＲＯＭ３２、およびＲＡＭ３３を備えており、これらはシステムバス３４により相互に接続されてマイクロコンピュータを構成している。また、コントローラ３は、システムバス３４に接続されるネットワークインタフェース３５を備え、このネットワークインタフェース３５を介してネットワーク１と接続されている。またコントローラ３はリモコン受光部３６を有しており、リモコン受光部３６はシステムバス３４に接続されている。このコントローラ３は、図１のシステム全体を制御する役割を担っており、具体的な制御は、システムの構成の設定や、システムの動作(コンテンツの再生)の制御等である。特に本実施の形態では、コンテンツ再生時におけるディレイの設定が機能と

して重要である。

【００１９】図５は、図１に示したコンテンツソース４の構成を説明するための図である。本実施の形態におけるコンテンツソース４は、ネットワーク１に接続された機器に、ネットワーク１を経由して、コンテンツストリームを供給している。図５に示すように、コンテンツソース４は、ＣＰＵ４１、ＲＯＭ４２、およびＲＡＭ４３を備えており、これらはシステムバス４４により相互に接続されマイクロコンピュータを構成している。また、コンテンツソース４は、ネットワークインタフェース４５を介してネットワーク１と接続されている。ネットワークインタフェース４５はシステムバス４４に接続されている。

【００２０】更に、本実施の形態におけるコンテンツソース４は、クロック発振器４６とカウンタ４７を備えている。クロック発振器４６の出力はカウンタ４７に入力されている。クロック発振器４６とカウンタ４７はそれぞれシステムバス４４に接続されている。また、コンテンツソース４は、ハードディスク装置４８、ビットストリーム解析器４９、バッファ５０を備えている。ハードディスク装置４８は、自身に格納されているコンテンツストリームをビットストリーム解析器４９に入力している。ビットストリーム解析器４９の出力はバッファ５０に入力され、その出力はシステムバス４４に供給されている。また、ハードディスク装置４８およびビットストリーム解析器４９は、それぞれシステムバス４４に接続されている。

【００２１】尚、コンテンツソース４は、デコーダを持たないためデコードディレイという値には意味は無いが、デコードディレイとして０を持っている。このデコードディレイ値は、ＲＯＭ４２に記録され、ＣＰＵ４１から読み出すことができるようになっている。また、コンテンツソース４は、自身の種別を知っており、その種別は「コンテンツソース」である。この種別はＲＯＭ４２に記録され、ＣＰＵ４１から読み出すことができるようになっている。本実施の形態における種別として、ディスプレイ６は「ディスプレイ型」、各スピーカ５は「モノラルスピーカ型」の種別を持つものとする。この他にも、例えば「ステレオスピーカ型」、「ディスプレイ＋ステレオスピーカ一体型」、「スーパーウーファー型」、「オーディオソース型」などの種別が考えられる。

【００２２】コンテンツソース４は、内部にクロックの再生機能を持ち、クロックソース２から供給される基準時刻信号からクロックを再生している。また、コントローラ３の制御を受けて、内部に蓄えているコンテンツストリームを供給している。このとき、コントローラ３から指定されるディレイ値と再生したクロックから、ストリームにタイムスタンプを付加する機能を備えている。

【００２３】図６は、図１に示したスピーカ５（５－１～５－５）の構成を説明するための図である。この図６に示す構成は、図１に示す５本のスピーカ５（５－１～５－５）の全てに共通しており、これらは同一の内部構造（および動作）を備えている。スピーカ５は、ＣＰＵ５１、ＲＯＭ５２、およびＲＡＭ５３を備えており、これらはシステムバス５４により相互に接続されマイクロコンピュータを構成している。また、スピーカ５は、システムバス５４に接続されているネットワークインタフェース５５を介して、ネットワーク１と接続されている。

【００２４】このスピーカ５は、システムバス５４に接続されるクロック発振器５６とカウンタ５７とを備えており、このクロック発振器５６の出力はカウンタ５７に入力されている。更に、スピーカ５は、バッファ５９、タイムスタンプ抽出器６０、デコーダ６１、アンプ６２、およびスピーカ６３を備えている。スピーカ５を流れるストリームデータは、システムバス５４からバッファ５９、タイムスタンプ抽出器６０を経てデコーダ６１に入力されてデコードされる。デコーダ６１の出力は、アンプ６２を経てスピーカ６３により、音声信号に変換されて出力される。タイムスタンプ抽出器６０、デコーダ６１、およびアンプ６２は、各々システムバス５４に接続されている。このデコーダ６１の内部は、後述するＭＰＥＧデコーダの構成をなしている。具体的にはデマルチプレクサ、オーディオバッファ、オーディオデコーダで構成されている。

【００２５】スピーカ５は、更に、ボタン/ＬＥＤ操作器５８を備えている。ＣＰＵ５１は、システムバス５４経由で、ボタン/ＬＥＤ操作器５８のボタンが押されたか否かの情報を読み出すことができる。また、ＣＰＵ５１は、システムバス５４経由でボタン/ＬＥＤ操作器５８におけるＬＥＤの点滅を制御することができる。

【００２６】ここで、スピーカ５は、自身のデコーダ６１におけるデコードディレイを知っている。ここでデコードディレイとは、デコーダ６１にデータが入力されてからデコードされたデータが出力されるまでの時間であり、各種のディレイである変換ディレイの中で代表的なものである。このデコードディレイ値はＲＯＭ５２に記録され、ＣＰＵ５１から読み出すことができるようになっている。また、スピーカ５は自身の種別を知っており、このスピーカ５の種別は「モノラルスピーカ」である。この種別はＲＯＭ５２に記録され、ＣＰＵ５１から読み出すことができるようになっている。

【００２７】スピーカ５にて実行される動作は大きく３つに分けられる。この３つの動作とは、（１）クロックの再生、（２）ユーザインタフェース対応、（３）ストリームの再生、である。これらの動作は、ＣＰＵ５１上で動作するタスクと、それぞれ必要な構成要素とが組み合わされて成り立っている。なお、ＣＰＵ５１上で動作するこれらの３つのタスクは、ＣＰＵ５１で動作するマルチタスクオペレーティングシステムの上で独立に動作してい

るものとする。

【００２８】上述した第１の動作である、(１)クロックの再生、では、クロックソース２から発信された基準時刻信号により、スピーカ５の内部にクロックが再生される。第２の動作である、(２)ユーザインタフェース対応、では、例えばコントローラ３の指示により内部のＬＥＤを発光させ、内部のボタンが押されたか否かの情報をコントローラ３に戻す動作がなされる。第３の動作である、(３)ストリームの再生、では、受け取ったコンテンツデータに必要なディレイが施され、デコードして出力される。

【００２９】図７は、図１に示したディスプレイ６の構成を説明するための図である。このディスプレイ６の内部構成は、図６で説明したスピーカ５の内部構成に、ＯＳＤ(On Screen Display)８２、表示装置８３を加えたものとして把握することができる。即ち、図７に示すディスプレイ６は、ＣＰＵ７１、ＲＯＭ７２、およびＲＡＭ７３を備えており、これらはシステムバス７４により相互に接続されマイクロコンピュータを構成している。また、ディスプレイ６は、システムバス７４に接続されているネットワークインタフェース７５を介して、ネットワーク１と接続されている。

【００３０】このディスプレイ６は、それぞれシステムバス７４に接続されているクロック発振器７６とカウンタ７７を備えており、クロック発振器７６の出力は、カウンタ７７に入力されている。更に、ディスプレイ６はバッファ７９、タイムスタンプ抽出器８０、デコーダ８１、ＯＳＤ８２、および表示装置８３を備えている。ディスプレイ６を流れるストリームデータは、システムバス７４からバッファ７９、タイムスタンプ抽出器８０を経てデコーダ８１に入力されてデコードされる。デコーダ８１からの出力は、ＯＳＤ８２の出力と混合された後に、表示装置８３に表示される。このタイムスタンプ抽出器８０、デコーダ８１およびＯＳＤ８２は、それぞれシステムバス７４に接続されている。尚、デコーダ８１の内部は、後述するＭＰＥＧデコーダの構成をなしている。具体的にはデマルチプレクサ、ビデオバッファ、ビデオデコーダで構成されている。

【００３１】また、ディスプレイ６は、ボタン/ＬＥＤ操作器７８を備えている。ＣＰＵ７１は、システムバス７４経由でボタンが押されたか否かの情報を読み出すことができる。また、ＣＰＵ７１は、システムバス７４経由でＬＥＤの点滅を制御することができる。ディスプレイ６は自身のデコーダ８１におけるデコードディレイを知っている。ここでデコードディレイとは、デコーダ８１にデータが入力されてから、デコードされたデータが出力されるまでの時間である。このデコードディレイ値はＲＯＭ７２に記録され、ＣＰＵ７１から読み出すことができるようになっている。また、ディスプレイ６は、自身の種別を知っている。ディスプレイ６の種別は「ディスプレイ」である。この種別はＲＯＭ７２に記録され、ＣＰＵ７１から読み出すことができるように構成されている。

【００３２】ディスプレイ６の行う動作は大きく３つに分けられる。この３つの動作とは、スピーカ５の動作と同様であり、(１)クロックの再生、(２)ユーザインタフェース対応、(３)ストリームの再生、である。これらの動作は、ＣＰＵ７１上で動作するタスクと、それぞれ必要な構成要素とが組み合わされて成り立っている。なお、ＣＰＵ７１上で動作するこれらの３つのタスクは、ＣＰＵ７１で動作するマルチタスクオペレーティングシステムの上で独立に動作しているとする。

【００３３】この(１)クロックの再生、とは、クロックソース２から発信された基準時刻信号により、ディスプレイ６の内部にクロックを再生することである。(２)ユーザインタフェース対応、とは、例えばコントローラ３の指示により内部のＬＥＤを発光させ、内部のボタンが押されたか否かの情報をコントローラ３に戻すことである。(３)ストリームの再生、とは、受け取ったコンテンツデータに必要なディレイを施した後、デコードして出力することである。

【００３４】次に、本実施の形態が適用されるネットワークシステムの全体の動作について説明する。図１に示すクロックソース２は、ネットワーク１で相互に接続された機器に対して、クロックを送信している。より具体的には、一定時間ごとの時刻情報(ここでは基準時刻信号'Ｔｓ')を送信している。但し、この時刻情報の送信は、必ずしも決まった時間間隔である必要はなく、所定時間ごとに送信されるものとすることができる。クロックの受信側(スピーカ５(５－１～５－５)、ディスプレイ６、コンテンツソース４)では、受け取った基準時刻情報と内部発振回路の出力の値とが比較されて、内部発振回路の発振周波数が調整され、クロックが再現される。

【００３５】図１に示すコントローラ３は、システム構成を設定している。この動作は、コントローラ３の指示に従いディスプレイ６にて表示されるメッセージや、コントローラ３の指示に従い各スピーカ５が発光させるＬＥＤに反応して、ユーザが各スピーカ５の持つスイッチを押す事により達成される。また、コントローラ３は、ストリーム再生を制御している。この動作は、コントローラ３の指示に従い、コンテンツソース４がストリームを出力し、各スピーカ５およびディスプレイ６がストリームをデコードすることにより行われる。また、コントローラ３の指示により各スピーカ５の音声出力を調整することも行われる。

【００３６】次に、クロックソース２が行うクロック配信について説明する。ここでは、図２および図３に示した２種類の構成のそれぞれに対して説明を行う。まず、図２に示すクロックソース２の動作を説明する。ここで

クロック発振器１１の発振周波数を'Ｕ'［Ｈｚ］とする。クロック発振器１１で発信されたクロックは、カウンタ１２とタイマ１３のそれぞれに入力される。カウンタ１２は、入力されたクロックにより自身のカウンタ１２をインクリメントし、時刻情報'Ｔ'を生成する。入力されるクロックが'Ｕ'［Ｈｚ］であるため、時刻情報'Ｔ'は'Ｕ'［Ｈｚ］の周波数でインクリメントされる。カウンタ１２の出力(時刻情報'Ｔ')は、ラッチ１４に入力される。ラッチ１４は、ラッチ信号が入力された時に入力信号をラッチして出力する。ラッチ信号が入力されない場合には、直前にラッチした信号が保持される。

【００３７】タイマ１３は、入力されたクロックを計数し、あらかじめ設定された設定値'Ｓ'ごとにトリガ信号を出力している。ここで、タイマ１３への入力クロックは'Ｕ'［Ｈｚ］であるので、トリガ信号は'Ｓ／Ｕ'［秒］ごとに出力される。言い換えると、トリガ信号の周期は'Ｓ／Ｕ'［秒］、あるいはトリガ信号の周波数は'Ｕ／Ｓ'［Ｈｚ］となる。タイマ１３から出力されるトリガ信号は、ラッチ１４のラッチ信号端子に入力されると同時に、ネットワークインタフェース１５の送信信号端子に入力される。

【００３８】ラッチ１４は、ラッチ信号端子にタイマ１３からのトリガ信号が入力された時点で、カウンタ１２から入力されている時刻信号'Ｔ'をホールドして出力する。ラッチ１４の出力は、次のラッチ信号が入力されるまで変化しない。この値を基準時刻信号'Ｔｓ'とする。送信信号端子にトリガ信号を受け取ったネットワークインタフェース１５は、入力されているラッチ１４の出力信号'Ｔｓ'を読み取り、ブロードキャストパケットに格納してネットワーク１に送り出す。ブロードキャストパケットは宛先を指定しないために、ネットワーク１に接続されている各機器に到着する。

【００３９】次に、図３に示すクロックソース２の動作を説明する。ここでクロック発振器２１の発振周波数を'Ｕ'［Ｈｚ］とする。クロック発振器２１で発信されたクロックは、カウンタ２２とタイマ２３のそれぞれに入力される。カウンタ２２は、入力されたクロックにより自身のカウンタをインクリメントして時刻情報'Ｔ'を生成する。入力されるクロックが'Ｕ'［Ｈｚ］であるため、時刻情報'Ｔ'は、'Ｕ'［Ｈｚ］の周波数でインクリメントされる。カウンタ２２の出力(時刻情報'Ｔ')はラッチ２４に入力される。ラッチ２４は、ラッチ信号が入力された時に入力信号をラッチして出力する。ラッチ信号が入力されない場合には、直前にラッチした信号が保持される。

【００４０】タイマ２３は、入力されたクロックを計数し、あらかじめ設定された設定値'Ｓ'ごとにトリガ信号を出力している。ここでタイマ２３への入力クロックは'Ｕ'［Ｈｚ］であるので、トリガ信号は'Ｓ／Ｕ'［秒］ごとに出力される。言い換えると、トリガ信号の周期は'Ｓ／Ｕ'［秒］、あるいはトリガ信号の周波数は'Ｕ／Ｓ'［Ｈｚ］となる。タイマ２３から出力されるトリガ信号は、ラッチ２４のラッチ信号端子に入力されると同時にＣＰＵ２５の割り込み端子に入力される。ラッチ２４は、ラッチ信号端子にタイマ２３からのトリガ信号が入力された時点で、カウンタ２２から入力されている時刻信号'Ｔ'をホールドして出力する。ラッチ２４の出力は、次のラッチ信号が入力されるまで変化しない。この値を基準時刻信号'Ｔｓ'とする。

【００４１】割り込み端子にトリガ信号を受け取ったＣＰＵ２５は、システムバス２０を介して、ラッチ２４の値、即ち基準時刻信号'Ｔｓ'を読み出す。さらにＣＰＵ２５は、システムバス２０経由でネットワークインタフェース２６に指令し、ネットワーク１に接続されている各機器に対して、ラッチ２４から読み出した基準時刻信号'Ｔｓ'を送信させる。ネットワークインタフェース２６は、ＣＰＵ２５から受け取った指令に従い、基準時刻信号'Ｔｓ'をブロードキャストパケットに格納して、ネットワーク１に送り出す。ブロードキャストパケットは宛先を指定しないために、ネットワーク１に接続されている各機器に到着する。尚、ＣＰＵ２５は、ＲＯＭ２７に蓄えられているプログラムによりＲＡＭ２８を利用して動作している。

【００４２】図８は、クロック配信におけるクロックソース２の動作を説明するための図である。ここでは、横軸に時刻をとり、縦軸はそれぞれの信号の値が割り当てられている。クロック発振器１１(あるいはクロック発振器２１)は、'Ｕ'［Ｈｚ］の周波数で発信している。カウンタ１２(あるいはカウンタ２２)の出力、即ち時刻情報'Ｔ'は、同じく'Ｕ'［Ｈｚ］の周波数でインクリメントされる。図８の横軸の原点'ｔ０'をタイマ１３(あるいはタイマ２３)のトリガ信号が出力された時刻にとり、その時点での時刻情報'Ｔ'を'Ｔ０'とする。

【００４３】次のトリガ信号は、時刻'ｔ０'からクロック'Ｕ'の周波数で周期'Ｓ'だけ計数した時刻'ｔ１'に出力され、この時の時刻情報'Ｔ'は'Ｔ０＋Ｓ'となる。このとき、同時に時刻情報'Ｔ'はラッチ１４(あるいはラッチ２４)によりラッチされ、基準時刻信号'Ｔｓ'は'Ｔ０＋Ｓ'となる。その後、トリガ信号は'Ｕ'の周波数でクロック周期'Ｓ'ごとに出力され、基準時刻情報'Ｔｓ'はそれぞれ'Ｔ０＋２Ｓ'、'Ｔ０＋３Ｓ'と変化して行く。基準時刻信号'Ｔｓ'は、トリガ信号が発生する度にネットワーク１に対してブロードキャストされる。

【００４４】次に、クロック再生の動作について、スピーカ５(５－１～５－５)を例に挙げて説明する。このクロック再生は、コンテンツソース４およびディスプレイ６においても実行される。ここで「クロックの再生」とは、クロックソース２から送られたクロックをクライアント機器内部で再び発生させることを示している。

【００４５】通常状態において、図６に示すクロック発

振器５６の出力はカウンタ５７に入力され、カウンタ５７は、入力されたクロックにより自身のカウンタをインクリメントして時刻情報'Ｔｔ'を生成する。ＣＰＵ５１は、バス５４を経由してカウンタ５７の値(時刻情報'Ｔｔ')を読み込むことができる。時刻情報'Ｔｔ'は、ＣＰＵ５１により読み出され、デコードのタイミング調整に使用される。クロック発振器５６は可変周波数発振器であり、ＣＰＵ５１がシステムバス５４を経由して指示を行うことにより、発振周波数を一定の範囲で変化させることができる。また、ＣＰＵ５１は、システムバス５４を経由して、カウンタ５７に所望の値を設定することができる。

【００４６】クロックソース２からの基準時刻信号'Ｔｓ'は、例えば一定時刻ごとに受信され、その際、ＣＰＵ５１は、以下で説明する時計合わせのタスクを実行する。カウンタ５７は、クロック発振器５６の自走クロックで時刻情報'Ｔｔ'を生成する。クロックソース２の発信した基準時刻信号'Ｔｓ'は、スピーカ５のネットワークインタフェース５５で受信される。ネットワークインタフェース５５は、ＣＰＵ５１に基準時刻信号'Ｔｓ'の到着を伝える。ＣＰＵ５１では、この信号により時計合わせのタスクが起動される。

【００４７】図９は、クロック再生における時刻合わせの処理を説明するためのフローチャートである。ＣＰＵ５１は、まず、ネットワークインタフェース５５に到着した基準時刻信号'Ｔｓ'を読み出す(ステップ１０１)。その後、カウンタ５７から時刻情報'Ｔｔ'を読み出す(ステップ１０２)。次に、基準時刻信号'Ｔｓ'と時刻信号'Ｔｔ'の差分を計算し('Ｔｓ'－'Ｔｔ')、変数 'ｄｉｆｆ'に代入する(ステップ１０３)。ＣＰＵ５１は、'ｄｉｆｆ'の絶対値と定数ｋを比較し(ステップ１０４)、'ｄｉｆｆ'の絶対値の方が大きいと判断された場合には、'Ｔｓ'の値をカウンタ５７に代入して(ステップ１０５)、処理を終了する。ステップ１０４にて'ｄｉｆｆ'の絶対値の方が大きいと判断されなかった場合には、ステップ１０６に進む。

【００４８】ステップ１０６にて、ＣＰＵ５１は、'ｄｉｆｆ'の値が０(ゼロ)であるか否かを確認し、'ｄｉｆｆ'の値が０でない場合にはステップ１０７へ進み、'ｄｉｆｆ'の値が０である場合には処理を終了する。ステップ１０７にて、ＣＰＵ５１は、'ｄｉｆｆ'の値と０(ゼロ)とを比較し、'ｄｉｆｆ'の値が０(ゼロ)より大きいと判断された場合('Ｔｓ'＞'Ｔｔ')にはステップ１０９に進み、そうでない('Ｔｓ'＜'Ｔｔ')と判断された場合にはステップ１０８へ進む。ステップ１０８において、ＣＰＵ５１は、クロック発振器５６に対して発振周波数を遅くするよう指令した後、処理を終了する。また、ステップ１０９においてＣＰＵ５１は、クロック発振器５６に対して発振周波数を速くするよう指令した後、処理を終了する。

【００４９】以上のようなＣＰＵ５１のタスクにより、受信した基準時刻信号'Ｔｓ'と内部時刻情報'Ｔｔ'の差が大きい場合には受信した基準時刻信号'Ｔｓ'を内部時刻カウンタに代入し、内部時刻情報'Ｔｔ'が少し進んでいる時にはクロックを遅くし、また内部時刻情報'Ｔｔ'が少し遅れている時にはクロックを速める。これらの動作により、基準時刻信号'Ｔｓ'にあわせた内部時刻情報'Ｔｔ'を再生することができる。尚、上記の例では、クロック配信に際してブロードキャスト通信を使用したが、ユニキャスト(一対一通信)の通信を用いることも可能である。

【００５０】次に、コントローラ３を中心に行われるシステム構成設定の動作について説明する。図１０および図１１は、ネットワーク１に各機器が接続された直後からの処理を示したフローチャートである。図１０に示すように、まず、コントローラ３は、物理的にネットワーク１に接続されて電源が投入されると、ネットワーク１への接続を確認する。即ち、図４に示すコントローラ３のＣＰＵ３１は、まず、ネットワークインタフェース３５に対してネットワーク１に接続したことを確認させる(ステップ２０１)。そして、ネットワーク１に接続したと判断されるとステップ２０３に進み、ネットワーク１に接続したと判断されない場合にはステップＳ２０１に戻って確認を繰り返す(ステップ２０２)。

【００５１】ステップ２０３において、ＣＰＵ３１は、「応答リクエストメッセージ」を作成し、ネットワークインタフェース３５に指令して、ネットワーク１に対して送信させる。この「応答リクエストメッセージ」は所定の文字列で構成されている。ネットワークインタフェース３５は、ＣＰＵ３１から受け取った指令に従い、「応答リクエストメッセージ」をブロードキャストパケットに格納してネットワーク１に送り出す。ブロードキャストパケットは、宛先を指定していないために、「応答リクエストメッセージ」はネットワーク１に接続されている各機器に到着する。「応答リクエストメッセージ」を受け取った各機器は、自身の種別およびデコードディレイ値を返答する。種別およびデコードディレイ値は各機器のＲＯＭに記録されている。例えばスピーカ５において、このメッセージはネットワークインタフェース５５に到着し、ＣＰＵ５１は、このメッセージを解読して、ＲＯＭ５２から種別とデコードディレイを読み取り、「応答」メッセージを作成し、ネットワークインタフェース５５に指令してコントローラ３に対して送信させる。各機器の返答したメッセージは、ネットワークインタフェース３５で受信され、ネットワークインタフェース３５は、ＣＰＵ３１にメッセージの到着を伝える。

【００５２】ステップ２０４においてメッセージが到着していると判断される場合には、コントローラ３のＣＰＵ３１は、ステップ２０５に進む。メッセージが到着していないと判断された時には、ステップ２０８に進む。

ステップ２０５において、ＣＰＵ３１は、ネットワークインタフェース３５に到着したメッセージを読み出す。その後、ＣＰＵ３１は、読み出したメッセージが各機器から返答された所定の「応答」であるかどうかを確認する（ステップ２０６）。所定の「応答」であった場合にはステップ２０７へ進む。所定の「応答」ではなかった場合には、ステップ２０８へ進む。

【００５３】ステップ２０７において、ＣＰＵ３１は、受け取った「応答」をＲＡＭ３３に置かれた機器データベースに記録する。図１２は、機器データベースの構造を示した図である。機器データベースの先頭には、登録されているエントリの数が記録され、引き続いて、エントリの数だけ各機器のレコードが記録されている。機器のレコードには、その機器をネットワーク上で特定するためのネットワークアドレス、種別、役割、デコードディレイがそれぞれ記録される。「応答」を機器データベースに記録する際には、まず、エントリ数を１だけインクリメントする。そして、新しいレコードをリストの終りに追加する。その後、新しく追加したレコードに対して、受け取った「応答」の送り先の機器を特定するためのネットワークアドレス、種別、デコードディレイを記録する。役割の項はシステム構成設定動作の後半で記録される。

【００５４】ステップ２０８において、ＣＰＵ３１では、ステップ２０３の「応答リクエストメッセージ」を送信してから、所定の時間が経過したかどうかが判断される。所定時間が経過したと判断された場合には、図１１に示すステップ２１０に進む。まだ所定時間が経過していないと判断された場合には、ステップ２０４に戻って、これまでの動作を繰り返す。

【００５５】図１１に示すステップ２１０において、ＲＡＭ３３に置かれている機器データベースの内容は、図１３に示すような内容になっており、そのエントリの数は7となっている。7つあるレコードの内訳としては、まず、ディスプレイ６を示すネットワークアドレスを持ち、種別がディスプレイとなっているレコードが置かれる。また、5つのスピーカ５（５－１～５－５）の各々を示すネットワークアドレスを持ち、種別がモノラルスピーカとなっているレコードが計5つ置かれている。最後に、コンテンツソース４を示すネットワークアドレスを持ち、種別がコンテンツソースとなっているレコードが置かれている。尚、レコードの順番は一例として示したものであり、図１３に示す順番になるものと決まっているものではない。

【００５６】ステップ２１０において、ＣＰＵ３１は、ＲＡＭ３３に置かれている機器データベースから種別がディスプレイとなっているものを探す。この例では、図１３に示すように、最初に記録されているレコードの種別がディスプレイとなっている。ＣＰＵ３１は、最初のレコードに記録されているネットワークアドレスを持つ機器がディスプレイであることを認識する。その後、ＣＰＵ３１は、「システムを設定開始する」旨のメッセージを作成し、ネットワーク１経由でディスプレイ６に指令して表示させる（ステップ２２０）。前のステップにてディスプレイ手段を持つ機器としてディスプレイ６を認識したので、ＣＰＵ３１は、直接、ディスプレイ６（を示すネットワークアドレスを持つ機器）に対して指令を行う。具体的には、ＣＰＵ３１は、メッセージを作成し、ネットワークインタフェース３５に指令してディスプレイ６に対して送信させる。ディスプレイ６では、図７に示すネットワークインタフェース７５がメッセージを受信し、ＣＰＵ７１によってそのメッセージが読み取られて解読され、ＯＳＤ８２に指令してメッセージを出力して表示装置８３に表示させる。

【００５７】次に、コントローラ３のＣＰＵ３１により「フロントレフト」が指定される（ステップ２３０）。即ち、「フロントレフトの位置に有るスピーカのスイッチを押して下さい」という旨のメッセージが作成され、ネットワーク１を経由してディスプレイ６に表示される。

【００５８】その後、コントローラ３は、「ＬＥＤ点滅／ボタン待ち」を指令するメッセージを作成し、ＲＡＭ３３に置かれている機器データベースから種別がモノラルスピーカになっている機器の全て（スピーカ５（５－１～５－５））に対して送出する。この「ＬＥＤ点滅／ボタン待ち」のメッセージを受信した機器は、ＬＥＤを点滅させ（ステップ２３１）、ボタンが押されることを待ち受ける。スピーカ５においては、このメッセージはネットワークインタフェース５５に到着し、ＣＰＵ５１はこのメッセージを解読して、ボタン/ＬＥＤ操作器５８に指令してＬＥＤを点滅させる。またＣＰＵ５１は、ボタン/ＬＥＤ操作器５８にてボタンが押されたことを探知したら、コントローラ３に返答を返すべく、待ち受ける。

【００５９】ユーザは、ＬＥＤの点滅しているスピーカ５（５－１～５－５）の中から、フロントレフトの位置にあたるスピーカ５－１のスイッチを押す。スピーカ５－１では、ボタンを押されたことをボタン/ＬＥＤ操作器５８が感知し、ＣＰＵ５１に伝える。ＣＰＵ５１は、「ボタン押下」のメッセージを作成し、ネットワークインタフェース５５に指令してコントローラ３に対して送信させる。コントローラ３のＣＰＵ３１は、ステップ２３２にて「ボタン押下」のメッセージが受信されるまで待ち（ステップ２３２）、受信するとスピーカ５－１のＬＥＤを消灯する（ステップ２３３）。

【００６０】その後、ＣＰＵ３１は、ネットワークインタフェース３５からメッセージを読み出し、「ボタン押下」メッセージ発信元のネットワークアドレスを読み取り、ＲＡＭ３３に記録されている機器データベースのレコードの中で、ネットワークアドレスが一致するもの、即ち、図１３の例では第２番目のレコードを探し出し、

該当するレコードの役割のフィールドに「フロントレフト」の属性を書き込んで設定する(ステップ２３４)。最後にＣＰＵ３１は、役割のフィールドの内容、ここでは「フロントレフト」の属性を記述した「役割メッセージ」を作成し、ネットワーク１を経由して「ボタン押下」メッセージの発信元に送り返す(ステップ２３５)。「ボタン押下」メッセージの発信元(ここではスピーカ５－１)は、「役割メッセージ」を受信し、役割(ここでは「フロントレフト」の属性)を取り出してＲＡＭ５３に格納する。即ち、このメッセージは、ネットワークインタフェース５５に到着し、ＣＰＵ５１は、このメッセージを解読して、ＲＡＭ５３に記録する。以上のようなステップ２３０からステップ２３５の操作によって、「フロントレフト」の位置に設置されたスピーカ５はスピーカ５－１であるということがコントローラ３の機器データベースに記録される。また、スピーカ５－１は、自身の役割(ここではフロントレフト)を知り、自身のＲＡＭ５３に記録する。

【００６１】以下同様に、ステップ２４０からステップ２４５の操作により「フロントライト」、ステップ２５０からステップ２５５の操作により「センター」、ステップ２６０からステップ２６５の操作により「リアレフト」、ステップ２７０からステップ２７５の操作により「リアライト」のそれぞれの位置に設置されたスピーカ５を知ることができる。更に、各スピーカ５(５－２～５　５)は、自身の役割を知り自身のＲＡＭ５３に記録する。

【００６２】最後に、コントローラ３のＣＰＵ３１は「システム設定が終了した」という旨のメッセージを作成し、ネットワーク１経由でディスプレイ６に指令して表示させる(ステップ２８０)。以上の手順により、物理的な配置とネットワーク１上でのアドレスの対応をとることができる。尚、ここでは説明しなかったが、ディスプレイ６が複数、存在する場合には、ディスプレイ６が持つボタン/ＬＥＤ操作器７８、あるいはリモコン７と画面表示などを用いて、上記と同様の設定を行うことができる。

【００６３】次に、スピーカ５(５－１～５－５)のデコーダ６１やディスプレイ６のデコーダ８１にて構成されるＭＰＥＧデコーダについて説明する。図１４は、ＭＰＥＧ２システム(ISO13818‐1)に準拠したデコーダシステムの概略を示した図である。図１４に示すデコーダシステムでは、ストリーム入力９１からストリームが入力され、デマルチプレクサ９２にてビデオストリームとオーディオストリームとに振り分けられる。ビデオストリームは、ビデオバッファ９３に入力され、所定のディレイ時間を経過した後にビデオデコーダ９４に入力され、デコードされてビデオ出力９７から出力される。オーディオストリームは、オーディオバッファ９５に入力され、所定のディレイ時間を経過した後にオーディオデコーダ９６に入力され、デコードされてオーディオ出力９８から出力される。

【００６４】図１５は、ＭＥＰＧデコーダに供給されるデータフォーマットの一例を示した図である。このフォーマットは、ＭＰＥＧ２の多重化ビットストリーム(ProgramStream)として規定されているものである。図１５に示すように、多重化ビットストリームは、１つ以上のパック(PACK)で構成され、各パックは１つ以上のパケット(PACKET)で構成される。パックの先頭にはパックヘッダ(PACK HEADER)が配置され、このパックヘッダにはパックの開始点を示すパックスタートコード(PACK START CODE)、ＳＣＲ(System Clock Reference)およびＭＵＸ ＲＡＴＥが配置されている。ＳＣＲは、その最後のバイトがデマルチプレクサ９２に入力される時刻を表わしている。ＭＵＸ ＲＡＴＥは、転送レートを示すものである。

【００６５】図１５に示す例においては、パックヘッダに続いて、ビデオパケット(VIDEO PACKET)とオーディオパケット(AUDIO PACKET)とが配置されている。これらのパケットの先頭にもパケットヘッダが配置され、これらのパケットヘッダは、ビデオまたはオーディオのパケットの開始点を表すビデオパケットスタートコード(VIDEO PACKET START CODE)またはオーディオパケットスタートコード(AUDIO PACKETSTART CODE)と、ビデオまたはオーディオデータのデコード(表示)開始時刻を示すＤＴＳ(V)(ＰＴＳＶ(ＰＴＳｖｉｄｅｏ))またはＤＴＳ(Ａ)(ＰＴＳＡ(ＰＴＳａｕｄｉｏ))が配置されている。そして、この各パケットヘッダの次に、ビデオデータあるいはオーディオデータが配置されている。尚、これらＳＣＲ、ＰＴＳ(ＰＴＳＶまたはＰＴＳＡ)などのタイミングデータは、９０ｋＨｚの周波数からなるクロックのカウント値で表わされ、３３ビットの有効数字を有している。更に、システム表現のための簡略モデル化を行っているため、デコードにかかる時間は０となっている。このためデコード開始時刻と表示開始時刻が等しい。

【００６６】図１６は、ＳＣＲとＰＴＳＶ(ＰＴＳｖｉｄｃｏ)、ＳＣＲとＰＴＳＡ(ＰＴＳａｕｄｉｏ)の関係を示した図である。パックヘッダがデマルチプレクサ９２を通過した時刻がｔ１であり、該当するパックに含まれるビデオとオーディオの表示時間がそれぞれｔ２とｔ３となる。ここでは、ＳＣＲからＰＴＳａｕｄｉｏまでの時間をΔＴａとし、ＳＣＲからＰＴＳｖｉｄｅｏまでの時間をΔＴｖとしている。ΔＴａとΔＴｖは、エンコード時に決まる任意の時間である。該当するパックがデマルチプレクサ９２に入力される時刻がＳＣＲであり、該当するパックに含まれるビデオデータとオーディオデータは、それぞれビデオバッファ９３とオーディオバッファ９５で所定の時間、遅延するために、そのデコード(表示)開始時刻はＳＣＲより遅くなる。また、一般的にビデオバッファ９３での遅延はオーディオバッファ９５

での遅延より大きいため、同じタイミングでデマルチプレクサ９２を通過した場合でも、ビデオのデコード(表示)時刻はオーディオのデコード(表示)開始時刻より遅くなる。尚、前述したように、デコードにかかる時間は０であるので、ここでの遅延は、それぞれビデオバッファ９３とオーディオバッファ９５で発生している。

【００６７】次に、本実施の形態におけるディレイ補償の考え方について説明する。コントローラ３は、コンテンツの再生を行うにあたり、Δｔを設定してコンテンツソース４に知らせる。コンテンツソース４は、コンテンツデータの送信時刻に対して、Δｔだけ遅らせた時刻にコンテンツの出力されるべき時刻を設定してコンテンツデータを送り出す。コンテンツのシンク機器群(スピーカ５(５－１～５－５)、ディスプレイ６)では、受け取ったコンテンツデータを指定された時刻に出力する。

【００６８】Δｔは、コンテンツソース４からコンテンツのシンク機器群(スピーカ５(５－１～５－５)、ディスプレイ６)に、コンテンツデータが送られた時のネットワークディレイ(通信ディレイ)のうち最大となる時間と、コンテンツのシンク機器群(スピーカ５(５－１～５－５)、ディスプレイ６)の実際のデコードディレイのうち最大となる時間とを加えたものである。場合によっては、マージンを増やすために幾分かの余裕を加える場合がある。また、この２つ以外の要因によるディレイを加える場合もある。例えば、サブネットをまたぐ場合のルーターのディレイ等が考えられる。

【００６９】更に、Δｔの値はシンク機器群の組み合わせによって変化することがある。即ち、コンテンツのシンク機器群の組み合わせが変化すると、その中での最大のネットワークディレイと最大のデコードディレイは変化する可能性がある。シンク機器群の組み合わせによってΔｔをその都度変えて制御する場合と、システム内で起こり得る最大のネットワークディレイと最大のデコードディレイとを用いて、即ち、Δｔの最大値を見積もってそれを使用する場合の二つの方法が考えられる。

【００７０】図１７は、本実施の形態におけるディレイを説明するための図である。コンテンツソース４がストリーム先頭に対する発信を開始した時刻をｔ１０とする。ｔ１０に最大の通信ディレイ(ネットワークディレイ)を加えた時刻をｔ１１とする。コンテンツソース４から発信されたコンテンツデータは、遅くともｔ１１までには全てのコンテンツのシンク機器群(スピーカ５(５－１～５－５)、ディスプレイ６)に到着している。即ち、通信のディレイに関しては、配信時刻から再生時刻を最大の通信ディレイだけ遅らせることで補償可能である。

【００７１】次に、ｔ１１に最大のデコードディレイを加えた時刻、即ち、ｔ１０にΔｔを加えた時刻をｔ１２とする。ＭＰＥＧ系の多重化ストリームの場合には、デコード開始とはビットストリームの処理開始、即ちデマルチプレクス開始ということになる。このため、時刻ｔ１２は、該当するビットストリームのＳＣＲが対応することになる。図１６に示した数値を当てはめて考えると、デコーダではｔ１２がＳＣＲの時刻となるから、オーディオの出力開始時刻ＰＴＳａｕｄｉｏは、ｔ１２にΔＴａを加えた時刻ｌ２０となり、ビデオの出力開始時刻ＰＴＳｖｉｄｅｏは、ｔ１２にΔＴｖを加えた時刻ｔ２２となる。

【００７２】ここで、実際のデコードディレイを考える。ディスプレイ６のデコーダ８１のデコードディレイをＤｖ、スピーカ５のデコーダ６１のデコードディレイをＤａとすると、ｔ２０にオーディオを出力するためには、ｌ２０よりＤａだけさかのぼった時刻ｌ２１にオーディオのデコードを開始する必要がある。また、同様に、ｔ２２にビデオを出力するためには、ｔ２２からＤｖだけさかのぼった時刻ｔ２３にビデオのデコードを開始する必要がある。

【００７３】このとき、ΔＴａあるいはΔＴｖが非常に小さく与えられたとすると、ｔ２０あるいはｔ２２はｔ１２に限りなく近付くので、ｔ２１あるいはｔ２３はｔ１２からデコードディレイ分さかのぼった時刻になる。このため、ｔ１２から最大のデコードディレイ分をさかのぼった時刻が、ｔ１１の後になっていれば良い。繰り返すと、ｔ１１までデコード開始を待てばネットワークによるディレイは無視できる。更に、見かけ上のデコード開始時刻(ｔ１２、ＳＣＲ)を、ｔ１１から実際のデコードディレイにおける最大のものだけ遅らせることにより、デコードディレイの影響も無視できるようになる。このために、Δｔには最大のデコードディレイが加算されている。

【００７４】次に、ストリーム再生動作について説明する。図１にて、ユーザがリモコン７を使用して、再生モードの指定をコントローラ３に対して与えた場合を考える。ここではディスプレイ６とスピーカ５(５－１～５－５)を使用するサラウンド再生を指定したものとする。このとき、リモコン信号は、コントローラ３のリモコン受光部３６で受信される。リモコン受光部３６は、ＣＰＵ３１に対して再生モード指定のコマンドを伝える。ＣＰＵ３１は、ＲＡＭ３３に記録されている機器データベースを参照し、シンク機器(スピーカ５(５－１～５－５)、ディスプレイ６)から、デコードディレイの最大値を選び出し、更に、予め決まっているネットワークのディレイをＲＯＭ３２から読み出して加算し、Δｔとする。

【００７５】ここで、ユーザがリモコン７を使用して、コンテンツ再生の命令をコントローラ３に対して与えた場合を考える。このリモコン信号は、コントローラ３のリモコン受光部３６で受信される。リモコン受光部３６は、ＣＰＵ３１にコンテンツ再生のコマンドを伝える。ＣＰＵ３１は、ＲＡＭ３３に記録されている機器データ

ベースを参照して、コンテンツソースを探し出す。ここではコンテンツソース４を発見し、ネットワークアドレスを得る。コントローラ３は、コンテンツソース４に対して「コンテンツ再生開始」のメッセージを送出する。このとき、Δｔの値も同時に与える。また、通常では、ここでコンテンツの指定がなされる。

【００７６】即ち、ＣＰＵ３１は、「コンテンツ再生開始メッセージ」を作成し、その中でΔｔを指定する。そして、ＣＰＵ３１は、ネットワークインタフェース３５に指令して、「コンテンツ再生開始メッセージ」をコンテンツソース４に対して送信させる。この「コンテンツ再生開始メッセージ」は、所定の文字列で構成される。

【００７７】コントローラ３は、今回のシンク機器群（スピーカ５（５－１～５－５）、ディスプレイ６）に対して、コンテンツソース４からのコンテンツデータを再生するように指示する。即ち、ＣＰＵ３１は、「コンテンツソース指定メッセージ」を作成し、その中でコンテンツソース４のネットワークアドレスを指定する。ＣＰＵ３１は、ネットワークインタフェース３５に指令して、今回のシンク機器群（スピーカ５（５－１～５－５）、ディスプレイ６）に対して、一つづつ順番に「コンテンツソース指定メッセージ」を送信させる。この「コンテンツソース指定メッセージ」は、所定の文字列で構成される。

【００７８】同時に、コントローラ３は、今回のシンク機器群の中で、スピーカ５群（スピーカ５　１～５　５）に対してボリュームを設定する。即ち、ＣＰＵ３１は、「ボリューム設定メッセージ」を作成し、その中でボリュームの値を指定する。ＣＰＵ３１は、ネットワークインタフェース３５に指令して、シンク機器群の中の指定されたシンク機器であるスピーカ５（５－１～５－５）に対して、一つづつ順番に「ボリューム設定メッセージ」を送信させる。この「ボリューム設定メッセージ」は、所定の文字列で構成されている。

【００７９】次に、コンテンツソース４の動作について説明する。コンテンツソース４は、「コンテンツ再生開始メッセージ」を受信する。即ち、ネットワークインタフェース４５にメッセージが到着し、ＣＰＵ４１はこのメッセージを解読し、Δｔの値を受け取ると共に、コンテンツの再生を開始する。まず、ＣＰＵ４１は、ハードディスク装置４８に対して所定のコンテンツの出力を指示する。ハードディスク装置４８から出力したストリームは、ビットストリーム解析器４９にて解析され、先頭のＳＣＲの値が読み出される。ＳＣＲの数値はＣＰＵ４１によって読み出され、ストリームはそのままバッファ５０へ進む。

【００８０】このとき、コンテンツソース４での時刻情報'Ｔｔ'が'Ｔ１０'、ストリーム先頭のＳＣＲの値が'Ｓ１０'であったとする。即ち、ビットストリームの'Ｓ１０'の値を持つタイムスタンプが、スピーカ５におけるデコーダ６１内の図示せぬデマルチプレクサ、あるいはディスプレイ６のデコーダ８１内の図示せぬデマルチプレクサにて処理される時刻'Ｔ１２'は、現在時刻'Ｔ１０'にコントローラ３より与えられた'Δｔ'を加えた時刻となる。

Ｔ１２＝Ｔ１０＋Δｔ

コンテンツソース４のＣＰＵ４１は、「タイムスタンプオフセットメッセージ」を作成し、'Ｓ１０'の値と'Ｔ１２'の値を指定する。ＣＰＵ４１は、ネットワークインタフェース４５に指令して、「タイムスタンプオフセットメッセージ」をネットワーク１にブロードキャストする。これにより、全てのシンク機器は、ＭＰＥＧ２のタイムスタンプと'クロック'の関係を知ることになる。

【００８１】コンテンツソース４のＣＰＵ４１は、ネットワークインタフェース４５に指定して、バッファ５０の内容（ストリーム）をネットワーク１にブロードキャストさせる。ここではネットワーク１に対してブロードキャストするが、ユニキャスト（一対一通信）でも構わない。その際には、コントローラ３が、シンク機器の一覧をコンテンツソース４に伝える必要が有り、コンテンツソース４は、その一覧データに従ってデータを送信する。

【００８２】次に、本実施の形態におけるスピーカ５の動作について説明する。スピーカ５のＣＰＵ５１は、コントローラ３の発信した「コンテンツソース指定メッセージ」のメッセージを受け取る。ＣＰＵ５１は、コンテンツソースとして指定されたネットワークアドレスをＲＡＭ５３に格納する。スピーカ５は、コントローラ３の指示により、コンテンツソース４からのコンテンツデータを待ち受ける。このメカニズムにより、同一ネットワーク上に、コンテンツソースとシンク機器の組み合わせが複数存在することができる。

【００８３】また、スピーカ５のＣＰＵ５１は、コントローラ３の発信した「ボリューム設定メッセージ」のメッセージを受け取る。ＣＰＵ５１はアンプ６２に対してボリュームを設定する。更にスピーカ５は、コンテンツソース４の発信した「タイムスタンプオフセットメッセージ」を受け取る。この「タイムスタンプオフセットメッセージ」は、コンテンツデータに伴って送られるものであり、スピーカ５は、予めコンテンツソース４からのコンテンツデータを待ち受けていたので、このメッセージを受け入れた。この「タイムスタンプオフセットメッセージ」からＭＰＥＧのタイムスタンプと時刻情報の関係が判る。スピーカ５のＣＰＵ５１は、差分オフセット値＝Ｓ１０－Ｔ１２を計算する。これによりスピーカ５での時刻情報'Ｔｔ'にオフセット値を加えることで、ＭＰＥＧの時計であるところのＳＴＣ（システムタイムクロック）の値を計算することができる。また、ＭＰＥＧ系のタイムスタンプからオフセットを減じることで、時刻情報'Ｔｔ'の値を求めることができる。

【００８４】スピーカ５は、コンテンツソース４の発信したストリームを受け取り、バッファ５９に入力する。バッファ５９に入ったストリームは、タイムスタンプ抽出器６０を経てデコーダ６１に入力される。ＣＰＵ５１は、タイムスタンプ抽出器６０からＳＣＲとＰＴＳａｕｄｉｏを読み出す。また、ＣＰＵ５１はデコードディレイＤａの値を知っている。デコーダ６１にはデコードディレイがあるために、名目上の出力開始時刻からＤａだけさかのぼって(前に)デコードを開始する必要がある。この調整を行うために、スピーカ５のＣＰＵ５１は以下の処理を行う。

【００８５】まず、Ｔ１２からＤａだけさかのぼった時刻をＴ１１(実際のデマルチプレクス開始時刻)とする。ストリームの最初のＰＴＳａｕｄｉｏ(名目上のオーディオデコード/表示開始)をＴ２０とし、Ｔ２０からＤａだけさかのぼった時刻をＴ２１とする。即ちＴ１２は名目上のデマルチプレクス開始時刻であり、Ｔ１１はデコードディレイの分、前倒しにした実際のデマルチプレクス開始時刻となる。また、Ｔ２０は、ストリームの最初のＰＴＳａｕｄｉｏにおける名目上のデコード(表示)開始時刻であり、Ｔ２１はデコードディレイの分、前倒しにした実際のデコード開始時刻となる。ＣＰＵ５１はカウンタ５７から再生された時刻情報'Ｔｔ'を読み出し、時刻がＴ１１に等しくなった時点でデマルチプレクスを開始する。引き続いて時刻がＴ２１になった時点でデコードを開始する。これにより、Ｔ２０にオーディオの出力を開始することができる。

【００８６】次に、本実施の形態におけるディスプレイ６の動作について説明する。ディスプレイ６のＣＰＵ７１は、コントローラ３の発信した「コンテンツソース指定メッセージ」のメッセージを受け取る。ＣＰＵ７１は、コンテンツソースとして指定されたネットワークアドレスをＲＡＭ７３に格納する。ディスプレイ６は、コントローラ３の指示により、コンテンツソース４からのコンテンツデータを待ち受ける。この仕掛けにより、同一ネットワーク上に、コンテンツソースとシンク機器の組み合わせが複数存在することができる。

【００８７】ディスプレイ６は、コンテンツソース４の発信した「タイムスタンプオフセットメッセージ」を受け取る。この「タイムスタンプオフセットメッセージ」は、コンテンツデータに伴って送られるものであり、ディスプレイ６は、予めコンテンツソース４からのコンテンツデータを待ち受けていたのでこのメッセージを受け入れた。この「タイムスタンプオフセットメッセージ」からＭＰＥＧのタイムスタンプと時刻情報の関係が判る。ディスプレイ６のＣＰＵ７１は、差分オフセット値＝Ｓ１０－Ｔ１２を計算する。これによりディスプレイでの時刻情報'Ｔｔ'にオフセット値を加えることで、ＭＰＥＧの時計であるところのＳＴＣ(システムタイムクロック)の値を計算することができる。また、ＭＰＥＧ系のタイムスタンプからオフセットを減じることで、時刻情報'Ｔｔ'の値を求めることができる。

【００８８】ディスプレイ６は、コンテンツソース４の発信したストリームを受け取り、バッファ７９に入力する。バッファ７９に入ったストリームは、タイムスタンプ抽出器８０を経てデコーダ８１に入力する。ＣＰＵ７１は、タイムスタンプ抽出器８０からＳＣＲとＰＴＳｖｉｄｅｏを読み出す。また、ＣＰＵ７１はデコードディレイＤｖの値を知っている。デコーダ８１にはデコードディレイがあるために、名目上の出力開始時刻からＤｖだけさかのぼって(前に)デコードを開始する必要がある。この調整を行うために、ディスプレイ６のＣＰＵ７１は以下の処理を行う。

【００８９】まず、Ｔ１２からＤｖだけさかのぼった時刻をＴ１１(実際のデマルチプレクス開始時刻)とする。ストリームの最初のＰＴＳｖｉｄｅｏ(名目上のビデオデコード/表示開始)をＴ２２とし、Ｔ２２からＤｖだけさかのぼった時刻をＴ２３とする。即ちＴ１２は名目上のデマルチプレクス開始時刻であり、Ｔ１１はデコードディレイの分、前倒しにした実際のデマルチプレクス開始時刻となる。また、Ｔ２２はストリームの最初のＰＴＳｖｉｄｅｏにおける名目上のデコード(表示)開始時刻であり、Ｔ２３はデコードディレイの分、前倒しにした実際のデコード開始時刻となる。ＣＰＵ７１は、カウンタ７７から再生された時刻情報'Ｔｔ'を読み出し、時刻がＴ１１に等しくなった時点でデマルチプレクスを開始する。引き続いて時刻がＴ２３になった時点でデコードを開始する。これにより、Ｔ２２にビデオの出力を開始することができる。

【００９０】最後に、位置補償の方法について説明する。上記の例では、Δｔの要素として、ネットワークによるディレイとデコードディレイの二つを考えたが、このほかにもスピーカ５(５－１～５－５)の位置によるディレイ調整が考えられる。例えば、スピーカ５－４がスピーカ５－５より１．７メートルだけリスニングポジション８に近いとする。音速はおよそ３４０メートル毎秒なので、１．７メートルの距離を進むのに約５ミリ秒かかる。即ち、スピーカ５－４とスピーカ５－５を同時に発した音は、スピーカ５－４からの音が５ミリ秒だけ早くリスニングポジション８に到達する。これを補償するために、スピーカ５－４に対して５ミリ秒のディレイを加えるという方法がある。言い換えれば、スピーカ５－４だけが１．７メートルだけリスニングポジション８に近い場合に、スピーカ５－４以外のスピーカに対して５ミリ秒だけ早く出力させるということになる。

【００９１】本実施の形態では、かかる方法にも対応することが可能である。具体的には、上記例においてΔｔを計算する際に、「ネットワークディレイの最大値」と「デコードディレイの最大値」に加えて「位置補償ディレイの最大値」を加算する。即ち、上記例よりもΔｔは

５ミリ秒だけ大きくなる。また、コントローラ３は、スピーカ５－４に対して０ミリ秒のディレイを指示し、スピーカ５－４以外のスピーカに対して５ミリ秒の前倒しを指示する。このとき、コントローラ３から発信される「コンテンツソース指定メッセージ」によって指定されたネットワークアドレスに基づいて、シンク機器毎にディレイの指示を受けることが可能である。また、コントローラ３のＣＰＵ３１により各スピーカ５の「位相補償ディレイ設定メッセージ」が作成され、各スピーカ５に送られることで、シンク機器毎にディレイを指示するように構成することも可能である。

【００９２】図１８は、位置補償ディレイの最大値を考慮した場合のディレイを説明するための図である。コンテンツソース４は、コンテンツの送り出し時刻ｔ１０に対してΔｔを加え、名目上のデマルチプレクス開始時刻としてｔ１２を指定する。リアレフトであるスピーカ５－４では、最初のＳＣＲにあたるｔ１２から”デコードディレイＤａ＋０ミリ秒”だけさかのぼった時刻ｔ１１にデマルチプレクスを開始し、最初のＰＴＳａｕｄｉｏにあたるｔ２０から”デコードディレイＤａ＋０ミリ秒”だけさかのぼった時刻ｔ２１にデコードを開始する。これによって、ｔ２０に音声出力することができる。

【００９３】スピーカ５－４以外のスピーカ５では、最初のＳＣＲにあたるｔ１２から”デコードディレイＤａ＋５ミリ秒”だけさかのぼった時刻ｔ３１にデマルチプレクスを開始し、最初のＰＴＳａｕｄｉｏにあたるｔ２０から”デコードディレイＤａ＋５ミリ秒”だけさかのぼった時刻ｔ３２にデコードを開始する。するとｔ２０から５ミリ秒前に音声出力することができる。即ち、これらの操作によりスピーカ５－４以外のスピーカ５からの出力は、スピーカ５－４からの出力よりも５ミリ秒だけ早く出力することができる。このようにしてスピーカ位置の補償が可能となる。

【００９４】以上、詳述したように、本実施の形態では、まず、ソース機器とシンク機器における動作クロックの関係(結びつき)について対策を施している。即ち、両者のクロックが非同期である場合、両者のデータ処理速度、即ち同一サンプル数に対する処理時間が異なるため、結果として二つの機器の間に存在するバッファには破綻(アンダーフロー、オーバーフロー)をきたすことになる。本実施の形態では、クロックの配信を行い、受信した基準時刻信号と内部時刻情報との差が大きい場合には受信した基準時刻信号を内部時刻カウンタに代入し、内部時刻情報が進んでいるときにはクロックを遅くし、遅れているときにはクロックを速めるように構成した。これによって、ディジタルネットワークに接続された機器におけるクロックの同期をとることが可能となる。

【００９５】また、ネットワークに接続された複数の入出力機器間の位相(タイミング)の関係について、ディジタルネットワークでは伝送遅延が固定ではないことから、十分に小さい位相誤差にて出力されるという保証が無い。また、アナログ接続に比べて遅延の絶対値が大きくなる。しかしながら、本実施の形態によれば、まず、上記のように、同じ時計を各機器で動作させ、入出力タイミングをその時計の時刻で指定することで、位相情報(入出力時刻/タイミング)配信を行っている。また、ネットワーク(通信)ディレイと末端機器におけるデコードディレイを考慮し、コンテンツの再生に際して生じるディレイ分を補正するように構成した。これによって、理想状態で位相を合わせた上で実際の位相調整を行うことが可能となる。

【００９６】尚、本実施の形態にて説明したデコードディレイに代表される変換ディレイには、例えば、サンプリング周波数変換、ノイズフィルタ、また、サラウンド機能を通した場合のディレイ等の各種ディレイを含むことができる。また、本実施の形態では、ストリームを再生して出力するデコーダを備える出力機器を中心に説明したが、例えば、ディジタルカメラやマイク、リモコン操作等を行うスイッチ等、各種入力機器がディジタルネットワークに接続された態様に対しても適用することができる。かかる態様の場合には、エンコードしてストリームを生成する際に生じるエンコードディレイを考慮し、位相の同期が図られる。

【００９７】

【発明の効果】以上説明したように、本発明によれば、ディジタルネットワークに接続された各機器の間において、同期をとることが可能となる。

【図面の簡単な説明】

【図１】　本実施の形態におけるネットワークシステムの全体像を示した構成図である。

【図２】　クロックソースにおける第１の構成例を示した図である。

【図３】　クロックソースにおける第２の構成例を示した図である。

【図４】　図１に示したコントローラの構成を説明するための図である。

【図５】　図１に示したコンテンツソースの構成を説明するための図である。

【図６】　図１に示したスピーカの構成を説明するための図である。

【図７】　図１に示したディスプレイの構成を説明するための図である。

【図８】　クロック配信におけるクロックソースの動作を説明するための図である。

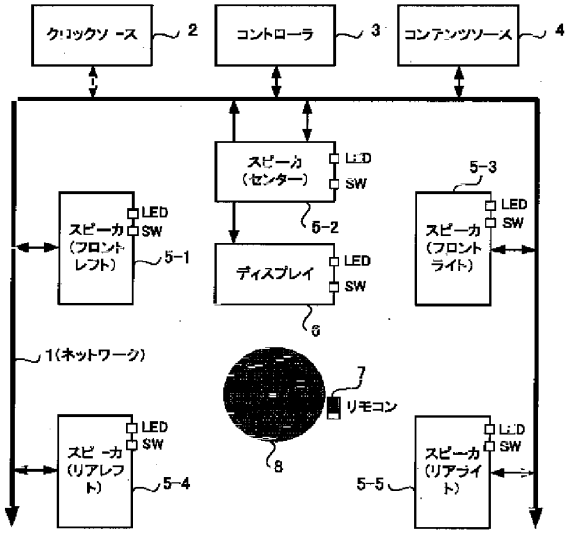
【図９】　クロック再生における時刻合わせの処理を説明するためのフローチャートである。

【図１０】　ネットワークに各機器が接続された直後からの処理を示したフローチャートである。

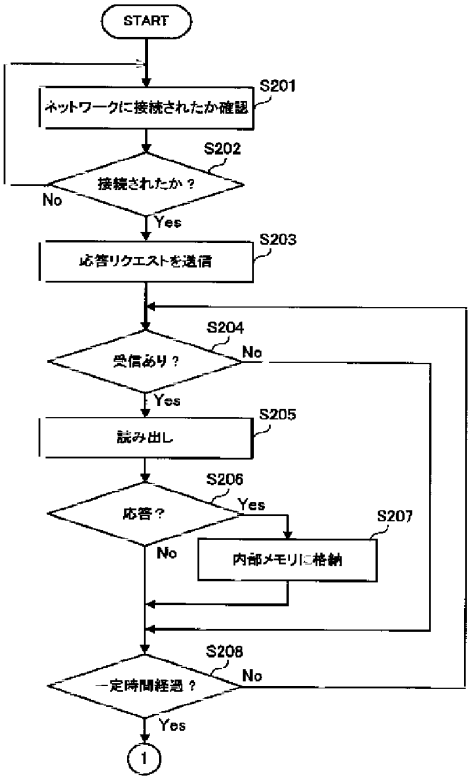
【図１１】　ネットワークに各機器が接続された直後からの処理を示したフローチャートである。

【図１２】　機器データベースの構造を示した図である。
【図１３】　コントローラのＲＡＭに置かれている機器データベースの内容を示した図である。
【図１４】　ＭＰＥＧ２システム(ISO13818－1)に準拠したデコーダシステムの概略を示した図である。
【図１５】　ＭＥＰＧデコーダに供給されるデータフォーマットの一例を示した図である。
【図１６】　ＳＣＲとＰＴＳＶ(ＰＴＳｖｉｄｅｏ)、ＳＣＲとＰＴＳＡ(ＰＴＳａｕｄｉｏ)の関係を示した図である。
【図１７】　本実施の形態におけるディレイを説明するための図である。
【図１８】　位置補償ディレイの最大値を考慮した場合のディレイを説明するための図である。
【符号の説明】
１…ネットワーク、２…クロックソース、３…コントローラ、４…コンテンツソース、５(５－１～５－５)…スピーカ、６…ディスプレイ、１１…クロック発振器、１２…カウンタ、１３…タイマ、１４…ラッチ、１５…ネットワークインタフェース、２０…システムバス、２１…クロック発振器、２２…カウンタ、２３…タイマ、２４…ラッチ、２５…ＣＰＵ、２６…ネットワークインタフェース、３１…ＣＰＵ、３２…ＲＯＭ、３３…ＲＡＭ、３４…システムバス、３５…ネットワークインタフェース、４１…ＣＰＵ、４４…システムバス、４５…ネットワークインタフェース、４６…クロック発振器、４７…カウンタ、４８…ハードディスク装置、４９…ビットストリーム解析器、５０…バッファ、５１…ＣＰＵ、５４…システムバス、５５…ネットワークインタフェース、５６…クロック発振器、５７…カウンタ、５９…バッファ、６０…タイムスタンプ抽出器、６１…デコーダ、７１…ＣＰＵ、７５…ネットワークインタフェース、７６…クロック発振器、７７…カウンタ、７９…バッファ、８０…タイムスタンプ抽出器、８１…デコーダ

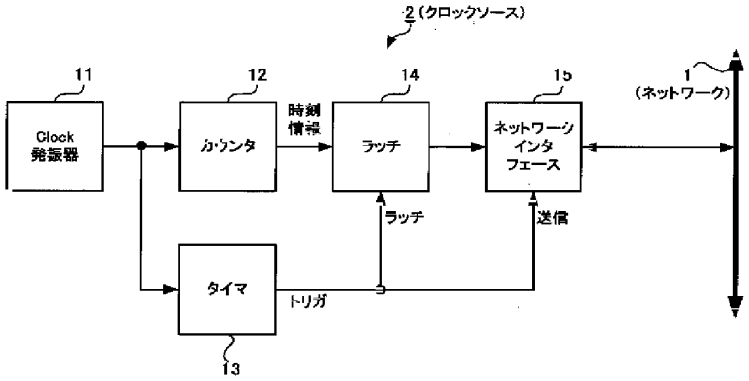
【図１】



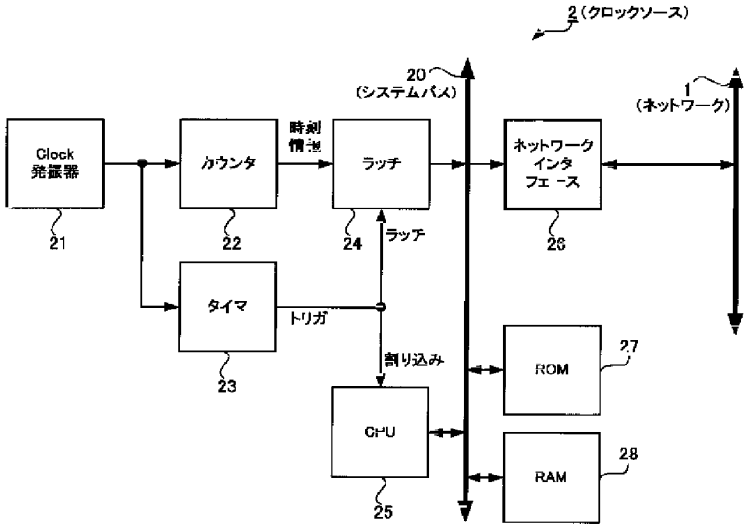
【図１０】



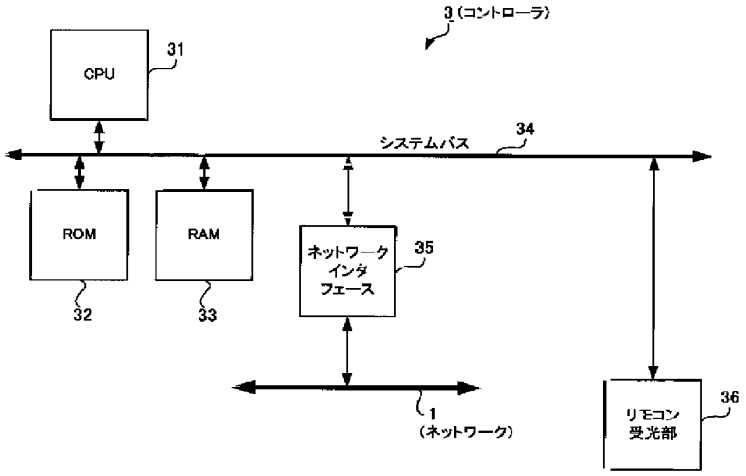
【図2】


2（クロックソース）
11
12
14
15
1
（ネットワーク）
Clock
発振器
カウンタ
時刻
情報
ラッチ
ネットワーク
インタ
フェース
ラッチ
送信
タイマ
トリガ
13

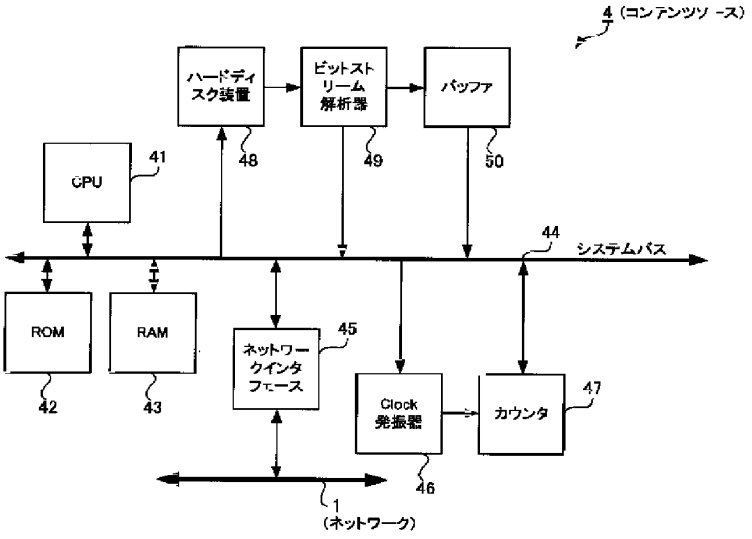
【図3】


2（クロックソース）
20
（システムバス）
1
（ネットワーク）
Clock
発振器
カウンタ
時刻
情報
ラッチ
ネットワーク
インタ
フェース
21
22
24
ラッチ
26
タイマ
トリガ
割り込み
ROM
27
CPU
RAM
28
23
25

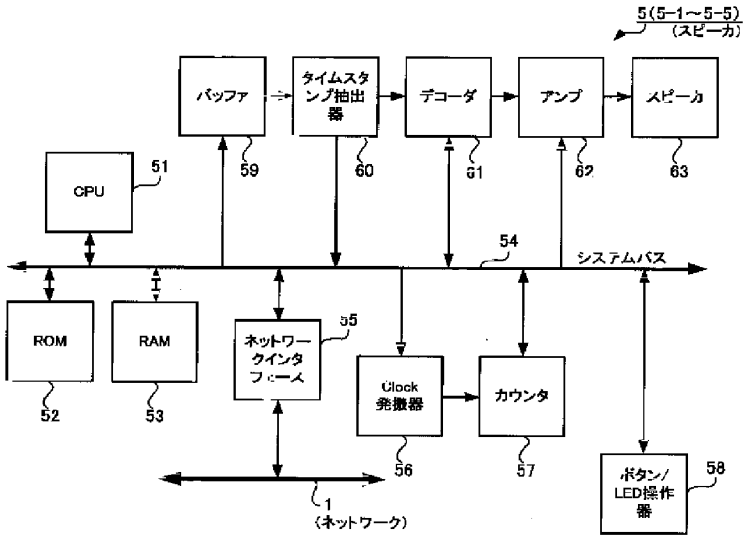
【図4】


3（コントローラ）
CPU
31
システムバス
34
ROM
RAM
ネットワークインタフェース
35
32
33
1
（ネットワーク）
リモコン受光部
36

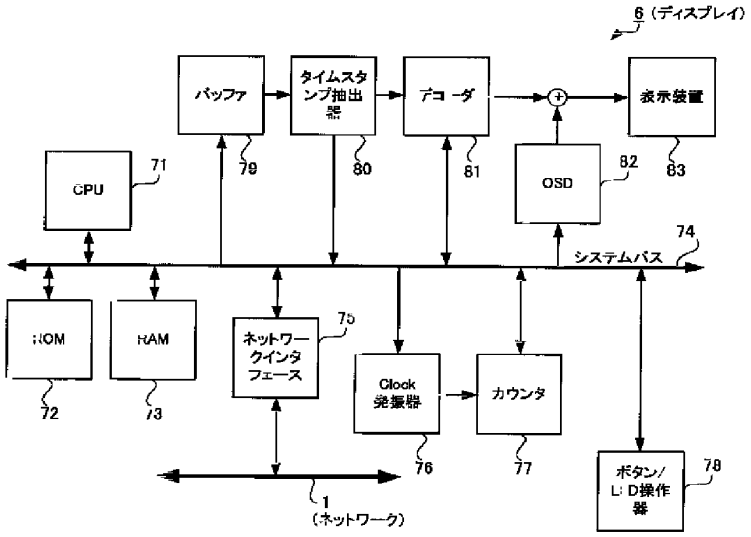
【図5】


4（コンテンツソース）
ハードディスク装置
ビットストリーム解析器
バッファ
48
49
50
CPU
41
44
システムバス
ROM
RAM
ネットワークインタフェース
45
42
43
Clock発振器
カウンタ
47
46
1
（ネットワーク）

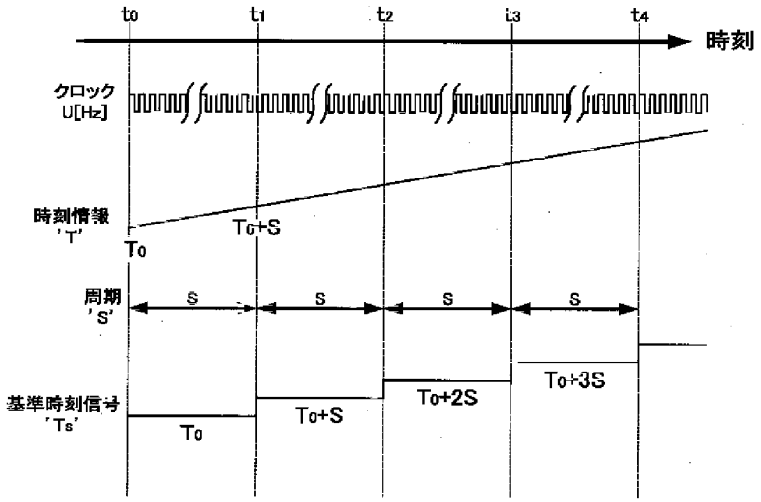
【図６】


5(5-1～5-5)
(スピーカ)
バッファ
タイムスタンプ抽出器
デコーダ
アンプ
スピーカ
59
60
61
62
63
CPU
51
54
システムバス
ROM
RAM
ネットワークインタフェース
55
Clock発振器
カウンタ
52
53
56
57
ボタン/LED操作器
58
1
(ネットワーク)

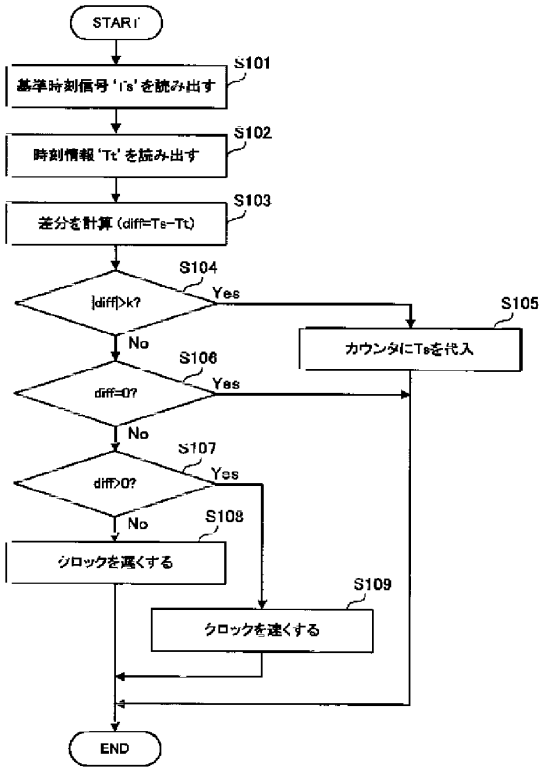
【図７】


6 (ディスプレイ)
バッファ
タイムスタンプ抽出器
デコーダ
表示装置
79
80
81
83
OSD
82
CPU
71
74
システムバス
ROM
RAM
ネットワークインタフェース
75
Clock発振器
カウンタ
72
73
76
77
ボタン/LED操作器
78
1
(ネットワーク)

【図８】



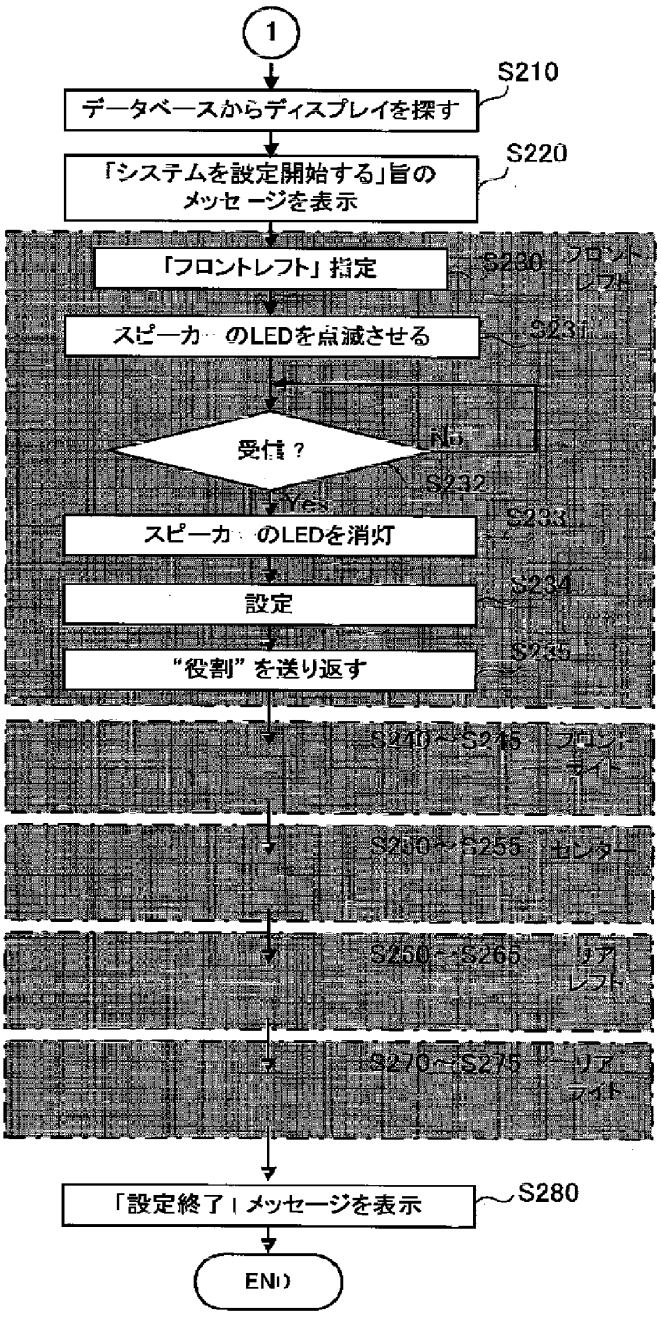
【図９】



【図１２】

エントリ数		N
第1番目	ネットワークアドレス	xxx.xxx.xxx.xxx
	種別	ディスプレイ
	役割	—
	デコードディレイ	xx.x
第2番目	ネットワークアドレス	xxx.xxx.xxx.xxx
	種別	モノラルスピーカ
	役割	—
	デコードディレイ	xx.x
第3番目	ネットワークアドレス	xxx.xxx.xxx.xxx
	種別	モノラルスピーカ
	役割	—
	デコードディレイ	xx.x

【図11】

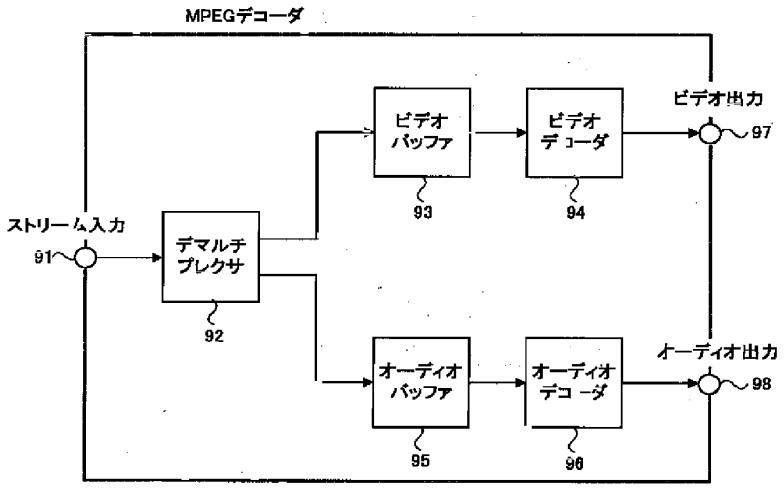

1
S210
データベースからディスプレイを探す
S220
「システムを設定開始する」旨の
メッセージを表示
「フロントレフト」指定
S230
フロント
レフト
スピーカのLEDを点滅させる
S231
受信？
No
S232
Yes
スピーカのLEDを消灯
S233
設定
S234
“役割”を送り返す
S235
S240～S245
フロント
ライト
S250～S255
センター
S260～S265
リア
レフト
S270～S275
リア
ライト
「設定終了」メッセージを表示
S280
END

【図１３】

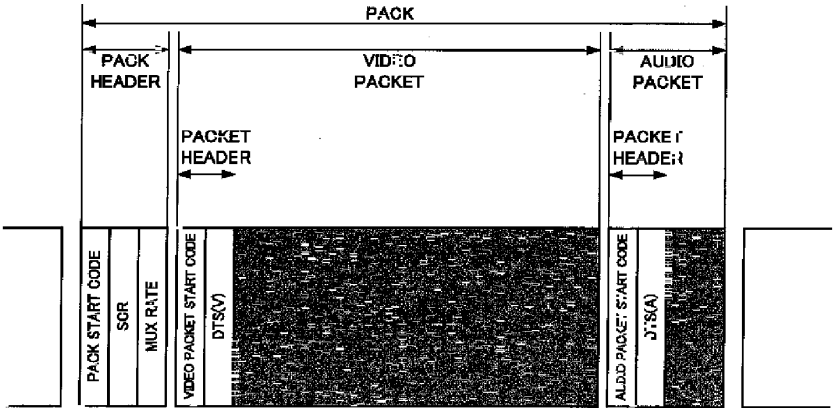
エントリ数	7

ネットワークアドレス	種別	役割	デコードディレイ	(参考)
xxx.xxx.xxx.xxx	ディスプレイ	—		(ディスプレイ 6)
xxx.xxx.xxx.xxx	モノラルスピーカ	—		(スピーカ 5-1)
xxx.xxx.xxx.xxx	モノラルスピーカ	—		(スピーカ 5-2)
xxx.xxx.xxx.xxx	モノラルスピーカ	—		(スピーカ 5-3)
xxx.xxx.xxx.xxx	モノラルスピーカ	—		(スピーカ 5-4)
xxx.xxx.xxx.xxx	モノラルスピーカ	—		(スピーカ 5-5)
xxx.xxx.xxx.xxx	コンテンツソース	—		(コンテンツソース 4)

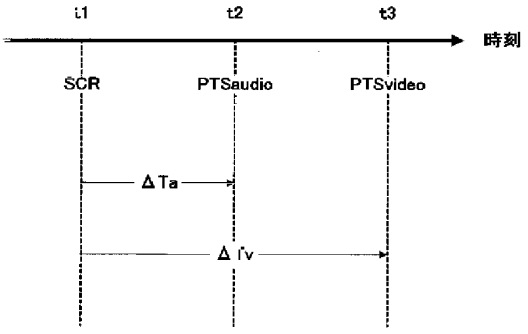
【図１４】



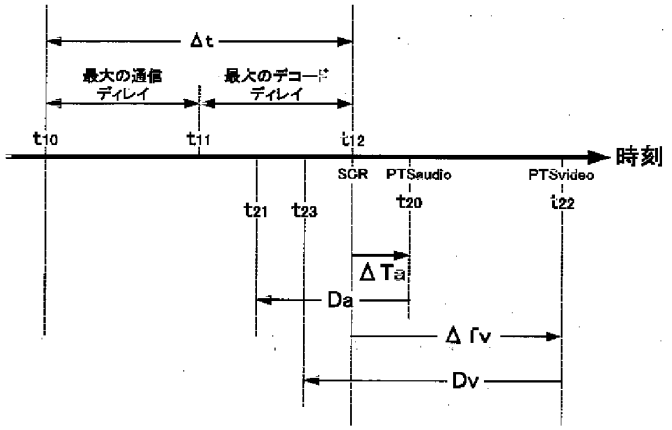
【図15】


PACK
PACK HEADER
VIDEO PACKET
AUDIO PACKET
PACKET HEADER
PACKET HEADER
PACK START CODE
SCR
MUX RATE
VIDEO PACKET START CODE
DTS(V)
AUDIO PACKET START CODE
DTS(A)

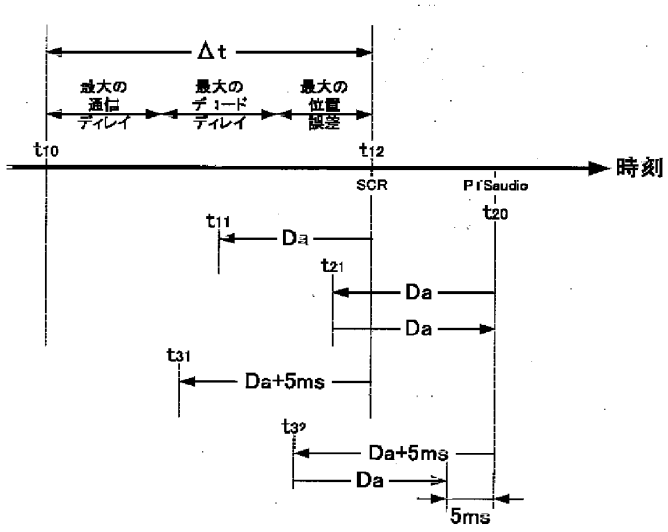
【図16】


t1
t2
t3
時刻
SCR
PTSaudio
PTSvideo
ΔTa
ΔTv

【図１７】


Δt
最大の通信ディレイ
最大のデコードディレイ
t10
t11
t12
時刻
SCR
PTSaudio
PTSvideo
t21
t23
t20
t22
ΔTa
Da
ΔTv
Dv

【図１８】


Δt
最大の通信ディレイ
最大のデコードディレイ
最大の位置誤差
t10
t12
時刻
SCR
PTSaudio
t20
t11
Da
t21
Da
Da
t31
Da+5ms
t32
Da+5ms
Da
5ms

Espacenet

Bibliographic data: JP2003101958 (A) — 2003-04-04

METHOD AND DEVICE FOR SYNCHRONOUS REPRODUCTION

Inventor(s): IDE KENICHI ± (IDE KENICHI)

Applicant(s): TOSHIBA CORP ± (TOSHIBA CORP)

Classification: **- international:** ***H04B1/16; H04L7/00; H04N5/93;***
(IPC1-7): H04B1/16; H04L7/00; H04N5/93
- cooperative:

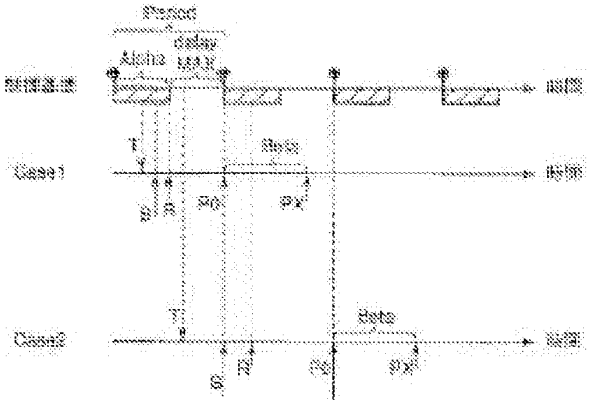
Application number: JP20010287375 20010920

Priority number(s): JP20010287375 20010920

Also published as: JP3544963 (B2)

Abstract of JP2003101958 (A)

PROBLEM TO BE SOLVED: To synchronize information synchronized with each other such as a moving picture and sound accompanying the moving picture in different devices and to reproduce the information without transmitting information about reproducing time. SOLUTION: A maximum estimation value of the total delay time including the transmission delay time of data and a processing delay time needed for decoding processing, etc., is defined as delayMax. A value Period larger than the delayMax is defined, and a reproduction start reference point is provided on a time base with an interval Period. An alpha is defined by (Period-delayMax), and the section between the reproduction start reference point and the Alpha is defined as a transmission commenceable section. A transmitting side device starts to transmit data within the transmission commenceable section.; After the data reach a receiving side device, the existence range of a time R when the data can be reproduced after being subjected to processing such as decoding is limited to a range existing after a transmission start time S and up to the first reproduction start reference point to come after passing the transmission commenceable section. The transmitting side device and the receiving side device individually start reproduction by taking the reproduction start reference point P0 as a reference.


Period
delay
Alpha
Case1
Beta
Case2

(19)日本国特許庁（ＪＰ）　(12) **公開特許公報**（Ａ）　(11)特許出願公開番号
特開2003－101958
（P2003－101958A）
(43)公開日　平成15年４月４日(2003.4.4)

(51)Int.Cl.[7]	識別記号	Ｆ Ｉ		テーマコード（参考）	
Ｈ０４Ｎ　5/93		Ｈ０４Ｂ	1/16	Ｍ	５Ｃ０５３
Ｈ０４Ｂ　1/16		Ｈ０４Ｌ	7/00	Ｚ	５Ｋ０４７
Ｈ０４Ｌ　7/00		Ｈ０４Ｎ	5/93	Ａ	５Ｋ０６１

審査請求　有　　請求項の数18　ＯＬ　（全 10 頁）

(21)出願番号　特願2001－287375(P2001－287375)

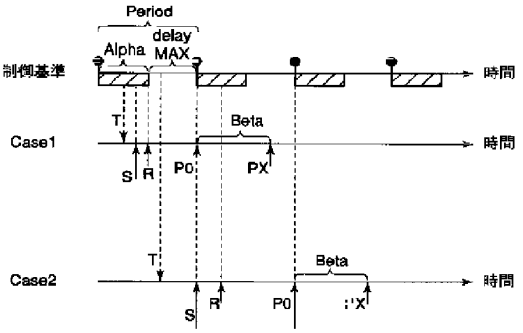
(22)出願日　平成13年９月20日(2001.9.20)

(71)出願人　000003078
株式会社東芝
東京都港区芝浦一丁目１番１号
(72)発明者　井手　賢一
東京都青梅市末広町２丁目９番地　株式会社東芝青梅工場内
(74)代理人　100058479
弁理士　鈴江　武彦　（外６名）

Ｆターム(参考)　5C053 FA24 GB21 JA01 JA22 JA26 KA24 LA05 LA06 LA20
5K047 AA03 AA18 BB01 DD01 DD02 JJ08 MM23
5K061 AA15 EF06

(54)【発明の名称】　同期再生のための方法および装置

(57)【要約】
【課題】再生時刻に関する情報を送信することなく、動画とそれに付随する音声のように互いに同期した情報を異なる機器で同期して再生する。
【解決手段】データの伝送遅延時間と復号処理などに要する処理遅延時間を含む合計の遅延時間の最大見積もり値をdelayMaxとする。そして、delayMaxよりも大きい値Periodを定め、間隔Periodで時間軸上に再生開始基準点を設ける。また、Alphaを（Period － delayMax）と定め、前記再生開始基準点からAlphaの区間を送信開始可能区間とする。送信側機器は送信開始可能区間内にデータ送信を開始する。このデータが受信側機器に届いた後、復号などの処理を経て再生可能になる時刻Ｒの存在範囲は送信開始時刻Ｓ以降で、送信開始可能区間経過後に最初に到来する再生開始基準点までの範囲に限定される。この再生開始基準点をＰＯを基準として、送信側機器および受信側機器がそれぞれ再生を開始する。



送信開始可能区間
再生開始基準点
T　再生ボタン押下時刻
S　音声データ送信開始時刻
R　最初の音声データが到着して再生可能となった時刻
P0　最も早い再生開始時刻
PX　一般の再生開始時刻(PX=P0+Beta)
Period　再生開始基準点の間隔
delayMax　(伝送遅延＋処理遅延)の最大値
Alpha　=Period－delayMax
Beta　0以上の定数。再生開始時刻をP0以降に遅延させる値

【特許請求の範囲】
【請求項１】　第一のデータの再生を行いかつ同期して再生すべき第二のデータを通信路を介して送信する第一の機器と前記第一の機器から送信された第二のデータを受信して再生する第二の機器との間で、前記第一および第二のデータの再生タイミングを同期させるための同期再生方法であって、
前記第二のデータの送信は、前記第二のデータを送信してから受信側で再生可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に開始し、
前記第一および第二のデータの再生は、前記第二のデータの送信を開始した送信開始可能区間後の最初の基準点を基に開始することを特徴とする同期再生方法。
【請求項２】　前記第二のデータの送信を開始する前に、前記第一の機器と前記第二の機器との間で、前記第一および第二の機器がそれぞれ前記基準点を定める時間軸の時刻基準となる時計を同期させる同期処理を実行することを特徴とする請求項１記載の同期再生方法。
【請求項３】　前記第一のデータは動画データであり、前記第二のデータは音声データであることを特徴とする請求項１記載の同期再生方法。
【請求項４】　前記第一のデータは音声データであり、前記第二のデータは動画データであることを特徴とする請求項１記載の同期再生方法。
【請求項５】　前記第一および第二のデータの再生は、前記第二のデータの送信を開始した送信開始可能区間後の最初の基準点から予め定められた時間が経過した時点で開始することを特徴とする請求項１記載の同期再生方法。
【請求項６】　前記通信路は、無線通信路であることを特徴とする請求項１記載の同期再生方法。
【請求項７】　第一のデータの再生を行いかつ同期して再生すべき第二のデータの再生開始を指示する再生開始コマンドを通信路を介して送信する第一の機器と前記第一の機器から送信された再生開始コマンドに応答して前記第二のデータの再生を行う第二の機器との間で、前記第一および第二のデータの再生タイミングを同期させるための同期再生方法であって、
前記再生開始コマンドの送信は、前記再生開始コマンドを送信してから受信側で前記再生開始コマンドに応答した再生処理を開始可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に開始し、
前記第一および第二のデータの再生は、前記再生開始コマンドの送信を開始した送信開始可能区間後の最初の基準点を基に開始することを特徴とする同期再生方法。
【請求項８】　前記再生開始コマンドの送信を開始する前に、前記通信路を介して前記第一の機器から前記第二の機器に前記第二のデータの送信を開始することを特徴とする請求項７記載の同期再生方法。
【請求項９】　第一のデータの再生を行いかつ同期して再生すべき第二のデータを通信路を介して送信する第一の機器と前記第一の機器から送信された第二のデータを受信して再生する第二の機器とを含む同期再生システムであって、
前記第一の機器は、
前記第二のデータを送信してから受信側で再生可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に、前記第二のデータの送信を開始する手段と、前記第二のデータの送信を開始した送信開始可能区間後の最初の基準点を基に、前記第一のデータの再生を開始する手段とを具備し、
前記第２の機器は、
前記第二のデータを受信する手段と、
前記時間間隔で時間軸上に一意に定められた前記複数の基準点を管理し、前記受信した第二のデータの再生開始可能後の最初の基準点を基に、前記第二のデータの再生を開始する手段とを具備することを特徴とする同期再生システム。
【請求項１０】　第一のデータの再生を行いかつ同期して再生すべき第二のデータの再生開始を指示する再生開始コマンドを通信路を介して送信する第一の機器と前記第一の機器から送信された再生開始コマンドに応答して前記第二のデータの再生を行う第二の機器とを含む同期再生システムであって、
前記第一の機器は、
前記再生開始コマンドを送信してから受信側で前記再生開始コマンドに応答した再生処理を開始可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に、前記再生開始コマンドの送信を開始する手段と、
前記再生開始コマンドの送信を開始した送信開始可能区間後の最初の基準点を基に、前記第一のデータの再生を開始する手段とを具備し、
前記第二の機器は、
前記再生開始コマンドを受信する手段と、
前記時間間隔で時間軸上に一意に定められる前記複数の基準点を管理し、前記再生開始コマンドに応答して再生処理が開始可能になった後の最初の基準点を基に、前記第二のデータの再生を開始する手段とを具備することを特徴とする同期再生システム。
【請求項１１】　第一のデータの再生を行いかつ同期して再生すべき第二のデータを通信路を介して受信側の装置に送信可能な再生装置であって、

前記第二のデータを送信してから受信側で再生可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に、前記第二のデータの送信を開始する手段と、
前記第二のデータの送信を開始した送信開始可能区間後の最初の基準点を基に、前記第一のデータの再生を開始する手段とを具備することを特徴とする再生装置。

【請求項１２】　第一のデータを再生する送信側装置から通信路を介して送信される第二のデータを受信して再生する再生装置であって、
前記第二のデータを送信してから受信側で再生可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に前記送信側装置によって送信が開始された前記第二のデータを受信する手段と、
前記時間間隔で時間軸上に一意に定められた前記複数の基準点を管理し、前記受信した第二のデータの再生開始可能後の最初の基準点を基に、前記第二のデータの再生を開始する手段とを具備することを特徴とする再生装置。

【請求項１３】　第一のデータの再生を行いかつ同期して再生すべき第二のデータの再生開始を指示する再生開始コマンドを通信路を介して受信側の装置に送信可能な再生装置であって、
前記再生開始コマンドを送信してから受信側で前記再生開始コマンドに応答した再生処理を開始可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に、前記再生開始コマンドの送信を開始する手段と、
前記再生開始コマンドの送信を開始した送信開始可能区間後の最初の基準点を基に、前記第一のデータの再生を開始する手段とを具備することを特徴とする再生装置。

【請求項１４】　第一のデータを再生する送信側装置から通信路を介して送信される再生開始コマンドに応答して、前記第一のデータと同期して再生すべき第二のデータを再生する再生装置であって、
前記再生開始コマンドを送信してから受信側で前記再生開始コマンドに応答した再生処理を開始可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に前記送信側装置が送信を開始した前記再生開始コマンドを受信する手段と、
前記時間間隔で時間軸上に一意に定められる前記複数の基準点を管理し、前記再生開始コマンドに応答して再生処理が開始可能になった後の最初の基準点を基に、前記第二のデータの再生を開始する手段とを具備することを特徴とする再生装置。

【請求項１５】　コンピュータに第一のデータを再生させ且つ前記第一のデータと同期して再生すべき第二のデータを通信路を介して受信側の装置に送信させるコンピュータプログラムであって、
前記第二のデータを送信してから受信側で再生可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に、前記第二のデータの送信を前記コンピュータに開始させる手順と、
前記第二のデータの送信を開始した送信開始可能区間後の最初の基準点を基に、前記第一のデータの再生を前記コンピュータに開始させる手段とを具備することを特徴とするコンピュータプログラム。

【請求項１６】　第一のデータを再生する送信側装置から通信路を介して送信される第二のデータをコンピュータに再生させるコンピュータプログラムであって、
前記第二のデータを送信してから受信側で再生可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に前記送信側装置によって送信が開始された前記第二のデータを前記コンピュータに受信させる手順と、
前記時間間隔で時間軸上に一意に定められた前記複数の基準点を前記コンピュータに管理させ、前記受信した第二のデータの再生開始可能後の最初の基準点を基に、前記第二のデータの再生を前記コンピュータに開始させる手順とを含むことを特徴とするコンピュータプログラム。

【請求項１７】　コンピュータに第一のデータを再生させ且つ前記第一のデータと同期して再生すべき第二のデータの再生開始を指示する再生開始コマンドを通信路を介して受信側の装置に送信させるコンピュータプログラムであって、
前記再生開始コマンドを送信してから受信側で前記再生開始コマンドに応答した再生処理を開始可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に、前記再生開始コマンドの送信を前記コンピュータに開始させる手順と、
前記再生開始コマンドの送信を開始した送信開始可能区間後の最初の基準点を基に、前記第一のデータの再生を前記コンピュータに開始させる手順とを具備することを特徴とするコンピュータプログラム。

【請求項１８】　第一のデータを再生する送信側装置から通信路を介して送信される再生開始コマンドに応答して、コンピュータに前記第一のデータと同期して再生すべき第二のデータを再生させるコンピュータプログラム

であって、
前記再生開始コマンドを送信してから受信側で前記再生開始コマンドに応答した再生処理を開始可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に前記送信側装置が送信を開始した前記再生開始コマンドを前記コンピュータに受信させる手順と、
前記時間間隔で時間軸上に一意に定められる前記複数の基準点を前記コンピュータに管理させ、前記再生開始コマンドに応答して再生処理が開始可能になった後の最初の基準点を基に、前記第二のデータの再生を前記コンピュータに開始させる手順とを具備することを特徴とするコンピュータプログラム。
【発明の詳細な説明】
【０００１】
【発明の属する技術分野】本発明は例えば動画データと音声データのように同期して再生する必要がある第一及び第二のデータを再生するための方法および装置に関し、特に第一及び第二のデータを互いに異なる機器で同期して再生するための方法および装置に関する。
【０００２】
【従来の技術】動画データ、音声データのようなストリームデータを送信側機器から受信側機器へ伝送して、予め定められた時刻に処理させる方法としては、例えば特開平１１　３１７７４８号公報に開示された手法に見られるように、送信するデータに処理されるべき時刻情報を付加して送信し、受信側では受信したデータに含まれる処理すべき時刻の到来を待ってデータを処理する方法が知られている。
【０００３】
【発明が解決しようとする課題】しかし、上記手法ではデータを処理すべき時刻を示すタイムスタンプをデータの他に送信する必要があった。このため、データを送信する手順やデータ形式とあわせて、データを処理すべき時刻を示すタイムスタンプについてもそれを送信する手順や形式が定められている必要がある。よって、これらが定められていない通信規約を用いる場合には適用できないという問題があった。
【０００４】本発明は上述の事情を考慮してなされたもので、再生時刻に関する情報を送信することなく、動画とそれに付随する音声のように互いに同期した情報を異なる機器で同期して再生することが可能な方法および装置を提供することを目的とする。
【０００５】
【課題を解決するための手段】上述の課題を解決するため、本発明は、第一のデータの再生を行いかつ同期して再生すべき第二のデータを通信路を介して送信する第一の機器と前記第一の機器から送信された第二のデータを受信して再生する第二の機器との間で、前記第一および第二のデータの再生タイミングを同期させるための同期再生方法であって、前記第二のデータの送信は、前記第二のデータを送信してから受信側で再生可能になるまでの最大遅延時間に所定の送信開始可能区間を加えた時間間隔で時間軸上に一意に定められる複数の基準点の内、いずれかの基準点から始まる送信開始可能区間内に開始し、前記第一および第二のデータの再生は、前記第二のデータの送信を開始した送信開始可能区間後の最初の基準点を基に開始することを特徴とする。
【０００６】このように、最大遅延時間に所定の送信開始可能区間を加えた時間間隔で予め時間軸上に複数の基準点を定め、これら複数の基準点の内のいずれかの基準点から始まる送信開始可能区間内に第二のデータの送信を開始することで、次の基準点以前に第二のデータが受信側で再生可能になることを保証することができる。よって、送信側と受信側が予め定められた同一の規則で複数の基準点それぞれの時刻を定めるだけで、再生時刻に関する情報を何等送信することなく、動画とそれに付随する音声のように互いに同期した情報を異なる機器で同期して再生することが可能となる。
【０００７】
【発明の実施の形態】以下、図面を参照して本発明の実施の形態を説明する。図１には、本発明の一実施形態に係る同期再生システムが示されている。この同期再生システムは音声が付随した動画データを再生するシステムであり、図１に示すように２台の機器から構成されている。第一の機器（送信側機器）１１は動画データおよびそれに付随する音声データが記録されている記憶装置を備え、動画データを再生してディスプレイ上に再生すると同時に、第二の機器（受信側機器）１４に通信路１３を介して音声データを送信する。この第一の機器１１は例えば携帯情報端末、ビデオプレーヤ、パーソナルコンピュータなどの電子機器から構成される。一方、第二の機器１４は、音声データの受信部と受信した音声データを再生する機構を備えており、第一の機器１１から受信した音声データを再生する。この第二の機器１４も第一の機器１１と同様に様々な電子機器として実現できるが、その典型的な一例としてはヘッドフォンである。
【０００８】このシステムでは通信路１３として無線通信路が用いられており、例えばＢｌｕｅｔｏｏｔｈ（２．４ＧＨｚ　ＩＳＭ帯を用いる短距離無線通信方式）を用いることによって動画表示を行う機器と音声の再生を行う機器とを無線接続することによって、便利な視聴システムを実現している。
【０００９】ここで、第一の機器１１で再生される動画データの呈示と第二の機器１４で再生される音声データの発音を同期させることによっていわゆるリップシンク(lipsync)を実現したいと考える。
【００１０】次に、上記の要求を実現するために必要な技術的要素について説明する。動画データとそれに付随

する音声データは、それらをディジタルデータとして取り込む際に、互いの時間関係を保存することが既に行われている。例えばパーソナルコンピュータ上の音声つき動画データの形式として用いられているＡＶＩ形式（Audio Video Interleave Format）では、動画および音声のメディア毎に、サンプリングレートが記録されており、かつ、記録されている各々のメディアの再生開始時刻が一致しているという仕組みになっている。このため、各々の再生を同時に開始し、かつ、各々のメディアの再生速度を、単一の時計あるいは正確に同期した時計とサンプリングレートに従って決めることによって、リップシンク（lip sync）を保つことができる仕組みになっている。

【００１１】また、複数のメディアの再生を同期させる他の方法としては、インターネット上のストリーミング伝送で用いられるＲＴＰ（Real time Transfer Protocol,インターネット上で動画,音声などのメディアデータを送受信するための方式。IETFのRFC1889等で規格化されている。）に見られるように、一群のメディアデータ毎にタイムスタンプを付与すると共にサンプリングレートに関する情報を付与するなどの方法も用いられている。

【００１２】このように、互いに同期して再生されるべき複数のメディアデータを同期して再生するためには、進み方が正確な時計と例えばメディアデータの呈示開始時刻を一致させる手段が必要である。

【００１３】ここで本発明が解決しようとする状況では、動画および音声データは異なる機器上で再生される。そこで、前述の要件を満たすためには、各々の機器の時計間で時刻の同期をとり、さらに、各々のメディアデータの再生開始時刻を２台の機器同上で示し合わせることが必要となる。

【００１４】まず、時刻の同期をとる件については、例えばＮＴＰ（Network Time Protocol、インターネット上で計算機同士の時刻あわせを行うための方式。IETFのRFC1305等で規格化されている。）と呼ばれる時刻合わせ方式を用いることで実現可能である。あるいは、例えば特開２０００－３３２８０２号公報に開示されているように、時刻情報を送信側機器から受信側機器へ送り、受信側機器の時計を送信側機器に合わせることによっても実現可能である。これら２例の他にも、方式によって実現可能な時刻あわせの精度が異なるものの、様々な時刻あわせ手法が知られており、これらの内の任意の手法を用いることが可能である。

【００１５】次に、再生開始時刻を示しあわせる件については、本実施形態では次のような方法を使用する。

【００１６】本実施形態では再生開始時刻に関する情報を機器間で伝送せずに、再生開始時刻を合わせることが特徴である。まず、前提条件として以下の条件を設ける。

前提条件：

・２台の機器１１，１４の時計は互いに同期している。

・送信側機器１１から受信側機器１４へ最初のメディアデータを送信する時刻Ｓを送信側機器１１は制御できる。

・送信側機器１１から受信側機器１４へのメディアデータ送信から、受信されたデータが再生されるまでの所要時間delayMax(伝送遅延、メディアデータのデコード処理などに必要な処理遅延時間を含む)の最大値を見積もり可能。

【００１７】上記の前提条件の下で、以下に示す方法によって再生開始時刻を同期させる。

同時刻に再生を開始する方法：送信側機器１１および受信側機器１４がメディアデータ、例えば動画と音声データを再生開始可能な時刻（再生開始基準点）は、予め定められた離散的な時刻に限る。各々の時刻の間隔はPeriod (= Alpha + delayMax)とする。Period, Alpha, delayMaxの定義は以下の通りである。

【００１８】

・delayMax　(伝送遅延+処理遅延)の最大見積もり値

・Period　再生開始基準点の間隔

・Alpha　(= Period - delayMax)

送信側機器１１がメディアデータを受信側機器１４へ送信開始する時刻Ｓは、各再生開始基準点以降Alphaの区間内に限る。すなわち、送信側機器１１および受信側機器１４ではそれぞれ同一の規則を用いて、互いに同期した時間軸上にPeriodの間隔で再生開始基準点を定める処理が行われる。そして、送信側機器１１から受信側機器１４への音声データの送信は、いずれかの再生開始基準点からAlphaの区間内に開始される。つまりAlphaが送信開始可能区間となる。

【００１９】図２は図１に示した各々の機器のシステム構成図である。まず、送信側機器１１について説明する。送信側機器１１は、図示のように、動画/音声記憶部２１、動画/音声分離部２２、動画復号部２３、動画一時記憶部２４、動画表示部２５、送信データ一時記憶部２６、音声データ送信部２７、制御部２８、時刻管理部２９、および制御データ送受信部２Ａから構成されている。

【００２０】動画/音声記憶部２１には送信側機器１１で再生される動画データと受信側機器１４で再生される音声データが記憶されている。これらのデータはたとえば前述のＡＶＩ形式で保存されており、動画と音声の対応関係を知るために必要な情報とともに記憶されている。また、一般に動画や音声のデータは圧縮されて保存されることが多く、本実施形態の場合も適当な形式で圧縮された上でＡＶＩ形式で保管されているものとする。

【００２１】動画/音声分離部２２は動画/音声記憶部２１から読み込まれたＡＶＩ形式の動画/音声データを動画データと音声データに分離する。動画復号部２３は圧

縮形式の動画データを復号して動画表示部２５で表示可能な形式に変換する。動画一時記憶部２４は復号された動画データを、表示される時点に達するまでの間、一時的に記憶する。一般に動画の復号処理は時間がかかるため、予め復号処理を行って動画一時記憶部２４に保存しておくことによって、表示すべき時刻の到来時に即座に表示することが可能になる。

【００２２】動画表示部２５は動画一時記憶部２４に記憶されている動画データを制御部２８からの指示に従って表示する。送信データ一時記憶部２６は音声データが送信されるべき時刻が来るまでの間、一時的に音声データを記憶する。音声データ送信部２７はＢｌｕｅｔｏｏｔｈなどの無線通信インタフェースを用いて受信側機器１４へ音声データを送信する。

【００２３】制御データ送受信部２ＡはＢｌｕｅｔｏｏｔｈなどの無線通信インタフェースを介して受信側機器１４との間で制御情報をやりとりする。時刻管理部２９は受信側機器１４の時刻と同期した時刻を刻み、動画再生や音声データの送信の基準となる時刻情報を提供する。受信側機器１４の時刻との同期は制御部２８および制御データ送受信部２Ａと連携して前述のＮＴＰ方式などの時刻合わせ手順を実行することで行う。

【００２４】制御部２８は送信側機器１１全体の連携制御を行う。制御部２８は例えば汎用ＣＰＵとＲＡＭ，ＲＯＭで構成されるマイクロコンピュータとして実現することが可能である。また、制御部２８以外の前述した機能モジュールも基本的にマイクロコンピュータにその動作を行わせる形での実現が可能である。

【００２５】受信側機器１４は、図示のように、音声データ受信部２Ｂ、音声復号部２Ｃ、音声データ一時記憶部２Ｄ、音声再生部２Ｅ、制御部２Ｆ、制御データ送受信部２Ｇ、時刻管理部２Ｈから構成されている。

【００２６】音声データ受信部２Ｂは送信側機器１１から送信された音声データを受信する。音声復号部２Ｃは圧縮されている音声データの復号を行い、音声再生部２Ｅで再生可能な形式に変換する。

【００２７】音声データ一時記憶部２Ｄは復号した音声データを再生すべき時刻に達するまでの間、一時的に保存する。一般に音声データは圧縮された形式で伝送される場合が多く、復号には時間がかかるため、予め復号を行って音声データ一時記憶部２Ｄに記憶しておくことによって、再生時刻の到来時に即座に発音させることが可能になる。

【００２８】時刻管理部２Ｈは送信側機器１１の時刻と同期した時刻を刻み、音声データの再生の基準となる時刻情報を提供する。送信側機器１１の時刻との同期は制御部２Ｆおよび制御データ送受信部２Ｇと連携して送信側機器１１と前述のＮＴＰ方式などの時刻合わせ手順を実行することで行う。制御部２Ｆは受信側機器１４全体の連携制御を行う。制御部２Ｆは例えば汎用ＣＰＵとＲＡＭ，ＲＯＭで構成されるマイクロコンピュータとして実現することが可能である。また、制御部２Ｆ以外の前述した機能モジュールも基本的にマイクロコンピュータにその動作を行わせる形での実現が可能である。

【００２９】次に、図３を用いて本実施形態のシステムの動作の概要を説明する。受信側機器１４は通常は待機モードで動作しており、送信側機器１１からの電源ＯＮ指示を受けて待機モードから通常の動作モードへ移行する。待機中は電力消費が非常に低い状態であり、電源ＯＮ指示のみに反応することが可能である。ユーザが送信側機器１１の操作パネル１２の電源オンボタンを操作することにより（Ｓ３８）、電源ＯＮ指示（Ｓ３１）が送信側機器１１から受信側機器１４に送信される。電源ＯＮの後、双方の機器１１，１４は通信路確立処理（Ｓ３２）を行う。

【００３０】次に、時刻あわせ処理（Ｓ３３）を前述のＮＴＰ方式などにより行う。送信側機器１１と受信側機器１４との間の時刻あわせが完了すると、送信側機器１１および受信側機器１４でそれぞれ上述の送信開始可能区間および再生開始基準点が設定されて同期再生を行うための準備が行われる。

【００３１】ユーザが送信側機器１１の操作パネル１２の再生ボタンを押下すると（Ｓ３９）、送信側機器１１では音声データの送信と動画データの再生が行われる。受信側機器１４では受信した音声データの再生が行われる。これらの動作は動画と音声の同期をとるために本発明で提案する手法に従って行われる。この処理の詳細については図４で後述するが、基本的には、送信側機器１１から受信側機器１４への音声データの送信は送信開始可能区間内でのみ開始され、そして、送信側機器１１および受信側機器１４それぞれにおける動画データおよび音声データの再生は、音声データが送信開始された送信開始可能区間後の最初の再生開始基準点を基準とすることにより、同時に開始されることになる。

【００３２】ユーザが操作パネル１２の電源ＯＦＦボタンを押下すると（Ｓ４０）、通信路切断および電源ＯＦＦ処理が行われ、受信側機器１３は待機状態に移行し、送信側機器１１は電源ＯＦＦ状態になる。

【００３３】次に図４を用いて本実施形態の特徴とする同期再生の方法を説明する。

【００３４】データの伝送遅延時間と復号処理などに要する処理遅延時間を含む合計の遅延時間の最大見積もり値をdelayMaxとする。このdelayMaxの値は、予め測定を行うことによって決められる値である。そして、delayMaxよりも大きい値Periodを定め、間隔Periodで時間軸上に再生開始基準点を設ける。また、Alphaを（Period－delayMax）と定め、前記再生開始基準点からAlphaの区間を送信開始可能区間とする。これらの情報は予め決めておくことにより、容易に双方の機器で再生開始基準点の周期を一致させることが可能である。例えば、delayM

axが100msecである場合にPeriodを200msecに定め、200msecの整数倍の時刻を再生開始基準点とするという具合である。時刻合わせを行っているので、双方の機器で各再生開始基準点の時刻は一致する。

【００３５】ここで、図４中のCase1およびCase2の場合を例にあげ、送信開始時刻および再生開始時刻の決め方を説明する。

【００３６】Case1は、送信開始可能区間内の時刻Tにユーザが再生ボタンを押下した場合を示している。この場合は、直ちにデータの送信を開始する。図では時刻Sで送信している。このデータが受信側機器１４に届いた後、復号などの処理を経て再生可能になる時刻をRとすれば、前述のdelayMaxの定義からRの存在範囲は時刻S以降で、前記送信開始可能区間経過後に最初に到来する再生開始基準点までの範囲に限定される。この再生開始基準点をＰ０とするとＰ０は送信側機器１１においても、受信側機器１４においても一意に定まる。送信側機器１１においては送信開始時刻を把握しているため自明であり、受信側機器１４においてはRの存在範囲から時刻R以降の最初の再生可能基準点をＰ０と判断すれば良い。要するに、受信側機器１４では受信データの再生が可能になってから最初の再生開始基準点がＰ０である。

【００３７】そこで、Ｐ０に予め定めた値Betaを加算した時刻ＰＸも一意に定まり、時刻ＰＸに両方の機器で再生を開始することによって、再生開始時刻を一致させることができる。Betaは０以上の値が可能であるが、０にした場合にユーザの再生ボタン押下に対する反応が一番速い。

【００３８】Case2は、送信可能区間外の時刻Tにユーザが再生ボタンを押下した場合を示している。この場合は、データの送信開始は時刻Tの直後に到来する送信開始可能区間まで待たされる。送信開始可能区間の到来と共に時刻Sにデータの送信を開始する。なお、データの送信開始はこの送信開始可能区間内であれば良い。受信側１４でデータが再生可能になる時刻Rの存在範囲は、時刻S以降で、時刻Sが含まれる送信開始可能区間経過後の最初の再生開始基準点までの範囲に限定される。この再生開始基準点をＰ０とするとP0は送信側機器１１においても受信側機器１４においても一意に定まる。送信側機器１１においては送信開始時刻を把握しているため自明であり、受信側機器１４においてはRの存在範囲から時刻R以降の最初の再生可能基準点をP0と判断すれば良い。そこで、Ｐ０に予め定めた値Betaを加算した時刻ＰＸも一意に定まり、時刻ＰＸに両方の機器で再生を開始することによって、再生開始時刻を一致させることができる。Betaの値はCase1の場合と同様、０の場合にユーザの再生ボタン押下に対する反応が一番速い。

【００３９】なお、Beta＝０とした場合には、Case1の場合のユーザの再生ボタン押下から再生開始までの最大遅延はPeriodとなり、Case2の場合のユーザの再生ボタン押下から再生開始までの最大遅延は(delayMax+Period)となる。

【００４０】次に図５のフローチャートを用いて、再生開始ボタン押下時の送信側機器１１の処理手順を説明する。

【００４１】ユーザが再生ボタンを押すと図５に示す処理が起動される。Step51では現在の時刻を取得する。これは送信側機器１１の時刻管理部２９に問い合わせを行うことにより行う。次にStep52に進み、現在、送信開始可能区間の内部であるか否かを判定する。送信開始可能区間に関する情報は時刻管理部２９が保持している。ここで、現在、送信開始可能区間の内部である場合にはStep53に進み、直ちに音声データの送信処理を起動する。引き続きStep54に進み、動画再生準備処理を起動した後、処理を終了する。動画再生準備処理では圧縮されている動画データの復号を表示に先行して行い、動画一時記憶部２４に蓄積する。この段階では表示は行われない。動画表示部２５への表示開始は、後述する動画データ再生開始判定処理の中で起動される。Step52で送信開始可能区間外と判定された場合はStep55に進み、既定の時間待機した後にStep51へ進み、再度、前述の処理を行う。

【００４２】次に図６のフローチャートを用いて、動画データ再生開始判定処理について説明する。

【００４３】この処理は図５のStep54から起動される。まず、Step61では現在の時刻を取得する。次にStep62に進み、現在、再生開始可能時刻（Ｐ０）に達しているか否かを判定する。再生開始可能時刻（Ｐ０）に関する情報は時刻管理部29が保持している。ここで、現在、再生開始可能時刻（Ｐ０）に達している場合はStep63へ進み、動画データの再生処理を起動する。Step62で再生開始時刻に達していないと判定された場合はStep64へ進み、既定の時間待機した後Step61へ進み、再度、前述の処理を行う。Step63で起動された動画データの再生処理では、時刻管理部２９の刻む時刻と動画データのフレームレートに従って再生を行う。

【００４４】次に、図７を用いて受信側機器１４における音声データの再生開始処理について説明する。この処理は例えば、最初の音声データを受信した時に起動される。Step71では現在の時刻を取得する。次にStep72に進み、現在、再生開始時刻（Ｐ０）に達しているかを判定する。再生開始時刻に達している場合はStep73へ進み、音声データの再生処理を起動する。Step72で再生開始時刻に達していないと判定された場合はStep74へ進み、既定の時間待機した後Step71へ進み、再度、前述の処理を行う。Step73で起動された音声データの再生処理では、時刻管理部２Ｈの刻む時刻と音声データのサンプリングレートに従って再生を行う。

【００４５】このように、本方式によれば送信開始時刻に関する情報を伝送路を介して送らなくとも、再生指示

から再生までの遅延時間が高々(delayMax+Period)で、同時に再生を開始することが可能になる。再生開始タイミングを揃えることで、後は、動画データおよび音声データそれぞれのサンプリングレートに従って再生処理を行うことで、リップシンク（lip sync）を保つことができる。

【００４６】なお、以上の説明では、音声データの送信開始を送信開始可能区間内に限定したが、これに代えて、送信側機器１１から受信側機器１４に対して再生開始を指示するための再生開始コマンドの送信開始時刻を、送信開始可能区間内に限定してもよい。つまり、音声データについては例えば再生開始ボタンが押された時点で送信し始めておき、再生開始コマンドの送信は、最初の送信開始可能区間で開始する。なお、この場合、delayMaxは、再生開始コマンドの送信を開始してから受信側機器１４がその再生開始コマンドに応答して再生処理を開始可能となるまでの最大時間となり、これは再生開始コマンドの伝送遅延時間とそのコマンド解釈処理などに要する処理遅延時間を含む合計の遅延時間の最大見積もり値として規定される。また、この場合には、音声データを記憶した記憶装置を受信側機器１４に用意しておくことも可能となる。他の点については音声データの送信開始タイミングを制御する場合と全く同じ方法が用いられる。

【００４７】また、上記の例では動画データを送信側機器１１で再生し、動画データと同期再生すべき音声データを送信側機器１１から受信側機器１４に送信したが、これとは逆に、音声データを送信側機器１１で再生し、音声データと同期再生すべき動画データを送信側機器１１から受信側機器１４に送信しても良い。この再生システムの一例を図８に示す。

【００４８】図８のシステムでは、携帯型オーディオプレイヤ１０１は、動画データおよびそれに付随する音声データが記録されている記憶装置を備えており、音声データを再生してイヤホン１０３から出力すると共に、携帯型ＬＣＤなどの表示装置１０２に無線通信路を介して動画データを送信する。

【００４９】また、本実施形態の同期再生方法は、同期再生する必要がある２つの動画データ同士を２つの異なる機器で再生する場合や、同期再生する必要がある２つの音声データ同士を２つの異なる機器で再生する場合にも適用可能である。２つの音声データを２つの異なる機器で同期再生するシステムの一例を図９に示す。

【００５０】図９のシステムは、ＤＶＤプレイヤ２０１を用いた視聴システムであり、ＤＶＤプレイヤ２０１によって再生される動画データはＴＶなどのモニタ２０３に表示される。またＤＶＤプレイヤ２０１から出力される複数チャネルの音声データはオーディオアンプ２０２に送られる。オーディオアンプ２０２はメインのＬ，Ｒチャネルの音声データを再生してＬ，Ｒのフロントスピーカ２０４，２０５から再生すると共に、サラウンド用のＬ，Ｒチャネルの音声データを無線通信路などを介してＬ，Ｒのバックスピーカシステム２０６，２０７に送信する。オーディオアンプ２０２とバックスピーカシステム２０６，２０７にそれぞれ本実施形態の方法を適用することで、音声データ同士の同期再生を容易に実現することが出来る。

【００５１】また、本実施形態の同期再生方法で使用する送信側機器および受信側機器それぞれの機能はすべてコンピュータプログラムによって実現することができるので、そのコンピュータプログラムが記憶されたコンピュータ読み取り可能な記憶媒体を通じてそのコンピュータプログラムを通常のコンピュータに導入して実行するだけで、本実施形態と同様の効果を容易に得ることが出来る。

【００５２】また、本発明は、上記実施形態に限定されるものではなく、実施段階ではその要旨を逸脱しない範囲で種々に変形することが可能である。例えばPeriodは常に一定である必要はなく、delayMaxよりも長ければ、予め決められたパターンで変化させるようにしても良い。同様に、Alphaの値も一定である必要はなく、予め決められたパターンで変化させるようにしても良い。また、delayMaxの値を通信環境に合わせて最適化する仕組みなどを用意することで、delayMaxの値も可変にすることが可能となる。

【００５３】更に、上記実施形態には種々の段階の発明が含まれており、開示される複数の構成要件における適宜な組み合わせにより種々の発明が抽出され得る。例えば、実施形態に示される全構成要件から幾つかの構成要件が削除されても、発明が解決しようとする課題の欄で述べた課題が解決でき、発明の効果の欄で述べられている効果が得られる場合には、この構成要件が削除された構成が発明として抽出され得る。

【００５４】

【発明の効果】以上説明したように、本発明によれば、再生時刻に関する情報を送信することなく、動画とそれに付随する音声のように互いに同期した情報を異なる機器で同期して再生することが可能となる。

【図面の簡単な説明】

【図１】本発明の一実施形態に係る再生システムの構成の一例を示すブロック図。

【図２】同実施形態の再生システムを構成する送信側機器と受信側機器それぞれの構成を示すブロック図。

【図３】同実施形態の再生システムの動作を示す図。

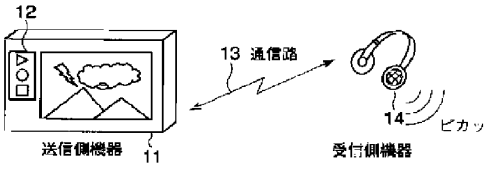
【図４】同実施形態の再生システムにおける同期再生処理のタイミングを説明するための図。

【図５】同実施形態の再生システムにおける再生開始ボタン押下時の送信側機器の処理手順を示すフローチャート。

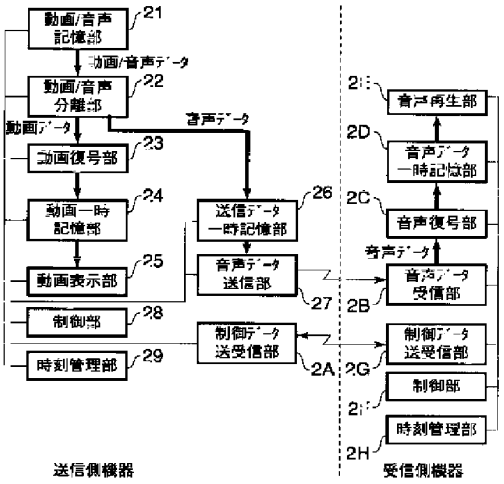
【図６】同実施形態の再生システムにおける送信側機器

の動画データ再生開始判定処理の手順を示すフローチャート。
【図７】同実施形態の再生システムにおける受信側機器の音声データ再生開始判定処理の手順を示すフローチャート。
【図８】本発明の一実施形態に係る再生システムの他の構成の一例を示すブロック図。
【図９】本発明の一実施形態に係る再生システムのさらに他の構成の一例を示すブロック図。
【符号の説明】
１１…送信側機器
１２…操作パネル
１３…通信路
１４…受信側機器
２１…動画/音声記憶部
２２…動画/音声分離部
２３…動画復号部
２４…動画一時記憶部
２５…動画表示部
２６…送信データ一時記憶部
２７…音声データ送信部
２８…制御部
２９…時刻管理部
２Ａ…制御データ送受信部
２Ｂ…音声データ受信部
２Ｃ…音声復号部
２Ｄ…音声データ一時記憶部
２Ｅ…音声再生部
２Ｆ…制御部
２Ｇ…制御データ送受信部
２Ｈ…時刻管理部

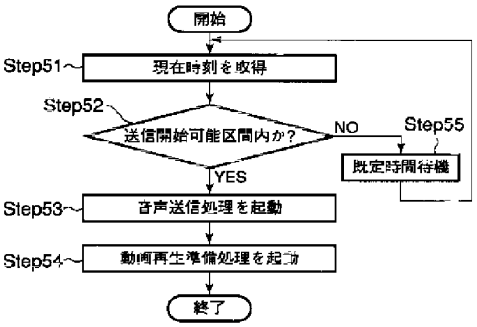
【図１】



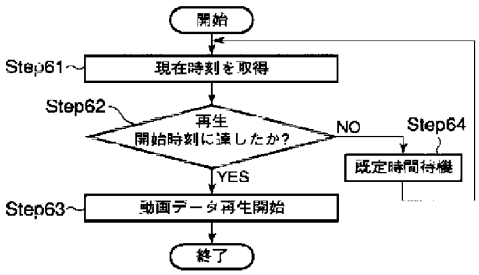
【図２】



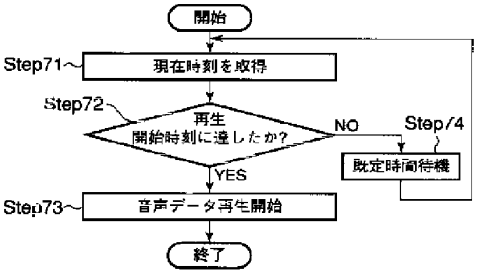
【図５】



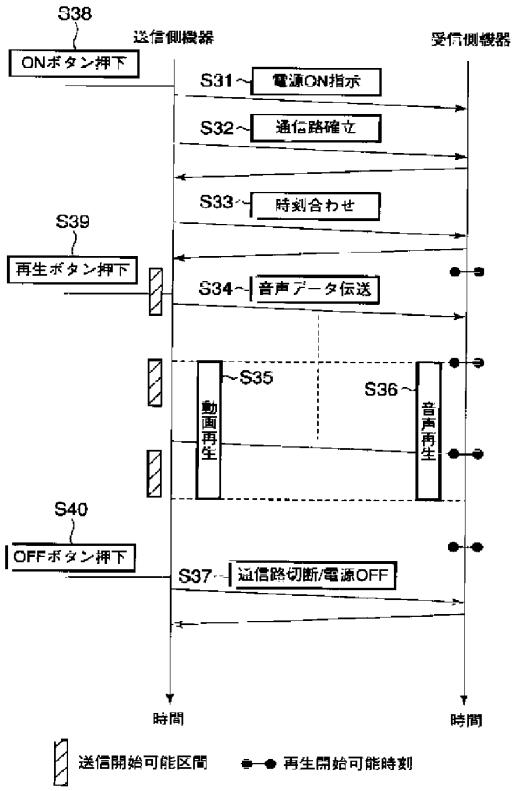
【図６】



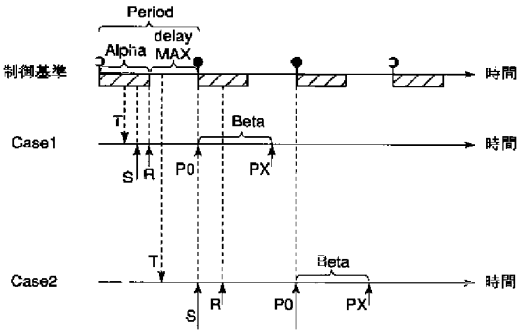
【図７】



【図3】

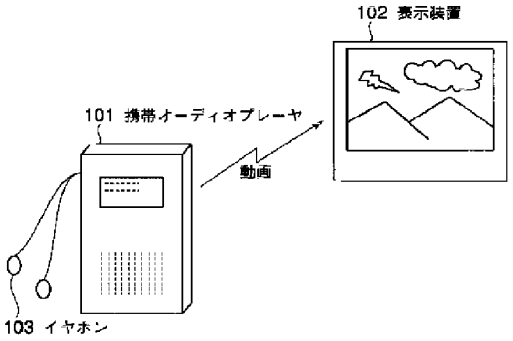


【図4】

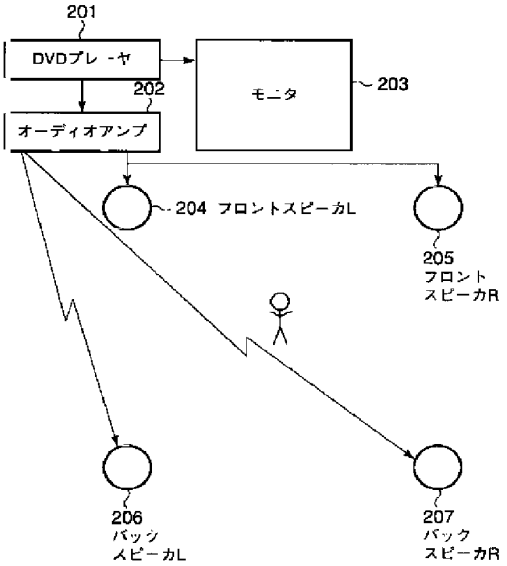


送信開始可能区間
? 再生開始基準点
T 再生ボタン押下時刻
S 音声データ送信開始時刻
R 最初の音声データが到着して再生可能となった時刻
P0 最も早い再生開始時刻
PX 一般の再生開始時刻(PX=P0+Beta)
Period 再生開始基準点の間隔
delayMax (伝送遅延+処理遅延)の最大値
Alpha =Period−delayMax
Beta 0以上の定数。再生開始時刻をP0以降に遅延させる値

【図8】



【図9】



(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
19 September 2002 (19.09.2002)

PCT

(10) International Publication Number
WO 02/073851 A1

(51) International Patent Classification[7]: H04J 3/06

(21) International Application Number: PCT/US02/07082

(22) International Filing Date: 7 March 2002 (07.03.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/805,735 13 March 2001 (13.03.2001) US

(71) Applicant: **PULSE-LINK, INC.** [US/US]; 9155 Brown Deer Road, Suite 8, San Diego, CA 92121 (US).

(72) Inventor: **TAYLOR, James, L.**; 8 Allee de la Grenouillere, 78680 Epone (FR).

(74) Agent: **BRISSON, Timothy A.**; SIERRA PATENT GROUP, LTD., P.O. Box 6149, Stateline, NV 89449 (US).

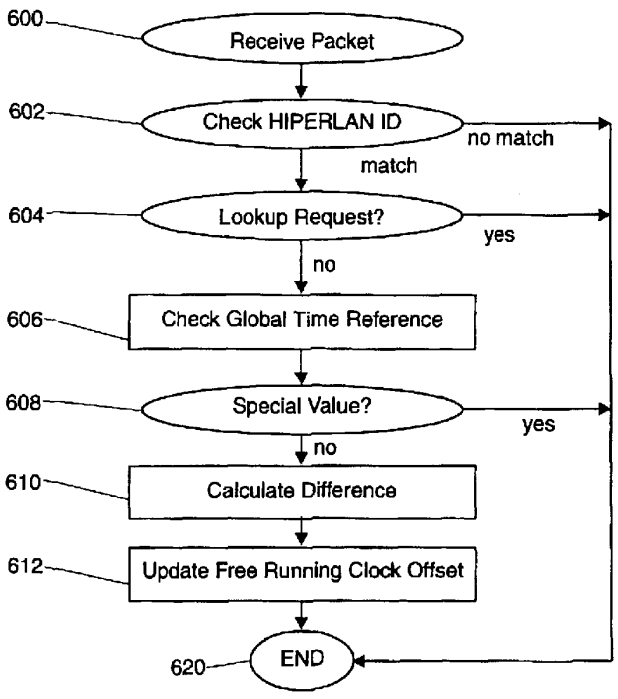
(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*

[Continued on next page]

(54) Title: MAINTAINING A GLOBAL TIME REFERENCE AMONG A GROUP OF NETWORKED DEVICES



(57) Abstract: Synchronization is maintained among a plurality of network devices (102,...) having local clocks. A first packet including global time reference (606) derived from local clock is broadcast from a first network device to other devices. The clocks of the network devices are adjusted. The difference between a free running clock and a first network global time reference is determined (610). A first network offset is calculated (612). Similarly, a second network offset is calculated to account for the difference between the free running clock and the second network global time reference.

WO 02/073851 A1

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

MAINTAINING A GLOBAL TIME REFERENCE AMONG A GROUP OF NETWORKED DEVICES

FIELD OF THE INVENTION

The present invention relates generally to networks of devices. More specifically, maintaining a global time reference in an ad hoc wireless network is disclosed.

BACKGROUND OF THE INVENTION

Wireless networks are becoming increasingly important as solutions for transferring data among devices that move, are located in a home, or are otherwise situated so that it is desirable to connect the devices without wires. Numerous physical and MAC layer solutions have been proposed. For example, a baseband ultra wideband wireless network is described in U.S. Patent Application No. 09/393,126 "Baseband Wireless Network for Isochronous Communication" by Aiello, et. al., which is herein incorporated by reference. An ultra wide band network transmits a signal that covers a very large portion of the frequency spectrum but transmits at a very low power in any individual frequency band.

Wireless networks are sometimes referred to as ad hoc networks because the network is comprised of devices that happen to be within the broadcast vicinity of the other devices in the network. Devices may join or leave the network as they are powered up or move out of range. A particular ad hoc network is described in the HIPERLAN Type 1, Functional Specification, ETS 300 652 published by the European Telecommunications Standards Institute. It is available at www.etsi.org.

Figure 1 is a block diagram illustrating a set of wireless devices that communicate together in a wireless ad hoc network. Video player 102 transmits an audio signal to speakers 106a, 106b, 106c and 106d. Video player 102 also transmits a video signal to screen 104. Screen 104 and speakers 106 may also transmit information to each other or to video player 102.

For an application such as the one shown in Figure 1, it is readily seen that synchronization among the various devices is important to achieve a desirable result. Different amounts of delay or jitter may exist in the signal paths between video player 102, speakers 106a and 106b, and screen 104. If the difference in delay is significant, then the

video presentation shown on screen 104 may not be properly synchronized with the audio being played on the speakers. Even worse, if speaker 106a and speaker 106b are not properly synchronized, then the left and right stereo audio channels could become separated in time, resulting in an unpleasant audio experience. Thus, it would be useful if a method of synchronizing the devices in such a wireless network could be developed.

SUMMARY OF THE INVENTION

Accordingly, a system and method for maintaining a global time reference among a group of wireless devices participating in a wireless network is disclosed. It should be noted that the techniques described herein may also be applied to other networks including certain networks that transmit signals partially or wholly using wires, fiberoptic cables or other means. For the purposes of example, this specification will describe in detail a wireless network and in particular a wireless ad hoc network that uses the HIPERLAN standard. It should be appreciated that the invention is not limited to such a network and that the techniques described are applicable to many different types of networks.

In one embodiment, each device in a network maintains its own clock. Instead of designating the clock of one particular device as the master clock, each device transmits its own clock signal or part of its own clock signal when it sends a packet to other devices in the network. The devices receiving the packet read the clock signal and adjust their own clocks to the transmitted value. This clock adjustment is made in a manner that causes the clocks to smoothly approach a common value without oscillating in an unstable manner.

It should be appreciated that the present invention can be implemented in numerous ways, including as a process, an apparatus, a system, a device, a method, or a computer readable medium such as a computer readable storage medium or a computer network wherein program instructions are sent over optical or electronic communication links. Several inventive embodiments of the present invention are described below.

In one embodiment, synchronization is maintained among a plurality of network devices having local clocks that participate in a network. A first packet is broadcast from a first network device to other network devices that participate in the network. The first packet includes a global time reference derived from the local clock of the first network device. The clocks of the network devices that receive the first packet are adjusted to be closer to the local clock of the first network device.

In one embodiment, a first network local time reference and a second network local time reference are maintained for a device that participates in a first network and a second network. A free running clock is maintained on the device. The difference between the free running clock and a first network global time reference is determined. A first network offset is calculated to account for the difference between the free running clock and the first network global time reference. The difference between the free running clock and a second network global time reference is determined. A second network offset is calculated to account for the difference between the free running clock and the second network global time reference.

These and other features and advantages of the present invention will be presented in more detail in the following detailed description and the accompanying figures which illustrate by way of example the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

Figure 1 is a block diagram illustrating a set of wireless devices that communicate together in a wireless ad hoc network.

Figure 2 is a diagram illustrating in some detail how a channel may introduce delay and jitter into the timing of packets sent from a transmitter to a receiver.

Figure 3 is a diagram illustrating how the timestamp, referred to as the global time reference is inserted into a packet header.

Figure 4A is a diagram illustrating how the global time reference is split into a most significant part and a least significant part.

Figure 4B is a diagram illustrating a memory register that stores the free running clock. The free running clock is a 32 bit clock.

Figure 4C is a diagram illustrating a memory register that stores the free running clock offset.

Figure 5 is a diagram illustrating an embodiment where devices belong to two different ad hoc networks and a new device is joining one of the networks.

Figure 6 is a flow chart illustrating a process that a network device executes when a

packet that includes a global time reference is received to determine whether to adjust the network device local time reference.

Figure 7A is a flow chart illustrating a process for including a global time reference in a packet.

Figure 7B is a flow chart illustrating a process executed by a receiver to prepare to update its free running clock offset.

DETAILED DESCRIPTION

A detailed description of a preferred embodiment of the invention is provided below. While the invention is described in conjunction with that preferred embodiment, it should be understood that the invention is not limited to any one embodiment. On the contrary, the scope of the invention is limited only by the appended claims and the invention encompasses numerous alternatives, modifications and equivalents. For the purpose of example, numerous specific details are set forth in the following description in order to provide a thorough understanding of the present invention. The present invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the present invention is not unnecessarily obscured.

Figure 2 is a diagram illustrating in some detail how a channel may introduce delay and jitter into the timing of packets sent from a transmitter to a receiver. Transmitter 200 sends packets in an isochronous manner. That is, each packet begins transmission at a regular interval. The difference in time between t_1, t_2, t_3, and t_4 is constant. Channel 202 introduces delay and jitter into the transmission so that Δ_1, Δ_2, Δ_3, and Δ_4 time delays are added to each of the times that the respective packets are received by the receiver 204. Thus, the evenly spaced transmitted packets are not evenly spaced when they are received.

So long as there is only one transmitter and one receiver, the inter packet timing can be corrected at the receiver if a timestamp is included at the beginning of each packet. The difference between each of the timestamps would indicate the correct relative time spacing of the packets enabling the receiver to reposition or perhaps even reorder the packets in time.

When there are more than two devices, then simply timestamping the packets will

not synchronize the display of data among the devices. It should also be noted that even when there are only two devices, the timestamp can adjust the relative spacing between packets but it does not facilitate making sure that events occur on the receiver at a specific time relative to an event occurring at the transmitter. Referring back to the example shown in Figure 1, even if packets sent from video player 102 to left and right speakers 106a and 106b are timestamped, there may still be a time offset between the left and right stereo channels. The timestamp on the packets only serves to enable each of the channels to correctly position in time the packets that it receives.

In order to synchronize the video or audio display of information received in packets over the wireless network or otherwise coordinate the use of information received from packets, a clock is maintained on each of the network devices that is in synchronization with the clocks on each of the other network devices.

In one embodiment, a clock is specified for each network device that needs to be correct only to about 5 parts per million. In order to maintain acceptable synchronization between all of the device clocks on the network, each time one of the devices transmits a packet, the device includes in the packet header at a well defined point the value or, as described below, part of the value of its clock. When the other devices receive that packet, they adjust each of their respective clocks toward the time indicated by the timestamp. As described below, in one embodiment, the time adjustment is implemented by calculating a clock offset and the clock offset is calculated by taking half the difference between the device clock and the timestamp.

Figure 3 is a diagram illustrating how the timestamp, referred to as the global time reference is inserted into a packet header. Packet header 300 includes a clock sync sequence 302. Clock sync sequence 302 is chosen to be an appropriate sequence to facilitate the receiver acquiring a clock signal used to sample the remainder of the packet. Frame sync sequence 304 is included next for the purpose of precisely identifying the beginning of encoded data in the packet. Thus, the end of the frame sync sequence precisely determines the beginning of a frame. Any number of well known frame sync sequences may be used. In one embodiment, a 13 bit Barker code is used. As soon as the frame sequence is sent, the transmitter records its local time reference.

It should be noted that ideally, the local time reference of the transmitter will agree with the local time reference kept at each of the devices participating in the network. If that

is the case, then each of the local time references actually is a global time reference because it is in agreement with every other local time reference in the system. For the purpose of this explanation, the local time reference that the transmitter records and then inserts in the packet will be referred to as the global time reference. That is because the transmitter asserts it to be the global time reference and it is the time that all receiving devices will adjust their own local time references toward. It should also be noted that each device considers its own local time reference to be the global time reference and asserts that its own local time reference is the global time reference it transmits a packet. Thus, each transmitting device asserts its local time reference as the global time reference when it includes its local time reference in a packet that it broadcasts to the other devices. Over time, as enough packets are sent, all of the local time references converge to the same global time and the broadcast global time reference results in only small, if any, changes by receiving devices.

Global time reference 306 is inserted in the packet after frame sync sequence 304. In the embodiment shown, the global time reference is inserted immediately after the frame sync sequence, but that is not required. The receiver notes its local time reference at the time when the frame sync sequence is detected and the frame sync sequence signal is asserted. The global time reference from the packet may be read immediately or it may be read much later, so long as it is compared to the time when the frame sync sequence signal is asserted. Thus, global time reference 306 may occur in the packet immediately following frame sync sequence 304 or may occur at a later point in the packet. The remainder of the packet 308 is then sent. The remainder of the packet will likely include further information for the packet header as well as finally the packet payload of data.

It should be noted that in one embodiment, the entire global time reference is not included in the packet. Instead, only the least significant bits, referred to as the least significant part, of the global time reference is included. It is sufficient to include just the least significant part of the global time reference so long as the difference between the clocks is less than the amount of time counted by the least significant part of the global time reference. This condition is met if the amount of drift between the clocks before an adjustment is made is not enough to move the clocks out of synchronization by more than the maximum amount of time that can be represented by the least significant part. In one embodiment, a 32 bit global time reference is used and the least significant part that is

transmitted with each packet is the least significant 12 bits of the 32 bit global time reference.

In the embodiment described above, the relative clock drift of the clocks being used is about 5 parts per million (ppm) and the clock rate is 1 MHz. The least significant part of the global time reference can count up to 2^{12} microseconds. Therefore, the clocks would need to be updated before a 2^{12} microsecond error could accumulate. For a 1 MHz clock with 5 ppm accuracy, the maximum time interval between updates would be on the order of about 5 minutes, depending on the method of updating used. Of course, the maximum time interval varies in different systems depending on the speed and accuracy of the clock and the number of bits transmitted with the least significant part of the clock. If more of the global time reference is provided in the least significant part, then the frequency of update is lower. For example, if 14 bits are used, then the least significant part of the global time reference can accurately represent a cumulative drift of up to 2^{14} microseconds. Assuming that the relative clock drift is unchanged, then updates would be required only one fourth as often.

One advantage of the described arrangement where the global time reference is included along with every packet sent by one of the devices participating in the network is that as the amount of network traffic increases, updates of the local time references occur more frequently. This decreases the variance among the local time references.

Figure 4A is a diagram illustrating how the global time reference is split into a most significant part and a least significant part. As mentioned above, the most significant part is, in one embodiment, 20 bits and the least significant part is 12 bits. The size of the least significant part may be increased to reduce the required frequency of updates at the cost of requiring more overhead for transmission. It should also be noted that, in addition to the least significant part being included in packets that are sent by each of the devices, periodically a MAC layer protocol prompts one or more of the devices to broadcast the entire global time reference, that is both the most significant part and the least significant part. In one embodiment, the global time reference is broadcast every 20 seconds. The reason for periodically broadcasting the entire global time reference is to allow for an error that occurred in one of the devices that changed the global time reference by more than the maximum time represented by the least significant part of the global time reference to be corrected.

In different embodiments, the task of broadcasting the entire global time reference may be performed by either one or all of the devices participating in the network. For example, each device may be directed to broadcast the global time reference periodically. The periodic broadcast may be staggered. Alternatively, the MAC layer protocol may direct a designated device to broadcast the entire global time reference at certain intervals.

Figure 4B is a diagram illustrating a memory register that stores the free running clock. The free running clock is a 32 bit clock. As mentioned previously, in one embodiment, the clock is a 1 megahertz clock. In other embodiments, the clock rate may be set as desired. Preferably, the free running clock is read only. That is, adjustments to the global time reference are not written to the clock. Rather, a free running clock offset is defined such that the global time reference is equal to the sum of the free running clock plus the free running clock offset.

Figure 4C is a diagram illustrating a memory register that stores the free running clock offset. Like the free running clock, the free running clock offset is a 32 bit register. However, the free running clock offset may be written as well as read. Each time a global time reference is received from one of the other devices, the free running clock is adjusted so that the sum of the free running clock and the free running clock offset converges toward the global time reference specified by the broadcasting device.

In different embodiments, adjustments to the local time reference based on the received global time reference are made according to different schemes. In one embodiment, a linear adjustment to the local time reference is made based on the difference between the received global time reference and the local time reference. The adjustment may be represented by the equation:

$$\Delta FRCO=(GTR-LTR) * L.$$

Where ΔFRCO is the change in the free running clock offset. GTR is the global time reference received with a packet. LTR is the current value of the local time reference stored by the device. As mentioned above, after several packets are received and adjustments are made, the local time references stored on each of the devices will converge toward a single value and that value will be the global time reference. L is a linear factor that adjusts the amount that the free running clock offset is adjusted compared to the difference between the global time reference and the local time reference that is computed. In one embodiment, L is equal to .5 and the adjustment to the free running clock offset is equal to half the

difference between the received global time reference and the stored local time reference.

Adjustment schemes are used in different embodiments that make non-linear adjustments to the free running clock offset based on the difference between the global time reference and the local time reference. In one embodiment, the change to the free running clock offset is equal to a scaling factor times the square of the difference between the global time reference and the local time reference. Squaring the difference between the two references has the advantage of making greater changes to the free running clock offset if the difference between the global time reference and the local time reference is large. If the square of the difference is used, then the sign of the difference between the global time reference and the local time reference is computed and applied to the difference so that the free running clock offset is adjusted in the proper direction.

In another embodiment, another non-linear adjustment scheme is implemented to allow the system to be more robust. The difference between the global time reference and the local time reference is computed. If the difference is greater than a threshold, then no adjustment is made to the local time reference. If the difference is less than a threshold, then the adjustment is made. Thus, if one of the devices in the network malfunctions or a new device signing on to the network attempts to assert a global time reference that is incorrect and would perturb the global time reference agreed upon by the current devices in the network, that device's assertion of the global time reference may be ignored.

In some embodiments, the threshold test can be turned off under certain circumstances. For example, if more than a certain number of changes are rejected by a device, then the threshold function may be temporarily turned off. This would be appropriate since a device that is receiving a large number of global time references that differ from its local time reference by a great amount most likely is actually incorrect by a large amount. A counter may be used to count the number of recent times that the difference between a received global time reference and the local time reference exceeds the threshold. In one embodiment, the counter is programmed to only count global time references from unique sources so that a single inaccurate source broadcasting repeatedly will be ignored. Both the threshold and the number of times that the threshold must be exceeded to turn off the threshold test may be programmed or may be preset.

Figure 5 is a diagram illustrating an embodiment where devices belong to two different ad hoc networks and a new device is joining one of the networks. Devices 502,

504, 506, and 508 all belong to a first ad hoc network 530. Device 508 along with devices 512 and 514 also belong to an ad hoc network 540. Device 522 is requesting to join network 530. As mentioned above, in one embodiment, the HIPERLAN protocol is utilized by the wireless ad hoc networks. In such an embodiment, each network has its own HIPERLAN ID. The HIPERLAN ID is included in packets sent by devices.

The two separate networks can have two separate independent clocks with a device such as device 508 that participates in both networks maintaining two separate clocks. More precisely, the device need only maintain a single free running clock and keep track of two free running clock offsets. Changes are applied to each of the respective free running clock offsets to keep track of the two different global times for the networks. Thus, device 508 may have a single free running clock but two registers used to calculate two different free running clock offsets to adjust the free running clock time to the global time of whichever network it is participating in at a given moment.

Other aspects of the HIPERLAN specification are used to facilitate the implementation of a global time reference shared among the devices of an ad hoc network. For example, the entire global time reference may be periodically broadcast using a Hello packet as described in the specification. The standard Hello procedure is modified to carry the full global time reference. The Hello procedure is described more fully in the HIPERLAN Type 1, Functional Specification, ETS 300 652 at 6.5.4 Neighbour Information Declaration. See also 6.1.14 Neighbour Information Base and 6.1.15 Hello Information Base.

As described above, a device may use the HIPERLAN ID to determine whether a packet belongs to its network and therefore whether to adjust its local time reference toward the global time reference specified in the packet. Also, a reserved network management ID may be used by devices joining the network. Network devices may be configured to ignore time references included in HIPERLAN lookup requests.

Various methods may be implemented to avoid a stable network global time reference from being perturbed by a new device. In one embodiment, new devices are instructed to provide a global time reference of all 0's which is a special value that is ignored by all network devices. Also, a new device is preferably configured to initially set its local time reference to equal a received global time reference instead of adjusting its local reference as described above. Specification of conventions for new devices, however,

does not protect against rogue devices that do not observe the correct conventions or devices that belong to other networks. For that reason, other procedures such as checking the HIPERLAN ID or rejecting certain large changes may also be implemented.

Figure 6 is a flow chart illustrating a process that a network device executes when a packet that includes a global time reference is received to determine whether to adjust the network device local time reference. The process starts at 600. In step 602, the HIPERLAN ID is checked. If there is no match, then the process ends at 620. If there is a match, then control is transferred to step 604 and it is determined whether the packet is a look up request. If the packet is a look up request, then the process ends at 620. If the packet is not a lookup request, then control is transferred to step 606 and the global time reference is checked. In step 608, it is determined whether the global time reference is a special value. If it is, then control is transferred to 620 and the process ends. If the global time reference is not a special value, then control is transferred to step 610 and the difference between the global time reference and the local time reference is calculated. Then, in step 612, the free running clock offset is updated. The free running clock may be updated according to any of the linear or nonlinear schemes described above or any other appropriate scheme. The process then ends at 620.

Figure 7A is a flow chart illustrating a process for including a global time reference in a packet. The process starts at 700 when a clock sync sequence is written. Next, in step 702, the frame sync sequence is written. As soon as the frame sync sequence is completed the local time reference of the transmitting device is read in step 704. This value will be inserted into the packet as the global time reference corresponding to the time that the frame sync sequence was finished. Next, in step 706, the global time reference is written into the packet. It should be noted that other steps may intervene between steps 704 and 706. The key point is that the global time reference is stored either immediately after or at some specific time after the frame sync sequence is completed so that the receiving device can compare its local time reference to the transmitted global time reference at a specific reference point.

Figure 7B is a flow chart illustrating a process executed by a receiver to prepare to update its free running clock offset. The process starts at 710 when the frame sync sequence is detected. The local time is immediately stored in step 712. It should be noted that although it is preferred that the local time be stored immediately, it is also possible to

store the local time at a set period of time after the assertion of the frame sync signal. Then, in step 714, the global time reference is read from the packet. In this manner, the transmitted global time and the local time are stored at a specific point in time referenced to the writing of the frame sync sequence and the detection of the frame sync sequence, respectively. In this manner, the local time reference and the global time reference can be compared and the free running clock offset can be appropriately adjusted.

Maintaining a global time reference in an ad hoc network has been disclosed. Each device maintains a local free running clock and a local free running clock offset, which together comprise a local time reference. Each free running clock is adjusted toward a received global time reference that is the local time reference of whatever device happens to be broadcasting. In this manner, eventually all of the local time references of devices in the network converge. A device may participate in more than one network by keeping track of one free running clock offset for each network.

Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. It should be noted that there are many alternative ways of implementing both the process and apparatus of the present invention. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

WHAT IS CLAIMED IS:

1. A method of maintaining synchronization among a plurality of network devices having local clocks that participate in a network comprising:

broadcasting a first packet from a first network device to other network devices that participate in the network wherein the first packet includes a global time reference derived from the local clock of the first network device; and

adjusting the clocks of the network devices that receive the first packet to be closer to the local clock of the first network device.

2. A method of maintaining synchronization among a plurality of network devices as recited in claim 1 further including:

broadcasting a second packet from a second network device to other network devices that participate in the network wherein the second packet includes a global time reference derived from the local clock of the second network device; and

adjusting the clocks of the other network devices network devices that receive the second packet to be closer to the local clock of the second network device.

3. A method of maintaining synchronization among a plurality of network devices as recited in claim 1 wherein the global time reference includes the least significant part of the local clock of the first network device and does not include the most significant part of the local clock of the first network device.

4. A method of maintaining synchronization among a plurality of network devices as recited in claim 3 wherein the most significant part of one of the network device local clocks is broadcast periodically and the most significant parts of all local clocks are conformed to the broadcast most significant part.

5. A method of maintaining synchronization among a plurality of network devices as recited in claim 1 wherein adjusting the clocks of the network devices that receive the first packet to be closer to the local clock of the first network device includes adjusting the clocks of the network devices according to a linear function of the difference between the unadjusted value of the clock being adjusted and the broadcast global time reference.

6. A method of maintaining synchronization among a plurality of network devices as recited in claim 1 wherein adjusting the clocks of the network devices that receive the first packet to be closer to the local clock of the first network device includes adjusting the

clocks of the network devices by approximately one half the difference between the broadcast global time reference and the unadjusted value of the clock being adjusted.

7. A method of maintaining synchronization among a plurality of network devices as recited in claim 1 wherein adjusting the clocks of the network devices that receive the first packet to be closer to the local clock of the first network device includes adjusting the clocks of the network devices according to a nonlinear function of the difference between the unadjusted value of the clock being adjusted and the broadcast global time reference.

8. A method of maintaining synchronization among a plurality of network devices as recited in claim 7 wherein the nonlinear function of the difference between the unadjusted value of the clock being adjusted and the broadcast global time reference causes substantially no adjustment to the clock being adjusted when the difference between the unadjusted value of the clock being adjusted and the broadcast global time reference is greater than a maximum adjustable difference.

9. A method of maintaining synchronization among a plurality of network devices as recited in claim 1 wherein the global time reference is derived from the local clock of the first network device immediately after a frame synchronization portion of a packet is transmitted.

10. A method of maintaining synchronization among a plurality of network devices as recited in claim 1 wherein adjusting the clocks of the network devices that receive the first packet to be closer to the local clock of the first network device includes comparing the local clock of a receiving network device at the instant that frame synchronization is detected to the transmitted global time reference.

11. A method of maintaining synchronization among a plurality of network devices as recited in claim 1 wherein all data packets broadcast by any network device that participates in the network includes a global time reference derived from the local clock of the broadcasting network device so that as more data is transmitted, the local clocks of the receiving network devices are updated more frequently.

12. A method of maintaining a first network local time reference and a second network local time reference for a device that participates in a first network and a second network comprising:

maintaining a free running clock on the device;

determining the difference between the free running clock and a first network global

time reference;

calculating a first network offset to account for the difference between the free running clock and the first network global time reference;

determining the difference between the free running clock and a second network global time reference; and

calculating a second network offset to account for the difference between the free running clock and the second network global time reference.

13. A receiving network device configured to maintain synchronization with a transmitting network device comprising:

a receiving network device local clock;

an interface configured to receive a first packet from the transmitting network device wherein the first packet includes a global time reference derived from a transmitting network device local clock; and

a processor configured to adjust the receiving network device local clock to be closer to the transmitting network device local clock.

14. A receiving network device configured to maintain synchronization with a transmitting network device as recited in claim 13 wherein the receiving network device local clock includes a free running clock and a free running clock offset.

15. A transmitting network device configured to maintain synchronization with a receiving network device comprising:

a transmitting network device local clock;

an interface configured to transmit a first packet wherein the first packet includes a global time reference derived from the transmitting network device local clock and wherein the global time reference is used by the receiving network device to adjust the receiving network device local clock to be closer to the transmitting network device local clock.

16. A network of devices comprising:

a transmitting network device configured to transmit a first packet wherein the first packet includes a global time reference derived from a transmitting network device local clock; and

a receiving network device configured to receive the first packet from the transmitting network device and to adjust a receiving network device local clock to be closer to the global time reference.

17. A network device comprising:

a free running clock;

a processor configured to:

determine the difference between the free running clock and a first network global time reference;

calculate a first network offset to account for the difference between the free running clock and the first network global time reference;

determine the difference between the free running clock and a second network global time reference; and

calculate a second network offset to account for the difference between the free running clock and the second network global time reference.

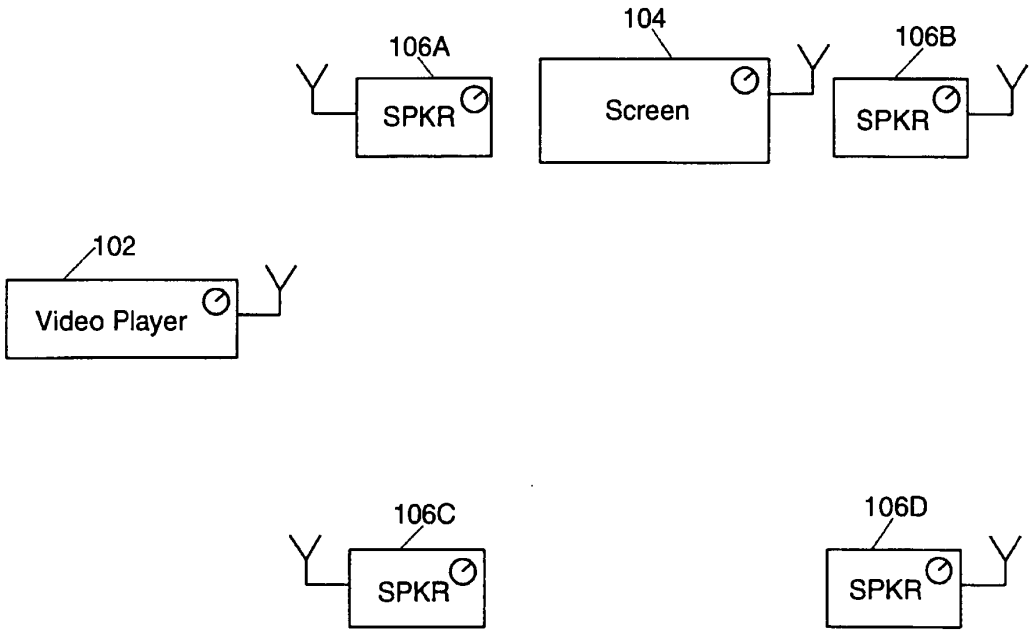

106A
SPKR
104
Screen
106B
SPKR
102
Video Player
106C
SPKR
106D
SPKR

FIG. 1

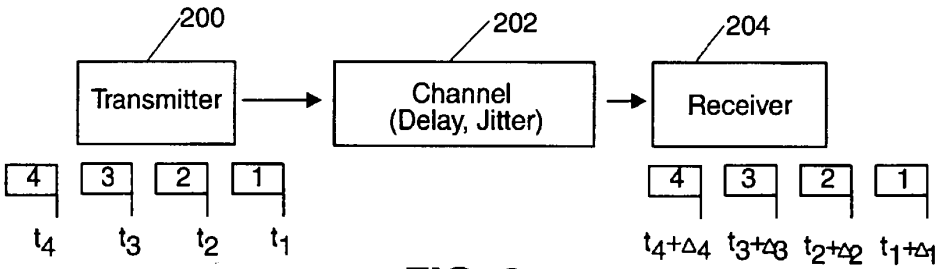

200
202
204
Transmitter
Channel
(Delay, Jitter)
Receiver
4
3
2
1
t_4 t_3 t_2 t_1
4
3
2
1
$t_4+\Delta_4$ $t_3+\Delta_3$ $t_2+\Delta_2$ $t_1+\Delta_1$

FIG. 2

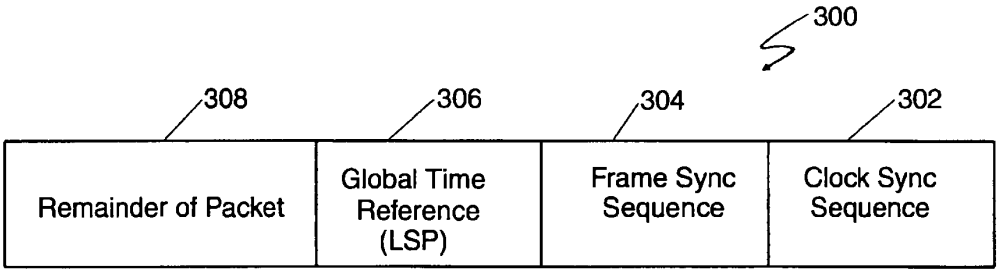

300
308
306
304
302
Remainder of Packet
Global Time Reference (LSP)
Frame Sync Sequence
Clock Sync Sequence

FIG. 3

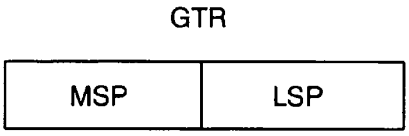


FIG. 4A

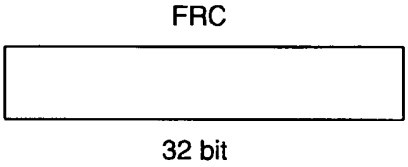
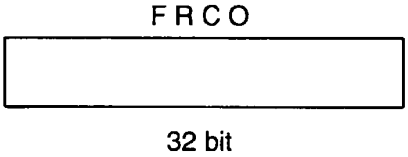


FIG. 4B



GTR = FRC + FRCO

FIG. 4C

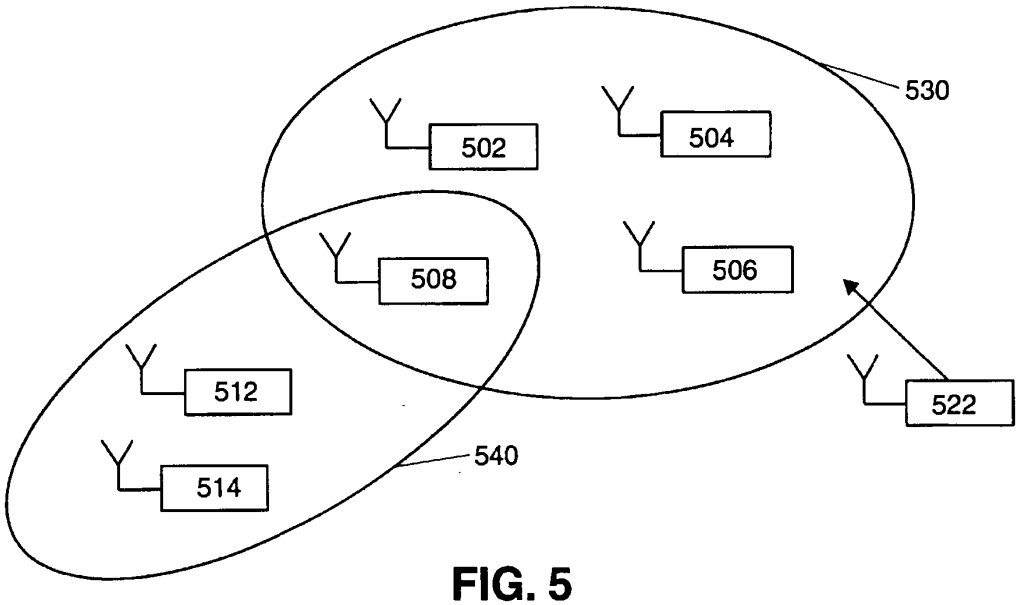


FIG. 5

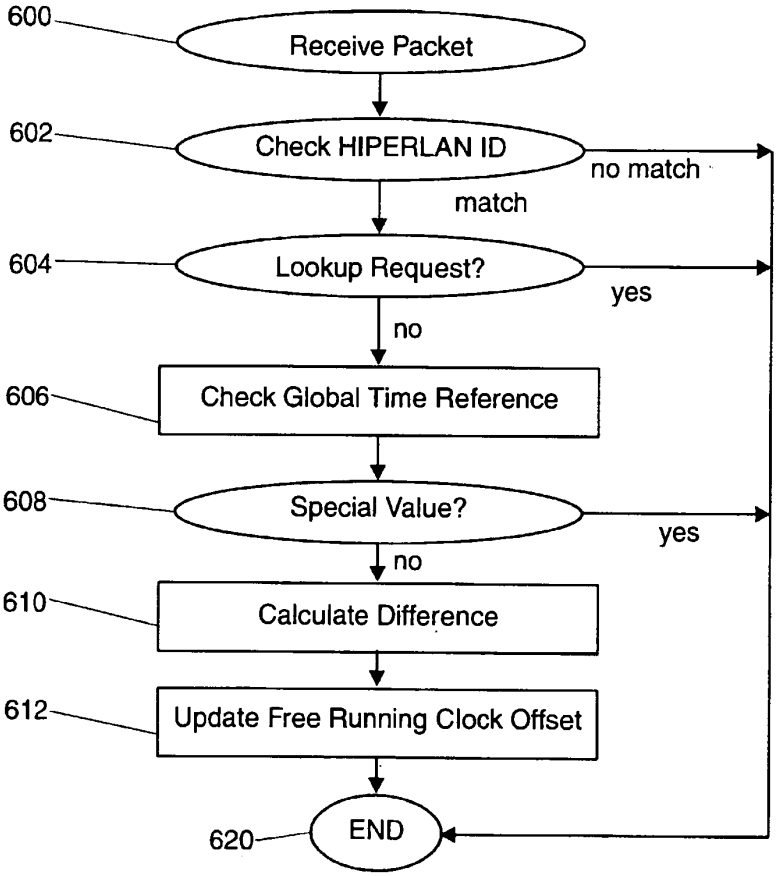

600
Receive Packet
602
Check HIPERLAN ID
no match
match
604
Lookup Request?
yes
no
606
Check Global Time Reference
608
Special Value?
yes
no
610
Calculate Difference
612
Update Free Running Clock Offset
620
END

FIG. 6

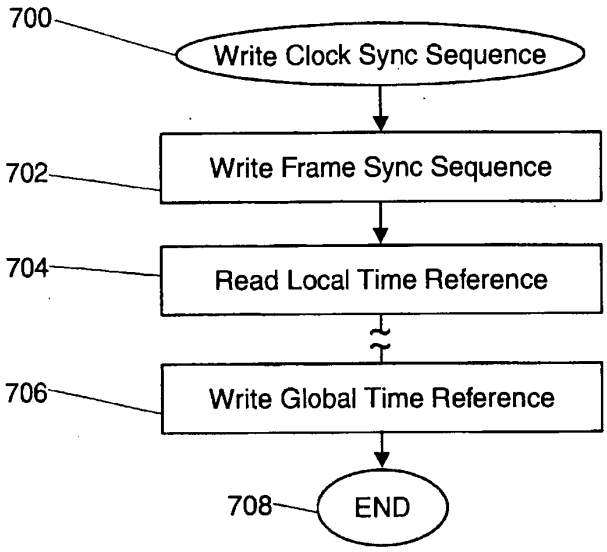

700
Write Clock Sync Sequence
702
Write Frame Sync Sequence
704
Read Local Time Reference
706
Write Global Time Reference
708
END

FIG. 7A

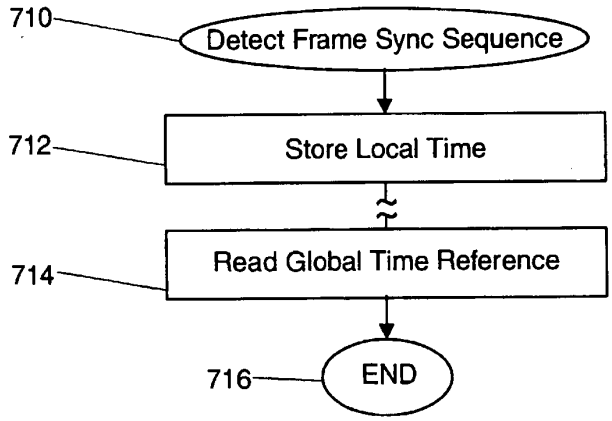

710
Detect Frame Sync Sequence
712
Store Local Time
714
Read Global Time Reference
716
END

FIG. 7B

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US02/07082

A. CLASSIFICATION OF SUBJECT MATTER
IPC(7) : H04J 3/06
US CL : 370/350, 432, 503, 516; 375/356; 455/265
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
U.S. : 370/350, 390, 432, 503, 516; 375/356; 455/208, 265; 709/248, 249

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EAST, WEST, IEEE, DIALOG

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y,P	US 5,334,975 A (WACHOB et al.) 02 August 1994, col. 1, lines 35-52, col. 2, lines 10-68, col. 3, lines 1-40.	1-17
Y,P	US 6,285,873 B1 (QUICK, Jr.) 04 September 2001, col. 4, lines 5-19, col. 6, lines 31-67, col. 7, lines 1-67, col. 8, lines 1-7, 54-63	1-17
Y,P	US 6,347,084 B1 (HULYALKAR et al.) 12 February 2002, col. 1, lines 37-48, 66-67, col. 2, lines 1-12, col. 6, lines 53-67, col. 7, lines 1-4, col. 8, lines 21-32, 43-64.	1-17

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
13 June 2002 (13.06.2002)	11 JUL 2002
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703)305-3230	Authorized officer Afsar M. Qureshi Telephone No. (703) 308 8522

Form PCT/ISA/210 (second sheet) (July 1998)

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	13864248
Filing Date	2013-04-17
First Named Inventor	MILLINGTON, Nicholas A. J.
Art Unit	2442
Examiner Name	SURVILLO, OLEG
Attorney Docket Number	SNS-71301US10 (13-0401)

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20020002565		2002-01-03	OHYAMA; Akira	
	2	20020062406		2002-05-23	CHANG; William Ho et al.	
	3	20100004983		2010-01-07	DICKERSON; Stephanie et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	MILLINGTON, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301US10 (13-0401)

If you wish to add additional Foreign Patent Document citation information please click the Add button			Add
NON-PATENT LITERATURE DOCUMENTS			Remove
Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T5
	1	"Final Office Action mailed 05 June 2014 for U.S. Application No. 13/907666 filed 31 May 2013".	☐
	2	"Non-Final Office Action mailed 05 January 2012 for U.S. Patent Application No. 13/298090 filed on 16 November 2011, 40 pages".	☐
	3	"Non-Final Office Action mailed 17 June 2014 for U.S. Application No. 14/176808 filed 10 February 2014".	☐
	4	"Non-Final Office Action mailed 18 December 2013 for U.S. Application No. 13/907666 filed 31 May 2013".	☐
	5	"Non-Final Office Action mailed 27 May 2014 for U.S. Application No. 14/186850 filed 21 February 2014".	☐
	6	"Yamaha DME Designer software manual: Copyright 2004", p.1-250.	☐
	7	"Yamaha DME Designer software manual: Copyright 2004", p.251-482.	☐
If you wish to add additional non-patent literature document citation information please click the Add button			Add

EXAMINER SIGNATURE			
Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	MILLINGTON, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301US10 (13-0401)

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	MILLINGTON, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301US10 (13-0401)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Robert J. Irvine III/	Date (YYYY-MM-DD)	2014-07-02
Name/Print	Robert J. Irvine III	Registration Number	41865

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	19483948
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 (13-0401)
Receipt Date:	07-JUL-2014
Filing Date:	17-APR-2013
Time Stamp:	21:36:40
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Non Patent Literature	13907666-FinalOA060514.pdf	407537 5cee0dff8d284ab99151703c816e7d2d532efd42	no	12

Warnings:	
Information:	

2	Non Patent Literature	13298090-Non-FinalOA010512.pdf	1401566 8a276db39b354e4c427f17b1fd0b422a7f7bf915	no	40
Warnings:					
Information:					
3	Non Patent Literature	14176808-FinalOA61714.pdf	230411 ad6e4f625c6d90b5e69ec1ddace354ad53b22705	no	6
Warnings:					
Information:					
4	Non Patent Literature	13907666-Non-FinalOA121813FIXED.pdf	2878882 6dbe0eebc708b6c0f5e1a9df2045902e4545b544	no	12
Warnings:					
Information:					
5	Non Patent Literature	14186850-NonfinalOA-052714.pdf	373755 af0dca4c73991941d17e5f87060228aa6b03e388	no	12
Warnings:					
Information:					
6	Non Patent Literature	Yamaha2004-1-250.pdf	8668316 7480c761eeb987ccaf98eddc6a844fc7b068245b	no	250
Warnings:					
Information:					
7	Non Patent Literature	Yamaha2004-251-482.pdf	10638326 aa8aea22efb9440f82e4bd8158b52225ddf1cc7d	no	232
Warnings:					
Information:					
8	Information Disclosure Statement (IDS) Form (SB08)	SNS-71301US10-SuppIDS-JUNE2014.pdf	706970 cc09bef4b6252a38222abb4623255100f993d837	no	5
Warnings:					
Information:					
		Total Files Size (in bytes):	25305763		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	13864248
Filing Date	2013-04-17
First Named Inventor	MILLINGTON, Nicholas A. J.
Art Unit	2442
Examiner Name	SURVILLO, OLEG
Attorney Docket Number	SNS-71301US10 (13-0401)

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20020095460		2002-07-18	BENSON; Michael	
	2	20050168630		2005-08-04	YAMADA; Norihiko et al.	
	3	20080126943		2008-05-29	PARASNIS; Shashank Mohan et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	MILLINGTON, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301US10 (13-0401)

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T5
	1	"Final Office Action mailed 27 March 2014 for U.S. Application No. 13/533105 filed 26 June 2012".	☐
	2	"Non-Final Office Action mailed 01 May 2014 for U.S. Application No. 14/184522 filed 19 February 2014".	☐
	3	"Non-Final Office Action mailed 06 May 2014 for U.S. Application No. 13/705176 filed 05 December 2012".	☐
	4	"Non-Final Office Action mailed 12 May 2014 for U.S. Application No. 14/184528 filed 19 February 2014".	☐
	5	"Non-Final Office Action mailed 14 May 2014 for U.S. Application No. 13/848932 filed 22 March 2013".	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	MILLINGTON, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301US10 (13-0401)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Robert J. Irvine III/	Date (YYYY-MM-DD)	2014-05-21
Name/Print	Robert J. Irvine III	Registration Number	41865

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	19088299
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 (13-0401)
Receipt Date:	23-MAY-2014
Filing Date:	17-APR-2013
Time Stamp:	21:44:32
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	SNS-71301US10SupplementalIDS-MAY2014.pdf	706931 de4453ecfb25855cf5ee399c676147c7dd3f0f15	no	4

Warnings:	
Information:	

2	Non Patent Literature	13533105FinalOA32714.pdf	951725 50db79854ca7b844590e839e6dd5cebef2e9a089	no	29
Warnings:					
Information:					
3	Non Patent Literature	14184522Non-FinalOA5114.pdf	1042059 3e0d22e91cd5d85ad3f946ad4a7897db11c7df09	no	31
Warnings:					
Information:					
4	Non Patent Literature	13705176Non-FinalOA5614.pdf	824173 c505030fa0f9e7ca2746c7bcf8c561a5a580308e	no	23
Warnings:					
Information:					
5	Non Patent Literature	14184528Non-FinalOA51214.pdf	784767 cedc7adc69fbf2021f105c90f100487cc3feb203	no	23
Warnings:					
Information:					
6	Non Patent Literature	13848932Non-FinalOA51414.pdf	461919 e2e7f62a13e4121e1416e89a9fd5fc441e0196d4	no	14
Warnings:					
Information:					
		Total Files Size (in bytes):	4771574		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	13864248
Filing Date	2013-04-17
First Named Inventor	Millington, Nicholas A. J.
Art Unit	2442
Examiner Name	SURVILLO, OLEG
Attorney Docket Number	SNS-71301.US.10 (13-0401)

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	5956088		1999-09-21	SHEN; Paul et al.	
	2	7143141		2006-11-28	MORGAN; Jeffrey A. et al.	
	3	7574274		2009-08-11	HOLMES; Stephen G.	
	4	7657644		2010-02-02	ZHENG; Qinghua	
	5	7676142		2010-03-09	HUNG; Andy Chao	
	6	8169938		2012-05-01	DUCHSCHER; Rob et al.	
	7	8290603		2012-10-16	LAMBOURNE; Robert A.	
	8	8588949		2013-11-19	LAMBOURNE; Robert A. et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Millington, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

U.S.PATENT APPLICATION PUBLICATIONS						Remove
Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20020049843		2002-04-25	BARONE; Samuel T. et al.	
	2	20020168938		2002-11-14	CHANG; Chin-Chi	
	3	20020194309		2002-12-19	CARTER; Harry Nick et al.	
	4	20030002609		2003-01-02	FALLER; Christof et al.	
	5	20030043924		2003-03-06	HADDAD; Semir S. et al.	
	6	20040032348		2004-02-19	LAI; Angela C. W. et al.	
	7	20040041836		2004-03-04	ZANER; Melora et al.	
	8	20040066736		2004-04-08	KROEGER; Brian William	
	9	20040075767		2004-04-22	NEUMAN; Darren et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Millington, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	10	20040078383		2004-04-22	MERCER; Ian Cameron et al.	
	11	20050021590		2005-01-27	DEBIQUE; Kirt A. et al.	
	12	20060119497		2006-06-08	MILLER; Kevin L. et al.	
	13	20070143493		2007-06-21	MULLIG; Richard M. et al.	
	14	20070180137		2007-08-02	RAJAPAKSE; Ravi	
	15	20080022320		2008-01-24	VER; Steeg William C.	
	16	20120290621		2012-11-15	HEITZ III; Geremy A. et al.	
	17	20130174100		2013-07-04	SEYMOUR; Eric T. et al.	
	18	20130197682		2013-08-01	MILLINGTON; Nicholas A. J.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Millington, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	1	1122931	EP	A2	2001-08-08			☐
	2	2284327	GB	A	1995-05-31			☐
	3	03093950	WO	A2	2003-11-13			☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T5
	1	BLAKOWSKI G. et al., "A Media Synchronization Survey: Reference Model, Specification, and Case Studies", January 1996, Vol. 14, No. 1, 5-35.	☐
	2	"European Extended Search Report for Application No. 13184747.7 mailed on 28 February 2014, 8 pages".	☐
	3	"Final Office Action mailed 3 July 2012 for U.S. Patent Application No. 13/298090 filed on 16 November 2011".	☐
	4	"Final Rejection mailed on 21 January 2010 for U.S. Application No. 11/906,702 filed on 2 October 2007".	☐
	5	"Maniactools, "Identify Duplicate Files by Sound," 09/28/2010, http://www.maniactools.com/soft/music-duplicate-remover/identify-duplicate-files-by-sound.shtml".	☐
	6	"Non-Final Office Action mailed 05 December 2013 for U.S. Application No. 13/827653 filed 14 March 2013".	☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Millington, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	7	"Non-Final Office Action mailed 19 March 2013 for U.S. Application No. 13/724048 filed 21 December 2012".	☐
	8	"Non-Final Office Action mailed 30 April 2012 for U.S. Patent Application No. 13/204511 filed 5 August 2011".	☐
	9	"Non-Final Office Action mailed 30 July 2013 for U.S. Application No. 13/724048 filed 21 December 2012".	☐
	10	"Non-Final Rejection mailed on 20 August 2009 for U.S. Application No. 11/906,702 filed on 2 October 2007".	☐
	11	"Notice of Allowance mailed 06 March 2014 for U.S. Application No. 13/827653 filed 14 March 2013".	☐
	12	"Notice of Allowance mailed 10 November 2011 for U.S. Patent Application No. 11/906,702 filed on 2 October 2007".	☐
	13	"Notice of Allowance mailed 13 November 2013 for U.S. Application No. 13/724048 filed 21 December 2012".	☐
	14	"Notice of Allowance mailed 21 January 2013 for U.S. Patent Application No. 13/298090 filed on 16 November 2011".	☐
	15	"Notice of Allowance mailed 5 October 2012 for U.S. Patent Application No. 13/204511 filed 5 August 2011".	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Millington, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Millington, Nicholas A. J.
	Art Unit	2442
	Examiner Name	SURVILLO, OLEG
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Robert J. Irvine III/	Date (YYYY-MM-DD)	2014-04-11
Name/Print	Robert J. Irvine III	Registration Number	41865

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt	
EFS ID:	18738292
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 (13-0401)
Receipt Date:	15-APR-2014
Filing Date:	17-APR-2013
Time Stamp:	11:34:16
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Foreign Reference	EP1122931A2.pdf	697874 c3e62931711416d6eda01a8838040629079cd066	no	14
Warnings:					
Information:					

2	Foreign Reference	GB2284327A.pdf	1407296 f9488162d246f7791a03621b11436cba00b17575	no	44
Warnings:					
Information:					
3	Foreign Reference	WO03093950A2.pdf	7403602 eb269cfb32fa43f1e41d3d3d5ebaa7acb084e908	no	147
Warnings:					
Information:					
4	Non Patent Literature	Blakowski_1996.pdf	2310463 07ba8fd2ba3e4aef6a9851d211a00b4e27cfc297	no	32
Warnings:					
Information:					
5	Non Patent Literature	EP_SearchRep_13184747_2-28-14.pdf	275479 7d25f006c97025c33c54a64411a2be59eaf6fd20	no	8
Warnings:					
Information:					
6	Non Patent Literature	13298090-FinalOA-7-3-12.pdf	1759081 9894548a9462602535e9dff4207dfaf4b0f19180	no	46
Warnings:					
Information:					
7	Non Patent Literature	11906702-FinalOA1-21-10.pdf	962578 34bc010829051d0d6ca94098397658a78f4abdc3	no	27
Warnings:					
Information:					
8	Non Patent Literature	Maniactools_2010.pdf	287189 5622a1d881bfedd96547dddbb8e2e1cc8b641c6f	no	2
Warnings:					
Information:					
9	Non Patent Literature	13827653-NonFinalOA-12-05-13.pdf	956176 7c049b2e0dfc529d725affa5d514e473e6f52596	no	28
Warnings:					
Information:					
10	Non Patent Literature	13724048Non-FinalOA-03-19-13.pdf	309790 314f1e1e01cd426e4e16c1c77aa7ec540362fd24	no	9
Warnings:					
Information:					

11	Non Patent Literature	13204511-NonFinalOA-4-30-12.pdf	1442596 f036adf0774552d726c437b14379230238fe49e7	no	30
Warnings:					
Information:					
12	Non Patent Literature	13724048-NonFinalOA-07-30-13.pdf	236255 b1c0978b73d6a67758f00be036778362f00956e3	no	7
Warnings:					
Information:					
13	Non Patent Literature	11960702-NonFinalOA8-20-09.pdf	921072 26d37ae4969a5883331fcf941387fd6908476ad7b	no	27
Warnings:					
Information:					
14	Non Patent Literature	13827653-NOA-03-06-14.pdf	774770 89016e0f3e27c11fc6068380333a15d78ebac595	no	17
Warnings:					
Information:					
15	Non Patent Literature	11906702NOA11-10-11.pdf	715784 2c85b9d0c66ef4de7be456bf5428ec93d0b15e4b	no	17
Warnings:					
Information:					
16	Non Patent Literature	13724048-NOA-11-13-13.pdf	360709 909256053537769f7d61a2fd71e06ab3e6d34aefa	no	7
Warnings:					
Information:					
17	Non Patent Literature	13298090-NOA-1-21-13.pdf	1640898 21642c65b28dc23dcb7700bf983c16d2dae4aa21	no	39
Warnings:					
Information:					
18	Non Patent Literature	13204511-NOA-10-5-12.pdf	489070 23da7a010281d341894893bf0287a182bac8f252	no	11
Warnings:					
Information:					
19	Information Disclosure Statement (IDS) Form (SB08)	SNS-71301US10-SuppIDS-APR2014.pdf	708274 07ee415d5a55c121a2e019d61551f92d6449af2d	no	8
Warnings:					
Information:					

Total Files Size (in bytes):	23658956

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 122 931 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
08.08.2001 Bulletin 2001/32

(51) Int Cl.[7]: **H04L 29/06**

(21) Application number: **01102060.9**

(22) Date of filing: **30.01.2001**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **07.02.2000 US 499701**

(71) Applicant: **SHARP KABUSHIKI KAISHA**
Osaka 545-8522 (JP)

(72) Inventor: **Ishii, Atsushi**
Vancouver, Washington 98665 (US)

(74) Representative: **MÜLLER & HOFFMANN**
Patentanwälte
Innere Wiener Strasse 17
81667 München (DE)

(54) Real-time media content synchronization and transmission in packet network apparatus and method

(57) A system for transmitting and synchronizing real-time multimedia content includes a multimedia server (62) for generating a multimedia packet (64) having multimedia audio/visual information therein in a play sequence, and for transmitting the multimedia packet; a packet-based communication network (66) connected to the multimedia server (62) for receiving the multimedia packets therefrom; the network having multiple routers (68, 70) therein to route the multimedia packets to plural destinations; means for inserting total delay information (TDI) into the multimedia packets, wherein the TDI includes total end-to-end delay (TED) and cumulative network delay (CND); and plural multimedia receivers (80, 82, 84) for receiving the multimedia packets having TDI therein, wherein each multimedia receiver (80, 82, 84) includes a buffer for storing received multimedia packets, a sequencing means for retrieving stored multimedia packets from the buffer and sequencing the multimedia packets in the play sequence, and a playing mechanism for playing the retrieved, sequenced multimedia packets at time TED after the multimedia packet has been transmitted by the multimedia server (62). A method of transmitting and synchronizing real-time multimedia content includes loading multimedia audio/visual into plural multimedia packets (64); inserting total delay information (TDI) into at least one of the multimedia packets (64), wherein TDI includes total end-to-end delay (TED) and cumulative network delay (CND); transmitting the multimedia packets to plural multimedia receivers (80, 82, 84) over a packet-based network (66) in a play sequence; buffering the multimedia packets in each multimedia receiver (80, 82, 84); recalling the multimedia packets in play sequence; and playing the multimedia packets (64) at each receiver (80, 82, 84) at time TED.

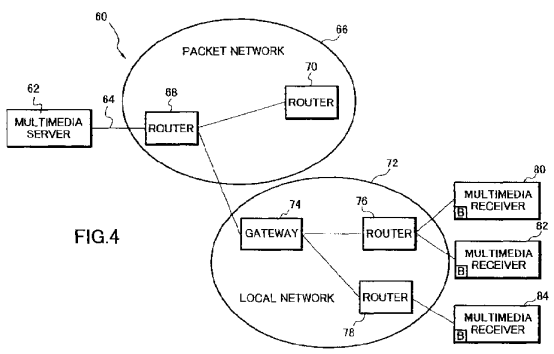


FIG.4

EP 1 122 931 A2

Printed by Jouve, 75001 PARIS (FR)

Description

Field of the Invention

[0001] This invention relates to multimedia networks, and specifically to the synchronization of multicast distributed media content when played by multiple multimedia devices, e.g., monitors and audio systems.

Background of the Invention

[0002] In some situations, such as at conventions and in exhibition halls, it is required to play the same video contents synchronized at multiple video monitors. Generally, it is also required to synchronize audio and video between separately located monitors and speakers. Current real-time media transmission technology allows transferring an audio/video stream over non-guaranteed quality of service (QoS) packet networks with reasonable quality and at a lower cost than provided by higher quality networks. Thus, there will be a demand for playing multimedia contents, delivered to a plurality of receivers over packet network, in synchronized manner.

[0003] Fig. 1 depicts a prior art multimedia content delivery system, generally at 10, which uses a packet network. A multimedia server 12 generates a stream of multimedia packets 14 that contain multimedia audio/video contents. Normally, server, or source, 12 generates multimedia packets at a constant rate (packet rate); *e.g.*, one packet every 20 ms. Packet size depends on the multimedia codec, which may result in fixed size or variable sized packets. The generated multimedia packets are then transferred through a packet network 16, having plural routers, or nodes, 18, 20, 22, 24, 26, 28 and 30 therein, and are delivered to a plurality of multimedia receivers, 32, 34, 36 and 38. The mechanism for delivery to multiple devices is well known as "multicast." All the receivers that wish to receive multimedia contents from the server join a multicast group. The routers in packet network 16 are aware of the multicast group, and transfer received multimedia packet to multiple routes. Each receiver is equipped with a speaker and/or a video display, and each plays the contents in the received multimedia packet in real-time. The multimedia receivers may buffer received multimedia packets for some duration in order to absorb the variance in delay caused by the network, and to playback multimedia packets in right order.

[0004] In conventional packet network, such as an IP network, packets are queued in each router. How long a packet stays in a queue depends on the amount of traffic that a particular router is handling. Hence, each multimedia packet delivered to each multimedia receiver has an unpredictable delay. As a result, the audio/video contents may be played at multiple locations without synchronization.

[0005] As previously noted, real-time multimedia transmission technology currently allows transferring an audio/video stream over non-guaranteed quality of service (QoS) packet networks with reasonable quality. As mobile Internet access increases, with its data transfer over bandlimited communication channels, such as wireless data networks, the technology to transmit real-time multimedia packets requiring minimum bandwidth becomes essential and of increases importance.

[0006] A packet-based real-time media transmission system is illustrated in Fig. 2, generally at 40. A multimedia source, or server, 42, periodically transmits multimedia packets 44 having audio/video multimedia contents to a packet network 46. In packet network 46, each multimedia packet is routed by a plurality of network nodes, or routers, 48, 50, 52 and 54, and delivered to a receiver 56. Each multimedia packet includes timing information that is used to set timing for playback at a receiver 56. The multiple routers and receivers of Fig. 1 are not shown in Fig. 2 for the sake of clarity.

[0007] Packet network 46 causes a variable length delay in the delivery of each multimedia packet. In addition, network 46 does not guarantee the delivery of the multimedia packets in the same order as transmitted. Thus, a receiver is normally equipped with a buffer to handle delay variances and/or disorders in multimedia packet receipt. However, due to the real-time property of the contents, the receiver will discard a received multimedia packet if the delay in its arrival exceeds a certain value, or if the multimedia packet does not arrive by the time it is to be played.

[0008] Fig. 3 shows an example of packets received and played back by the receiver, where each packet contains the timing information. Line I depicts seven multimedia packets transmitted in sequential order by server 42. Line II depicts the order of receipt of the multimedia packets at a receiver. Line III depicts a play order by the receiver. The figure depicts that the multimedia packet with timing information 3 is discarded due to an excessive delay. Basically, the routers in the prior art packet network are not aware of delays in packets transmission and receipt, nor are they aware of disordering of packet transmission. Thus, each router re-directs every received real-time multimedia packet to the next router or the receiver, even if the multimedia packet may not have any chance of being played at the receiver. If a physical link between some two routers, or the physical link between the receiver and the adjacent router has limited bandwidth, the transmission of the multimedia packets which will be discarded at the receiver possibly wastes substantial portion of valuable bandwidth on the link. Delay and disordering is illustrated by the variable routing of multimedia packets from router 48 through router 50, router 52 and router 54, or directly from router 48 to router 54.

[0009] U.S. Patent No. 4,506,358, for *Time stamping for a packet switching system*, to Montgomery, granted March

19, 1985, describes a method of measuring the network delay in a packet network by including an information field in each packet that indicates the cumulative network delay, time stamping each packet's arrival and re-transmission in each network node, and updating the information field in each network node for each packet by adding the delay occurring in each node, *i.e.*, re-transmission time - arrival time.

[0010] U.S. Patent No. 4,530,091, for *Synchronization of real-time clocks in a packet switching system*, to Crockett, granted July 16, 1985, describes a means to synchronize a remote real-time clock to a standard real-time clock via an X.25 (CCITT) network, but does not teach or suggest how to synchronize a multimedia stream, and does not consider network delay caused in a packet network.

[0011] U.S. Patent No. 4,918,687 for *Digital packet switching networks*, to Bustini *et al.*, granted April 17, 1990, describes a method to avoid consecutive packet loss by clipping voice packets whose contents indicate whether the packet is eligible for clipping, if required, in a node queue within the network.

[0012] U.S. Patent No. 5,432,790 for *Method for allocating internodal link bandwidth in a packet oriented communication network to guarantee delay quality-of-service*, to Hluchyj *et al.*, granted July 11, 1995, describes a method to allocate bandwidth on internodal links such that worst case maximum delays and worst case average delays are guaranteed while providing statistical gains in maintaining predetermined end-to-end delay QoS.

[0013] U.S. Patent No. 5,444,709 for *Protocol for transporting real time data*, to Riddle, granted August 22, 1995, describes a method to maintain the isochronous nature of the real-time data transmission by discarding data that is not timely received.

[0014] U.S. Patent No. 5,450,394 for *Delay monitoring of telecommunication networks*, to Gruber *et al.*, granted September 12, 1995, describes a means to measure a round-trip delay of network transit time, however, the delay which must be known for multimedia transmission and synchronization is a one-way network delay.

[0015] U.S. Patent No. 5,557,724 for *User interface, method, and apparatus selecting and playing channels having video, audio, and/or text streams*, to Sampat *et al.*, granted September 17, 1996, describes a user interface in a computer system to playback multimedia streams.

[0016] U.S. Patent No. 5,594,660 for *Programmable audio-video synchronization method and apparatus for multimedia systems*, to Sung *et al.*, granted January 14, 1997, describes a means to playback an audio and video stream synchronously to a single network-based receiver.

[0017] U.S. Patent No. 5,598,352 for *Method and apparatus for audio and video synchronizing in MPEG playback systems,* to Rosenau *et al.*, granted January 28, 1997, describes a means to playback an audio and video stream synchronously to a single network-based receiver.

[0018] U.S. Patent No. 5,623,483 for *Synchronization system for networked multimedia streams,* to Agrawal *et al.*, granted April 22, 1997, describes a means for a receiving device to determine the total end-to-end delay from a statistical network delay distribution and the synchronization of the two media streams from different sources through the provision of a synchronization clock located in each server and in each receiver.

[0019] U.S. Patent No. 5,629,936 for *Control of consecutive packet loss in a packet buffer*, to Lee *et al.*, granted May 13, 1997, describes a method for dropping packets to control consecutive packet loss based on a desired packet gap.

[0020] U.S. Patent No. 5,675,573 for *Delay-minimizing system with guaranteed bandwidth delivery for real-time traffic*, to Karol *et al.*, granted October 7, 1997, describes a bandwidth allocation mechanism based on individual guaranteed bandwidth requirements associated with each traffic flow.

[0021] U.S. Patent No. 5,682,384 for *Apparatus and methods achieving multiparty synchronization for real-time network application,* to Zarros, granted October 28, 1997, describes a means to playback contents from multiple sources synchronously to a single receiver on a network.

[0022] U.S. Patent No. 5,768,527 for *Device, system and method of real-time multimedia streaming,* to Zhu *et al.*, granted June 16, 1998, describes a method to recover lost packets by retransrnission.

[0023] U.S. Patent No. 5,802,060 for *Data network and a process for real time transmission of data packets in a data network*, to Fischbach *et al.*, granted September 1, 1998, describes a method to prioritize packet transmissions for stations whose network access is controlled by token exchanges.

[0024] U.S. Patent No. 5,815,634 for *Stream synchronization method and apparatus for MPEG playback system,* to Daum *et al.*, granted September 29, 1998, describes a means to playback an audio and video stream synchronously to a single network-based receiver.

[0025] U.S. Patent No. 5,844,600 for *Methods, apparatus, and systems for transporting multimedia conference data streams through a transport network,* to Kerr, granted December 1, 1998, describes a means for audio/video conferencing over an ATM network. Audio streams are combined in a network-based mixer and distributed to terminals. At each terminal, the received audio and video streams are played synchronously, however, there is no provision for synchronous play at multiple terminals on a network.

[0026] U.S. Patent No. 5,874,997 for *Measuring and regulating synchronization of merged video and audio data*, to Haigh, granted February 23, 1999, describes a means to playback an audio and video stream synchronously to a single, network-based receiver.

[0027] The Internet Engineering Task Force (IETF) publishes requests for comments (RFC) and comments filed in response thereto.

[0028] RFC 1889: A real-time end-to-end protocol utilizing existing transport layers for data that has real-time properties.

[0029] RFC 2508: Compressing IP/UDP/RTP Headers for Low-Speed Serial Links

[0030] RFC 1144: Compressing TCP/IP Headers for Low-Speed Serial Links,

[0031] Similar to RFC 2508, but for TCP/IP packets.

Summary of the Invention

[0032] A system for transmitting and synchronizing real-time multimedia content includes a multimedia server for generating a multimedia packet having multimedia audio/visual information therein in a play sequence, and for transmitting the multimedia packet; a packet-based communication network connected to the multimedia server for receiving the multimedia packets therefrom; the network having multiple routers therein to route the multimedia packets to plural destinations; means for inserting total delay information (TDI) into the multimedia packets, wherein the TDI includes total end-to-end delay (TED) and cumulative network delay (CND); and plural multimedia receivers for receiving the multimedia packets having TDI therein, wherein each multimedia receiver includes a buffer for storing received multimedia packets, a sequencing means for retrieving stored multimedia packets from the buffer and sequencing the multimedia packets in the play sequence, and a playing mechanism for playing the retrieved, sequenced multimedia packets at time TED after the multimedia packet has been transmitted by the multimedia server. A method of transmitting and synchronizing real-time multimedia content includes loading multimedia audio/visual into plural multimedia packets; inserting total delay information (TDI) into at least one of the multimedia packets, wherein TDI includes total end-to-end delay (TED) and cumulative network delay (CND); transmitting the multimedia packets to plural multimedia receivers over a packet-based network in a play sequence; buffering the multimedia packets in each multimedia receiver, recalling the multimedia packets in play sequence; and playing the multimedia packets at each receiver at time TED.

[0033] An object of the invention is to provide synchronized A/V output at a plurality of receivers over and multimedia carrying network.

[0034] Another object of the invention is to prevent transmissions of any multimedia packets which will not be used by a receiver on the network.

[0035] These and other objects and advantages of the invention will become more fully apparent as the description which follows is read in conjunction with the drawings.

Brief Description of the Drawings

[0036] Fig. 1 depicts a prior art multimedia content delivery system.

[0037] Fig. 2 depicts a packet-based real-time multimedia content transmission system.

[0038] Fig. 3 is a block diagram of packet receipt and playback.

[0039] Fig. 4 depicts a LAN multimedia content delivery system constructed according to the invention.

[0040] Fig. 5 depicts a multimedia content delivery system wherein the video and audio streams are separately routed.

Detailed Description of the Preferred Embodiment

[0041] Synchronization and transmission of multicast multimedia, particularly over packet-based systems is an efficient and inexpensive techniques for distributing information. Prior art systems simply do not consider the problems associated with synchronizing and transmitting a multimedia data stream, from one or more sources, over a packet-based system, to multiple receivers. Although there exist a number of technologies that enable a plurality of correlated media streams, transmitted from different locations, to be played at single receiver in a synchronized manner, the known prior art does not allow media streams generated by a single source, or multiple sources, to be multicast over a packet-based network to be played synchronously at multiple locations. This invention provides a system and method to transmit and synchronously play the same audio/video stream at the different locations, and to eliminate transmission of a multimedia packet when it is apparent that the multimedia packet will not be played at its destination.

[0042] Referring now to Fig. 4, a content synchronized multimedia transmission system constructed according to the invention is depicted generally at 60. System 60 includes a multimedia server 62, which has the means to insert TED and CND information, which is defined and described later herein, into each multimedia packet 64 that passes therethrough. A packet network 66 includes two routers, 68, 70, although many more routers are likely present in such a generalized network. Router 68 is connected to a local area network (LAN) 72, which includes a gateway 74 and

routers 76, 78. Multimedia receivers 80, 82 and 84 are connecter to LAN 72.

[0043] The method of the invention includes loading multimedia information into a series of packets 64, referred to herein as multimedia packets. Multimedia server 62 inserts fields of total end-to-end delay (TED) and cumulative network delay (CND) in each multimedia packet, referred to herein as total delay information (TDI). TED indicates the time from when the multimedia packet is transmitted from the server to when the multimedia packet is to be played at a receiver, whereas CND indicates the actual latency caused by routers in the communications network before the multimedia packet is delivered to a receiver. Multimedia server 62 determines the value of TED from network statistics. When transmitting a multimedia packet, server 62 sets the value of jitter to CND. The jitter is obtained by

$$\text{time interval of transmissions of two consequent packets - packet rate.} \quad (1)$$

For example, if the packet rate is 20 ms and a packet is actually transmitted 30 ms after the previous packet, jitter is 10 ms and the CND of the packet is set to 10 ms.

[0044] Each router in the network, regardless of whether the router is located in packet network 66 or LAN 72, measures the queuing delay of each multimedia packet caused by that router. The queuing delay is obtained from time stamps of reception and transmission of the multimedia packet. When transmitting the multimedia packet, each router updates CND by adding the queuing delay to the current CND value. If the router needs to route a received multimedia packet to multiple routes, *i.e.,* multicast, such as router 68 multicasting to router 70 and gateway 74, the multimedia packet is duplicated, and the queuing delay for the duplicated multimedia packet is handled separately from the original. There may be more than one duplicated packet, and the queuing delay is placed in each duplicated packet.

[0045] Each multimedia receiver includes a buffer "B" to store incoming multimedia packets, and recalls a multimedia packet from the buffer to play the content of a multimedia packet at (TED-CND) after its reception. Thus, a multimedia packet delivered to the multiple multimedia receivers is played at TED after transmitted at the server, regardless of the actual network delay caused in the route to each receiver.

[0046] As described in the previous example, CND contains only a queuing delay, and does not include physical latencies caused by the physical transmission between two routers. The physical latency is assumed to be small and is therefore negligible. Alternatively, a router may be provided with a known physical latency, which may be added to the queuing delay for increased accuracy.

[0047] If all the receivers are aware of the packet rate, performing the procedures described above for the first multimedia packet of the media stream is sufficient, because a receiver receiving a multimedia packet with TED and CND knows when to play the multimedia packets which follow, however, where the first multimedia packet containing TED and CND is "lost," *i.e.*, not received by one or more receivers, the server may be required to include TED and CND in the first several multimedia packets. Alternatively, the server may be required to include TED and CND until it receives an acknowledgement that a multimedia packet having TED and CND has been received by all of the receivers belonging to the multicast group.

[0048] As mentioned earlier, multimedia server 62 determines the value of TED, based on network statistics. TED may be provided to the server by a network administrator, or measured dynamically. The actual determination of TED is beyond of the scope of this invention, however, U.S. Patent No. 5,623,483 provides a technique to determine TED from the statistical distribution of network delay. The overall solution of U.S. Patent No. 5,623,483, however, is that the server and the receivers have synchronized clocks. The server specifies, in each packet, the time at which the receiver is to play the packet content. However, it is unrealistic to expect that the servers and receivers of a packet network maintain precisely the same time on their synchronized clocks.

[0049] It is foreseeable that situations may arise where queuing delays and physical latencies add up to the point where a multimedia packet reaches a receiver, and is "late." In the case of prior art systems, audible and visual discontinuities are perceptible to a user, without any means to correct this situation. The system and method of the invention have provisions to function in the case where one of the receivers detects a CND in a multimedia packet which exceeds the TED for the multimedia packet. In this case, the multimedia packet cannot be played synchronously, because the time for play of the contents of the multimedia packet has already passed. The receiver includes means to request a larger TDI, and sends a message to the server with the value CND-TED. The server then increases the value of TED, and thus TDI, for following multimedia packets.

[0050] In the embodiment shown in Fig. 4, server 62 and all the routers included in packet network 66 and LAN 72 for the multicast group supports TED and CND. As depicted in Fig. 4, all of the multimedia receivers in the multicast group are in a local area network (LAN), however, a multimedia receiver may also be connected to a router in packet network 66. A multimedia server does not have to include special hardware or software to insert TDI into the multimedia packets, as described above. Instead, a gateway, which connects the LAN to the (public) packet network, generates TDI from TED and CND for multimedia packets received from the server. The gateway must be capable of generating the value of TED, and must be supplied with the packet rate. The routers and the receivers in the LAN may perform

the same function. Consequently, the invention may be implemented even if the entire packet network does not support the method presented here. Put another way, the means for inserting TDI may be provided by a multimedia server, a router, or a gateway.

[0051] Another embodiment, illustrated in Fig. 5, generally at 90, includes a multimedia server 92, which generates a video stream 94 and an audio stream 96, a packet network 98, having routers 100, 102, 104, 106 and 108, and a video monitor 110 and a speaker 112. Video stream 94 and audio stream 96 are to be played separately yet synchronously at video monitor 110 and at speaker 112, respectively, which are connected to packet network 98. The destination addresses of video packets and audio packets are the addresses for the video monitor and for the speaker, respectively. Server 92 includes means to insert TDI, including TED and CND, in both video and audio packets. The routers perform the same procedure as shown in the other embodiments.

[0052] The method of U.S. Patent No. 4,506,358 is different from the invention described herein in that the present invention employs an information field in each packet of a type neither taught nor suggested in the prior art, that indicates the time when the audio/video content in the packet should be played. This enables the same content to be played synchronously at multiple receivers, without the necessity of synchronized server/receiver clocks.

[0053] This invention also maximizes the efficiency of bandwidth on physical links among routers and receivers by preventing a router from transmitting multimedia packets that will not be played at the final destination, *i.e.*, a media player, or receiver, as in the case where CND > TED.

[0054] Non-viable multimedia packets may be discarded if it is determined that they are not need, or if it is determined that they will not be used. The following may be used as a stand-alone technique, or in conjunction with the synchronization technique previously described herein. A router places a time stamp and sequence number in every received multimedia packet transferred between the media source and the receiver. Let:

$t_{ts}[n]$: the time stamp of the multimedia packet with sequence number n, and

$t_p[n]$: the timing information included in the multimedia packet with sequence number n, indicating when the content in this multimedia packet is to be played at the receiver. This information provides relative timing information among series of multimedia packets and does not indicate the absolute time.

[0055] At the receipt of the first multimedia packet, having sequence number = n_0, the router stores $t_{ts}[n_0]$ and $t_p[n_0]$ in memory and re-directs the multimedia packet to the next router or the receiver. Now, let

$$T_u = t_u[n_0] \text{ and} \qquad (2)$$

$$T_p = t_p[n_0]. \qquad (3)$$

[0056] At the receipt of every following received multimedia packet with sequence number n, the router examines the values of $t_{ts}[n]$ and $t_p[n]$, and, if

$$t_{ts}[n] - T_{ts} < tp[n] - T_p + D, \qquad (4)$$

where D denotes the maximum TED allowance, including the buffering delay in the receiver, the router re-directs the multimedia packet and sets

$$T_{ts} = t_{ts}[n] \text{ and} \qquad (5)$$

$$T_p = t_p[n]; \qquad (6)$$

and determines that the multimedia packet is a viable packet. Otherwise, the router discards the multimedia packet.

[0057] The value $t_{ts}[n] - T_{ts}$ is the difference between the arrival time of the multimedia packet of interest and the latest successfully re-directed multimedia packet at the router. Similarly, $t_p[n] - T_p$ indicates the time interval at play back between these two multimedia packets. Thus, if

$$t_{ts}[n] - T_{ts} > t_p[n] - T_p + D, \quad (7)$$

the multimedia packet with sequence number *n* has already consumed the entire TED allowance. In this case, there is no chance for the multimedia packet to be played at the receiver, because the multimedia packet will not arrive at the receiver before or at the time it is to be played. By eliminating unnecessary transmission of such a multimedia packet, referred to herein as a non-viable packet, the bandwidth on the physical link between this router and the next entity, regardless of whether the next entity is a router or a receiver, may be utilized efficiently. The router thus includes means to determine whether a multimedia packet is a viable multimedia packet, in which case the multimedia packet is transmitted; or whether the multimedia packet is a non-viable multimedia packet, in which case the multimedia packet is discarded.

[0058] The value of *D* at any router may be provided by: (1) being hard-coded on the node's software or hardware; (2) being manually configured by a system administrator; or (3) being provided in every real-time multimedia packet by the multimedia server. The value is determined by the required QoS and must be greater than the buffering delay in the receiver. *D* will be a small value for delay-sensitive applications, such as two-way voice conversation. The receiver may negotiate the value of *D* with the media source using common communication protocols.

[0059] The transmission method of the invention functions with IP/UDP/RTP header compressor, specified in RFC 2508, wherein the header information of real-time multimedia packets is compressed by a differential coding scheme. The compressor, which may be implemented in a transmitting portion of a router, provides a means to identify a specific real-time media stream between a media source and a receiver, among other traffic. Because of differential coding, the compressor is required to examine the timing information in the header of real-time multimedia packets in any packet-based network, and the extra processing power required to implement the method of the invention is minimal.

[0060] Thus, a system and method for transmitting and synchronizing real-time, multimedia packet-based content over a packet-based network has been disclosed. Although several variations of the system and method of the invention have been disclosed, it will be appreciated that further variations and modification may be made thereto without departing from the scope of the invention as defined in the appended claims.

Claims

1. A system (60) for transmitting and synchronizing real-time multimedia content, comprising:

   a multimedia server (62) for generating a multimedia packet (64) having multimedia audio/visual information therein in a play sequence, and for transmitting said multimedia packet (64);
   a packet-based communication network (66) connected to said multimedia server (62) for receiving said multimedia packets therefrom; said network having multiple routers (68, 70) therein to route said multimedia packets (64) to plural destinations;
   means for inserting total delay information (TDI) into said multimedia packets (64), wherein said TDI includes total end-to-end delay (TED) and cumulative network delay (CND); and
   plural multimedia receivers (80, 82, 84) for receiving said multimedia packets (64) having TDI therein, wherein each multimedia receiver (80, 82, 84) includes a buffer for storing received multimedia packets, a sequencing means for retrieving stored multimedia packets from said buffer and sequencing said multimedia packets in said play sequence, and a playing mechanism for playing the retrieved, sequenced multimedia packets at time TED after said multimedia packet (64) has been transmitted by said multimedia server.

2. The system (60) of claim 1 wherein said means for inserting is located in said multimedia server (62).

3. The system (60) of claim 1 wherein said means for inserting is located in a router (68, 70).

4. The system (60) of claim 1 which further includes a local area network (LAN) (72) located between said packet-based network (66) and at least some of said multimedia receivers (80, 82, 84), wherein said LAN (72) includes a gateway (74).

5. The system (60) of claim 4 wherein said means for inserting is located in said gateway (74).

6. The system (60) of claim 1 wherein at least one of said routers (68, 70) includes duplication means for duplicating multimedia packets (64), wherein said means for inserting is located in said router (68, 70), and wherein said router

(68, 70) inserts TDI into each duplicated multimedia packet.

7. The system (60) of claim 1 which further includes means to determine whether a multimedia packet (64) is a viable packet or whether a multimedia packet is a non-viable packet, and wherein said non-viable packets are discarded.

8. The system (60) of claim 1 wherein said multimedia receivers (80, 82, 84) includes means to request a larger TDI from said means for inserting.

9. A method of transmitting and synchronizing real-time multimedia content, comprising:

loading multimedia audio/visual into plural multimedia packets (64);
inserting total delay information (TDI) into at least one of the multimedia packets (64), wherein TDI includes total end-to-end delay (TED) and cumulative network delay (CND);
transmitting the multimedia packets (64) to plural multimedia receivers (80, 82, 84) over a packet-based network (66) in a play sequence;
buffering the multimedia packets (64) in each multimedia receiver (80, 82, 84);
recalling the multimedia packets (64) in play sequence; and
playing the multimedia packets (64) at each receiver (80, 82, 84) at time TED.

10. The method of claim 9 wherein said inserting includes inserting TDI in a multimedia server (62).

11. The method of claim 9 wherein said inserting includes inserting TDI in a router (68, 70) located in the packer-based network (66).

12. The method of claim 9 wherein said inserting includes inserting TDI in a gateway (74) located in a local area network (LAN) (72) which is located between the packet-based network (66) and a multimedia receiver (80, 82, 84).

13. The method of claim 9 which includes duplicating a multimedia packet (64) for multicasting to plural multimedia receivers (80, 82, 84), and wherein said duplicating includes inserting TDI into at some of the duplicated multimedia packets.

14. The method of claim 9 which includes discarding a multimedia packet (64) if at any time during said transmitting CND>TED.

15. The method of claim 14 which includes increasing the value of TED if CND>TED.

16. The method of claim 9 which includes determining whether a multimedia packet (64) is a viable packet or a non-viable packet, and which further include discarding any non-viable packets.

17. A method of transmitting and synchronizing real-time multimedia content, comprising:

determining whether a multimedia packet (64) is a viable packet or a non-viable packet; and
discarding any non-viable packets.

18. The method of claim 17 which further includes:

loading multimedia audio/visual into plural multimedia packets (64);
inserting total delay information (TDI) into at least one of the multimedia packets (64), wherein TDI includes total end-to-end delay (TED) and cumulative network delay (CND);
transmitting the multimedia packets (64) to plural multimedia receivers (80, 82, 84) over a packet-based network (66) in a play sequence;
discarding a multimedia packet (64) if at any time during said transmitting CND>TED;
buffering the multimedia packets (64) in each multimedia receiver (80, 82, 84);
recalling the multimedia packets (64) in play sequence; and
playing the multimedia packets (64) at each receiver (80, 82, 84) at time TED.

19. The method of claim 18 wherein said inserting includes inserting TDI in a multimedia server (62).

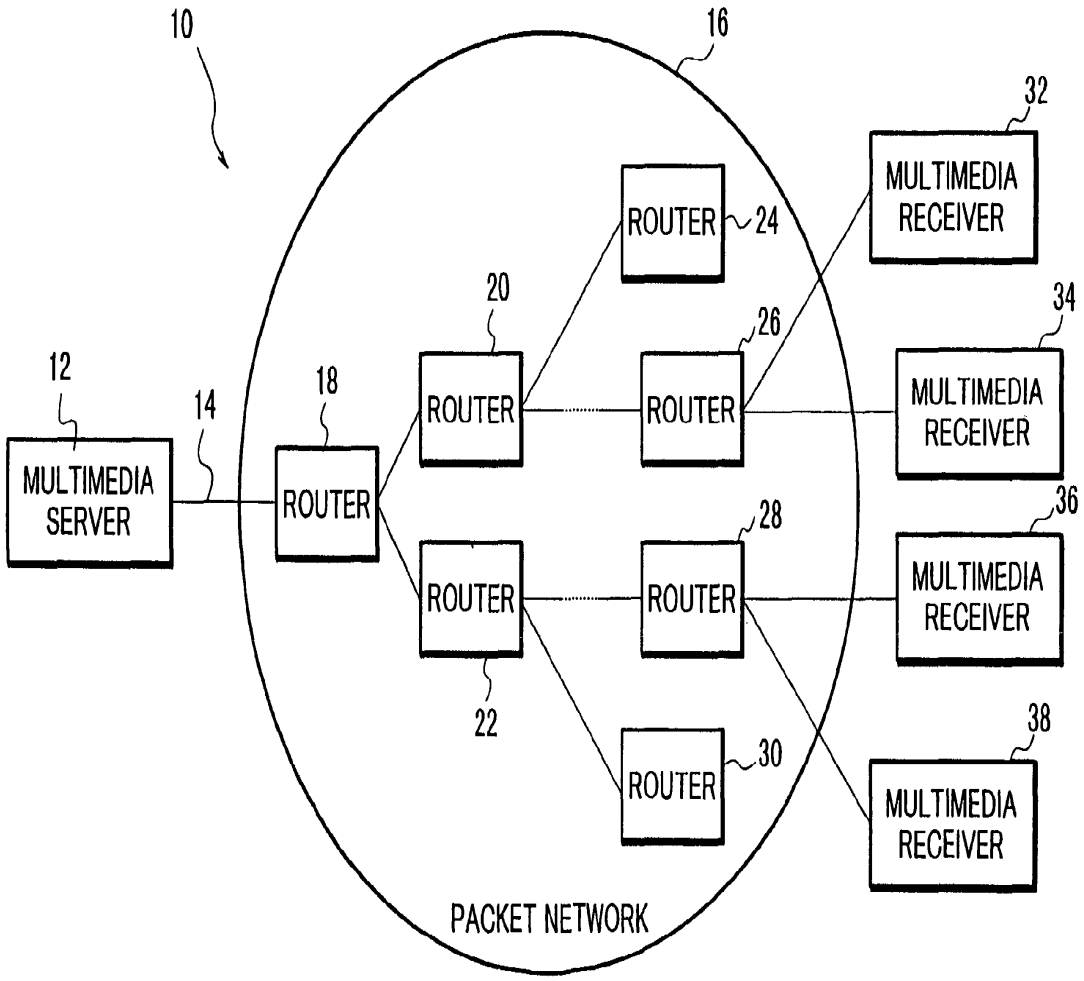
20. The method of claim 18 wherein said inserting includes inserting TDI in a router (68, 70) located in the packet-based network (66).

21. The method of claim 18 wherein said inserting includes inserting TDI in a gateway (74) located in a local area network (LAN) (72) which is located between the packet-based network (66) and a multimedia receiver (80, 82, 84).

22. The method of claim 18 which includes duplicating a multimedia packet (64) for multicasting to plural multimedia receivers (80, 82, 84), and wherein said duplicating includes inserting TDI into at some of the duplicated multimedia packets.

23. The method of claim 18 which includes increasing the value of TED if CND>TED.

FIG.1
(PRIOR ART)



EP 1 122 931 A2

FIG.2

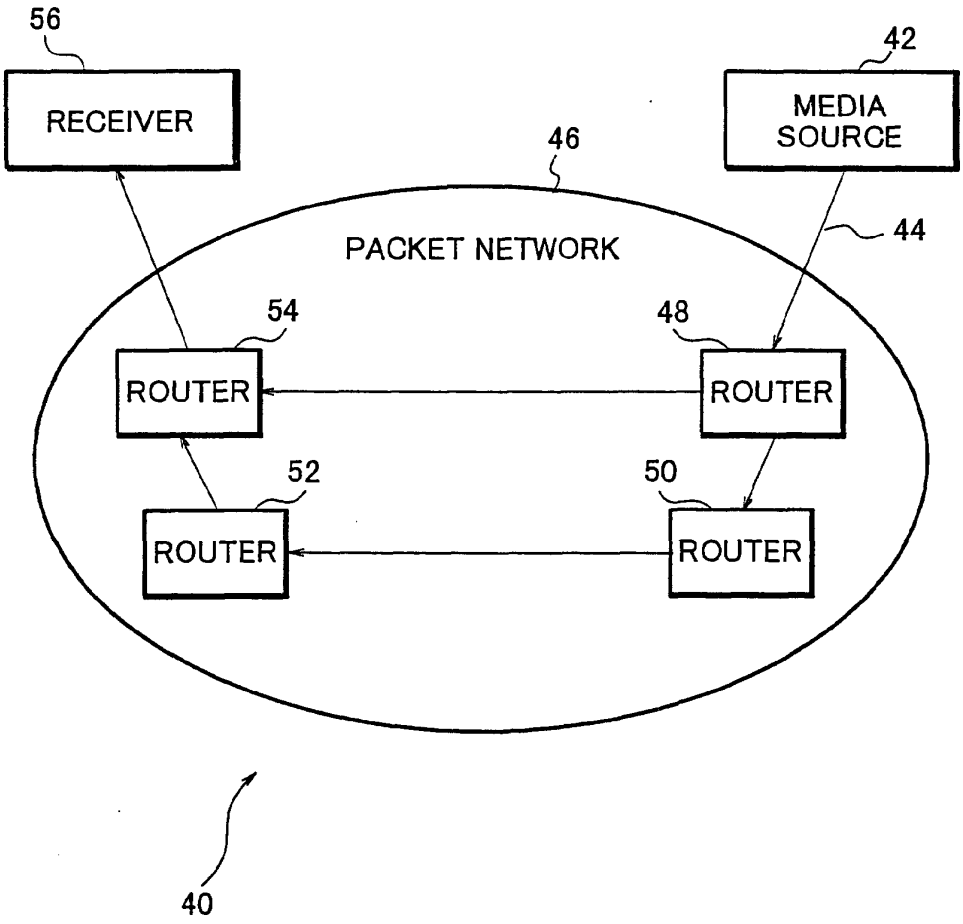
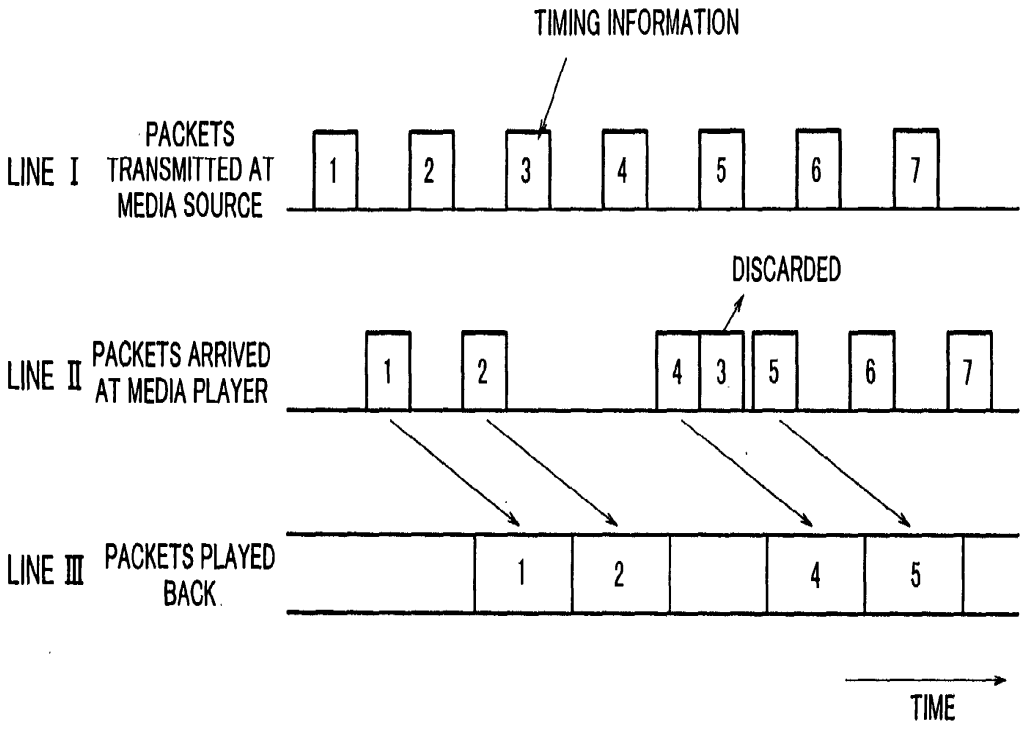

56
RECEIVER
42
MEDIA SOURCE
46
PACKET NETWORK
44
54
48
ROUTER
ROUTER
52
50
ROUTER
ROUTER
40

FIG.3



12

EP 1 122 931 A2

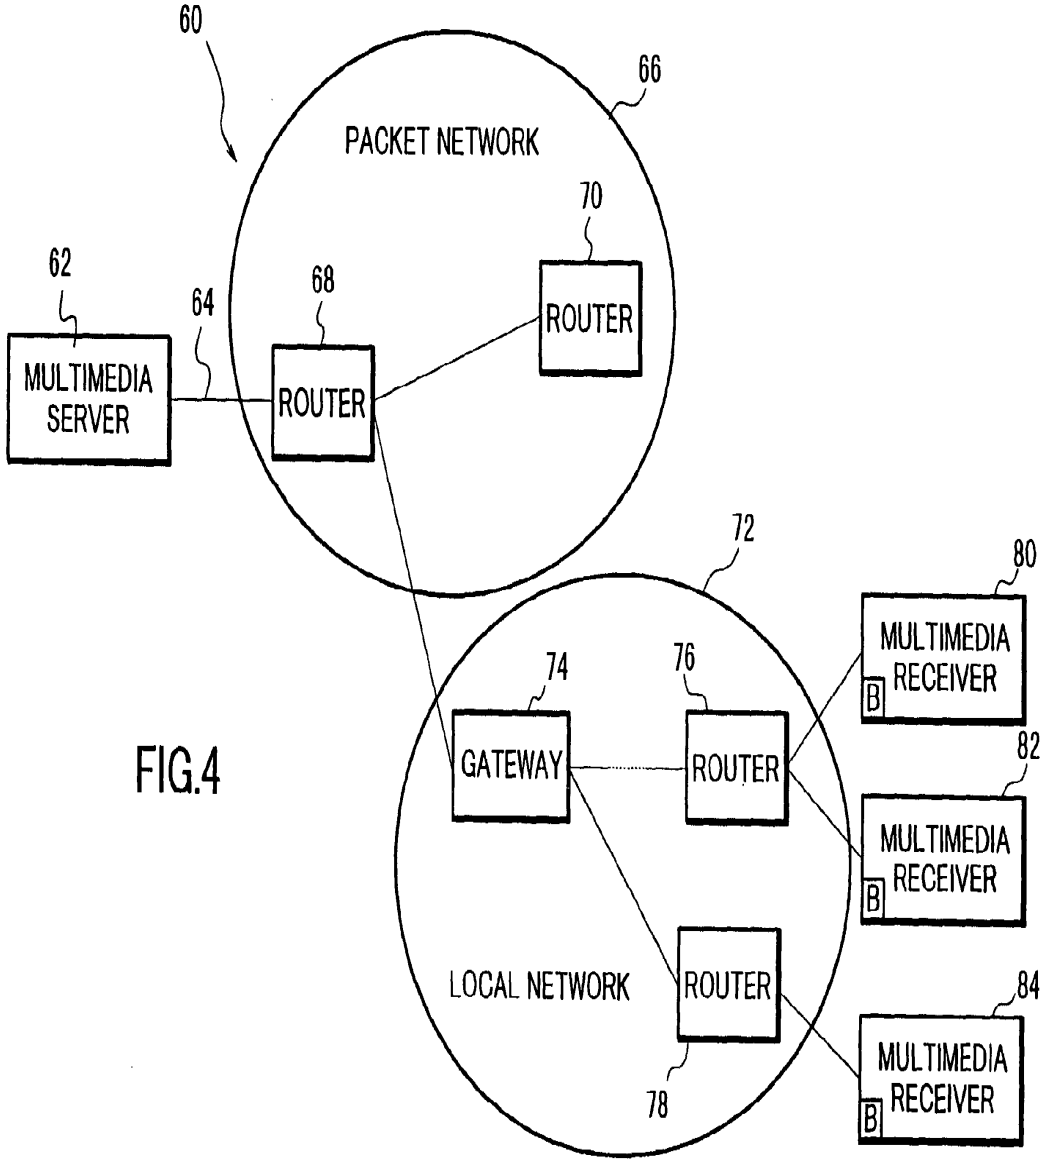
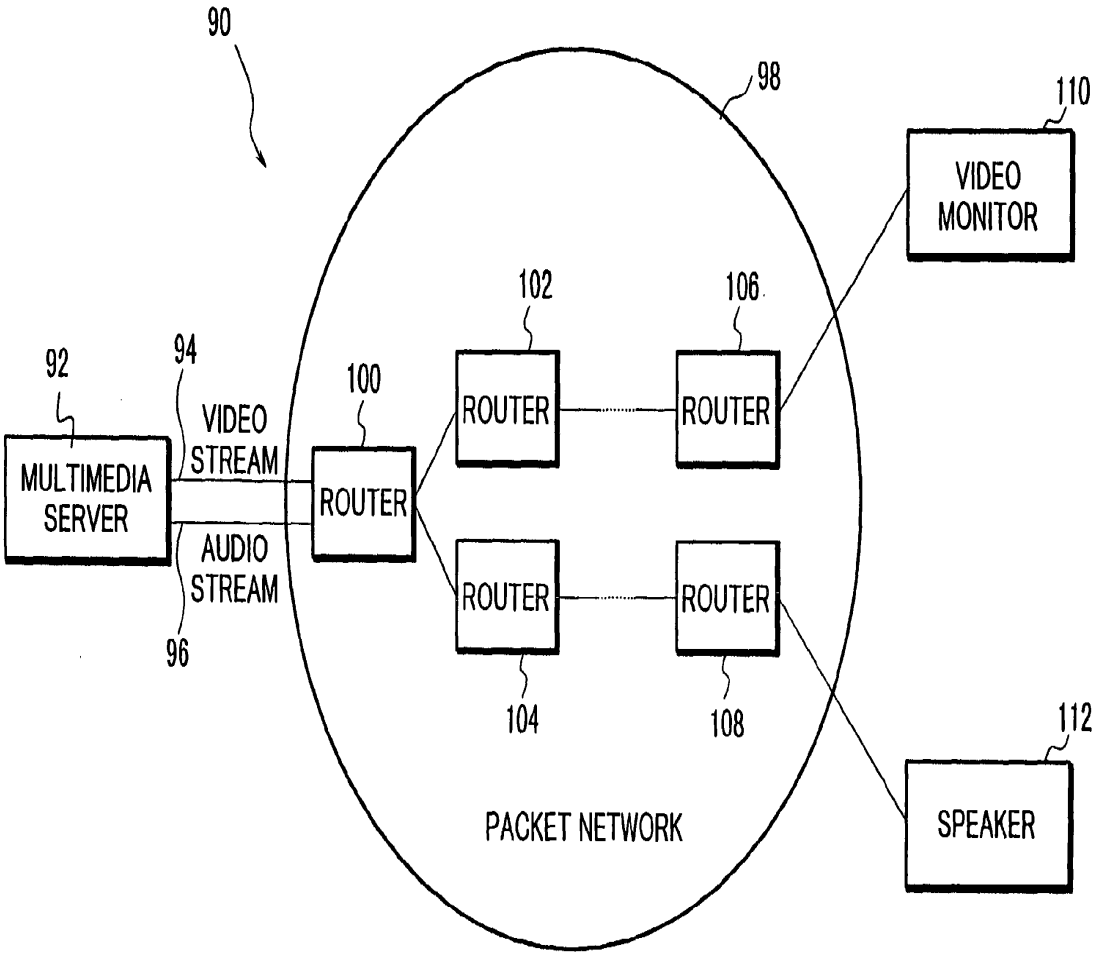

60
66
PACKET NETWORK
70
62
64
68
ROUTER
MULTIMEDIA SERVER
ROUTER
72
80
MULTIMEDIA RECEIVER
B
74
76
GATEWAY
ROUTER
82
MULTIMEDIA RECEIVER
B
LOCAL NETWORK
ROUTER
84
78
MULTIMEDIA RECEIVER
B

FIG.4

FIG.5


90
98
110
VIDEO MONITOR
102
106
92
94
100
ROUTER
ROUTER
MULTIMEDIA SERVER
VIDEO STREAM
ROUTER
AUDIO STREAM
96
ROUTER
ROUTER
104
108
112
PACKET NETWORK
SPEAKER

(12) UK Patent Application (19) GB (11) 2 284 327 (13) A

(43) Date of A Publication 31.05.1995

(21) Application No 9416276.5

(22) Date of Filing 08.08.1994

(30) Priority Data
(31) 08159857 (32) 29.11.1993 (33) US

(51) INT CL[6]
H04L 7/10

(52) UK CL (Edition N)
H4P PPF

(56) Documents Cited
EP 0234861 A2 EP 0059821 A1

(58) Field of Search
UK CL (Edition M) H4P PPF PPS PSB
INT CL[5] H04L 7/00 7/10 12/56
ONLINE DATABASES : WPI,INSPEC.

(71) Applicant(s)
Intel Corporation

(Incorporated in USA - Delaware)

2200 Mission College Boulevard, Santa Clara, California 95052, United States of America

(72) Inventor(s)
Gunner Danneels

(74) Agent and/or Address for Service
Potts, Kerr & Co
15 Hamilton Square, BIRKENHEAD, Merseyside, L41 6BR, United Kingdom

(54) **Synchronizing multiple independent data streams in a networked computer system**

(57) Multiple independent data streams 318 transferred across a communication medium from a server computer system 310, 312 to a client computer system 320, 322 are synchronised by 1) capturing information into a data packet, the capturing step being performed by the server system; 2) storing time information into the data packet, the time information indicating when the capturing step is complete; 3) transferring the data packet to the client system across the communication medium; 4) initializing a virtual server system clock if the data packet is a first data packet in the data stream; 5) retrieving the time information from the data packet; 6) comparing the time information with the virtual server system clock; and 7) processing the information in the data packet if the time information indicates a time equal to a time indicated by the virtual server system clock. The present invention further includes the steps of: 1) scheduling processing of the information in the data packet if the time information indicates a time later than a time indicated by the virtual server system clock; and 2) discarding the data packet if the time information indicates a time earlier than a time indicated by the virtual server system clock.

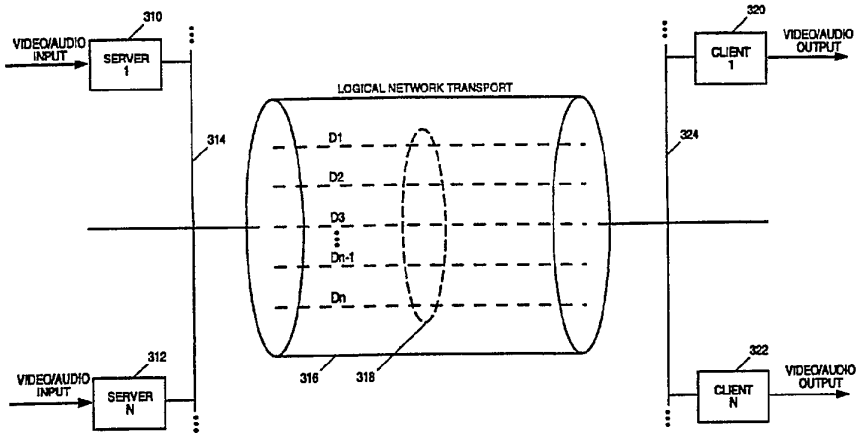


FIG.3

At least one drawing originally filed was informal and the print reproduced here is taken from a later filed formal copy.

GB 2 284 327 A

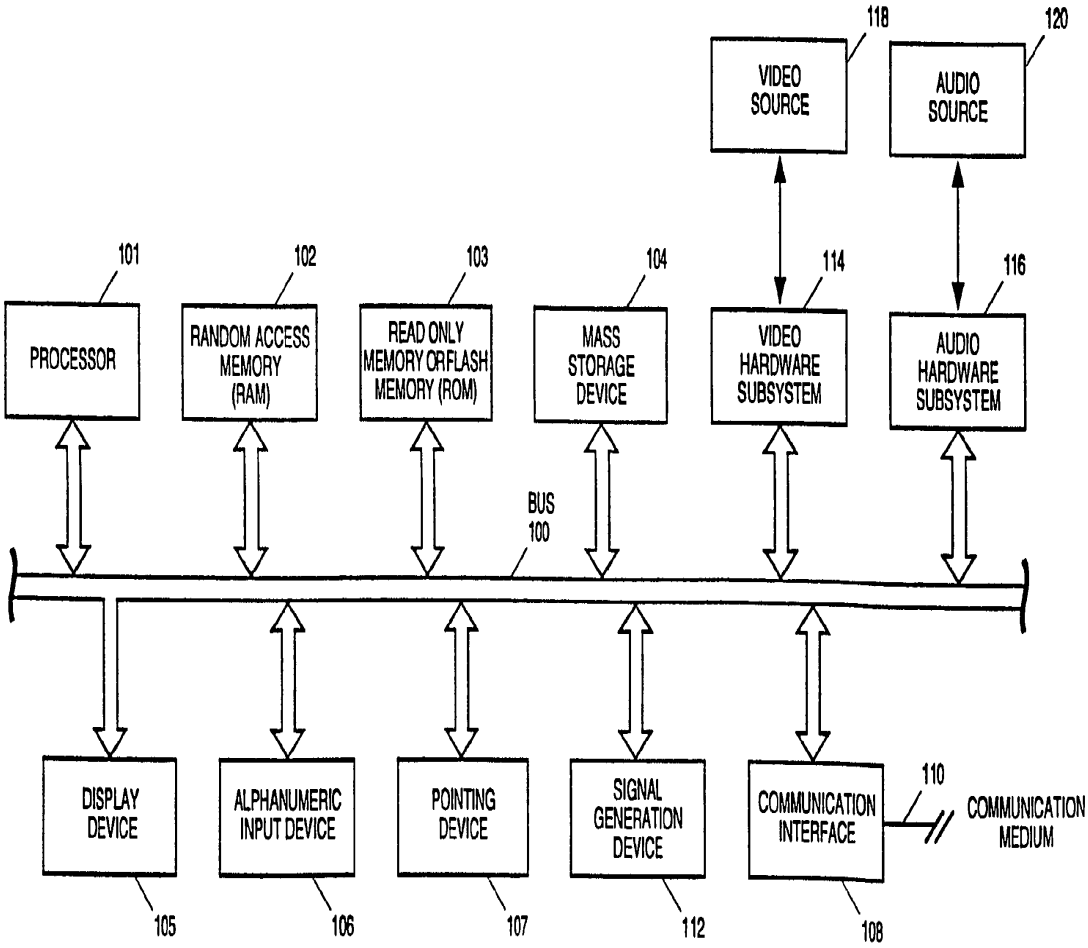

VIDEO SOURCE
118
AUDIO SOURCE
120
PROCESSOR
101
RANDOM ACCESS MEMORY (RAM)
102
READ ONLY MEMORY OR FLASH MEMORY (ROM)
103
MASS STORAGE DEVICE
104
VIDEO HARDWARE SUBSYSTEM
114
AUDIO HARDWARE SUBSYSTEM
116
BUS
100
DISPLAY DEVICE
105
ALPHANUMERIC INPUT DEVICE
106
POINTING DEVICE
107
SIGNAL GENERATION DEVICE
112
COMMUNICATION INTERFACE
108
110
COMMUNICATION MEDIUM

FIG.1

1/9

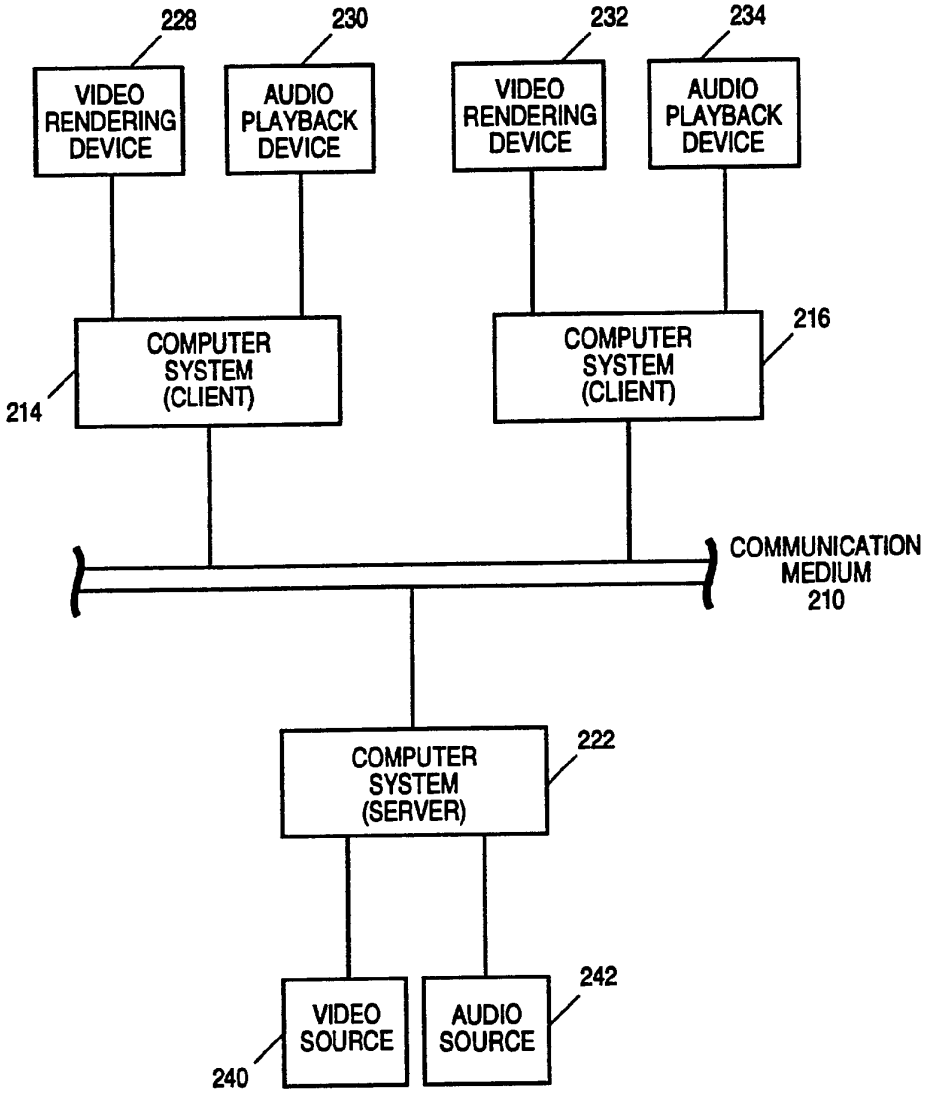

228
230
232
234
VIDEO RENDERING DEVICE
AUDIO PLAYBACK DEVICE
VIDEO RENDERING DEVICE
AUDIO PLAYBACK DEVICE
COMPUTER SYSTEM (CLIENT)
214
COMPUTER SYSTEM (CLIENT)
216
COMMUNICATION MEDIUM 210
COMPUTER SYSTEM (SERVER)
222
VIDEO SOURCE
240
AUDIO SOURCE
242

FIG.2

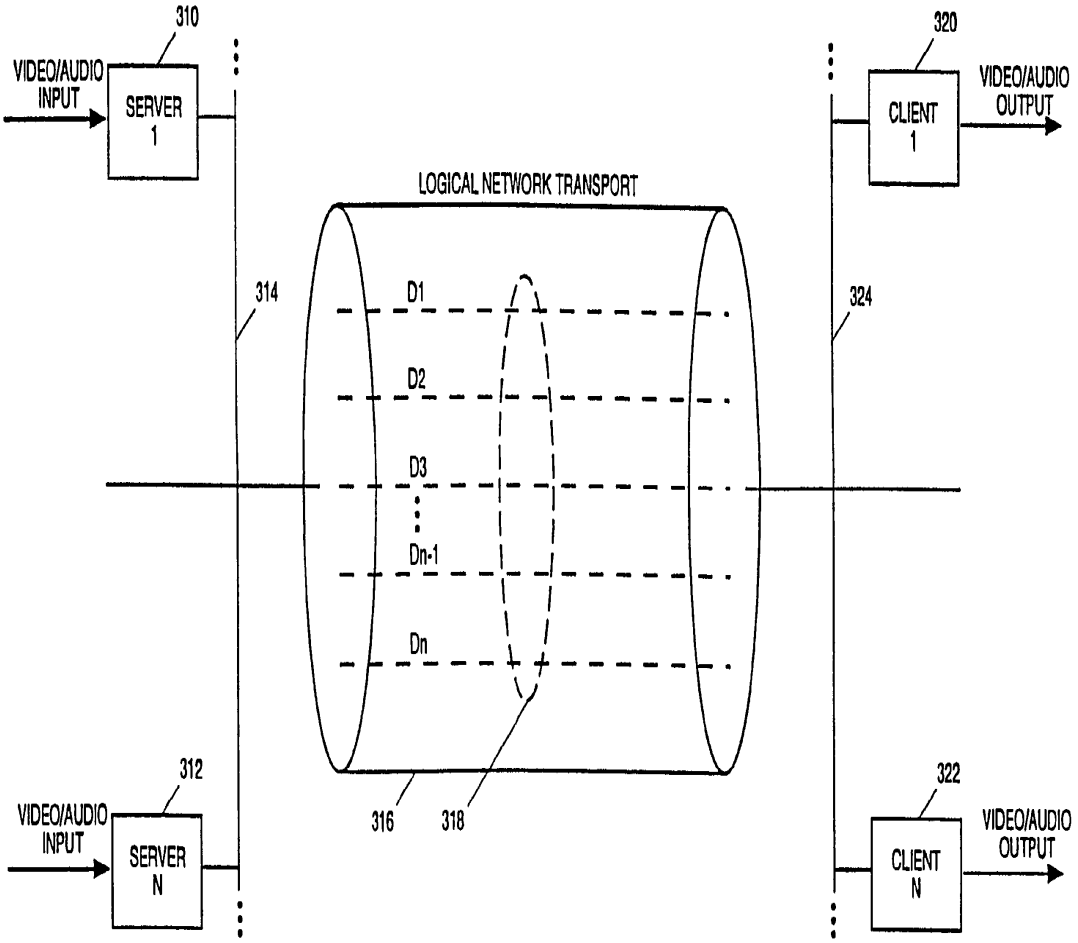

310
VIDEO/AUDIO INPUT
SERVER 1
320
CLIENT 1
VIDEO/AUDIO OUTPUT
LOGICAL NETWORK TRANSPORT
314
324
D1
D2
D3
Dn-1
Dn
316
318
312
VIDEO/AUDIO INPUT
SERVER N
322
CLIENT N
VIDEO/AUDIO OUTPUT

FIG.3

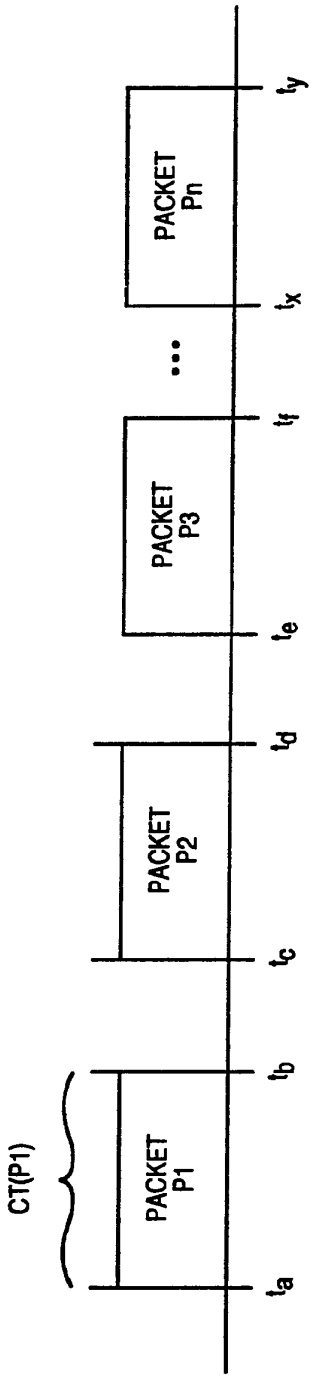

CT(P1)
PACKET P1
PACKET P2
PACKET P3
PACKET Pn
t_a
t_b
t_c
t_d
t_e
t_f
t_x
t_y
$t_b = t_a + CT(P1)$
$t_a = t_b - CT(P1)$

FIG.4

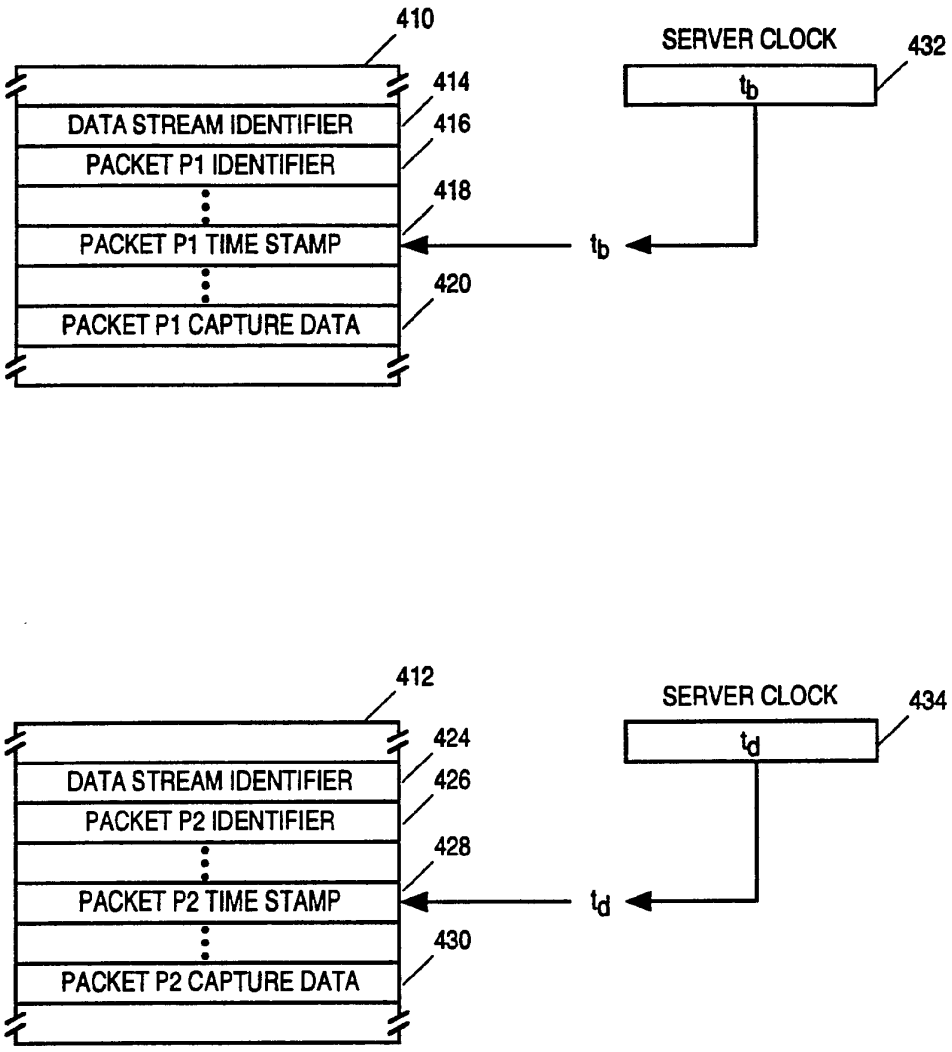

410
414
DATA STREAM IDENTIFIER
416
PACKET P1 IDENTIFIER
418
PACKET P1 TIME STAMP
420
PACKET P1 CAPTURE DATA
SERVER CLOCK
432
t_b
t_b
412
424
DATA STREAM IDENTIFIER
426
PACKET P2 IDENTIFIER
428
PACKET P2 TIME STAMP
430
PACKET P2 CAPTURE DATA
SERVER CLOCK
434
t_d
t_d

FIG.5

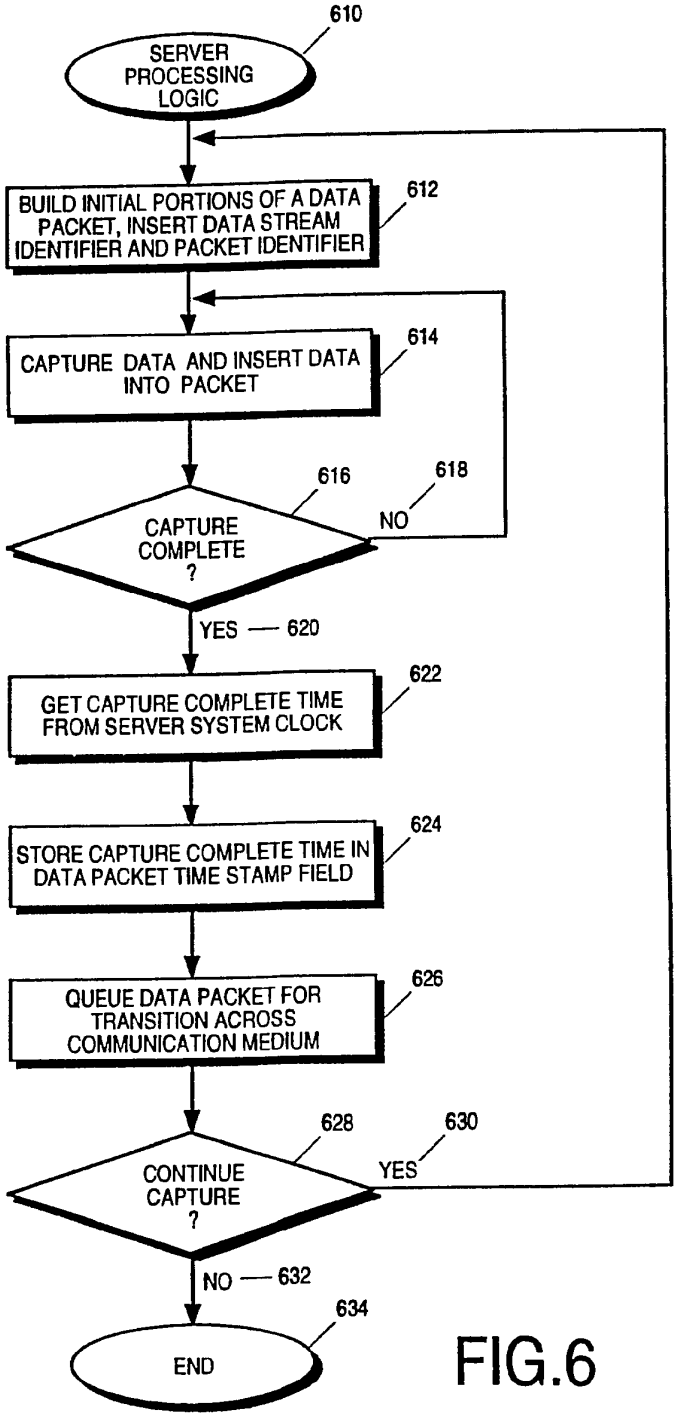

610
SERVER PROCESSING LOGIC
BUILD INITIAL PORTIONS OF A DATA PACKET, INSERT DATA STREAM IDENTIFIER AND PACKET IDENTIFIER
612
CAPTURE DATA AND INSERT DATA INTO PACKET
614
CAPTURE COMPLETE ?
616
NO
618
YES — 620
GET CAPTURE COMPLETE TIME FROM SERVER SYSTEM CLOCK
622
STORE CAPTURE COMPLETE TIME IN DATA PACKET TIME STAMP FIELD
624
QUEUE DATA PACKET FOR TRANSITION ACROSS COMMUNICATION MEDIUM
626
CONTINUE CAPTURE ?
628
YES
630
NO — 632
END
634

FIG.6

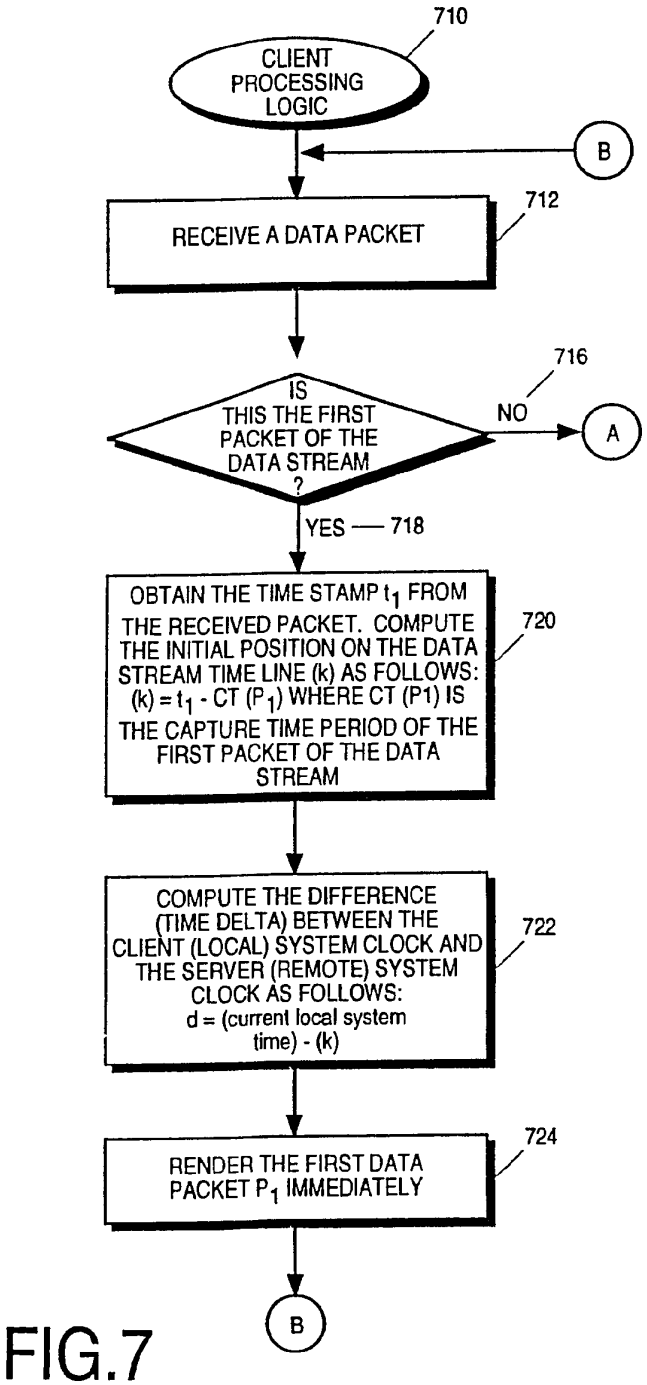

710
CLIENT PROCESSING LOGIC
B
712
RECEIVE A DATA PACKET
716
IS THIS THE FIRST PACKET OF THE DATA STREAM ?
NO
A
YES — 718
OBTAIN THE TIME STAMP t_1 FROM THE RECEIVED PACKET. COMPUTE THE INITIAL POSITION ON THE DATA STREAM TIME LINE (k) AS FOLLOWS: (k) = t_1 - CT (P_1) WHERE CT (P1) IS THE CAPTURE TIME PERIOD OF THE FIRST PACKET OF THE DATA STREAM
720
COMPUTE THE DIFFERENCE (TIME DELTA) BETWEEN THE CLIENT (LOCAL) SYSTEM CLOCK AND THE SERVER (REMOTE) SYSTEM CLOCK AS FOLLOWS: d = (current local system time) - (k)
722
RENDER THE FIRST DATA PACKET P_1 IMMEDIATELY
724
B

FIG.7

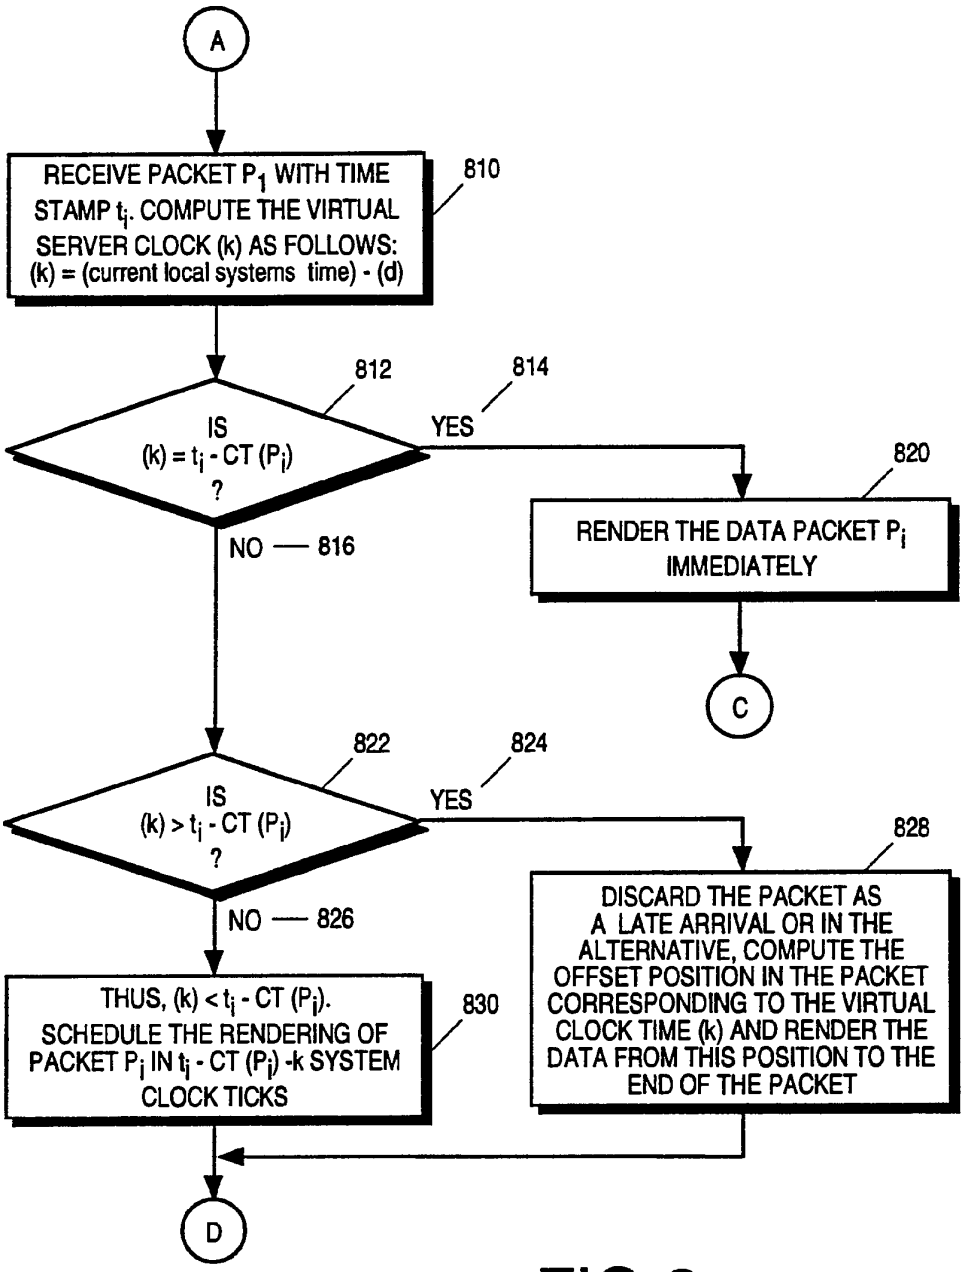

A
RECEIVE PACKET P_1 WITH TIME STAMP t_i. COMPUTE THE VIRTUAL SERVER CLOCK (k) AS FOLLOWS: (k) = (current local systems time) - (d)
810
812
IS (k) = t_i - CT (P_i) ?
814
YES
820
RENDER THE DATA PACKET P_i IMMEDIATELY
C
NO — 816
822
IS (k) > t_i - CT (P_i) ?
824
YES
828
DISCARD THE PACKET AS A LATE ARRIVAL OR IN THE ALTERNATIVE, COMPUTE THE OFFSET POSITION IN THE PACKET CORRESPONDING TO THE VIRTUAL CLOCK TIME (k) AND RENDER THE DATA FROM THIS POSITION TO THE END OF THE PACKET
NO — 826
THUS, (k) < t_i - CT (P_i). SCHEDULE THE RENDERING OF PACKET P_i IN t_i - CT (P_i) -k SYSTEM CLOCK TICKS
830
D

FIG.8

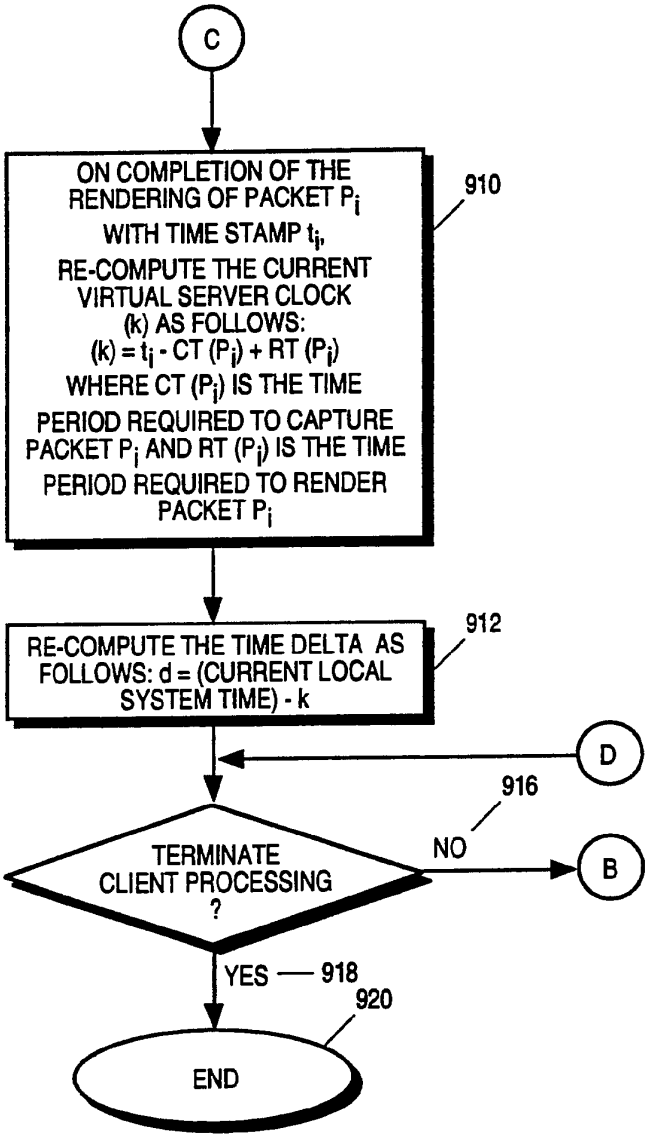

C
ON COMPLETION OF THE RENDERING OF PACKET P_i WITH TIME STAMP t_i, RE-COMPUTE THE CURRENT VIRTUAL SERVER CLOCK (k) AS FOLLOWS: $(k) = t_i - CT(P_i) + RT(P_i)$ WHERE $CT(P_i)$ IS THE TIME PERIOD REQUIRED TO CAPTURE PACKET P_i AND $RT(P_i)$ IS THE TIME PERIOD REQUIRED TO RENDER PACKET P_i
910
RE-COMPUTE THE TIME DELTA AS FOLLOWS: d = (CURRENT LOCAL SYSTEM TIME) - k
912
D
916
TERMINATE CLIENT PROCESSING ?
NO
B
YES
918
920
END

FIG.9

A METHOD AND APPARATUS FOR SYNCHRONIZING MULTIPLE INDEPENDENT DATA STREAMS IN A NETWORKED COMPUTER SYSTEM

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to the field of a networked computer system. Specifically, the present invention relates to the transfer and synchronization of multiple data streams across a computer network or various types of communication media.

2. Description of Related Art

In a networked computer system, it is often necessary to send multiple streams of data to a remote network site. Typically, the streams of data must be processed in the correct order and synchronized with respect to each data stream. For example, audio and video streams in a teleconferencing application must be synchronized to provide the correct relationship between the audio and video data. Conventional systems implement data streams as a plurality of data packets that can be independently transferred across a network. Currently, packetized data streams may be synchronized by interleaving the data packets that make up the streams in such a way that the packets are sent to the network in synchronized order. This interleaving method has several drawbacks. First, not all networks can guarantee an in-order delivery of data packets. The synchronization between data streams may be lost by a network that randomly delays the transfer of a data packet. Secondly, even if in-order delivery of packets is guaranteed, the interleaving method does not allow synchronization of data streams where the data

stream cannot be sent in a single interleaved stream. For example, where the data streams are sent on different network channels, two different network addresses, or on different networks, the data streams cannot be sent in a single interleaved stream. In other situations, multiple originators or producers of data streams (i.e. servers) may exist on a network. In this case, data streams from a single server must be synchronized with respect to each other; but, synchronization from one server to another may not be required. In other situations it may be necessary to mix a set of data streams originated by multiple servers into one unified data stream for delivery to a client network node. For example, instead of two servers each sending an audio and video stream to a client node for processing, it may be more efficient to mix the audio signals into a unified stream which then must be synchronized with both video streams. Conventional systems are incapable of providing these multiple levels of data stream synchronization.

Thus, a better method for synchronizing multiple independent data streams in a networked computer system is needed.

SUMMARY OF THE INVENTION

The present invention is an apparatus and method for synchronizing multiple independent data streams in a networked computer system. In a computer network having a plurality of computer systems coupled by a communication medium, one of the plurality of computer systems being a server system, a different one of the plurality of computer systems being a client system, a method for synchronizing a data stream transferred from the server system to the client system, the method comprising the steps of: 1) capturing information into a data packet, the capturing step being performed by the server system; 2) storing time information into the data packet, the time information indicating when the capturing step is complete; 3) transferring the data packet to the client system across the communication medium; 4) initializing a virtual server system clock if the data packet is a first data packet is the data stream; 5) retrieving the time information from the data packet; 6) comparing the time information with the virtual server system clock; and 7) processing the information in the data packet if the time information indicates a time equal to a time indicated by the virtual server system clock. The present invention further includes the steps of: 1) scheduling processing of the information in the data packet if the time information indicates a time later than a time indicated by the virtual server system clock; and 2) discarding the data packet if the time

information indicates a time earlier than a time indicated by the virtual server system clock.

It is therefore an advantage of the present invention to provide a means and method for time stamping packets of a data stream. It is a further advantage of the present invention to provide a means and method for synchronizing data streams arriving from one server and providing them to a client in the same time sequence as they were captured. It is a further advantage of the present invention to provide a data stream synchronization means and method wherein data streams may take independent paths across a network. It is a further advantage of the present invention to provide a means and method for mixing a plurality of data streams for playback in one unified stream.

These and other advantages of the present invention are fully described in the following detailed description of the preferred embodiment.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram illustrating the computer system hardware used in the preferred embodiment.

Figure 2 illustrates a typical network environment in which the present invention operates.

Figure 3 illustrates a multiple server, multiple client network environment.

Figure 4 illustrates a time line with associated data packets of a data stream.

Figure 5 illustrates the data structure of a data packet.

Figures 6 through 9 are flow charts illustrating the processing logic of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention is an apparatus and method for synchronizing multiple independent data streams in a networked computer system. In the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the invention. However, it will be apparent to one of ordinary skill in the art that these specific details need not be used to practice the present invention. In other circumstances, well known structures, materials, circuits, and interfaces have not been shown or described in detail in order to not unnecessarily obscure the present invention.

Referring now to Figure 1, a block diagram of the computer system hardware used in the preferred embodiment is illustrated. The computer system used in the preferred embodiment comprises a Bus 100 for communicating information between computer system components. These components coupled to Bus 100 include Processor 101. Processor 101 may be implemented as any of a variety of commonly available microprocessors such as the 386 brand, i486 brand, or Pentium brand microprocessors manufactured by Intel Corporation, Santa Clara, California. The Intel, 386, i486, and Pentium marks are trademarks of Intel Corporation. It will be apparent of ordinary skill in the art that other equivalent processors may be used with the present invention. Other system components include Random Access Memory (RAM)

102, Read Only Memory (ROM) 103, and mass storage device or disk drive 104. The computer system of the preferred embodiment also includes a display device 105 coupled to the Bus 100 for displaying information to a computer user, an alphanumeric input device 106 coupled to the Bus 100 for communicating information and command selections to the Processor 101, a pointing device or cursor control device 107 coupled to the Bus 100 for communicating information and command selections to Processor 101, and a signal generation device 112 coupled to the Bus 100 for communicating command selections to the Processor 101. Display device 105 may be a liquid crystal device, a cathode ray tube, or other suitable display device. Alphanumeric input device 106 is typically an alphanumeric or function key keyboard. Pointing device or cursor control device 107 is typically a mouse, trackball, or stylus device allowing the computer user to dynamically signal the two dimensional movement of a visible symbol (i.e., cursor) on the display screen of display device 105. Many implementations of the cursor control device 107 are well known in the art including trackball systems, mouse systems, joy sticks, stylus, or special keys on the alphanumeric input device 106 capable of signaling movement in a given direction. Signal generation device 112 is typically implemented as a button or switch located on or near the cursor control device 107.

The computer system used in the preferred embodiment includes a communication interface 108 coupled to Bus 100. Communication

interface 108 is also coupled to a communication medium 110. Communication interface 108 provides a means for the computer system to communicate with a network of other computer systems. A variety of networks or communication media are well known in the art. These communication media 110 include well known Local Area Networks (LAN) such as Ethernet, Wide Area Networks (WAN), telephone modem links, or other well known and conventional computer networking technology.

The computer system of the preferred embodiment also includes video hardware subsystem 114 and audio hardware subsystem 116 both coupled to Bus 100. A video source 118, such as a video camera or video tape player, is coupled to video hardware subsystem 114. An audio source 120, such as a microphone or audio tape player, is coupled to audio hardware subsystem 116. Video hardware subsystem 114 provides a means by which the computer system of the preferred embodiment may receive video signals and convert the video signals for transfer to other system components as a stream of digital video information via Bus 100. Similarly, audio hardware subsystem 116 provides a means by which the present invention may receive audio signals and convert the audio signals to a stream of digital audio information for transfer to other system components via Bus 100. In a similar manner, the video data streams provided by video hardware subsystem 114 and the audio data streams provided by audio hardware

subsystem 116 may be transferred under control of Processor 101 to other networked computer systems via communication interface 108 and communication medium 110. The present invention provides a means and method for synchronizing these data streams in a networked computer system. Except as described herein, the configuration and interconnection of computer system components as illustrated in Figure 1 are well known to those of ordinary skill in the art.

Referring now to Figure 2, a typical environment in which the present invention operates is illustrated. The present invention operates in an environment comprising a set of computer systems linked in a network topology by a communication medium 210. It will be apparent to those of ordinary skill in the art that any arbitrary number of computer systems may be so linked. In the example network environment illustrated in Figure 2, each computer system 214, 216 and 222 is shown coupled to communication medium 210. Each of these computer systems comprise components such as those illustrated in Figure 1. Computer systems 214 and 216 may be coupled to video rendering devices, such as a display, and audio playback devices. Computer system 222 may be coupled to a video source and an audio source. For example, computer system 214 is coupled to video rendering device 228 and audio playback device 230. Similarly, computer system 222 is coupled to video source 240 and audio source 242. It is not necessary for each computer system coupled to

communication medium 210 to have video rendering devices, video sources, audio playback devices, and/or audio sources connected to them. In a variety of configurations, computer systems coupled to communication medium 210 may independently be coupled to a video source or rendering device and an audio source or playback device, a video source only, an audio source only, or neither a video source or audio source. A single video source or audio source may also be shared between two computer systems. In any of these configurations, the video source provides a stream of video data which a computer system may provide to other computer systems via communication medium 210. Similarly, an audio source provides a steam of audio data which may be provided to other computer systems via communication medium 210. It will be apparent to one of ordinary skill in the art that data types other than video or audio may equivalently be supported.

A computer system on which video or audio data steams originate is called a server for the data streams. In the example of Figure 2, computer system 222 is a server. When acting as a server, computer system 222 receives video signals via video source 240 and audio signals via audio source 242 and converts these signals to video and audio data streams, which are provided to other computer systems via communication medium 210. These data streams are transferred on communication medium 210 by server 222 and transported to remote computer systems such as computer system 214 and 216. Any remote

computer system receiving data streams from a server and rendering or playing the data to a local user is known as a client system for the data streams. Thus, in general, a server system is an originator of data streams and a client system is a receiver or consumer of data streams. The client system is responsible for synchronizing the data streams arriving from a server and for rendering the data streams to a local user in the same time sequence as the data streams were captured by the server system. It is not essential to the present invention that the data streams from the server arrive on the same network media. It is only important that the data streams are provided in some way to a client computer system via a communication medium. The client system may then synchronize the data streams and render them to a local user using the technique taught herein.

Although the example illustrated in Figure 2 shows a single server system 222, the present invention is applicable to network environments wherein multiple servers and multiple clients exist on the network at one time. A multiple server, multiple client network configuration is illustrated in Figure 3.

Referring to Figure 3, servers 310 and 312 are coupled to clients 320 and 322 via communication medium 314, 316 and 324. A portion of network medium 316 is blown up in Figure 3 to illustrate the multiple independent data streams 318 concurrently being transferred across network medium 316 from servers 310 and 312 to clients 320 and 322.

Each data stream of data streams 318 consist of a sequence of data packets. Each data packet comprises a portion of information captured by a server system. This information may represent a captured video image, or portion thereof, a captured audio representation or any other collection of information that may be represented as binary information. In a typical implementation, the video hardware subsystem of a server system captures a video image and creates a sequence of data packets that are transferred across network 316 as one of data streams 318. Similarly, other data streams of data steams 318 are generated by other server systems. In the synchronization process of the present invention, the client system seeks to establish the same time line or synchronization for rendering the data packets as was present during the capturing of the data packets performed by the corresponding server system. In this manner, the client system is able to render information in a manner corresponding to the way the information was captured by a server system. This synchronization can occur regardless of the manner by which the information is transferred between the server and client systems. Moreover, the present invention can be used to synchronize data streams with respect to each other. In this manner, a video stream may be synchronized with an audio stream and rendered by a client system appropriately. Further details of the synchronization process of the present invention follow.

Referring now to Figure 4, data packets P_1 to P_n are shown on a time line. Data packets P_1 through P_n represent the data packets of a data stream such as data stream D_1 illustrated in Figure 3. Each packet comprises information collected or captured during a time frame defined by a beginning capture time and an ending capture time. For example, referring to Figure 4, packet P_1 comprises information captured between time t_a and time t_b. A function CT defines the capture time for a particular packet of a data stream. As illustrated in Figure 4, the capture time for packet P_1, CT (P_1) is defined as the time period between t_a and t_b. In one embodiment of the present invention, the capture time for each data packet of a data stream is identical and preset to some fixed value. In an alternative embodiment of the present invention, the capture time for each data packet of a data stream may be different. Thus, function CT is used to determine the capture time for a particular data packet. As indicated by the equations in Figure 4, the ending capture time, t_b, may be determined by knowing the beginning capture time, t_a, and the capture time CT (P_1). Similarly, the beginning capture time, t_a, may be determined if the ending capture time, t_b, and the capture time CT (P_1) is known. Similar computations may be performed for each data packet P_1 through P_n of a data stream.

In support of the synchronization process of the present invention, a time stamp is applied to each data packet of a data stream. When a server computer system captures data into data packets, the server

stamps its system clock time into each packet of the data stream. In the preferred embodiment, this time stamp represents the system clock time at the data packet ending capture time. Thus, for the example illustrated in Figure 4, the time stamp applied to packet P_1 is t_b. The computer system clock is assumed to be a monotonically increasing count of the system clock ticks. Thus, CT (P_1) equals the number of system clock ticks between t_a and t_b. Each data packet generated by a server is stamped with the ending capture time of the packet. Thus, packet P_2 is stamped with time t_d, packet P_3 is stamped with time t_f, and packet P_n is time stamped with time t_y. In this manner, each data packet of a data stream is referenced to a unique position on a data stream time line, such as the one illustrated in Figure 4. Each data packet is thereby associated with a fixed moment in time on the server system clock.

Referring now to Figure 5, the data structure of two data packets 410 and 412 are illustrated. In the example of Figure 4, data packet 410 corresponds to packet P_1 and data packet 412 corresponds to packet P_2.

Each data packet in a data stream comprises a data steam identifier 414, a packet identifier 416, a packet time stamp 418, and packet capture data 420 as illustrated in Figure 5. Data stream identifier 414 uniquely identifies the data stream of which the data packet is a member. The data stream identifier 414 may include an encoded value representing the source of capture data 420 or the identity of the server that generated the data packet 410. It will be apparent to one of ordinary

skill in the art that other means for uniquely identifying a particular data stream may be employed and encoded into data stream identifier 414. Packet identifier 416 is used to uniquely identify a particular data packet within a data stream identified by data stream identifier 414. In the preferred embodiment, packet identifier 416 is a consecutively numbered sequence integer that defines a particular data packet's relative position in a data stream relative to other data packets in the data stream. Packet time stamp 418 is a binary value corresponding to the time at which the capture of capture data 420 was completed by a server system. For the example illustrated in Figure 4, capture for packet P_1 was completed at time t_b. Referring again to Figure 5, the time t_b is present on the server system clock 432 at the time capture is complete for packet P_1. During the process of building packet P_1, a server system captures data into region 420 of packet 410. When capture for this packet is complete, the time from server clock 432 is retrieved and stored in time stamp field 418. For the example of data packet P_1, the time in server clock 432 will be t_b. Having completed the generation of packet 410, the server system queues packet 410 for transmission across a communication medium and begins to generate the next data packet 412 in the data stream. In this case, the data stream identifier 424 will be the same as data stream identifier 414 in the previous data packet 410. Packet identifier 426 will uniquely identify data packet 412. A new set of capture data will be stored in capture data

region 430. When capture for data packet 412 is complete, the server clock 434 is read and the capture complete time (t_d) is stored in time stamp field 428. In this manner, a server system generates a data stream of data packets with a time stamp in each data packet that defines a capture data time line for the data stream.

Referring now to Figure 6, a flowchart illustrates the processing logic of the server system. The server processing starts by building the initial portions of a data packet (processing block 612). The data stream identifier, the packet stream identifier, and any other initial fields of the data packet are filled in processing block 612. The server system captures data and inserts the data into the data packet in processing block 614. For example, digitized video information may be stored into the capture data area 420 of the data packet in processing block 614. The capture and storage of data into the data packet continues until either a predetermined time period has elapsed or the capture data area 420 of the data packet is full. When this occurs, capture is complete and processing path 620 is taken. If capture is not complete, processing path 618 is taken and the data capture process continues at processing block 614.

Once capture is complete, the capture complete time is read from the server system clock in processing block 622. The capture complete time is stored in the data packet in time stamp field 418 in processing block 624. The data packet is queued for transmission across the

communication medium in processing block 626. As long as the capture process is not terminated, processing path 630 is taken to processing block 612 where a new data packet is generated. Processing for the server system terminates through processing path 632 at end bubble 634.

In the synchronization process of the present invention, the client system seeks to establish the same data packet time line for rendering the data packet that was established by the server system during the capture process. Because the transfer of data packets across the communication medium may occur at any arbitrary time or in any arbitrary sequence, the client system must rebuild the data stream into the same form and synchronization as the data stream generated by the server system.

Referring now to Figures 7 through 9, the processing logic for the client system is illustrated. Referring now to Figure 7, the client system receives a data packet in processing block 712. If this data packet is the first packet received by the client system for this data stream, processing path 718 is taken to processing block 720. Otherwise, processing path 716 is taken to the bubble labeled A and illustrated in Figure 8.

In processing block 720, the time stamp t_1 is retrieved from the received first data packet. The virtual server clock or the initial position on the data stream time line (k) is computed in processing block 720. The value (k) is a variable residing in and maintained by the client system that represents the virtual clock of the server system. The

variable (k) represents the value of the server system's clock at the current position in the rendered data stream. The value of (k) is initially set to the time stamp of the first data packet (t_1) less the capture time period of the first data packet of the data stream, CT (P_1). Another value (d) is computed in processing block 722. The variable (d) contains a time delta representing the difference between the client system clock and the virtual server system clock represented by the variable (k). The value of (d) is computed as the difference between the client's current local system time and the virtual server system clock (k). The values in variables (k) and (d) are used extensively during the synchronization process of the present invention as will be described in more detail below.

Once the initial processing for the first data packet is completed, the first data packet P_1 is rendered by the client's system in processing block 724. Rendering a data packet comprises extracting the capture data contained in the data packet and converting the data to a form suitable for transfer to an appropriate output device, such as a display device, a print device, an audio device, or other device compatible with the type of information contained in the data packet. Means and methods for rendering data from a data packet are well known to those of ordinary skill in the art.

On completion of the rendering step in processing block 724, processing control loops back to the top of the processing flow

illustrated in Figure 7 at the bubble labeled B. The next data packet is received in processing block 712. In this case, this data packet will not be the first data packet in the data stream. Therefore, processing path 716 is taken to the bubble labeled A illustrated in Figure 8.

Referring now to Figure 8, the capture system processing logic continues at the bubble labeled A. The next packet received in processing block 712 (P_i) contains time stamp t_i. The value of the virtual server clock (k) is computed in processing block 810. The virtual server clock (k) is set to the client current local system time less the time delta (d). The value of (k) now represents the current position on the data stream time line. The time stamp t_i is used to determine where on this time line the received packet P_i should be placed. The beginning capture time for packet P_i is computed and used for this process. The beginning capture time of packet P_i is the time stamp value t_i less the capture time CT of packet P_i. If the virtual server clock (k) is equal to the starting capture time of packet P_i, processing path 814 is taken to processing block 820 where packet P_i is immediately rendered. Processing for this packet then continues at the bubble labeled C illustrated in Figure 9. If, however, the virtual server clock is not equal to the starting capture time of packet P_i, processing path 816 is taken to decision block 822. If the virtual server clock (k) is greater than the starting capture time of packet P_i, processing path 824 is taken to processing block 828. In this case, the current position in the data

stream time line has already passed the position corresponding to this data packet. Thus, the packet P_i is discarded as a late arrival.

In an alternative embodiment, a portion of a single packet is discarded while the remainder of the packet is rendered. The starting capture time of packet P_i can be used to generate an offset into the packet that corresponds to the virtual sever clock time (k). If the number of samples in the packet is known or if the starting capture time of a subsequent packet is known, the offset into the current packet is generated as an interpolation of these values with the virtual server clock time (k). In this manner, the offset into the packet corresponding to (k) is computed. The portion of the packet P_i from the starting capture time to the offset value is discarded. The portion of packet P_i from the offset value to the and of the packet is rendered.

Processing then continues at the bubble labeled D illustrated in Figure 9. If the virtual server clock is less than the start capture time of packet P_i, processing path 826 is taken to processing block 830 where the rendering of packet P_i is scheduled for a future time corresponding to a number of client system clock ticks in advance of the current time. This future time is computed as a number of client system clock ticks equal to time stamp t_i less the capture time CT of packet P_i less virtual server clock value (k). Means and methods for scheduling the execution of an event at a particular time in a computer is well known to those of

ordinary skill in the art. Processing for the client system continues at the bubble labeled D illustrated in Figure 9.

Referring now to Figure 9, client processing logic continues at the bubble labeled C. On completion of the rendering of packet P_i with a time stamp t_i, the current virtual server clock (k) is recomputed in processing block 910. The value of the virtual server clock (k) is set to the time stamp t_i less the capture time CT of packet P_i plus the time required to render RT packet P_i. The time required to render a packet is represented by the function RT (P_1). The time delta (d) is recomputed in processing block 912. Time delta (d) is set to the client current local system time less the value of a virtual server clock (k). If client processing is selected to continue, processing path 916 is taken to the bubble labeled B illustrated in Figure 7 where the next data packet of the data stream is received. If client processing is selected for termination, processing path 918 is taken and capture processing terminates at end bubble 920.

The process of the present invention described above maintains the virtual server clock (k) and ensures that data packets are rendered by the client system at the same time with respect to the virtual server clock as they were captured. The virtual server clock variable (k) is efficiently maintained in the present invention in that (k) is not updated on every local system clock tick. Rather, a time delta (d) is maintained as the difference between the client local system clock and the virtual server

clock (k). Therefore, as the local system clock is updated, the virtual server clock (k) is also updated. The offset between the clocks is recalculated for each correctly rendered packet to avoid clock skew introduced if the local client clock runs slightly faster or slower than the server system clock.

The process of the present invention described above details an effective method of synchronization of a data stream arriving from one server. The same basic process may be used if there are multiple servers supplying data streams to a single client system. In one case, the streams from one server must be synchronized, but no synchronization is required between streams from different servers. In this case, the client system maintains a set of virtual clock offsets d_1 to d_n, one for each server. In addition, a virtual server clock value k_i is maintained for each server system. When a data packet arrives from a particular server, that server's virtual clock, k_i, is calculated from the correct offset d_i. When a packet is finished rendering, the virtual clock offset d_i for that particular server is updated in a manner similar to the process described above.

In some environments, it is necessary to mix a set of data streams into one unified stream for playback. For example, instead of two servers each sending an audio and video stream to a client for rendering, it may be more efficient to mix the audio into a unified stream which then must be synchronized with both video streams. In this case, the

client system must maintain a set of virtual clock offsets for more than one server. The mixed audio stream must preserve the time stamps from all of the originating servers. Therefore, the time stamp of a packet in the mixed stream is a set of server identifiers and time stamp pairs. The decision to render a packet is now made as a consensus of whether the time stamps for each of the servers is ahead of or behind the virtual clock value for that server. A mixed data packet is scheduled for rendering at a time which satisfies one or more of the virtual clocks. When a mixed packet has been rendered, the virtual clock offsets for all of the servers contributing to the mixed packet are updated.

Thus, an apparatus and method for synchronizing multiple independent data streams in a networked computer system is disclosed. Although the present invention is described herein with reference to a specific preferred embodiment, many modifications and variations therein will readily occur to those of ordinary skill in the art. Accordingly, all such variations and modifications are included within the intended scope of the present invention as defined by the following claims.

CLAIMS

1. In a computer network having a plurality of computer systems coupled by a communication medium, one of said plurality of computer systems being a server system, a different one of said plurality of computer systems being a client system, a method for synchronizing a data stream transferred from said server system to said client system, said method comprising the steps of:

capturing information into a data packet, said capturing step being performed by said server system;

storing time information into said data packet, said time information indicating when said capturing step is complete;

transferring said data packet to said client system across said communication medium;

initializing a virtual server system clock if said data packet is a first data packet is said data stream;

retrieving said time information from said data packet;

comparing said time information with said virtual server system clock; and

processing said information in said data packet if said time information indicates a time equal to a time indicated by said virtual server system clock.

2. The method as claimed in claim 1 further including the step of:

scheduling processing of said information in said data packet if said time information indicates a time later than a time indicated by said virtual server system clock.

3. The method as claimed in claim 1 further including the step of:

discarding said data packet if said time information indicates a time earlier than a time indicated by said virtual server system clock.

4. The method as claimed in claim 1 further including the steps of:

calculating an index into said information in said data packet corresponding to a value in said virtual server system clock;

discarding a first portion of said data packet from a beginning position in said data packet to a position in said data packet indicated by said index; and

processing a second portion of said data packet from a position in said data packet indicated by said index to an ending position of said data packet.

5. In a computer network having a plurality of computer systems coupled by a communication medium, one of said plurality of computer systems being a server system, a different one of said plurality of computer systems being a client system, a method for synchronizing

data streams transferred from said server system to said client system, said method comprising the steps of:

capturing information into a first set of data packets, said capturing step being performed by said server system;

capturing information into a second set of data packets, said capturing step being performed by said server system;

storing time information into said first set of data packets, said time information indicating when said capturing step is complete;

storing time information into said second set of data packets, said time information indicating when said capturing step is complete;

transferring said first set of data packets to said client system across said communication medium in a first data stream;

transferring said second set of data packets to said client system across said communication medium in a second data stream;

initializing a virtual server system clock if a data packet is a first data packet is a data stream;

retrieving said time information from data packets of said first set or said second set of data packets;

comparing said time information with said virtual server system clock; and

processing said information from data packets of said first set or said second set of data packets if said time information indicates a time equal to a time indicated by said virtual server system clock.

6. The method as claimed in claim 5 further including the step of:

scheduling processing of said information from data packets of said first set or said second set of data packets if said time information indicates a time later than a time indicated by said virtual server system clock.

7. The method as claimed in claim 5 further including the step of:

discarding data packets of said first set or said second set of data packets if said time information indicates a time earlier than a time indicated by said virtual server system clock.

8. The method as claimed in claim 5 further including the steps of:

calculating an index into said information in data packets of said first set or said second set of data packets corresponding to a value in said virtual server system clock;

discarding a first portion of data packets of said first set or said second set of data packets from a beginning position in said data packets to a position in said data packets indicated by said index; and

processing a second portion of data packets of said first set or said second set of data packets from a position in said data packets indicated by said index to an ending position of said data packets.

9. The method as claimed in claim 1 wherein said information captured in said data packet comprises either video or audio information.

10. The method as claimed in claim 5 wherein said information captured in said first set of data packets comprises video information and said information captured in said second set of data packets comprises audio information.

11. In a computer network having a plurality of computer systems coupled by a communication medium, one of said plurality of computer systems being a server system, a different one of said plurality of computer systems being a client system, an apparatus for synchronizing a data stream transferred from said server system to said client system, said apparatus comprising:

means for capturing information into a data packet, said means for capturing residing in said server system;

means for storing time information into said data packet, said time information indicating when said means for capturing is finished;

means for transferring said data packet to said client system across said communication medium;

means for initializing a virtual server system clock if said data packet is a first data packet is said data stream;

means for retrieving said time information from said data packet;

means for comparing said time information with said virtual server system clock; and

means for processing said information in said data packet if said time information indicates a time equal to a time indicated by said virtual server system clock.

12. The apparatus as claimed in claim 11 further including:

means for scheduling processing of said information in said data packet if said time information indicates a time later than a time indicated by said virtual server system clock.

13. The apparatus as claimed in claim 11 further including:

means for discarding said data packet if said time information indicates a time earlier than a time indicated by said virtual server system clock.

14. The apparatus as claimed in claim 11 further including:

means for calculating an index into said information in said data packet corresponding to a value in said virtual server system clock;

means for discarding a first portion of said data packet from a beginning position in said data packet to a position in said data packet indicated by said index; and

means for processing a second portion of said data packet from a position in said data packet indicated by said index to an ending position of said data packet.

15. In a computer network having a plurality of computer systems coupled by a communication medium, one of said plurality of computer systems being a server system, a different one of said plurality of computer systems being a client system, an apparatus for synchronizing data streams transferred from said server system to said client system, said apparatus comprising:

means for capturing information into a first set of data packets, said means for capturing residing in said server system;

means for capturing information into a second set of data packets, said means for capturing residing in said server system;

means for storing time information into said first set of data packets, said time information indicating when said means for capturing is finished;

means for storing time information into said second set of data packets, said time information indicating when said means for capturing is finished;

means for transferring said first set of data packets to said client system across said communication medium in a first data stream;

means for transferring said second set of data packets to said client system across said communication medium in a second data stream;

means for initializing a virtual server system clock if a data packet is a first data packet is a data stream;

means for retrieving said time information from data packets of said first set or said second set of data packets;

means for comparing said time information with said virtual server system clock; and

means for processing said information from data packets of said first set or said second set of data packets if said time information indicates a time equal to a time indicated by said virtual server system clock.

16. The apparatus as claimed in claim 15 further including:

means for scheduling processing of said information from data packets of said first set or said second set of data packets if said time information indicates a time later than a time indicated by said virtual server system clock.

17. The apparatus as claimed in claim 15 further including:

means for discarding data packets of said first set or said second set of data packets if said time information indicates a time earlier than a time indicated by said virtual server system clock.

18. The apparatus as claimed in claim 15 further including:

means for calculating an index into said information in data packets of said first set or said second set of data packets corresponding to a value in said virtual server system clock;

means for discarding a first portion of data packets of said first set or said second set of data packets from a beginning position in said data packets to a position in said data packets indicated by said index; and

means for processing a second portion of data packets of said first set or said second set of data packets from a position in said data packets indicated by said index to an ending position of said data packets.

19. The apparatus as claimed in claim 11 wherein said information captured in said data packet comprises either video or audio information.

20. The apparatus as claimed in claim 15 wherein said information captured in said first set of data packets comprises video information and said information captured in said second set of data packets comprises audio information.

21. A computer network having a plurality of computer systems, substantially as hereinbefore described, with reference to the accompanying drawings.

Patents Act 1977 Examiner's report to the Comptroller under Section 17 (The Search report) 34	Application number GB 9416276.5
Relevant Technical Fields (i) UK Cl (Ed.M) H4P (PPF, PSB, PPS) (ii) Int Cl (Ed.5) H04L 7/10, 7/00, 12/56	Search Examiner J P Coules
	Date of completion of Search 21 September 1994
Databases (see below) (i) UK Patent Office collections of GB, EP, WO and US patent specifications. (ii) ONLINE DATABASES : WPI, INSPEC	Documents considered relevant following a search in respect of Claims :- 1-21

Categories of documents

X:	Document indicating lack of novelty or of inventive step.	**P:**	Document published on or after the declared priority date but before the filing date of the present application.
Y:	Document indicating lack of inventive step if combined with one or more other documents of the same category.	**E:**	Patent document published on or after, but with priority date earlier than, the filing date of the present application.
A:	Document indicating technological background and/or state of the art.	**&:**	Member of the same patent family; corresponding document.

Category	Identity of document and relevant passages		Relevant to claim(s)
A	EP 0234861 A2	(AMERICAN TELEPHONE AND TELEGRAPH CO) see Figure 11, page 28	1,5,11,15
A	EP 0059821 A1	(FOUR DEE AKTIEBOLOG) see time marker clock 20	1,5,11,15

Databases: The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
13 November 2003 (13.11.2003)

PCT

(10) International Publication Number
WO 03/093950 A2

(51) International Patent Classification[7]: G06F

(21) International Application Number: PCT/US03/14154

(22) International Filing Date: 6 May 2003 (06.05.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

60/378,415	6 May 2002 (06.05.2002)	US
60/388,887	14 June 2002 (14.06.2002)	US
60/452,230	4 March 2003 (04.03.2003)	US

(71) Applicants and

(72) Inventors: **GOLDBERG, David** [US/US]; 4055 Pinion Drive, Boulder, CO 80303 (US). **SIMON, Neil, R.** [US/US]; 3084 4th Street, Boulder, CO 80304 (US). **GOLDBERG, Martha, B.** [US/US]; 4055 Pinion Drive, Boulder, CO 80303 (US). **GOLDBERG, Miriam, D.** [US/US]; 4055 Pinion Drive, Boulder, CO 80303 (US). **GOLDBERG, Benjamin, M.** [US/US]; 4055 Pinion Drive, Boulder, CO 80303 (US).

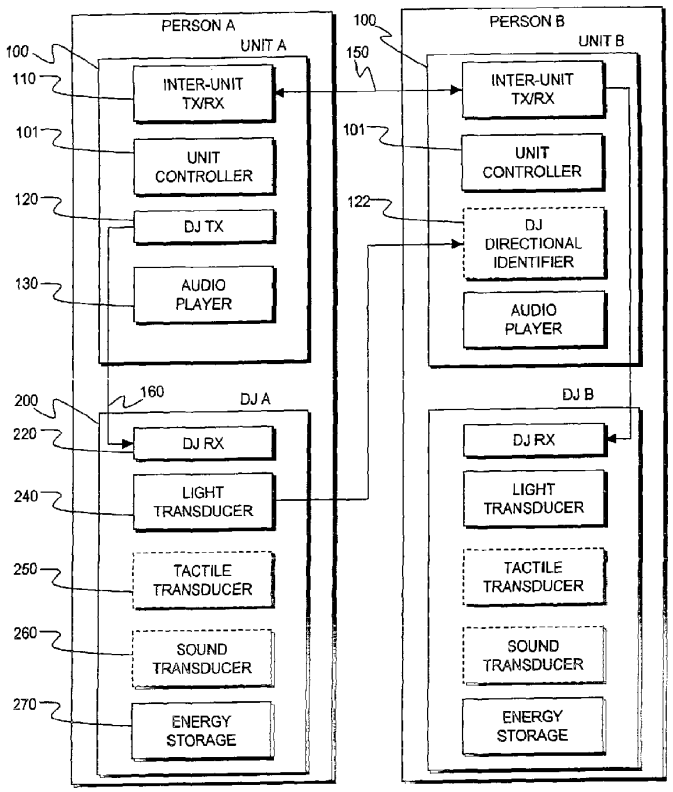
(74) Agent: **McINTOSH, Michael, D.**; 8000 E. Prentice Avenue, Suite B-6, Greenwood Village, CO 80111 (US).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,

[Continued on next page]

(54) Title: LOCALIZED AUDIO NETWORKS AND ASSOCIATED DIGITAL ACCESSORIES



(57) Abstract: Individuals share music through portable audio devices that synchronously play music by transmitting music signals from one device to the other. Upon request, other individuals can join such a sharing group as a member, and these new members can either be accepted automatically, or through voting by existing group members. So that members of a group can converse, the fraction of ambient sound admitted by the earphones of the portable device can be manually set. In order to enhance the enjoyment of the music, wearable transducers, such as light emitting diodes, respond to signals related to the beat of the music. Signals can be automatically generated, or manually specified, and signals thereafter stored for playback with the music to which it is related. Wearable transducers associated with individuals at a large music event can be set to receive common signals such that the transducers respond in synchrony with one another

WO 03/093950 A2

ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declaration under Rule 4.17:

— *of inventorship (Rule 4.17(iv)) for US only*

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

LOCALIZED AUDIO NETWORKS AND ASSOCIATED DIGITAL ACCESSORIES

Cross Reference To Related Patent Applications

This application is related to and claims priority from Provisional Patent Application No. 60/378,415, filed May 6, 2002, titled "Localized Audio Networks and Associated Digital Accessories," and from Provisional Patent Application No. 60/388,887, filed June 14, 2002, titled " Localized Audio Networks and Associated Digital Accessories," and from Provisional Patent Application No. 60/452,230, filed March 4, 2003, titled " Localized Audio Networks and Associated Digital Accessories," the contents of each of which are incorporated herein by reference.

Technical Field

The present invention relates to localized wireless audio networks for shared listening of recorded music, and wearable digital accessories for public music-related display, which can be used in conjunction with one another.

Background

Portable audio players are popular consumer electronic products, and come in a variety of device formats, from cassette tape "boom boxes" to portable CD players to digital flash-memory and hard-disk MP3 players. While boom boxes are meant to make music to be shared among people, most of the portable audio players are designed for single person use. While some of this orientation to personal music listening is due to personal preference, other important considerations are the technical difficulties of reproducing music for open area listening with small, portable devices, as well as the social imposition of listening to music in public places with other people who do not wish to listen to the same music, or who are listening to different music that would interfere with one's own music.

There are numerous audio devices that are designed to allow the transfer of music from one portable audio device to another, especially through those that store music in the MP3 audio format. These devices suffer from two main difficulties: firstly, listeners still do not hear the music simultaneously, which is the optical manner to share music, and secondly, there are serious copyright issues associated with the transfer of music files. Thus, it would be preferable for the transfer of the music for simultaneous enjoyment, and which did not

result in a permanent transfer of the music files between the devices, so as not to infringe on the intellectual property rights of the music owners.

Given the sharing of music, listeners will on occasion want to purchase the music for themselves. In such case, it would be beneficial for the user to have a way to obtain the music with minimal effort. It would further be desirable for there to be a way to keep track of the person from whom the listener heard the music, so that the person could be in some way encouraged or compensated.

The earphones associated with a portable music player admit a relatively constant fraction of ambient sound. If listening to music with a shared portable music device, however, one might at times want to talk with a friend, and at times listen to music without outside audible distraction. In such case, it would be desirable to have an earphone for which the amount of external ambient sound could be manually set.

Furthermore, many people like to show their individual preferences, to exhibit themselves, and to demonstrate their group membership. Furthermore, music preferences and listening to music together are among the more important means by which individuals express their individual and group identities. It would be beneficial for there to be a way for individuals to express themselves through their music, and for groups of individuals listening to music together, to be able to demonstrate their group enjoyment of the music.

One means for a person to express their identity through motion would be through having wearable transducers wherein the transduction signal is related to the music. If the transducer were a light transducer, this would result in a display of light related to the music that was being listened to. It would be further beneficial if there were means by which a person could generate control signals for the transducer so that instead of a wholly artificial response to the music, the transducer showed a humanly interpreted display. It would be preferable if these signals could be shared between people along with music files, so that others could entertain or appreciate the light display so produced.

At popular music concerts, there is often a "light show" that pulsates in rough relation to the music. In contrast to the generally vigorous light show, the patrons at the concerts often have light bracelets or other such static displays which are used to join with the displays on the stage. It would be beneficial for there to be a way in which patrons could participate in the light show in order to enhance their enjoyment of the concert.

It is to the solution of these and other problems that the present invention is directed.

Summary of the Invention

It is an object of the present invention to provide users a means of listening to music together using mobile devices.

It is also an object of the present invention to provide users a means of choosing with whom to listen to music.

It is additionally an object of the present invention to provide users the ability to monitor the people that are listening together.

It is furthermore an object of the present invention to provide users a means of expressing their enjoyment of the music they are listening to through visual displays of wearable accessories.

It is yet another object of the present invention to provide users a means of demonstrating their identity with other people they are listening to music with.

It is still further an object of the present invention to provide users to provide users with means to choreograph the visual displays.

Additional objects, advantages and novel features of this invention shall be set forth in part in the description that follows, and will become apparent to those skilled in the art upon examination of the following specification or may be learned through the practice of the invention. The objects and advantages of the invention may be realized and attained by means of the instrumentalities, combinations, and methods particularly pointed out in the appended claims.

To achieve the foregoing and other objects and in accordance with the purposes of the present invention, as embodied and broadly described therein, the present invention is directed to a method for sharing music from stored musical signals between a first user with a first music player device and at least one second user with at least one second music player device. The method includes the step of playing the musical signals for the first user on the first music player device while essentially simultaneously wirelessly transmitting the musical signals from the first music player device to the at least one second music player device. The method additionally includes receiving the musical signals by the at least one second player device, such that the musical signals can be played on the at least one second player device essentially simultaneously with the playing of the musical signals on the first music player device. In this method, the first and the at least one second users are mobile and maintain less than a predetermined distance.

The present invention is also related to a system of music sharing for a plurality of users. The system includes a first sharing device and at least one second sharing device, each comprising a musical signal store, a musical signal transmitter, a musical signal receiver, and a musical signal player. Furthermore, the system comprises a broadcast user operating the first sharing device and at least one member user operating the at least one second sharing device. The broadcast user plays the musical signal for his own enjoyment on the first sharing device and simultaneously transmits the musical signal to the receiver of the at least one second sharing device of the at least one member user, on which the musical signal is played for the at least one member user. The broadcast user and the at least one member user hear the musical signal substantially simultaneously.

The present invention yet further is related to a wireless communications system for sharing audio entertainment between a first mobile device and a second mobile device in the presence of a non-participating third mobile device. The system includes an announcement signal transmitted by the first mobile device for which the second mobile device and the third mobile device are receptive. In addition, the system includes a response signal transmitted by the second mobile device in response to the announcement signal for which the first mobile device is receptive and for which the third mobile device is not receptive. Also, the system includes an identifier signal transmitted by the first mobile device to the second mobile device in response to the response signal, and which is not receptive to the third mobile device. Finally, the system includes a broadcast signal comprising audio entertainment that is transmitted by the first mobile device, and which is receptive by the second mobile device on the basis of the reception of the identifier signal.

The present invention additionally is related to an audio entertainment device. The device includes a signal store that stores an audio entertainment signal, a transmitter that can transmit the stored audio entertainment signal, a receiver that can receive the transmitted audio entertainment signal from a transmitter of another such device, and a player that can play audio entertainment from a member selected from the group of stored audio entertainment signals or audio entertainment signals transmitted from the transmitter of another such device.

The present invention yet still is related to a system for identifying a first device that introduces a music selection to a second device. The system includes a mobile music transmitter operated by the first device and a mobile music receiver operated by the second device. In addition, the system includes a music signal comprising the music selection transmitted by the transmitter and received by the receiver, an individual musical identifier

that is associated with the music selection, and an individual transmitter identifier that identifies the transmitter. The transmitter identifier and the individual music identifier are stored in association with each other in the receiver.

The present invention is still further related to an audio entertainment device. The device includes a wireless transmitter for the transmission of audio entertainment signals and a wireless receiver for the reception of the transmitted audio entertainment signals from a transmitter of audio entertainment signals. A first manually-separable connector for electrically connecting with an audio player allows transfer of audio entertainment signals from the player to the device. The device also includes a second connector for connecting with a speaker and a control to manually switch between at least three states. In the first state the speaker plays audio entertainment signals from the audio player and the transmitter does not transmit the audio entertainment signals. In the second state the speaker plays audio entertainment signals from the audio player and the transmitter essentially simultaneously transmits the audio entertainment signals. In the third state the speaker plays audio entertainment signals received by the receiver.

The present invention also still is related to a system for the sharing of stored music between a first user and a second user. The system includes a first device for playing music to the first user, comprising a store of musical signals. A first controller prepares musical signals from the first store for transmission and playing, and a first player takes musical signals from the first controller and plays the signals for the first user. A transmitter is capable of taking the musical signals from the controller and transmitting the musical signals via wireless broadcast. A second device for playing music to the second user comprises a receiver receptive of the transmissions from the transmitter of the first device, a second controller that prepares musical signals from the receiver for playing, and a second player that takes musical signal from the second controller and plays the signals for the second user. The first user and the second user hear the musical signals at substantially the same time.

The present invention also is related to an earphone for listening to audio entertainment allowing for the controlled reception of ambient sound by a user. The earphone includes a speaker that is oriented towards the user's ear and an enclosure that reduces the amount of ambient noise perceptive to the user. In addition, a manually-adjustable characteristic of the enclosure adjusts the amount of ambient sound perceptive to the user.

The present invention is further related to a mobile device for the transmission of audio entertainment signals. The mobile device includes an audio signal store for the storage

of the audio entertainment signals, and an audio signal player for the playing of the audio entertainment signals. The device also includes a wireless transmitter for the transmission of the audio entertainment signals and a transmitter control to manually switch between two states consisting of the operation and the non-operation of the audio transmitter.

The present invention yet still is related to a mobile device for the reception of digital audio entertainment signals. The mobile device includes an audio signal store for the storage of the digital audio entertainment signals and an audio receiver for the reception of external digital audio entertainment signals from a mobile audio signal transmitter located within a predetermined distance of the audio receiver. The device also includes a receiver control with at least a first state and a second state. An audio signal player plays digital audio entertainment signals from the audio signal store when the receiver control is in the first state, and plays digital audio entertainment signals from the audio receiver when the receiver control is in the second state.

The present invention furthermore relates to a method for the shared enjoyment of music from stored musical signals between a first user with a first music player device and at least one second user with at least one second music player device. The method includes the step of playing the musical signals for the first user on the first music player device while essentially simultaneously wirelessly transmitting synchronization signals from the first music player device to the at least one second music player device. The method also includes receiving the synchronization signals by the at least one second player device. The synchronization signals allow the musical signals on the at least one second player device to be played essentially simultaneously with the playing of the musical signals on the first music player device. The first and the at least one second users are mobile.

The present invention yet furthermore relates to a wireless communications system for sharing audio entertainment between a first mobile device and a second mobile device. The system includes a broadcast identifier signal transmitted by the first mobile device to the second mobile device. A personal identifier signal is transmitted by the second mobile device to the first mobile device. A broadcast signal comprising audio entertainment is transmitted by the first mobile device of which the second device is receptive. The first mobile device and the second mobile device have displays which can display the identifier signal that they receive and the second mobile device can play the audio entertainment from the broadcast signal that it receives.

The present invention also relates to a method for enhancing enjoyment of a musical selection. The method includes the steps of obtaining control signals related to the musical

selection, transmitting the control signals wirelessly, receiving the control signals, and converting the control signals to a humanly-perceptible form.

The present invention further yet relates to a method for generating and storing control signals corresponding to musical signals. The method includes the steps of playing musical signals for a user and receiving manual input signals from the user that are produced substantially in synchrony with the music. The method also includes the steps of generating control signals from the input signals, and storing the control signals so that they can be retrieved with the musical signals.

The present invention still additionally relates to a wearable personal accessory. The accessory includes an input transducer taken from the group consisting of a microphone and an accelerometer. The transducer generates a time-varying input transduction signal. The accessory also includes a controller that accepts the input transduction signal, and generates an output transducer signal whose signal varies in amplitude with time. An output transducer receptive of the output transducer signal provides a humanly-perceptible signal. An energy source powers the input transducer, controller and output transducer.

The present invention also still relates to a wearable personal accessory controlled via wireless communications. The accessory includes a wireless communications receiver that is receptive of an external control signal. The accessory also includes a controller that accepts the external control signal and that generates a time-varying visual output transducer signal. A visual output transducer is receptive of the output transducer signal, and provides a humanly-perceptible visual signal. An energy store powers the receiver, controller and output transducer. The visual output transducer generates visually-perceptive output.

The present invention still further relates to a device for converting user tactile responses to stored music into a stored control signal. The device includes a player that plays stored music audible to the user and a manually-operated transducer that outputs an electrical signal. The transducer is actuated by the user in response to the music. A controller receives the electrical signal and outputs a control signal and a store receives the control signal and stores it.

The present invention furthermore relates to a music player that wirelessly transmits control signals related to the music, wherein the control signals control a wearable electronic accessory. The music player includes a store of music signal files and a controller that reads a musical signal file from the store and generates audio signals. The controller further generates the control signals. A transducer converts the audio signals into sound audible to

the user and a wireless transmitter transmits the control signal to the wearable electronic accessory.

The present invention yet relates to a music player that wirelessly transmits control signals related to the music, wherein the control signals control a wearable electronic accessory. The music player includes a store of music signal files and a second store of control signal files associated with the music signal files. A controller reads a musical signal file from the store and generates audio signals. The controller further reads an associated control signal file. A transducer converts the audio signals into sound audible to the user, and a wireless transmitter transmits the control signals from the associated control signal file to the wearable electronic accessory.

The present invention also relates to a system for exhibition of music enjoyment. The system includes a source of music signals, a controller that generates control signals from the music signals, and a transmitter of the control signals. The transmission of the control signals is synchronized with the playing of the music signals. In addition, the system includes a receiver of the control signals and a transducer that responds to the control signals.

The present invention further relates to a method for transferring a wearable-accessory control file stored on a first device to a second device in which an associated music file is stored. The method includes the steps of storing on the first device the name of the music file in conjunction with the control file with which it is associated and requesting by the second device of the first device for a control file stored in conjunction with the name of the music file. In addition, the method includes the step of transferring the control file from the first device to the second device. The control file is stored on the second device in conjunction with the name of the associated music file.

The present invention also relates to a device for transmitting control signals to a wearable accessory receptive of such control signals. The device includes a manually-separable input connector for connecting to an output port of an audio player. Audio signals are conveyed from the audio player to the device across the connector. The device also includes a controller for generating control signals from the audio signals and a transmitter for transmitting the control signals.

Brief Description Of The Drawings

Figure 1 is a schematic block diagram of a local audio network comprised of two linked audio units operated by two persons, and associated digital jewelry conveyed by the two persons.

Figure 2A is a schematic block diagram of a DJ with multiple independently controlled LED arrays.

Figure 2B is a schematic block diagram of a DJ with an LED array with independently controlled LEDs.

Figures 3A-C are schematic block diagrams of unit elements used in inter-unit communications.

Figure 4 is a schematic flow diagram of DJ entraining.

Figures 5A-B are schematic block diagrams of DJs associated with multiple people bound to the same master unit.

Figure 6 is a schematic block diagram of a cluster comprising a broadcast unit and multiple receive units, with an external search unit.

Figure 7 is a schematic diagram of a broadcast unit transmission.

Figure 8A is a schematic block diagram of audio units with self-broadcast so that audio output is highly synchronized.

Figure 8B is a schematic flow diagram for synchronous audio playing with multiple rebroadcast.

Figures 9A and 9B are schematic block diagrams of hierarchically-related clusters.

Figure 10 is a top perspective view of an earphone with manually adjustable external sound ports.

Figures 11 A and B are cross-sectional diagrams of a earpiece with an extender to admit additional ambient sound.

Figure 12A is a schematic diagram of a modular audio unit.

Figure 12B is a schematic diagram of modular digital jewelry.

Figure 12C is a schematic block diagram of a modular transmitter that generates and transmits control signals for digital jewelry from an audio player.

Figure 13A is a schematic cross-section through a search unit and a broadcast unit in which communications are provided via visible or infrared LED emission in search transmission mode.

Figure 13B is a schematic cross-section through a search unit and a broadcast unit n which communications are provided via a visible or infrared laser in search transmission mode.

Figure 13C is a schematic cross-section through a search unit and a broadcast unit in which communications are provided via visible or infrared emission from a digital jewelry element in broadcast transmission mode.

Figure 13D is a schematic cross-section through a search unit and a broadcast unit in which communications are provided via contact in mutual transmission mode.

Figure 13E is a schematic cross-section through a search unit and a broadcast unit in which communications are provided via sonic transmissions in broadcast transmission mode.

Figure 13F is a schematic cross-section through a search unit and a broadcast unit in which communications are provided via radio frequency transmissions in broadcast transmission mode.

Figure 14A is a schematic block diagram of the socket configurations on the broadcast unit and the receive unit.

Figure 14B is a schematic block flow diagram of using IP sockets for establishing and maintaining communications between a broadcast unit and the receive unit, according to the socket diagram of Figure 14A.

Figure 15 is a schematic block diagram of the IP socket organization used with clusters comprising multiple members.

Figure 16 is a schematic block flow diagram of transfer of control between the broadcast unit and the first receive unit.

Figure 17 is a matrix of DJ and searcher preferences and characteristics, illustrating the matching of DJ and searcher in admitting a search to a cluster.

Figure 18A is a screenshot of an LCD display of a unit, taken during normal operation.

Figure 18B is a screenshot of an LCD display of a unit, taken during voting for a new member.

Figure 19 is a table of voting schemes for the acceptance of new members into a cluster.

Figure 20 is a time-amplitude trace of an audio signal automatically separated into beats.

Figure 21A is a block flow diagram of a neural network method of creating DJ transducer control signals from an audio signal as shown in Figure 20.

Figure 21B is a block flow diagram of a deterministic signal analysis method of creating DJ transducer control signals from an audio signal as shown in Figure 20.

Figure 21C is a schematic flow diagram of a method to extract fundamental musical patterns from an audio signal to create DJ control signals.

Figure 21D is a schematic flow diagram of an algorithm to identify a music model resulting in a time signature.

Figure 22A is a top-view diagram of an audio unit user interface, demonstrating the use of buttons to create DJ control signals.

Figure 22B is a top-view diagram of a hand-pad for creating DJ control signals.

Figure 22C is a schematic block diagram of a set of drums used for creating DJ control signals.

Figure 23A is a schematic block flow diagram of the synchronized playback of an audio signal file with a DJ control signal file, using transmission of both audio and control signal information.

Figure 24 is a schematic block diagram of a DJ unit with associated input transducers.

Figure 25 is a schematic flow diagram indicating music sharing using audio devices, providing new means of distributing music to customers.

Figure 26 is a schematic diagram of people at a concert, in which DJs conveyed by multiple individuals are commonly controlled.

Figure 27 is a schematic block flow diagram of using a prospective new member's previous associations to determine whether the person should be added to an existing cluster.

Figure 28 is a block flow diagram indicating the steps used to maintain physical proximity between the broadcast unit and the receive unit via feedback to the receive unit user.

Figures 29A is a schematic block diagram of the connection of an Internet-enabled audio unit with an Internet device through the Internet cloud, using an Internet access point.

Figures 29B is a schematic block diagram of the connection of an Internet-enabled audio unit with an Internet device through the Internet cloud, with an audio unit directly connected to the Internet cloud.

Figure 30 comprises tables of ratings of audio unit users.

Figure 31 is tables of DJ, song and transaction information according to the methods of Figure 25.

Figure 32A is a schematic block diagram of maintaining privacy in open transmission communications.

Figure 32B is a schematic block diagram of maintaining privacy in closed transmission communication.

Figure 33 is a schematic block diagram of a hierarchical cluster, as in Figure 9A, in which communications between different units is cryptographically or otherwise restricted to a subset of the cluster members.

Figure 34A is a schematic block flow diagram of the synchronization of music playing from music files present on the units 100.

Figure 34B is a schematic layout of a synchronization record according to Figure 34A.

Figure 35 is a schematic block diagram of DJ switch control for both entraining and wide-area broadcast.

Figure 36 is a schematic block diagram of mode switching between peer-to-peer and infrastructure modes.

Best Mode for Carrying-Out the Invention

Overview

Figure 1 is a schematic block diagram of a local audio network comprised of two linked audio units 100 operated by two persons, and associated digital jewelry 200 conveyed by the two persons. The persons are designated Person A and Person B, their audio units 100 are respectively Unit A and Unit B, and their digital jewelry 200 are denoted respectively DJ A and DJ B. In this patent specification, "DJ" is used to denote either the singular "digital jewel" or the plural "digital jewelry".

Each unit 100 is comprised of an audio player 130, and an inter-unit transmitter/receiver 110. In addition, each unit 100 comprises a means of communication with the digital jewelry, which can be either a separate DJ transmitter 120 (Unit A), or which can be part of the inter-unit transmitter/receiver 110 (Unit B). Furthermore, unit 100 can optionally comprise a DJ directional identifier 122, whose operation will be described below. Also, unit 100 will generally comprise a unit controller 101, which performs various operational and executive functions of intra-unit coordination, computation, and data transfers. The many functions of the controller 101 will not be discussed separately below, but will be described with respect to the general functioning of the unit 100.

In operation, Unit A audio player 130 is playing recorded music under the control of a person to be designated User A. This music can derive from a variety of different sources and storage types, including tape cassettes, CDs, DVDs, magneto-optical disks, flash memory, removable disks, hard-disk drives or other hard storage media. Alternatively, the audio signals can be received from broadcasts using analog (e.g. AM or FM) or digital radio receivers. Unit A is additionally broadcasting a signal through DJ transmitter 120, which is received by DJ 200 through a DJ receiver 220 that is worn or otherwise conveyed by User A.

It should be noted that the audio signals can be of any sound type, and can include spoken text, symphonic music, popular music or other art forms. In this specification, the terms audio signal and music will be used interchangeably.

The DJ 200 transduces the signal received by the DJ receiver 220 to a form perceptible to the User A or other people near to him. This transduced form can include audio, visual or tactile elements, which are converted to their perceptible forms via a light transducer 240, and optionally a tactile transducer 250 or an audio transducer 260. The transducers 240, 250 and 260 can either directly generate the perceptible forms directly from the signals received by the DJ receiver 220, or can alternatively incorporate elements to filter or modify the signals prior to their use by the transducers.

When a second individual, User B, perceives the transduced forms produced by User A DJ 200, he can then share the audio signal generated by the audio player 130 of Unit A, by use of the inter-unit transmitter/receiver 110 of Unit A and a compatible receiver 110 of Unit B. Audio signals received by Unit B from Unit A are played using the Unit B audio player 130, so that User A and User B hear the audio signals roughly simultaneously. There are a variety of means by which the Unit B can select the signal of Unit A, but a preferred method is for there to be a DJ directional identifier 122 in Unit B, which can be pointed at the DJ of User A and which receives information needed to select the Unit A signal from the User A DJ, whose transduced signal is perceptible to User B.

Given the audio signal now being exchanged between Unit A and Unit B, User A and User B can experience the same audio signal roughly simultaneously. Within the spirit of the present invention, it is preferable for the two users to hear the audio signals within 1 second of one another, and more preferable for the users to hear the audio signals within 200 milliseconds of one another, and most preferable for the users to hear the audio signals within 50 milliseconds of one another. Furthermore, DJs 200 being worn by User A and User B can receive signals from their respective units, each emitting perceptible forms of their signals. Preferably, the transduced forms expressed by the DJs 200 are such as to enhance the personal or social experience of the audio being played.

Unit 100 Structure

Units 100 comprise a device, preferably of a size and weight that is suited for personal wearing or transport, which is preferably of a size and format similar to that of a conventional portable MP3 player. The unit can be designed on a "base" of consumer electronics products

such as cell phones, portable MP3 players, or personal digital assistants (PDAs), and indeed can be configured as an add-on module to any of these devices.

In general, the unit 100 will comprise, in addition to those elements described in Figure 1, other elements such as a user interface (e.g. an LCD or OLED screen, which can be combined with a touch-sensitive screen, keypad and/or keyboard), communications interfaces (e.g. Firewire, USB, or other serial communications ports), permanent or removable digital storage, and other components.

The audio player 130 can comprise one or more modes of audio storage, which can include CDs, tape, DVDs, removable or fixed magnetic drives, flash memory, or other means. Alternatively, the audio can be configured for wireless transmission, including AM/FM radio, digital radio, or other such means. Output of the audio signal so generated can comprise wireless or wired headphones or wired or wireless external speakers.

It is also within the spirit of the present invention that the unit 100 can have only receive capabilities, without having separate audio information storage or broadcast capabilities. In concept, such a device can have as little user interface as an on/off button, a button to cause the unit 100 to receive signals from a new "host", and a volume control. Such devices can be very small and be built very inexpensively.

Unit 100 Audio Output

One of the goals of the present invention is to assist communications between groups of people. In general, with mobile audio devices, the music is listened to through headphones. Many headphones are designed so as to reduce to the extent possible the amount of sound which is heard from outside of the headphones. This, however, will have the general effect of reducing the verbal communications between individuals.

In order to avoid this potential problem, it is within the teachings of the present invention that headphones or earphones be provided that allow ambient sound, including a friend's voice, to be easily perceptible to the wearer of the headphones, and that such headphones can be provided that variably allow such sound to be accessible for the headphone's wearer. Such arrangement of the headphones can be obtained either through physical or electronic means. If through electronic means, the headphones can have a microphone associated with them, through which signals received are played back in proportion through the headphone speakers, said proportion being adjustable from substantially all sound being from the microphone to substantially no sound being from the microphone. This microphone can also be a part of a noise cancellation system, such that the

phase of the playback is adjustable – if the phase is inverted relative to the ambient sound signal, then the external noise is reduced, whereas if the phase is coincident with the ambient sound signal, then the ambient sounds are enhanced.

Figure 10 is a top perspective view of an earphone 900 with adjustable external sound ports. A speaker element 940 is centrally located, and the outside circumferential surface is a rotatable sound shield 910 in which sound ports 930 are placed. The sound ports 930 are open holes to admit sound. Beneath the sound shield 930 is a non-rotatable sound shield in which fixed sound ports 920 are placed in a similar arrangement. As the sound shield 910 is rotated manually by the user, the sound ports 930 and the fixed sound ports 920 come into registration, so that open ports between sources of ambient noise and the outer ear chamber is created, increasing the amount of ambient sound that the user perceives.

Figures 11 A and B are cross-sectional diagrams of an earpiece with an extender 980 that admits additional ambient sound. In Figure 11A, the face of a speaker 960 with a cord 970 is covered with a porous foam block 950 that fits snugly into the ear. While some ambient sound is accessible to the ear through the foam block 950, the majority of the sound is input is impeded. In Figure 11B, the foam extender 980 is placed over the foam block 950 so that a formed shape at the distal end of the extender 980 fits snugly into the ear. A hollow cavity 982 can be allowed in the extender 980 so as to reduce the sound impedance from the speaker 960 to the ear. Ambient sound is allowed into the space between the speaker 960 and the distal end of the extender 980 (shown by the arrows).

Many other arrangements are allowed within the spirit of the present invention to allow ambient sound to more easily access the user's ear, including adjustable headphones or earplugs as in Figure 10, or accessories that can modify the structure of existing earphones and headphones, as in Figure 11B. Such effects can include increasing the number of apertures admitting ambient sound, increasing the size of an aperture (e.g. by adjusting the overlap between two larger apertures), changing the thickness or number of layers in the enclosure, or by placing a manually detachable cup that covers the earphone and ear channel so as to reduce ambient sound.

DJ 200 Transducers

DJs 200 will have a number of common elements, including communications elements, energy storage elements, and control elements (e.g. a manual ON/OFF switch or a switch to signal DJ entraining, as will be described below). In this section, the structure and function of transducers will be described.

Light Transducers 240

The DJ 200 transducers are used to create perceptible forms of the signals received by the receiver 220. Light transduction can include the use of one or more light-emitting devices, which can conveniently be colored LEDs, OLEDs, LCDs, or electroluminescent displays, which can be supplemented with optical elements, including mirrors, lenses, gratings, and optical fibers. Additionally, motors, electrostatic elements or other mechanical actuators can be used to mechanically alter the directionality or other properties of the light transducers 240. There can be either a single device or an array of devices, and if more than a single device, can display in synchrony, or can be "choreographed" to display in a temporal and/or spatial pattern.

Figure 2A is a schematic block diagram of a DJ 200 with multiple independently controlled LED arrays, wherein the number of LED arrays is preferably between 2 and 8, and is even more preferably between 2 and 4. The signal received from unit 100 via the DJ receiver 220 is passed to a multi-port controller 242 with two ports 294 and 296 connected respectively with two separate arrays 290 and 292 of LEDs 246. These arrays 290 and 292 can be distinguished by spatial placement, color of emitted light, or the temporal pattern of LED illumination. The signal is converted via analog or digital conversion into control signals for the two arrays 290 and 292, which are illuminated in distinct temporal patterns.

It should be noted that the signal received by receiver 220 from the unit 100 can comprise either a signal already in the form required to specify the array and temporal pattern of LED 246 activity, or it can alternatively be converted from a differently formatted signal into temporal pattern signals. For example, the unit 100 can transmit a modulated signal whose amplitude specifies the intensity of LED light amplitude. For multiple LED arrays, signals for the different arrays can be sent together and decoded by the DJ receiver 220, such as through using time multiplexing, or transmission on different frequencies.

Alternatively, the signal could be not directly related to the transduction intensity, such as in the direct transmission of the audio signal being played by the unit 100. In such case, the controller 242 can modify the signal so as to generate appropriate light transduction signals. For example, low frequency bandpass filters could provide the signals for the first array 290, whereas high-frequency bandpass filters could provide the signals for the second array 292. Such filtering could be accomplished by either analog circuitry or digital software within a microprocessor in the controller 242. It is also within the spirit of the present

invention for the different arrays to respond differently to the amplitude of the signal within a frequency band or the total signal.

An alternative control of LED arrays is presented in Figure 2B, a schematic block diagram of a DJ 200 with an LED array with independently controlled LEDs. In this case, the control signal received by the receiver 200 is passed through a single-port, multiple ID controller 243 to a single array of LEDs, each responsive only to signals with a particular characteristic or identifier. One or more of the LEDs 246 can have the same identifier or be responsive to the same characteristic so as to constitute a virtual array of LEDs.

As mentioned above, the transduced light signal can alternatively or additionally comprise multi-element arrays, such as an LED screen. In such case, the signal received by the receiver 220 can be either a specification of image elements to be displayed on the LED screen, or can be as before, a signal unrelated to the light transduction output. For example, many audio players on computers (e.g. Windows Media player) come with pattern generators that are responsive to the frequency and amplitude of the audio signal. Such pattern generators could be incorporated into the controllers 242 or 243.

Alternatively, the light transducer 240 can be a single color illuminated panel, whose temporal pattern of illumination was similar to that of the LEDs of Figures 2A and 2B. In such case, users can partially cover the panel with opaque or translucent patterns, such as a dog or a skull or a representation of a favorite entertainer.

Whereas the receiver 220 and the light controllers 242 or 243 can be hidden from view, either behind the light transducers or separated from the transducers by a wire, for example, the light transducers are meant to be perceptible to other people. For this purpose, the light transducers can be fashioned into fashion accoutrements such as bracelets, brooches, necklaces, pendants, earrings, rings, hair clips (e.g. barrettes), ornamental pins, netting to be worn over clothing, belts, belt buckles, straps, watches, masks, or other objects. Additionally, the light transducers can be fashioned into clothing, such as arrays of lighting elements sewn onto the outside of articles of clothing such as backpacks, wallets, purses, hats, or shoes. For those articles of clothing that are normally washed, however, the lighting transducers and associated electronics will preferably be able to withstand cleaning agents (e.g. water or dry cleaning chemicals), or will be used in clothing such as scarves and hats that do not need to be washable.

It is also convenient for there to be modular lighting arrangements in which the configuration can easily changed by a user. One example of such a modular arrangement is a light pipe made of a flexible plastic cable or rod, at one or both ends of which is positioned

a light source that directs light into the rod. At predetermined locations along the rod, the rod surface can be roughened so as to allow a certain amount of light to escape, on which transparent glass or plastic pieces can be clipped, and that are lighted when the pipe is lighted. Alternatively, the light can be uniformly smooth, and transparent pieces of roughly index of refraction matching material can be clipped onto the rod, allowing some fraction of the light to be diverted from the rod into the pieces. The light sources and associated energy sources used in such an arrangement can be relatively bulky and be carried in a backpack, pouch or other carrying case, and can brightly illuminate a number of separate items.

It should be noted that the transducers require an energy store 270, which is conveniently in the form of a battery. The size of the battery will be highly dependent on the transduction requirements, but can conveniently be a small "watch battery". It is also convenient for the energy store 270 to be rechargeable. Indeed, all of the electric devices of the present invention will need energy stores or generators of some sort, which can comprise non-rechargeable batteries, rechargeable batteries, motion generators that can convert energy from the motion of the user into electrical energy that can be used or stored, fuel cells or other such energy stores or converters as are convenient.

Sound Transducers 260

Sound transducers 260 can supplement or be the primary output of the audio player of the unit 100. For example, the unit 100 can wirelessly transmit the audio signal to DJ 200 comprising a wireless headphone sound transducer. This would allow a user to listen to the audio from the audio player without the need for wires connecting the headphones to the unit 100. Such sound transducers can comprise, for example, electromagnetic or piezoelectric elements.

Alternative to headphone or earphone audio production, external speakers, which can be associated with light transducers 240 or tactile transducers 250, can be used to enhance audio reproduction from external speakers associated with the unit 100. In addition or alternative to simple reproduction of the audio signal output by the audio player 130, the sound transducers 260 can play modified or accompanying signals. For example, frequency filters can be used to select various frequency elements from the music (for low bass), so as to emphasize certain aspects of the music. Alternatively, musical elements not directly output from the audio player 130 can be output to complete all instrumental channels of a piece of music, for example.

Tactile Transducers 250

DJs 200 can be configured with tactile transducers, which can provide vibrational, rubbing, or pressure sensation. As before, signals of a format that control these transducers can be sent directly from the DJ transmitter 120, or can be filtered, modified or generated from signals of an unrelated format that are sent from the transmitter 120. As before, the signal can be the audio signal from the audio player 130, which can, for example, be frequency filtered and possibly frequency converted so that the frequency of tactile stimulation is compatible with the tactile transducer. Alternatively, signals that are of the sort meant for light transduction can be modified so as to be appropriate for tactile transduction. For example, signals for light of a particular color can be used to provide vibrational transduction of a particular frequency, or light amplitudes can be converted into pressure values.

The tactile transducer can comprise a pressure cuff encircling a finger, wrist, ankle, arm, leg, throat, forehead, torso, or other body part. The tactile transducer can alternatively comprise a rubbing device, with an actuator that propels a tactile element tangentially across the skin. The tactile transducer can also alternatively comprise a vibrational device, with an actuator that drives an element normally to the skin. The tactile transducer can further alternatively comprise elements that are held fixed in relation to the skin, and which comprise moving internal elements that cause the skin to vibrate or flex in response to the movement of the internal element.

The tactile transducer can lack any moveable element, and can confer tactile sensation through direct electrical stimulation. Such tactile elements are best used where skin conductivity is high, which can include areas with mucus membranes.

Tactile transduction can take place on any part of the body surface with tactile sensation. In addition, tactile transduction elements can be held against the skin overlying bony structures (skull, backbone, hips, knees, wrists), or swallowed and conveyed through the digestive tract, where they can be perceived by the user.

Input Transducers

It should also be understood that the DJ 200 can comprise input transducers in order to create control signals from information or stimuli in the local environment. Figure 24 is a schematic block diagram of a DJ unit 200 with associated input transducers. The input-enabled DJ 1320 comprises energy storage 270, a controller 1322, output transducers 1324, a

DJ receiver 220 and input transducers 1326. The input transducers 1326 can comprise one or more of a microphone 1328 and an accelerometer 1330.

In operation, the energy storage 270 provides energy for all other functions in the DJ 1320. The controller 1322 provides control signals for the output transducers 1324, which can comprise tactile transducers 250, sound transducers 260, and/or light transducers 240. Input to the controller can be provided via the input transducers 1326, optionally along with input from the DJ receiver 220.

For example, on a dance floor, the microphone 1328 can provide electrical signals corresponding to the ambient music. These signals can be converted into transducer 1324 control signals in a manner similar to that described below for the automatic generation of control signals according to Figures 21 A-C, as will be described below. This allows the use of the DJ functionality in the absence of an accompanying audio unit 100, expanding the applications of the DJ 200. An automatic gain filter can be applied so as to compensate for the average volume level – because the user can be close or far from the sources of ambient music and the music can vary in volume, the strength of the DJ 200 transduction can be normalized. In addition, it can also be preferable for there to be a manual amplitude control 1323, such as a dial or two position rocker switch, by which the average intensity of the DJ 200 control signals can be varied to suit the taste of the user. The amplitude control 1323 can operate through modulating the input transducer 1326 output or as an input to the controller 1322 as it generates the signals for the output transducers 1324.

Alternatively, the accelerometer 1330 can track the movement of the person wearing the DJ 100, such that a signal indicating acceleration in one direction can be converted by the controller 1322 into signals for a channel of output transducers 1324. The accelerometer 1330 can be outfitted with sensors for monitoring only a single axis of motion, or alternatively for up to three independent directions of acceleration. Thus, the controller 1322 can convert sensed acceleration in each direction into a separate channel, horizontal axes of acceleration could be combined into a single channel and the vertical axis into a second channel, or other such linear or non-linear combination of sensed acceleration can be combined in aesthetic fashion.

It is also within the spirit of the present invention that multiple input signals be combined by the controller 1322 to create control signals for aesthetic output from the output transducers 1324. For example, one channel can be reserved for control signals generated from accelerometer signals, another channel for control signals generated from microphone signals, and yet a third channel from control signals generated from DJ receiver 220 input. In

general, the information from the DJ receiver 220 and from the microphone 1328 will be of the same type (i.e. generated from audio signals), so that the most common configurations will be control signals from a combination of the microphone 1328 and accelerometer 1330, and signals from a combination of the DJ receiver 220 and the accelerometer 1330.

The input transducers 1326 can further comprise a light sensor, such that the DJ would mimic light displays in its environment, making it appear that the DJ is part of the activity that surrounds it. In this case, the controller 1322 would preferably generate control signals based on rapid changes in the ambient lighting, since it would be less aesthetic to have the DJ transducers provide constant illumination. Furthermore, slowly changing light (on the order of tens or hundreds of milliseconds) will be created naturally by the movement of the user, whereas changes in the lighting (e.g. strobes, laser lights, disco balls) will be of much faster change (on the order of milliseconds). Thus, to match the ambient dance lighting, it is aesthetic for the DJ 200 to respond most actively to ambient light that is changing in intensity a predetermined percentage in a predetermined time, wherein the predetermined percentage is at least 20% and the predetermined time is 20 milliseconds or less, and even more preferably for the predetermined percentage to be at least 40% and the predetermined time is 5 milliseconds or less.

Unit to Unit Communication

Units 100 transfer audio signals from the audio player in one unit 100 to the audio player 130 of another unit 100. Figures 3A-C are schematic block diagrams of unit 100 elements used in inter-unit communications. Each diagram presents communications between a Unit A and a Unit B, with Unit A transmitting audio signals to Unit B. Dashed connectors and elements indicate elements or transfers that are not being utilized in that unit 100, but are placed to indicate the equivalence of the transmitting and receiving units 100.

In Figure 3A, compressed audio signals (e.g. in MP3 format or MPEG4 format for video transfers, as described below) stored in a compressed audio storage 310 are transferred to a signal decompressor 302, where the compressed audio signal is converted into an uncompressed form suitable for audio output. In Unit A, this decompressed signal is passed both to the local speaker 300, as well as to the inter-unit transmitter/receiver 110. The Unit B inter-unit transmitter-receiver 110 receives the uncompressed audio signal, which is sent to its local speaker for output. Thus, both Unit A and Unit B play the same audio from the Unit A storage, in which uncompressed audio is transferred between the two units 100.

In Figure 3B, compressed audio signals from the Unit A compressed audio storage 310 are sent both to the local signal decompressor 302 and to the inter-unit transmitter/receiver 110. The Unit A decompressor 302 conditions the audio signal so that it is suitable for output through the Unit A speaker 300. The compressed audio signal is sent via Unit A transmitter-receiver 110 to the Unit B transmitter/receiver 110, where it is passed to the Unit B decompressor 302 and thence to the Unit B speaker 300. In this embodiment, because compressed audio signals are transmitted between the units 100 transmitter/receivers 302, lower bandwidth communications means can be used in comparison with the embodiment of Figure 3A.

In Figure 3C, compressed audio signals from the Unit A compressed audio storage 310 are sent to the Unit A signal decompressor 302. These decompressed signals are sent to both the local speaker 300 as well as to a local compressor 330, which recompresses the audio signal to a custom format. In addition to decompressed audio signal input, the compressor also optionally utilizes information from a DJ signal generator 320, which generates signals to control DJ transducers 240, 250 and 260, which can be sent in conjunction with the audio signal. The signal generator 320 can include analog and/or digital filtering or other algorithms that analyze or modify the audio signals, or can alternatively take manually input transducer control signals input as described below. The custom compression can include multiplexing of the audio signals with the transducer control signals.

The custom compressed audio signals, are then passed to the Unit A inter-unit transmitter/receiver 110, which are then transferred to the Unit B inter-unit transmitter/receiver 110, and thence to the Unit B signal decompressor 302 and speaker 300.

Given the time delays in signal transfer between the units 100, custom compression that takes place in the sending unit, and any subsequent decompression that takes place in the receiving unit 100, it can be convenient to place a delay on the local (i.e. Unit A) speaker output of tens of milliseconds, so that both units 100 play the audio through their speakers at roughly the same time. This delay can include limited local digital storage between the local signal decompression and speaker 300 output.

Various hardware communications protocols will be discussed below with respect to unit-to-unit communications, but in general it is required that the distance between the units that must be maintained be preferably at least 40 feet, and more preferably at least 100 feet, and most preferably 500 feet, in order to allow units 100 sharing music to be able to move reasonably with respect to one another (e.g. for a user to go to the bathroom without losing contact), or to find each other in a large venue such as a shopping mall.

Communications Protocols

Communication between the inter-unit transmitter/receivers 110 can involve a variety of protocols within the teachings of the present invention, and can include IP protocol-based transmissions mediated by such physical link layers as 802.11a, b or g, WDCT, HiperLAN, ultra-wideband, 2.5 or 3G wireless telephony communications, custom digital protocols such as Bluetooth or Millennial Net i-Beans. Indeed, it is not even necessary for the transmissions to be based on Internet protocol, and conventional analog radio-frequency or non-IP infrared transmissions are also within the spirit of the present invention. Each unit 100 will generally have both transmission and reception capabilities, though it is possible for a unit to have only reception capabilities. While the bandwidth of the broadcast is dependent on the compression of the audio signal, it is preferable for the transmission bandwidth to be larger than 100 kb/sec, and even more preferable for the transmission bandwidth to be greater than 250 kb/sec.

While the distance of transmission/reception is not bounded within the teachings of the present invention, it will generally be less than a few hundred meters, and often less than 50 meters. The distance of communication is limited in general by the amount of power required to support the transmission, the size of antennae supported by portable devices, and the amount of power allowed by national regulators of broadcast frequencies. Preferably, however, the range of transmission will be at least 10 meters, and even more preferably at least 30 meters, in order to allow people sharing communications to move some distance from one another without communications being lost.

The unit 100 is characterized generally by four sets of roughly independent characteristics: playing audio or not playing audio, transmitting or not transmitting, receiving or not receiving, searching or not searching.

Units 100 will often function in conditions with large numbers of other units 100 within the communications range. For example, in a subway car, a classroom, bicycling, or at a party, a unit 100 can potentially be within range of dozens of other units. A unit 100 that is playing audio from local compressed audio storage 310 can, at the user's prerogative, choose to broadcast this audio to other units 100. A unit 100 that is currently "listening" to a broadcast or is searching for a broadcast to "listen" to will require a specific identifier roughly unique to a broadcaster in order to select that broadcaster signal from among the other possible broadcasters. Some of the communications protocols listed above, such as those based on IP protocols, 2.5G or 3G wireless, or Bluetooth communications, have such

identifiers as part of the protocols. Custom radio frequency based protocols will require protocols to allow signals to be tagged with specific identifiers.

A unit 100 that is transmitting signals can, within the spirit of the present invention, be prevented from simultaneously receiving signals. Preferably, however, units 100 can both transmit and receive simultaneously. One example of the use of simultaneous transmission and reception is for a unit 100 that is receiving a signal to send a signal indicating its reception to the transmitting unit 100. This allows the transmitting unit to determine the number of units 100 that are currently receiving its broadcast. In return, this information could be sent, along with the audio signal, so that all of the users with units 100 receiving the broadcast can know the size of the current reception group. Alternatively, a user with a unit 100 that is currently broadcasting can be searching for other broadcasting units, so that the user can decide whether to continue broadcasting or whether to listen to the broadcast of another unit.

Unit to DJ Communication

Communication between the unit 100 and the DJ 200 can be either through the inter-unit transmitter/receiver 110, or through a separate system. In general, the requirement of the DJ 200 is for reception only, although it is permissible for the DJ 200 to include transmission capabilities (e.g. to indicate to the unit 100 when the DJ 200 energy storage 270 is near depletion).

The signals for which the DJ 200 is receptive is dependent on how the transduction control signals are generated. For example, for a controller 242 that incorporates a filter or modifier that takes the audio signal as its input, the DJ receiver 220 would receive all or a large fraction of the audio signal. In this case, the communication between the unit 100 and the DJ 200 would require a bandwidth comparable to that of inter-unit communication, as described above.

However, if the signals are either generated in the unit 100, or pre-stored along with the stored compressed audio signal, then the communications bandwidth can be quite modest. Consider a DJ 200 with 2 arrays 290 and 292 of LEDs 246, which flash with a frequency of no more than 10 Hertz, and that the LEDs are in either an ON or an OFF state, without intermediate amplitudes. In such case, the maximum bandwidth required would be only 20 bits/second, in addition to the DJ control signals.

The range of unit to DJ communications need not be far. In general, the unit 100 and the DJ 200 will be carried by the same user, so communications ranges of 10 feet can be

adequate for many applications. Some applications (see below) can require, however, somewhat larger ranges. On the other hand, longer communications ranges will tend to confer the possibility of overlap and interference between two different units 100 to their respective DJs 200. In general, for the application of unit to DJ communications, it is preferable for the minimum range of communications to be at least 1 foot, and more preferably for the minimum range of communications to be at least 10 feet, and most preferably for the minimum range of communications to be at least 20 feet. Also, for the application of unit to DJ communications, it is preferable for the maximum range of communications to be no more than 500 feet, and more preferably for the maximum range of communications to be no more than 100 feet, and most preferably for the maximum range of communications to be no more than 40 feet. It should be noted that these communications ranges refer primarily to the transmission distance of the units 100, especially with regard to the maximum transmission distance.

Because there can be multiple unit 100/DJ 200 ensembles within a relatively short distance, communications between a unit 100 and a DJ 200 preferably comprise both a control signal as well as a unit identification signal, so that each DJ 200 receives its control signals from the correct unit 100. Because the unit 100 and the DJ 200 will not, in general, be purchased together, or that a user can buy a new unit 100 to be compatible with already owned DJs 200, it is highly useful to have a means of "entraining" a DJ 200 to a particular unit 100, called its "master unit", and a DJ 200 entrained to a master unit is "bound" to that unit.

Figure 4 is a schematic flow diagram of DJ entraining. To entrain a DJ 200, the DJ is set into entraining mode, preferably by a physical switch on the DJ 200. The master unit 100 to which the DJ 200 is to be entrained is then placed within communications range, and the unit 100 transmits through the DJ transmitter 120 an entraining signal that includes the master unit 100 identifier. Even should there be other units 100 transmitting in the vicinity, it is unlikely that they would be transmitting the entraining signal, so that entraining can often take place in a location with other active units 100. Verification that the entraining took place can involve a characteristic sequence of light output (for light transduction), audio output (for sound transduction) or motion (for tactile transduction). After verification, the DJ 100 is reset to its normal mode of operation, and will respond only to control signals accompanied by the identifier of its master unit 200.

It should be noted that there can be multiple DJ's 200 bound to the same master unit 100. Thus, a single person can have multiple light transducing DJs 200, or DJs 200 of various modes (light, sound, tactile) transduction.

While DJs 200 will generally be bound to a master unit associated with the same person, this is not a requirement of the present invention. Figures 5A-B are schematic block diagrams of DJs 200 associated with multiple people bound to the same master unit. In Figure 5A, DJ A 200and DJ B 200 are both bound to the same DJ transmitter 120, even though DJ A 200 and DJ B 200 are carried by different persons. This is particularly useful if the control signals are choreographed manually or through custom means by one person, so that multiple people can then share the same control signals. Such a means of synchronization is less necessary if the DJ 200 control signals are transmitted between units 100 through the inter-unit transmitter/receiver 110 along with the audio signals. Furthermore, in this case, it is better for the range of unit-to-DJ communication to be in the range of the inter-unit communication described above.

In the case of sound transducers 260, the DJ B 200 can comprise a wireless audio earpiece, allowing users to share music, played on a single unit 100, privately. Consider Figure 5A, configured with sound transducers 260 (see, for example, Figure 1) in DJ A 200 and DJ B 200. Signals from the audio player 130 are transmitted by the DJ transmitter 120, where they are received by DJs 200 – DJ A and DJ B – that are carried by Person A and Person B, respectively. In this case, both persons can listen to the same music.

Figure 5B shows the operation of a wide-area broadcast unit 360, which is used primarily to synchronize control of a large number of DJs 200, such as might happen at a concert, party or rave. In this case, the audio player 130 is used to play audio to a large audience, many of whom are wearing DJs 200. In order to synchronize the DJ output, a relatively high-power broadcast transmitter 125 broadcasts control signals to a number of different DJs 200 carried by Person A, Person B and other undesignated persons. The entraining signal can be automatically sent on a regular basis (e.g. whenever music is not being played, such as between songs, or interspersed within compressed or decompressed songs) so that patrons or partygoers could entrain their DJs 200 to the broadcast unit 360. The broadcast unit 360 can also transmit inter-unit audio signals, or can only play the audio through some public output speaker that both Person A and Person B can enjoy.

Figure 26 is a schematic diagram of people at a concert, in which DJs 200 conveyed by multiple individuals are commonly controlled. At a concert venue 1370, music is produced on a stage 1372, and concert patrons 1376 are located on the floor of the venue.

Many of the patrons have DJs 200 which are receptive to signals generated by a broadcast DJ controller 1374. The broadcast DJ controller creates signals as described below, in which the music is automatically converted into beats, where microphones are used to pick up percussive instruments, and/or where individuals use a hand-pad to tap out control signals. These control signals are either broadcast directly from the area of the broadcast DJ controller 1374, or alternatively are broadcast from a plurality of transmitters 1380 placed around the venue 1370, and which are connected by wires 1378 to the controller 1374 (although the connection can also be wireless within the spirit of the present invention). It should be understood that the protocol for transmitting DJ control signals can be limited either by hardware requirements or by regulatory standards to a certain distance of reception. Thus, to cover a sufficiently large venue, multiple transmitters can be necessary to provide complete coverage over the venue 1370. In general, it is preferable for the maximum transmission distance of transmission from the transmitters to be at least 100 feet, and more preferably at least 200 feet, and most preferably at least 500 feet, so as to be able to cover a reasonable venue 1370 size without needing too many transmitters 1380.

An alternative embodiment of unit 100 to DJ 200 communications is the use of radio frequency transmitters and receivers, such as those used in model airplane control, which comprise multi-channel FM or AM transmitters and receivers. These components can be very small (e.g. the RX72 receivers from Sky Hooks and Riggings, Oakville, Ontario, Canada), and are defined by the crystal oscillators that determine the frequency of RF communications. Each channel can serve for a separate channel of DJ control signals. In such cases, an individual can place a specific crystal in their audio unit 100, and entraining the DJ 200 is then carried out through the use of the same crystal in the DJ 200. Because of the large number of crystals that are available (e.g. comprising approximately 50 channels in the model aircraft FM control band), interference with other audio units 100 can be minimized. Furthermore, control of many DJs 200 within a venue, as described above, can take place by simultaneously transmitting over a large number of frequencies.

As described above, the wide-area broadcast transmitter 125 can transmit entriaing signals to which the DJs 200 can be set to respond. However, there are a number of other preferred means by which DJs 200 can be used to respond to control signals to which they have not been entrained. For example, the DJs 200 can be set to respond to controls signals to which they have not been entrained should there be no entrained control signals present (e.g. the corresponding unit 100 is not turned on).

Figure 35 is a schematic block diagram of DJ 200 switch control for both entraining and wide-area broadcast. The DJ 200 comprises a three-way switch 1920. In a first state 1922, the DJ 200 is entrained to the current control signal as described above. Thereafter, in a second state 1924, the DJ 200 responds to control signals corresponding to the entraining signal encountered in the step 1922. In a third state 1926, the DJ 200 responds to any control signal for which its receiver is receptive, and can therefore respond to a wide-area broadcast, thereby providing the user with manual control over the operational state of the DJ 200. It should be noted that the switch 1920 can be any physical switch with at least three discreet positions, or can alternatively be any manual mechanism by which the user can specify at least three states, including a button presses that have a visible user interface or a voice menu.

Figure 12B is a schematic drawing of modular digital jewelry 201. The modular jewelry 201 is comprised of two components: an electronics module 1934 and a display module 1932. These modules 1934 and 1932 can be electrically joined or separated through an electronics module connector 1936 and a display module connector 1938. The value of the modular arrangement is that the electronics module 1934 comprises, in general, relatively expensive components, whose combined price can be many-fold that of the display module 1932. Thus, if a user wants to change the appearance of the jewelry 201 without having to incur the cost of additional electronics components such as the energy storage 270, receiver 220 or controller 1322, they can simply replace the display module 1932 with its arrangement of output transducers 1324 with an alternative display module 1933 with a different arrangement of output transducers 1325.

The transmitter for DJ 200 control signals has been previously discussed primarily in terms of its incorporation within a unit 100. It should be understood, however, that the transmitter can be used in conjunction with a standard audio player unrelated to unit-to-unit communications. Figure 12C is a schematic block diagram of a modular digital jewelry transmitter 143 that generates and transmits control signals from an audio player 131. The modular transmitter 143 is connected to the audio player 131 via audio output port 136 through the cable 134 to the audio input port 138 of the modular transmitter 143. The modular transmitter 143 comprises the the DJ transmitter 120, which can send unit-to-DJ communications. The output audio port 142 is connected to the earphone 901 via cable 146. The earphone 901 can also be a wireless earphone, perhaps connected via the DJ transmitter 120.

The audio output from the player 131 is split both to the earphone 901 and to the controller 241 (except, perhaps where the DJ transmitter transmits to a wireless earphone).

The controller 241 automatically generates control signals for the DJ 200 in a manner to be described in detail below. These signals are then conveyed to the DJ transmitter 120. It should be understood that this arrangement has the advantage that the digital jewelry functionality can be obtained without the const of the components for the audio player 131, and in addition, that the modular transmitter 143 can then be used in conjunction with multiple audio players 131 (either of different types or as the audio players are lost or broken).

Inter-unit Audio Sharing

Overview

Inter-unit communication involves the interactions of multiple users, who may or may not be acquaintances of each other. That is, the users can be friends who specifically decide to listen to music together, or it can be strangers who share a transient experience on a subway train. The present invention supports both types of social interaction.

An important aspect of the present invention is the means by which groups of individuals join together. Figure 6 is a schematic block diagram of a cluster 700 of units 100, indicating the nomenclature to be used. The cluster 700 is comprised of a single broadcast unit 710, and its associated broadcast DJ 720, as well as one or more receive units 730 and their associated DJs 740. The broadcast unit 710 transmits music, while the receive unit 730 receives the broadcasted music. A search unit 750 and its associated search DJ 760 are not part of the cluster 700, and comprise a unit 100 that is searching for a broadcast unit 710 to listen to or a cluster 700 to become associated with.

It should be noted that many communications systems can be operated alternatively in two modes: one that supports peer-to-peer communications and one that requires a fixed infrastructure such as an access point. Figure 35 is a schematic block diagram of mode switching between peer-to-peer and infrastructure modes. A mode switch 1950 is made by the user, either manually, or automatically – for example, that the user chooses between different functions (listening or broadcasting, file transfers, browsing the Internet) and the system determine the optimal mode to use. A peer-to-peer mode 1952 is well configured for mutual communications between mobile units 100 that are within a predetermined distance, and is well-suited for short-range wireless communications and audio data streaming 1954. Alternatively, the mode switch 1950 enables an infrastructure mode 1956, which is of particular usefulness in gaining access to a wide area network such as the Internet, through

which remote file transfer 1958 (e.g. downloading and uploading) and remote communications such as Internet browsing can be made through access points to the fixed network.local wireless audio streaming.

It should be noted, however, that certain communications systems, such as many modes of telephony, do not distinguish between mobile communications and communications through fixed access points, and that both file transfer 1958 and audio streaming 1954 can be available through the same mode. Even in those cases, however, it can be convenient to have two modes in order to make optimal use of the advantages of the different modes. In such cases, however, the two modes can alternatively be supported by multiple hardware and software systems within the same device – for example, for remote communications to be made through a telephony system (e.g. GSM or CDMA), while the local audio streaming 1954 can be made through a parallel communications system (e.g. Bluetooth or 802.11) – indeed, the two systems can operate simultaneously with one another.

Inter-unit Transmission Segmentation

Preferably, the broadcast unit 710 and the receive units 730 exchange information in addition to the audio signal. For example, each user preferably has indications as to the number of total units (broadcast units 710 and receive units 730) within a cluster, since the knowledge of cluster 700 sizes is an important aspect of the social bond between the users. This also will help search units 750 that are not part of the cluster determine which of the clusters 700 that might be within their range are the most popular.

The additional information shared between members of a cluster 700 would include personal characteristics that a person might allow to be shared (images, names, addresses, other contact information, or nicknames). For example, the broadcast unit 710 will preferably, along with the music, transmit their nickname, so that other users will be able to identify the broadcast unit 710 for subsequent interactions, and a nickname is significantly easier to remember than a numerical identifier (however, such numerical identifier can be stored in the unit 100 for subsequent searching).

Such additional information can be multiplexed along with the audio signal. For example, if the audio signal is transferred as an MP3 file, assuming that there is additional bandwidth beyond that of the MP3 file itself, the file can be broken into pieces, and can be interspersed with other information. Figure 7 is a schematic diagram of a broadcast unit 710 transmission 820. The transmission is comprised of separate blocks of information, each represented in the figure as a separate line. In the first line, a block code 800 is transmitted,

which is a distinctive digital code indicating the beginning of a block, so that a search unit 750 receiving from the broadcast unit 710 for the first time can effectively synchronize itself to the beginning of a digital block. Following the block code 800 is a MP3 block header 802, which indicates that the next signal to be sent will be from a music file (in this case an MP3 file). The MP3 block header 802 includes such information as is needed to interpret the following block of MP3 file block 804, including such information as the length of the MP3 block 804, and characteristics of the music (e.g. compression, song ID, song length, etc.) that are normally located at the beginning of a MP3 file. By interspersing this file header information at regular intervals, a user can properly handle music files that are first received in the middle of the transmission of an MP3 file. Next, the MP3 block 804 containing a segment of a compressed music file is received.

Dependent on the amount of music compression and the bandwidth of the inter-unit communications, other information can be sent, such as user contact information, images (e.g. of the user), and personal information that can be used to determine the "social compatibility" of the user with the broadcast unit 710 and the receive unit 730. This information can be sent between segments of MP3 files or during "idle" time, and is generally preceded by a block code 800, that is used to synchronize transmission and reception. Next, a header file is transmitted, which indicates the type of information to follow, as well as characteristics that will aid in its interpretation. Such characteristics could include the length of information, descriptions of the data, parsing information, etc. In Figure 7, an ID header 806 is followed by an ID block 808, which includes nicknames, contact information, favorite recording artists, etc. Later, an image header 810 can be followed by an image block with an image of the user. The image header 810 includes the number of rows and columns for the image, as well as the form of image compression.

It should be understood that the communications format described in Figure 7 is only illustrative of a single format, and that a large number of different formats are possible within the present invention. Also, the use of MP3 encoding is just an example, and other forms of digital music encoding are within the spirit of the present invention, and can alternatively comprise streaming audio formats such as Real Audio, Windows Media Audio, Shockwave streaming audio, QuickTime audio or even streaming MP3 and others. Furthermore, these streaming audio formats can be modified so as to incorporate means for transmitting DJ 200 control signals and other information.

Transmitting Dynamic Data and Control Information

As described above, there are benefits to two-way communications between the broadcast unit 710 and the receive unit 730. There are many methods of carrying out this communication, even if the inter-unit transmitter/receiver 110 does not permit simultaneous transmission and reception. For example, additional transmission and reception hardware could be included in each unit 100. Alternatively, in the transmission 820 above, specific synchronization signals such as the block code 800 can be followed by specific intervals during which the inter-unit transmitter/receiver 110 that is transmitting switches into receive mode, while the inter-unit transmitter/receiver 110 that was receiving switches to transmit mode. This switch in communications direction can be for a specific interval, or can be mediated through conventional handshake methods of prior art communications protocols.

It should be noted that in addition to transfer of static information (e.g. identifiers, contact information, or images), dynamic information and control information can also be transferred. For example, the user at the receive unit 730 can be presented with a set of positive and negative comments (e.g. "Cool!" "This is awful!") that can be passed back to the broadcast unit 710 with the press of a button. Such information can be presented to the user of the broadcast unit 710 either by visual icon on, for example, an LCD screen, by a text message on this screen, or by artificial voice synthesis generated by the broadcast unit 710 and presented to the user in conjunction with the music.

Alternatively, the user of the receive unit 730 can speak into a microphone that is integrated into the receive unit 730, and the user voice can be sent back to the broadcast unit 710. Indeed, the inter-unit communications can serve as a two-way or multi-way communications method between all units 100 within range of one another. This two-way or multi-way voice communication can be coincident with that of the playing of the audio entertainment, and as such, it is convenient for there to be separate amplitude control over the audio entertainment and the voice communication. This can be implemented either as two separate amplitude controls, or alternatively as an overall amplitude control, with a second control that sets the voice communications amplitude as a ratio to that of the audio entertainment. In this latter mode, the overall level of audio output by the unit is relatively constant, and the user then selects only the ability to hear the voice communication over the audio entertainment.

In order to express their feelings and appreciation about the music they are hearing, users within a cluster 700 can also press buttons on their units 100 that will interrupt or

supplement the control signals being sent to their respective DJs 200, providing light shows that can be made to reflect their feelings. For example, it can be that all lights flashing together (and not in synchrony with the music) can express dislike for music, whereas intricate light displays could indicate pleasure.

It is also possible to send control requests between units 710. For example, a receive unit 730 can make song requests (e.g. "play again", "another by this artist") that can show on the broadcast unit 710 user interface. Alternatively, the user of a receive unit 730 can request that control be switched, so that the receive unit 730 becomes the broadcast unit 710, and the broadcast unit 710 becomes a receive unit 730. Such requests, if accepted by the initial broadcast unit 710 user, will result in the memory storage of the identifier of the broadcast unit 710 being set in all units in the cluster 700 to that of the new broadcast unit 730. Descriptions of the communications resulting in such a transfer of control will be provided below.

Additionally, it is also possible for users of units 100 to privately "chat" with other users while they are concurrently receiving their audio broadcasts. Such chat can be comprised of input methods including keyboard typing, stylus free-form writing/sketching, and quickly selectable icons.

It should be understood that within the spirit of the present invention that the functional configuration can be supported by the extension of certain existing devices. For example, the addition of certain wireless transmitter and receiver, as well as various control and possibly display functionality to a portable audio player would satisfy some embodiments of the present invention. Alternatively, by the addition of music storage and some wireless transmitter and receiver functionality, a mobile telephone would also allow certain embodiments of the present invention. In such case, the normal telephony communications, perhaps supported by expanded 3G telephony capabilities, could serve to replace aspects of the IP communications described elsewhere in this specification.

IP Socket Communication Embodiments

A standard set of protocols for inter-unit communications is provided through IP socket communications, which is widely supported by available wireless communications hardware, including 820.11a, b and g (Wi-Fi). An embodiment of inter-unit communications is provided in Figures 14A-B. Figure 14A is a schematic block diagram of the socket configurations on the broadcast unit 710 and the receive unit 730.

In the discussion below, transfer of the different messages and audio information are provided, generally but not always, through an Internet protocol. At the transport layer of such protocols, there will generally be used either a connectionless protocol or a connection-oriented protocol. Among the most common of these protocols are respectively the User Datagram Protocol (UDP) and the Transmission Control Protocol (TCP), and wherever these protocols are used below, it should be noted that any like protocol (connectionless or connection-oriented), or the entire class of protocol can generally be substituted in the discussion.

The broadcast unit 710, prior to the membership of the receive unit 730, broadcasts the availability of the broadcast on a broadcast 1050, which is generally a TCP socket. The annunciator 1050 broadcasts on a broadcast address with a predetermined IP address and port. The receive unit 730 has a client message handler 1060 that is also a TCP socket that is looking for broadcasts on the predetermined IP address and port. When it receives the broadcast, a handshake creates a private server message handler 1070 on a socket with a new address and port on the broadcast unit 710. The broadcast unit 710 and the receive unit 730 can now exchange a variety of different messages using the TCP protocol between the server message handler 1070 and the client message handler 1060. This information can comprise personal information about the users of the broadcaster unit 710 and the receive unit 730. Alternatively or additionally, the broadcast unit 710 can transfer a section of the audio signal that is currently being played, so that the user of the receive unit 730 can "sample" the music that is being played on the broadcast unit 710. It should be noted that, in general, the broadcast unit 710 continues its broadcast on the broadcast annunciator 1050 for other new members.

Once it is established that the broadcast unit 710 and the receiver unit 730 are mutually desirous of providing and receiving an audio broadcast, respectively, sockets optimized for broadcast audio are created both on the broadcast unit 710 and the receiver unit 730. These sockets will often be UDP sockets – on the broadcast unit 710, a multicast out socket 1080 and on the receiver unit 730, a multicast in socket 1090.

Figure 14B is a schematic block flow diagram of using IP sockets for establishing and maintaining communications between a broadcast unit 710 and the receive unit 730, according to the socket diagram of Figure 14A. In a step 1100, the broadcast annunciator 1050 broadcasts the availability of audio signals. In a step 1102, the receiver unit 730 searches for a broadcast annunciator 1050 on the client message handler 1060 socket. Once a connection is initiated in a step 1104, the broadcast unit 710 creates the message handler

socket 1070 in a step 1106, and the receiver unit 730 retasks the message handler socket 1060 for messaging with the broadcast unit 730. The broadcast annunciator 1050 continues to broadcast availability through the step 1100.

In a step 1110, the broadcast unit 710 and the receiver unit 730 exchange TCP messages in order to establish the mutual interest in audio broadcasting and reception. Should there not be mutual acceptance, then the system returns to the original state in which the broadcast unit 710 is transmitting the broadcast annunciation in the step 1100, and the receive unit 730 searches for broadcasts in the step 1102. Given that the receive unit 730 and the broadcast unit 710 will be within communications distance, and that the broadcast unit 710 is transmitting an annunciation for which the receive unit 730 is receptive, the broadcast unit 710 will be set into a state where it will not establish communications with the receive unit 730 in the step 1106. This can occur either by not creating the message socket in the step 1106 when connection is made with the receiver unit 730, or that the annunciator 1050 remains silent for a predetermined period, perhaps for a period of seconds.

If the broadcast unit 710 and the receiver unit 730 do mutually accept a multicasting relationship, the broadcast unit 710 creates the multicast out UDP socket 1080 in a step 1112 and the receiver unit 730 creates the multicast in UDP socket 1090 in the step 1114, and multicast audio transmission and reception is initiated in a step 1116. It should be noted that should the broadcast unit 710 already be multicasting audio to a receiver unit 730 prior to the step 1112, the multicast out socket 1080 is not created, but that the address of this existing socket 1080 is communicated to the new cluster member.

Given that a cluster can comprise many members, the system of Figures 14A-B must be able to expand to include multiple members. Figure 15 is a schematic block diagram of the IP socket organization used with clusters comprising multiple members. The broadcast unit 710 includes a broadcast annunciator 1050 indicating broadcast availability for new members. For each member in the cluster, the broadcast unit further comprises a message handler 1070 dedicated to the specific member, whose receive unit 730 in turn comprises a message handler 1060, generally in a one-to-one relationship. The broadcast unit comprises N messaging sockets 1070 for the N receive units of the cluster, while each member has only a single socket 1060 connected to the broadcast unit. Thus, when a member wishes to send a message to the other members of the cluster, the message is sent via the receive unit message handler 1060 to the broadcast unit message handler 1070, and which is then multiply sent to the other receive unit message handlers 1060. It is also within the teachings of the present invention for each member of the cluster to have direct messaging capabilities with each

other member, assisted in the creation of the communications by the broadcast unit 710, which can share the socket addresses of each member of the cluster, such that each member can assure that it is making connections with other members of the cluster rather than units of non-members. The broadcast unit 710 also comprises a multicast out socket 1080 which transfers audio to individual receiver sockets 1090 on each of the members of the cluster.

Members of the cluster may come and go, especially since members will frequently move physically outside of the transmission range of the broadcast unit 710. In order for the broadcast unit 710 to determine the current number of members of its cluster, it is within the teachings of the present invention for the broadcast unit 710 to use the messaging sockets 1060 and 1070 to "ping" the receive units 730 from time to time, or otherwise attempt to establish contact with each member of the cluster 700. Such communications attempts will generally be done at a predetermined rate, which will generally be more frequent than once every ten seconds. Information about the number of members of a cluster can be sent by the broadcast unit 710 to the other members of the cluster, so that the users can know how many members there are. Such information is conveniently placed on a display on the unit (see, for example, Figures 18A-B).

Music Synchronization

It will be generally desirable that the synchronicity of the audio playback on the broadcast unit 710 and the receive units 730 be highly synchronized, preferably within 1 second (i.e. this provides a low level functionality of listening to music together), more preferably within 100 milliseconds (i.e. near-simultaneous sharing of music, but an observer would be able to hear – or see through DJ 200 visible cues – the non-synchronicity), and most preferably within 20 milliseconds of one another. In a simple embodiment of the present invention, all members of a cluster 700 must communicate directly with the broadcast unit 710, without any rebroadcast. In such cases, making playback on the two units 710 and 730 as similar as possible will tend to synchronize their audio production.

Figure 8A is a schematic block diagram of audio units 100 with self-broadcast so that audio output is highly synchronized. Two audio units 100 are depicted, including a broadcast unit 710 and a receive unit 730. The organization of audio unit 100 elements is chosen to highlight the self-broadcast architecture. The audio media 1500, which can be compressed audio storage 310, stores the audio signals for broadcast. The output port 1502, which can comprise the inter-unit transmitter/receiver 110, transmits a broadcast audio signal, provided by the audio media 1500. The audio media comprise a variety of different storage protocols

and media, including mp3 files, .wav files, or .au files which are either compressed or uncompressed, monoaural or stereo, 8-bit, 16-bit or 24-bit, and stored on tapes, magnetic disks, or flash media. It should be understood that the spirit of the present invention is applicable to a wide variety of different audio formats, characteristics, and media, of which the ones listed above are given only by way of example. This broadcast audio signal transmitted from the output port 1502 is received at the input port 1504, which can also comprise aspects of the inter-unit transmitter/receiver 110. The signal so received is then played to the associated user via the audio output 1508.

It should be noted that the audio output is normally connected to the audio media 1500 for audio playing when the unit 710 is not broadcasting to a receive unit 730. In such case, there is no need for the audio signals to go to the output port 1502 and thence to the input port 1504. Indeed, even when broadcasting, the audio signal within the broadcast unit 710 can go both directly to the audio output 1508 as well as to be broadcast from the output port 1502.

However, in order to assure the synchronicity of the audio output on the broadcast unit 710 and the receive unit 730, the broadcast unit 710 can present all audio signal from the audio media 1500 for output on the output port 1502. The signal will be received not only on the receiver 730 input port 1504, but also on the input port 1504 of the broadcast unit 710. This can take place either through the physical reception of the broadcast audio signal on a radio frequency receiver, or through local feedback loops within the audio unit 100 (e.g. through employment of IP loopback addresses).

In the receive unit 730, the audio signal received at the input port 1504 goes directly to the audio output 1508, and the other elements of the unit 100 depicted are not active. In the broadcast unit 710, however, if means are used to transfer audio signal between the output port 1502 and the input port 1504 are utilized, and if such transfer means requires less time than that taken for transmitting signal from the output port 1502 of the broadcast unit 710 to the input port 1504 of the receive unit 730, then a delay means 1506 is introduced to provide a constant delay between the input port 1504 and the audio output 1508. This delay can comprise a digital buffer if the signal is digitally encoded, or an analog delay circuit if the signal is analog. Generally, the delay introduced into the audio playback will be a predetermined amount based on the characteristics of the unit hardware and software.

Alternatively, in the case of a digital signal, the delay can be variably set according to the characteristics of the communications system. For example, if there are IP-based communications between the units, the units can "ping" one another in order to establish the

time needed for a "round-trip" communications between the systems. Alternatively, each receive unit 730 of a cluster 700 can transmit to the broadcast unit 710 a known latency of the unit based on its hardware and transmission characteristics. It should be noted that in order to handle different delays between multiple members of a cluster, a delay can be introduced into both the broadcast unit 710 and the receive unit 730, should a new member to the cluster have a very long latency in communications.

Note that the delay 1506 can serve a second purpose, which is to buffer the music should there be natural interruptions in the connections between the members of the cluster 700 (for example, should the receive units 730 move temporarily outside of the range of the broadcaster unit 710). In such case, should enough audio signal be buffered in the delay 1506, there would not be interruption of audio signal in the receive unit 730. Even in such cases, however, in order to accommodate the differences in time to play audio between units and within a unit, the delays in the broadcast unit 710 can be larger than those in the receive unit 730.

If the music compression and the bandwidth of the inter-unit communications are large enough, it can be that the broadcast unit 710 will broadcast less than half of the time. This will generally allow the receive unit 730 to rebroadcast the information from an internal memory store, allowing the effective range of the broadcast signal to potentially double. This can allow, through multiple rebroadcasts, for a very large range even if each individual unit 100 has a small range, and therefore for a potentially large number of users to listen to the same music.

In order to synchronize those that listen to the music through first, second and Nth rebroadcast, a scheme for multi-broadcast synchronization is presented in Figure 8B, a schematic flow diagram for synchronous audio playing with multiple rebroadcast. In such a case, the cluster 700 is considered to be all units 100 that synchronize their music, whether from an original broadcast or through multiple rebroadcasts. In a first step 780, a unit 100 receives a music broadcast along with two additional data. The first data is the current "N", or "hop" of the broadcast it receives, where "N" represents the number of rebroadcasts from the original broadcast unit 710. Thus, a unit 100 receiving music from the original broadcast unit 710 would have an "N" of "1" (i.e. 1 hop), while a unit 100 that received from that receiving unit 100 would have an "N" of "2" (2 hops), and so on. A second piece of information would be the "largest N" that was known to a unit 100. That is, a unit 100 is in contact generally with all units 100 with which it either receives or transmits music, and each send the "largest N" with which it has been in contact.

In a second step 782, the unit 100 determines the duration between signals in the broadcasts it is receiving. Then, two actions are taken. In a step 786, the unit 100 rebroadcasts the music it has received, marking the music with both its "N" and the largest "N" it knows of (either from the unit from which it received its broadcast or from a unit to which it has broadcast).

Also, in a step 784, the music that has been received is played after a time equal to the duration between signals and the "largest N" minus the unit's "N". This will allow for all units 100 to play the music simultaneously. Consider, for example the original broadcast unit 710. It's "N" is "0", and its "largest N" is the maximum number of rebroadcasts in the network. It will store music for a period of "largest N" (equals "largest N" minus "0") times the duration of a rebroadcast cycle, and then play it. For a unit 100 at the furthest rebroadcast, it's "N" and "largest N" will be equal to one another, so that it will store music for no time (i.e. "largest N" minus "N" = 0), but will play it immediately. This will allow all units 100 in the cluster to play music simultaneously. The limitation, however, is that there is memory in each unit 100 to store the music for a sufficient period of time. The units 100 on the system, however, can transfer the amount of storage that is available with the other information, and the number of rebroadcasts can be limited to the amount of memory available within the units 100 that comprise the cluster 700.

As the size of this multi-broadcast cluster 700 changes, the "largest N" can vary, and it will take generally on the order of "largest N" steps for the system to register "largest N". In such cases, there can be temporary gaps in the music on the order of the duration between signals, which will generally be on the order of tens of milliseconds, but which can be longer.

It should be noted that the synchronization of music does not need to accompany the transfer of an actual music signal. Figure 34A is a schematic block flow diagram of the synchronization of music playing from music files present on the units 100. In this embodiment, in a step 1900, the broadcast unit establishes the presence or absence of the music file comprising the music signals to be played on the receive unit. The music file can be referenced either with respect to the name of the file (e.g. "Ooops.mp3"), or a digital identifier that is associated with the music file.

If the music file is not present, then transfer of the music file from the broadcast unit to the receive units can automatically proceed through a file transfer mechanism such as peer-to-peer transfer in a step 1904. If the file was already present, or if the file has been transferred, or alternatively, if the file transfer has begun and enough of the file is present to

allow the simultaneous playing of music between the two units 100, transmission of synchronization signals between the two units 100 can commence in a step 1902.

These synchronization signals can comprise many different forms. For example, the synchronization signal can be the time stamp from the beginning of the music file to the current position of the music file being played on the broadcast unit. Alternatively, the broadcast unit can send the sample number that is currently being played on the broadcast unit 100. In order to allow receiving units to begin synchronous playing in the middle of a transmission from a broadcast unit, the synchronization signals will preferably include information about the song being played, such as the name of the file or the digital identifier associated with the file.

Transmission of this synchronization signal continues until the termination of the song, or until a manual termination (e.g. by actuating a Pause or Stop key) is caused (the frequency of transmission of the synchronization signal will be discussed below). At this point, the broadcast unit can send a termination, pause or other signal in a step 1906. Note that this method of synchronization can operate when the receiving unit establishes connection with the broadcast unit even in the middle of a song.

Figure 34B is a schematic layout of a synchronization signal record 1910 according to Figure 34A. The order and composition of the fields can vary according to the types of music files used, the means of establishing position, the use of digital jewelry, the desire for privacy, and more.

The position field 1912 (SAMPLE#) which contains an indicator of position in a music file – in this case the sample number within the file. The music file identifier field 1914 (SONGID) comprises a textual or numerical identifier of the song being played. The third field is the sample rate field 1916 (SAMPLERATE), and is primarily relevant if the position field 1912 is given in samples, which allows a conversion into time. Given that the same audio entertainment can be recorded or saved at different sample rates, this allows the conversion from a potentially relative position key (samples) to one independent of sample rate (time). The jewelry signal field 1918 (JEWELSIGNAL) is used to encode a digital jewelry 200 control signal for controlling the output of the digital jewelry 200, should the receiver unit be associated with jewelry 200. The order and composition of the fields can vary according to the types of music files used, the means of establishing position, the use of digital jewelry, the desire for privacy, and more.

The frequency with which the record 1910 is broadcast can vary. The time of reception of the record 1910 sets a current time within the song that can adjust the position of

the music playing on the receiver unit. It is possible for the record to be broadcast only once, at the beginning of the song, to establish synchronization. This, however, will not allow others to join in the middle of the music file. Furthermore, if the record 1910 is received or processed at different times for the single record, the music can be poorly synchronized. With multiple synchronization signals, the timing can be adjusted to account for the most advanced reception of the signal – that is, the music playing will be adjusted forward for the most advanced signal, but not be adjusted back for a more laggard signal.

If the record further contains a jewelry signal field 1918, the frequency with which the record 1910 should be sent should be comparable or faster than the rate with which these signals change, and should be preferably at least 6 times a second, and even more preferably at least 12 times a second. If less frequent record 1910 transmission is desired, then multiple jewel signal fields 1918 can be included in a single record 1910.

It should be noted that given units 100 of different design or manufacture, there can be different intrinsic delays between reception of music and/or synchronization signals and the playing of the music. Such delays can result from different speeds of MP3 decompression, different sizes of delay buffers (such as delay 1506), different speeds of handling wireless transmission, differing modes of handling music (e.g. directly from audio media 1500 to audio output 1508 on the broadcast unit, but requiring transmission through an output port 1502 and input port 1504 for the receiver unit), and more. In such cases, it is preferable for receiver units to further comprise a manual delay switch that can adjust the amount of delay on the receiver unit. This switch will generally have two settings: to increase the delay and to decrease the delay, and can conveniently be structured as two independent switches, a rocker switch, a dial switch or equivalent. It is useful for the increments of delay determined by the switch be adjustable so as to allow users to sense the music from the broadcast unit and the receiver unit as being synchronous, and it is preferable for the increments of delay to be less than 50 milliseconds, and even more preferable for the increments of delay to be less than 20 milliseconds, and most preferable for the units of delay to be less than 5 milliseconds.

Creation and Maintenance of Clusters

Search units 750 can be playing music themselves, or can be scanning for broadcast units 710. Indeed, search units 750 can be members of another cluster 700, either as broadcast unit 710 or receive unit 730. To detect a different cluster 700 in which it might desire membership, the search unit 750 can either play the music of the broadcast unit 710 to

the search unit 750 user, or it can scan for personal characteristics of the broadcast unit 710 user that are transmitted in the ID block 808. For example, a user can establish personal characteristic search criteria, comprising such criteria as age, favorite recording artists, and interest in skateboarding, and respond when someone who satisfies these criteria approaches.

Alternatively, the search unit 750 user can also identify a person whose cluster he wishes to join through visual contact (e.g. through perceiving the output of the person's light transducer 240).

Before a search unit 750 user can establish contact, it is preferable for a broadcast unit 710 user, or a receive unit 730 user, to provide permissions for others to join the cluster. For example, each unit 100 will generally be able to changeably set whether no one can join with their unit 100, whether anyone can join with their unit 100, or whether permission is manually granted for each user who wishes to join with their unit into a cluster. For a cluster 700, membership in the cluster can be provided either if any one member of the cluster 700 permits a search unit 750 user to join, or it can be set that all members of a cluster 700 need to permit other users to join, or through a variety of voting schemes. The permissions desired by each member will generally be sent between units 100 in a cluster as part of the ID block 808 or other inter-unit communications. Furthermore, these permissions can be used to establish the degree to which others can eavesdrop on a unit 100 transmission. This can be enforced either through the use of cryptography, which can only provide decryption keys as part of becoming a cluster 700 member, through provision of a private IP socket address or password, through standards agreed by manufacturers of unit 100 hardware and software, or by unit 100 users limiting the information that is sent through the ID block 808 through software control.

The search unit 750 user can then establish membership in the group in a variety of ways. For example, if the search unit 750 is scanning music or personal characteristics of the unit 100 user, it can alert the search unit 750 user about the presence of the unit 100. The search unit 750 user can then interact with the search unit 750 interface to send the unit 100 user a message requesting membership in the cluster 700, which can be granted or not. This type of request to join a cluster 700 does not require visual contact, and can be done even if the search unit 750 and cluster are separated by walls, floors, or ceilings.

Another method of establishing contact between a search unit 750 user and a cluster 700 member is for the search unit 750 user to make visual contact with the cluster 700 member. In such case that physical contact or physical proximity is easily made between the unit 100 of the cluster member and the search unit 750, digital exchange can be easily made

either through direct unit 100 contact through electrical conductors, or through directional signals through infra-red LEDs, for example. For example, the search unit 750 user can point his unit 100 at the cluster 700 member unit, and then if the cluster member wishes the search unit 750 user to join the cluster, could point his unit 100 at the search unit 100, and with both pressing buttons, effect the transfer of IDs, cryptography keys, IP socket addresses or other information that allows the search unit 750 user to join the cluster 700.

Alternatively, the broadcast DJ 720 (or the receive DJ 740) can present digital signals through the light transducer. For example, most DJ 720 light transduction will be modulated at frequencies of 1-10 Hz, with human vision not being able to distinguish modulation at 50 Hz or faster. This means that digital signals can be displayed through the light transducer 240 at much higher frequencies (kHz) that will not perceived by the human eye, even while lower frequency signals are being displayed for human appreciation. Thus, the broadcast DJ 720 can receive a signal from the broadcast unit 710 DJ transmitter 120 containing information needed for a search unit 750 to connect to the broadcast unit's cluster 700. This information will be expressed by the light transducer 240 of the broadcast DJ 720 in digital format. The search unit 750 can have an optical sensor, preferably with significant directionality, that will detect the signal from the light transducer 240, so that the search unit 750 is pointed in the direction of the broadcast DJ 720, and the identifier information required for search unit 750 to become a member of cluster 700. This optical sensor serves as the DJ directional identifier 122 of Figure 1. At this point, if desired, the broadcast unit 710 user can determine if they want the search unit 750 user to join the cluster 700.

A summary of means to effect joining of a cluster is provided in Figures 13A through E, which display means for a search unit 750 to exchange information prior to joining a cluster 700 via a broadcast unit 710. It is also within the teachings of the present invention for the search unit 750 to institute communications with a receive unit 730 for the purposes of joining a cluster in a similar fashion, particularly since it may be difficult for a person outside of the cluster 700 to determine which of the cluster 700 members is the broadcast unit 710, and which is a receiver unit 730.

If should be noted in the Figures 13A-G that limited range and directionality are preferred. That is, there can be a number of broadcast units 710 within an area, and being able to select that one broadcast unit 710 whose cluster one wishes to join requires some means to allow the search unit 750 user to select a single broadcast unit 710 among many. This functionality is generally provided either by making a very directional communication between the two devices, or by depending on the physical proximity of the search unit 750

and the desired broadcast unit 710 (i.e. in a greatly restricted range, there will be fewer competing broadcast units 710). In the following description, the "broadcaster" denotes the user using the broadcast unit 710, and the "searcher" denotes the user using the search unit 750.

In the Figures 13A-G, the selection of the cluster by the searcher occurs in three ways, that will referred to as "search transmission mode", "broadcast transmission mode", and "mutual transmission mode", according to the entity that is conveying information. In search transmission mode, the searcher sends an ID via the search unit 750 to the broadcast unit 710. This ID can comprise a unique identifier, or specific means of communication (e.g. an IP address and port for IP-based communication). With this ID, the broadcast unit can either request the searcher to join, or can be receptive to the searcher when the searcher makes an undifferentiated request to join local units within its wireless range. In broadcast transmission mode, the broadcaster sends an ID via the broadcast unit 710 to the search unit 750. With this ID, the searcher unit can then make an attempt to connect with the broadcast unit 710 (e.g. if the ID is an IP address and port), or the search unit can respond positively to a broadcast from the broadcast unit 710 (e.g. from a broadcast annunciator 1050), wherein the ID is passed and checked between the units early in the communications process. Mutual transmission mode comprises a combination of broadcast transmission mode and search transmission mode, in that information and communication is two way between the broadcaster and the searcher.

Figure 13A is a schematic cross-section through a search unit 750 and a broadcast unit 710 in which communications are provided via visible or infrared LED emission in search transmission mode. On the right of the figure, a LED 1044 with an associated lens 1046 (the two of which can be integrated) transmits a directional signal from the unit case 1000. This light can optionally pass thorough a window 1048 that is transparent to the light. On the left of the figure, a lens element 1040 collects light through a broad solid angle and directs it onto a light sensing element 1042, which is conveniently a light-sensing diode or resistor. The directionality of the communication is conferred by the transmitting lens 1046 and the collecting lens 1040.

Alternatively, the LED 1044 can be replaced by a visible laser. Figure 13B is a schematic cross-section through a search unit 750 and a broadcast unit 710 in which communications are provided via a visible or infrared laser in search transmission mode. The search unit 750 comprises a diode laser 1041 that is conditioned by a lens 1043 to form a beam that is sensed by the light sensing element 1042 on the broadcast unit 710. Because a

collimated laser beam can be difficult to aim with precision at a photosensing element carried by a person, the optics can comprise a two focus lens 1043 that has a portion that produces a collimated beam 1045, and a second portion that produces a diverging beam 1047. The collimated beam is used by the user of the search unit 750 as a guide beam to direct the pointing of the unit 750, while the divergent beam provides a spread of beam so that the human pointing accuracy can be relatively low. The means for creating the two focus lens 1043 can include the use of a lens with two different patterns of curvature across its surface, or the use of an initial diverging lens whose output intersects a converging lens across only a part of its diameter, where the light that encounters the second lens is collimated, and the light that does not encounter the second lens remains diverging. It is also within the teachings of the present invention for the lens to be slowly diverging without a collimating portion, such that the user does not get visible feedback of their pointing accuracy. In such case, the laser can emit infrared rather than visible wavelengths.

Figure 13C is a schematic cross-section through a search unit 750 and a broadcast unit 710 in which communications are provided via visible or infrared emission from a digital jewelry element 200 in broadcast transmission mode. The digital jewelry 200 is carried by the broadcaster on a chain 1033, with the digital jewelry 200 visible. The digital jewelry is emitting through a light transducer 1031 a high frequency signal multiplexed within the visible low frequency signal. The search unit 750 is pointed in the direction of the digital jewelry 200, and receives a signal through the light-sensing element 1042. This manner of communication is convenient because the searcher knows, via the presence of the visible signal on the digital jewelry 200, that the broadcaster is receptive to cluster formation.

Figure 13D is a schematic cross-section through a search unit 750 and a broadcast unit 710 in which communications are provided via contact in mutual transmission mode. In this case, the broadcast unit 710 and the search unit 750 both comprise a contact transmission terminus 1030, and electronic means by which contact transmission is performed. This means can operate either inductively (via an alternative current circuit), through direct electrical contact with alternating or direct current means, or other such means that involves a direct physical contact (indicated by the movement of the search unit 750 to the position of the unit depicted in dotted lines). The search unit 750 or the broadcast unit 710 can, via automatic sensing of the contact or manual control, initiate communications transfer. Given the mutuality of contact as well as the physical equivalence of the two units 710 and 750, information transfer is possible in both directions. It should be noted that in the case of direct

current connection, the termini 1030 will comprise two contact points, both of which must make electrical contact in order for communications to occur.

Figure 13E is a schematic cross-section through a search unit 750 and a broadcast unit 710 in which communications are provided via sonic transmissions in broadcast transmission mode. The broadcaster (or receivers) will be listening to the audio information generally through headphones 1020 or earphones, all of which comprise speakers 1022 that, to one extent or another, leak sonic energy. The use of audio output devices as depicted in Figure 10 and Figures 11A and 11B that admit external sound, will also increase the amount of sound energy lost. This sound energy can be detected by the searcher via a directional speaker comprising a sound collector 1024 and a microphone 1026. This system requires that the sound output of the broadcast unit 710 and the receiver unit 750 also output an ID encoded in the sound. Such sound can be conveniently output at inaudible frequencies, such as 3000-5000 Hz, which carry sufficient bandwidth to encode short messages or identifiers (e.g. an IP address and port number can be carried in 5 bytes). Sound energy, especially at higher frequencies, can be quite directional, depending on the shape of the collector 1024 and the structure of the microphone 1024, allowing good directional selection by the searcher.

Figure 13F is a schematic cross-section through a search unit 750 and a broadcast unit 710 in which communications are provided via radio frequency transmissions in broadcast transmission mode. The radio frequency transmissions are not strongly directional (and for the purposes of the broadcast of audio information, are designed to be as directionless as possible). In order to distinguish a desired cluster 700 to join and an undesired cluster 700, a number of strategies can be employed. For example, the strengths of the various signals can be measured and the strongest chosen for connection. Alternatively, if there are multiple broadcast connections available, the search unit 750 can sequentially attempt a connection with each broadcast unit 710. When the attempt is made, the broadcast unit 710 can, prior to alerting the broadcaster of the attempted joining by a new member, cause the digital jewelry 200 associated with the broadcast unit 710 to visibly flash a characteristic signal. The searcher can then verify by pressing the appropriate button on the search unit 750 his desire to join the cluster 700 of the broadcast digital jewelry 200 that had just flashed. If the searcher decided not to join that cluster 700, the search unit 750 could search for yet another unit broadcast unit 750 within range, and attempt to join.

At any time, the members of a cluster 700 can share personal characteristics (nickname, real name, address, contact information, face or tattoo images, favorite recording artists, etc.) through selection of choices of the unit 100 interface, with all such characteristics

or a subset thereof to be stored on the units 100. In order to assist cluster 700 members in determining whether or not to accept a person into their cluster 700, a search unit 750 member can display either the total number of people with whom he has shared personal characteristics, or he can alternatively allow the cluster members to probe his store of persons with whom personal characteristics have been stored to see whether a particular trusted person or group of common acquaintances are present therein. It is also within the spirit of the present invention for individuals to rate other individual members of their cluster, and such ratings can be collated and passed from person to person or cluster to cluster, and can be used for a cluster 700 to determine whether a search unit 750 person should be added to the cluster 700.

Figure 17 is a matrix of broadcaster and searcher preferences and characteristics, illustrating the matching of broadcaster and searcher in admitting a searcher to a cluster. A broadcaster preference table 1160 includes those characteristics that the broadcaster wishes to see in a new member of a cluster. These characteristics can include gender, age, musical "likes" and "dislikes", the school attended, and more. The searcher similarly has a preference table 1166. The searcher preference table 1166 and broadcaster preference table 1160 are not different in form, as the searcher will at another time function as the broadcaster for another group, and his preference table 1166 will then serve as the broadcaster preference table.

The broadcaster preference table 1160 can be automatically matched with a searcher characteristics table 1162. This table 1162 comprises characteristics of the searcher, wherein there will be characteristics that overlap in type (e.g. age, gender, etc.) which can then be compared with the parameters in the broadcaster preference table. This matching occurs during the period when the searcher is interrogating the cluster with interest in joining. Similarly, there is a broadcaster characteristics table 1164 indicating the characteristics of the broadcaster, which can be matched against the searcher preferences tale 1166.

The algorithm used in approving or disapproving of an accord between a preference table and a characteristics table can be varied and set by the user – whether by the broadcaster to accept new members into a cluster, or by a searcher to join a new cluster. For example, the user could require that the gender be an exact match, the age within a year, and the musical preferences might not matter. The user can additionally specify that an accord is acceptable if any one parameter matches, specify that an accord be unacceptable if any one parameter does not match, specify an accord be acceptable based on the overlap of a majority of the individual matches, or other such specification.

It should be noted that the broadcaster preferences table 1160 and the broadcaster characteristics table 1164 (and likewise with the searcher tables 1162 and 1166) can be a single table, according to the notion that a person will prefer people who are like themselves. Each user could then express the acceptable range of characteristics of people with which to join as a difference from their own values. For example, the parameter "same" could mean that the person needs to match closely, whereas "similar" could indicate a range (e.g. within a year) and "different" could mean anyone. In this way, there would not be the burden on the user to define the preference table 1160 or 1166 in a very detailed manner.

In the case of a cluster, the transfer of information between the searcher and the cluster can, as mentioned above, involve not only the broadcaster, but also other members of the cluster (especially since the searcher may not know the identity of a cluster's broadcaster from external observation). The cluster can also make communal decisions about accepting a new member. That is, if there are 4 members of a cluster, and a searcher indicates an interest in joining the cluster, there can be voting among the members of a cluster regarding the acceptance of the new member. The procedure of voting will normally be done by messaging among the members, which can be assisted by structured information transfer as will be described below.

A number of such voting schemes are described in Figure 19, a table of voting schemes for the acceptance of new members into a cluster. The first column is the name of the rule, and the second column describes the algorithm for evaluation according to the rule. In the "BROADCASTER" rule, the broadcaster decides whether or not the new member will be accepted. The new member is accepted when the broadcaster indicates "yes" and is otherwise rejected.

In the "Majority" rule, the members are polled, and whenever a majority of the members vote either acceptance or rejection, the new member is accordingly accepted or rejected. It should be noted that this rule (as well as the rules to follow) depends on the broadcaster or other member of the cluster having knowledge of the number of members in the cluster, which will generally be the case (e.g. in an IP socket based system, the broadcaster can simply count the number of socket connections). Thus, if the number of members in a cluster is given as N_{mem}, as soon as $(N_{mem}/2) + 1$ members have indicated the same result, that result is then communicated to the broadcaster, the members and the prospective new member. If the number of members is even, and there is a split vote, the result goes according to the broadcaster's vote.

According to the "Unanimous" rule, a new member is accepted only on unanimous decision of the members. Thus, the prospective new member is rejected as soon as the first "no" vote is received, and is accepted only when the votes of all members of the cluster are received, and all of the votes are positive.

The "Timed Majority" rule is similar to that of the "Majority" rule, except that a timer is started when the vote is announced, the timer being of a predetermined duration, and in a preferred embodiment, is indicated as a count down timer on the unit 100 of each member of the cluster 700. The vote is completed when $(N_{mem}/2)+1$ members vote with the same indication ("yes" or "no") if the timer has not completed its predetermined duration. If all of the members have voted, and the vote is a tie, the result goes in accordance with that of the broadcaster. If the timer has expired, and the vote has not been decided, the number of members that have voted is considered a quorum of number Q. If (Q/2)+1 members have voted in some fashion, that is the result of the vote. Otherwise, in the case of a tie, the result goes according to the vote of the broadcaster. If the broadcaster has not voted, the vote goes according to the first vote received.

The "Synchronized Majority" rule is similar to the Timed Majority rule, but instead of initiating the vote, and then waiting a predetermined period for members to vote, the vote is announced, and then there is a predetermined countdown period to the beginning of voting. The voting itself is very limited in time, generally for less than 10 seconds, and preferably for less than 3 seconds. Counting votes is performed only for the quorum of members that vote, and is performed according to the rules for the Timed Majority.

There are many different voting schemes consistent with creating, growing and maintaining clusters within the spirit of the present invention. For instance, in cases where there are close votes, the voting can be reopened for individuals to change their vote. In cases, members can request a new round of voting. Furthermore, the voting can be closed ballot, in which the votes of individuals are not known to the other members, or open voting, in which the identity of each member's vote is publicly displayed on each unit 100.

In addition, the voting can be supported and enhanced by information made available to each member through displays on the units 100. Figure 18A is a screenshot of an LCD display 1170 of a unit 100, taken during normal operation. The display 1170 is comprised of two different areas, an audio area 1172 and a broadcaster area 1174. The audio area 1172 includes information about the status of the audio output and the unit 100 operation, which can include battery status, the name of the performer, the title of the piece of music, the time the audio has been playing, the track number and more. The broadcaster area 1174 comprises

information about the status of the cluster 700. In the example given, the broadcaster area includes the number "5", which represents the number of people current in the cluster, the text "DJ", which indicates that the unit 100 on which the display 1170 is shown is currently the broadcaster of the cluster 700, and the text "OPEN", which indicates that the cluster is open for new members to join (the text "CLOSED" would indicate that no new members are being solicited or allowed).

Figure 18B is a screenshot of an LCD display 1170 of a unit 100, taken during voting for a new member. The audio area 1172 is replaced by a new member characteristics area 1176, in which characteristics of the prospective new member are displayed. Such characteristics can include the name (or nickname) of the prospective new member, their age, and their likes (hearts) and dislikes (bolts). In the broadcaster area 1174, the digit "3" indicates that there are three current members of the cluster 700, and an ear icon indicates that the current unit 100 is being used to receive from the broadcaster rather than being a broadcaster, and the name [ALI] indicates the name of the current broadcaster. The text "VOTE-MAJ" indicates that the current vote is being done according to the Majority rule. The broadcaster area 1174 and the new member characteristics areas 1176 provide the information needed by the existing member to make a decision about whether to allow the prospective new member to join.

The displays 1170 of Figures 18A-B are indicative only of the types of information that can be placed on a display 1170, but it should be appreciated that there are many pieces of information that can be placed onto the displays 1170 and that the format of the display can be very widely varied. Furthermore, there need not be distinct audio areas 1172 and broadcaster areas 1174, but the information can be mixed together. Alternatively, especially with very small displays 1170, the display 1170 can be made to cycle between different types of information.

It is also within the spirit of the present invention for individuals to rate other individual members of their cluster, and such ratings can be collated and passed from person to person or cluster to cluster, and can be used for a cluster 700 to determine whether a search unit 750 person should be added to the cluster 700. Figure 27 is a schematic block flow diagram of using a prospective new member's previous associations to determine whether the person should be added to an existing cluster.

In a step 1400, from a search unit 750, the prospective new member places an external communication request with an operational broadcast annunciator 1050 by a broadcast unit 710. In a step 1402, a temporary message connection is established through which

information can be passed mutually between the search unit 750 and the broadcast unit 710. The broadcast unit 710 requests personal and cluster ID's from the search unit 750. The personal ID is a unique identifier that can be optionally provided to every audio unit 100, and which can further be optionally hard-encoded into the hardware of the unit 100. The cluster IDs represent the personal ID's of other units 100 with which the search unit 750 has been previously associated in a cluster. In a step 1406, the broadcast unit 710 matches the incoming personal IDs and cluster IDs with personal ID's and cluster IDs that are stored in the memory of the broadcast unit 710. If there exist a sufficient number of matches, which can be computed as a minimum number or as a minimum fraction of the IDs stored in the broadcast unit 710, the new member of the search unit 750 can be accepted into the cluster. In a step 1412, the search unit 750 can then store the ID of the broadcast unit 710 and the other members of the existing cluster 700 into his cluster IDs, and the broadcast unit 750 and the other receive units 730 of the cluster can then store the personal ID of the search unit 750 into their cluster IDs. If there does not exist a sufficient number or quality of matches, the broadcast unit 710 will reject the prospective new member, optionally send a message of rejection, and then close the socket connection (or other connection that had been created) between the broadcast unit 710 and the search unit 750. No new IDs are stored on either unit 710 or 750.

It is also within the spirit of the present invention for other information associated with the personal and cluster IDs to be shared and used in the algorithm for determining whether to accept or reject a prospective new member into a cluster 700. This information can include rating information, the duration of association with another cluster 700 (i.e. the longer the association, the more suitable the social connection of that person with the cluster 700 would have been), the size of the cluster 700 when the searcher was a member of a particular cluster 700, the popularity of a cluster 700 (measured by the number of cluster IDs carried by the broadcast unit 710), and more. The matching program, likewise, would weight the existence of a match by some of these quality factors in order to determine the suitability of the searcher to join the cluster.

While the comparisons can be made between a search unit 750 personal and cluster IDs and those from the broadcast unit 710, representing the personal experience of the owners of the respective units, it is also possible that the reputation or desirability of individuals with a given personal ID can be posted to or retrieved from trusted people. For example, two friends can swap the information of which IDs are to be trusted or not between two units 100, or alternatively, can be posted onto or retrieved from the Internet. For

example, after a bad personal experience with a unit 100 with a personal ID of 524329102, a person could post that ID on the Internet to share with friends, so that the friends could avoid allowing that person to join, or avoid joining a cluster with that person.

It should be noted that publishing a list of personal IDs allows people to establish the breadth of their contacts. By posting their contacts on web sites, people can demonstrate their activity and popularity. This also encourages people to join clusters, in order to expand the number of people with whom they have been associated. Furthermore, the personal ID serves as a "handle" by which people can further communicate with one another. For example, on the Internet, a person can divulge a limited amount of information (e.g. an email address) that would allow other people with whom they have been in a cluster together to contact them.

It should be noted that the formation and maintenance of a cluster 700 requires the initial and continued physical proximity of the broadcast unit 710 and the receive unit 730. In order to help maintain such physical proximity conducive to cluster maintenance, feedback mechanisms can be used to alert the users to help them maintain the required physical proximity, as will be discussed below.

Figure 28 is a block flow diagram indicating the steps used to maintain physical proximity between the broadcast unit 710 and the receive unit 730 via feedback to the receive unit user. In a step 1530, the wireless connection between the broadcast unit 710 and the receive unit 730 is established. In a step 1532, the connection between the two units 710 and 730 is tested. There are a number of different means by which this testing can take place. For example, in IP-based communications, the receive unit 730 can from time to time – though generally less than every 10 seconds, and even more preferably less than every 1 second – use the "ping" function to test the presence and speed of connection with the broadcast unit 710. Alternatively, the receive unit 730 will be receiving audio signals wirelessly almost continuously from the broadcast unit 710, and a callback alert function can be instituted such that loss of this signal determined at a predetermined repeat time – which is conveniently less than 5 seconds, and even more preferably less than every 1 second – and which is then reported to the system.

While the methods above determine the absolute loss of a signal, they do not anticipate loss of signal. A method that does anticipate signal issues prior to loss is the measurement of signal strength. This can be done directly in the signal reception hardware by measuring the wireless signal induced current or voltage.

In a step 1534, the results of the connection testing performed in the step 1532 is analyzed in order to determine whether the signal is adequate. It should be noted that a temporary loss of signal, lasting even seconds, may or may not be of importance. For example, the broadcast unit 710 user and receive unit 730 users could walk on opposite sides of a metallic structure, enter a building at different times, change their body posture such that the antennae are not optimally situated with respect to one another, etc. Thus, an algorithm is generally used to time average the results of the step 1532, with the results conveniently time averaged over a matter of seconds.

Whatever the results of the signal test of the step 1534, the step 1532 is continuously repeated as long as the connection between the broadcast unit 710 and the receive unit 730 is present. If the signal is deemed inadequate, however, feedback to that effect is provided to the receive unit 730 user in a step 1536. The user feedback can occur through a variety of mechanisms, including visual (flashing lights) and tactile (vibration) transducers, emanating either from the audio unit 100 or the digital jewelry 200. For example, the receiver unit 730 can send a signal to the associated digital jewelry 200 to effect a special sequence of light transducer output.

It is most convenient, however, for the audio output of the receiver unit 730 as heard by the user to be interrupted or overlain with an audio signal to alert the user to the imminent or possible loss of audio signal. This audio signal can include clicks, beeps, animal sounds, closed doors, or other predetermined or user selected signals heard over silence or the pre-existing signal, with the signal possibly being somewhat reduced in volume such that the combination of the pre-existing signal and the feedback signal is not unpleasantly loud.

It should be noted that the flow diagram of Figure 28 refers specifically to alerting the receive unit 730 user of potential communications issues. Such alerting can also be usefully transferred to or used by the broadcast unit 710. For example, with knowledge of the communications issues, the broadcast unit 710 user can move more slowly, make sure that the unit is not heavily shielding, that any changes in posture that could relate to the problems are reversed, etc. The broadcast unit 710 can perform communications tests (as in the step 1532) or analyze the tests to determine if the communications are adequate (as in the step 1534) – particularly through use of the messaging TCP channels. Given that there can be multiple receive units 730 connected to s single broadcast unit 710, it is generally preferable for the tests to be performed on the receive units 730, and problems to be communicated to the broadcast unit 710 – provided, however, that communications still exist for such communication.

In order to overcome this deficiency, it is possible for the receiver unit 730 to communicate potential problems in communications to the broadcast unit 710 at an early indication. The broadcast unit 710 then starts a timer of predetermined length. If the broadcast unit 710 does not receive a "release" from the receive unit 730 before the timer has completed its countdown, it can then assume that communications with the receive unit 730 have been terminated, and it can then send feedback to the broadcast unit 710 user.

It is also within the teachings of the present invention for both the broadcast unit 710 and the receive unit 730 to independently monitor the connections with each other, and alert their respective users of communications problems.

It should be noted that the use of audio alerts can be used more generally within the user interface of the audio units 100. Thus, audio alerts can be conveniently used to inform the user of the joining of new members to the cluster 700, the initiation of communications with search units 750 outside of the group, the leaving of the group by existing cluster 700 members, the request by a receive unit 730 to become the broadcast unit 710, the transfer of cluster control from a broadcast unit 710 to a receive unit 730, and more. These alerts can be either predetermined by the hardware (e.g. stored on ROM), or can be specified by the user. Furthermore, it can be convenient for the broadcast unit 710 to temporarily transfer to new members of the cluster custom alerts, so that the alerts are part of the experience that the broadcast unit 710 user shares with the other members of the cluster. Such alerts would be active only as long as the receive units were members of the cluster 700, and then would revert back to the alerts present before becoming cluster members.

Cluster Hierarchy

A receive unit 730 can also be the broadcast unit 710 of a separate cluster 700 from the cluster 700 of which it is a member. This receive unit is called a broadcasting receiver 770. In such case, it is convenient for the receive units 730 that are associated with the broadcasting receiver 770 to become associated with the cluster 700 of which the broadcasting receiver 770 is a member. This can conveniently be accomplished in two different ways. In a first manner, the receive units 730 that are associated with the broadcasting receiver 770 can become directly associated with the broadcast unit 710, so that they are members only of the cluster 700, and are no longer associated with the broadcasting receiver 770. In a second manner, the receive units 730 associated with the broadcasting receiver 770 can remain primarily associated with the broadcasting receiver 770, as shown in Figures 9A and 9B, which are schematic block diagrams of hierarchically-related clusters. In

Figure 9A, the receive units 730 that are members of a sub-cluster 701 of which the broadcast unit is a broadcast receive unit 770, can receive music directly from the broadcast receive unit 710, while retaining their identification with the broadcasting receiver 770, such that if the broadcasting receiver 770 removes itself or is removed from the cluster 700, these receive units 730 similarly are removed from the cluster 700. In order to provide this form of hierarchical control, the sub-cluster 701 receive units 730 can obtain an identifier, which can be an IP socket address, from the broadcast receive unit 770, indicating the desired link to the broadcast unit 710. The sub-cluster receive units 730, however, maintain direct communications with the broadcast receive unit 770, such that on directive from the unit 770, they break their communications with the unit 710, and reestablish normal inter-unit audio signal communications with the broadcast receive unit 770. In an embodiment using IP addressing and communications, this can involve the maintenance of TCP messaging communications between the sub-cluster 701 receive units 730 with the broadcast receive unit 770, during the time that the sub-cluster 701 is associated with the cluster 700.

In Figure 9B, the receive units 730 of the sub-cluster 701 receive music directly from the broadcasting receiver 770, which itself receives the music from the broadcast unit 710. In such case, as the broadcasting receiver 770 is removed from the cluster 700, the receive units 730 of the sub-cluster 701 would also not be able to hear music from the cluster 700.

It would be apparent that such an arrangement can be hierarchically arranged, such that the receive unit 730 of the sub-cluster 701 can itself be the broadcast receiver 770 of another sub-cluster 701, and so forth. The advantage of this arrangement is that people that are associated with one another, forming a cluster 700, can move as a group from cluster to cluster, maintaining a separate identity.

It should be also noted that the configuration of communications between members of a hierarchical cluster can be variously arranged, not only as shown in Figures 9A and 9B. For example, every member of the cluster 700 can have a direct link between every other member of the cluster 700, such that no re-broadcast of messages needs to take place. Furthermore, given that there are different inter-unit communications (for example, messaging versus audio broadcast), it is within the teachings of the present invention that the configuration for the different modes of communication can be different – for example direct communications between the broadcast unit 710 for audio broadcast, but peer-to-peer communications between individual units for messaging.

Maintaining Private Communications

In order to restrict membership in a cluster 700, either the information transfer must be restricted, such as by keeping private the socket IP addresses or passwords or other information that is required for a member to receive the signal, or the signal can be transmitted openly in encrypted form, such that only those members having been provided with the encryption key can properly decode the signal so sent. Both of these mechanisms are taught within the present invention, and are described at various points within this specification.

Figure 32A is a schematic block diagram of maintaining privacy in open transmission communications. In this case, the transmission is freely available to search units 750 in a step 1830, such as would occur with a digital RF broadcast, or through a multicast with open a fixed, public socket IP address available in certain transmission protocols. In this case, the broadcast audio signal or information signal is made in encrypted form, and membership in the cluster is granted through transfer of a decoding key in a step 1832.

Figure 32B is a schematic block diagram of maintaining privacy in closed transmission communication. In a step 1834, the broadcast unit 710 makes a closed transmission broadcast, such as through a socket IP address, that is not publicly available. In a step 1836, the broadcast unit 710 provides the private address to the search unit 750, which can now hear the closed transmission from the step 1834, which is not encrypted. Alternatively, or in addition to the provision of the private address in the step 1836, the establishment of the connection through the private, closed transmission is effected via a password provided in a step 1838. This password can, for example, be used in the step 1110 (e.g. see Figure 14B) to determine whether the broadcast unit 710 accepts the search unit 750 for audio multicasting.

In this section, the encryption of the musical signal and/or associated information about personal characteristics of members of the cluster 700 is described. The custom compressor 330 of the unit 100 can perform the encryption. In such a case, before joining a cluster, the search unit 750 can only receive some limited information, such as characteristics of the music being heard or some limited characteristics of the users in the cluster 700. If the search unit 750 user requests permission to join the cluster 700 and it is granted, the broadcast unit 710 can then provide a decryption key to the search unit 750 that can be used to decrypt the music or provide a private IP address for multicasting, as well as supply additional information about the current members of cluster 700.

It should be noted that in certain cases, it can be useful to have multiple forms of privacy protection. For example, a broadcast unit 710 can provide a search unit 750 access to audio signals and information for the cluster 700, but can reserve certain information based on encryption to only some members of the cluster 700. For example, if a group of friends comprise a cluster 700, and accept some new members into the cluster 700, access to more private information about the friends, or communications between friends, can be restricted on the basis of shared decryption keys.

Figure 33 is a schematic block diagram of a hierarchical cluster, as in Figure 9A, in which communications between different units is cryptographically or otherwise restricted to a subset of the cluster members. Thus, there are three types of communication that are used in the communication: channel A, which takes place between the members of the original cluster; channel B, which takes place between the members of the original cluster (mediated through the broadcast unit 710) and members of the sub-cluster 701; and channel C, which takes place between the members of the sub-cluster 701. Thus, a communications originating from the broadcast unit 710 can be directed either through channel A or channel B, and likewise, a communications originating from the broadcast receive unit 770 can be directed either at members only of the sub-cluster 701 through channel C, or to all members of the cluster 700 through both channels C and B, which is then communicated trough channel A.

A number of means can be used to maintain such independent channels. For example, separate socket communications can be established, and the originators of the communications can determine that information which is carried on each separate channel. For example, given an open transmission scheme such as digital RF signal, the information can be encoded with separate keys for the different channels of communication – thus, the cryptographic encoding determines each channel. A given unit 100 can respond to more than one encoding. Indeed, a channel identifier can be sent with each piece of information indicating the ID of the decoding key. If a unit 100 does not have the appropriate decoding key, then it is not privy to that channel communications.

Alternatively, if the communications is IP socket based, then each channel is determined by IP socket addresses. Furthermore, access to those addresses can be, for example, password controlled. Also, the socket communications can be broadcast so that any unit 100 can receive such broadcast, but that decoding of the broadcast can be mediated through cryptographic decoding keys.

It should be noted that there can be multiple forms of communication, which can comprise messaging communications using the TCP/IP protocols, versus multicasting using

UDP protocols, and also DJ 200 control signals using yet another protocol. The access to each of these communications can be controlled via different privacy hierarchies and techniques. For example, the audio multicasting will be available to all members within a cluster, while the messaging may retain different groupings of privacy (e.g. hierarchical), while the DJ control signals will generally be limited to communications between a given unit 100 and its corresponding DJs 200.

Broadcast Control Transfer

The dynamics of cluster 700 can be such that it will be desirable for a receive unit 730 to become the broadcast unit for the cluster. Such a transfer of broadcast control will generally require the acquiescence of the broadcast unit 710 user. To effect such a transfer, the user of the receiver unit 730 desiring such control will send a signal to the broadcast unit 710 expressing such intention. If the user of the broadcast unit 710 agrees, a signal is sent to all of the members of the cluster indicating the transfer of broadcast control, and providing the identifier associated with the receive unit 730 that is to become the broadcast unit 710. The broadcast unit 710 that is relinquishing broadcast control now becomes a receive unit 730 of the cluster 700.

It should be noted that the transfer of control as described above requires the manual transfer of control, such as actuation of a DJ switch. This switch can be limited to this function, or can be part of a menu system, in which the switch is shared between different functions. It is also within the spirit of the present invention that there be voice-activated control of the unit 100, in which the unit 100 further comprises a microphone for input of voice signals to a suitable controller within the unit 100, wherein the controller has voice-recognition capabilities.

In the case of a cluster 700 whose broadcast unit 710 is no longer broadcasting (e.g. it is out of range of the receive units 730, or it is turned off), the cluster can maintain its remaining membership by selecting one of the receive units 730 to become the new broadcast unit 710. Such a choice can happen automatically, for example by random choice, by a voting scheme, or by choosing the first receive unit 730 to have become associated with the broadcast unit 710. If the users of the cluster-associated units deem this choice to be wrong, then they can change the broadcast unit 710 manually as described above.

The receive unit 730 that is chosen to become the broadcast unit 710 of the cluster 700 will generally prompt its user of the new status, so that the newly designated broadcast unit 710 can make certain that it is playing music to the rest of the cluster 700. It can be

further arranged so that a newly-designated broadcast unit 710 will play music at random, from the beginning, or a designated musical piece in such case.

An embodiment of a transfer of broadcast control using IP socket communications protocols is described here. Figure 16 is a schematic block flow diagram of transfer of control between the broadcast unit 710 and the first receive unit 730. In a step 1130, the receive unit 730 requests broadcast control (designated here as "DJ" control). In a step 1132, the user of the broadcast unit 710 decides whether control will be transferred. The decision is then transferred back to the first receive unit 730 via the TCP messaging socket. If the decision is affirmative, the first receive unit 730 severs its UDP connection to the broadcast unit 710 multicast. The reason for this is to allow the receive unit 730 opportunity to prepare the beginning of its broadcast, if such time is required, and the user cannot both listen to the multicast as well as prepare its own audio selections, which occurs in a step 1136. In a step 1138, the receive unit 730 creates a multicast UDP socket with which it will later broadcast audio to other members of the cluster, while in a step 1140, the receive unit 730 creates a broadcast annunciator TCP socket with which to announce availability of the cluster, as well as to accept transfers of members from the broadcast unit 710 to itself as the new broadcast unit.

When the two new sockets (multicast and annunciator) are created, the receive unit 730 transmits the new socket addresses to the broadcast unit 710 in a step 1142. Since the other members of the cluster are guaranteed to be in contact with the broadcast unit, they can get addresses of the new, soon-to-be broadcast unit from the existing broadcast unit. In a step 1144, the original broadcast unit 710 transmits to the other cluster members (receive units 730 numbers 2-N) the addresses of the sockets on the receive 1 unit 730 that is now the new broadcast unit 710, and terminates its own multicast. The termination is performed here because the other receive units will be transferring to the new multicast, and because the original broadcast unit 710 is now becoming a receive unit 730 in the reconstituted cluster. In the step 1148, multicast of audio is now provided by the receive 1 unit 730 that has now become the new broadcast unit 710), and the original broadcast unit is listening to audio provided not by itself, but rather by the new broadcast unit.

In a step 1146, performed roughly synchronously with the step 1144, the original broadcast unit 710 transmits the socket addresses of the message handler TCP sockets of the other members of the cluster 700 (i.e. the receive units 730 numbers 2-N). In the subsequent step 1150, the original broadcast unit 710 and the receive units 730 numbers 2-N establish new messaging connections with the receive 1 unit 730 that is now the new broadcast unit

710. While there can be a set of criteria for the acceptance of a new member to a cluster, because the receive 1 unit 730 has received the message socket addresses of the other members of the cluster in the step 1144, the receive 1 unit 730 accepts new members with the socket addresses received. It should be noted that instead of socket addresses being the identifiers passed, the identifiers can also be unique machine IDs, random numbers, cryptograpically encoded numbers, or other such identifiers that can be transmitted from one member of the cluster to another.

It should be noted in certain embodiments, that there can be insufficient time for the new broadcast unit 710 to determine a set of music to broadcast to the members of its cluster. It is within the spirit of the present invention for a user to set a default collection of music that is broadcast when no other music has been chosen. This set of music can comprise one or more discrete audio files.

Audio and DJ Choreography

One of the attractions of the present invention is that it allows users to express themselves and share their expressions with others in public or semi-public fashion. Thus, it is highly desirable for users to be able to personalize aspects of both the audio programming as well as the displays of their DJs 200.

Audio

Audio personalization comprises the creation of temporally linked collections of separate musical elements in "sets." These sets can be called up by name or other identifier, and can comprise overlapping selections of music, and can be created either on the unit 100 through a visual or audio interface, or can be created on a computer or other music-enabled device for downloading to the unit 100.

In addition, the unit 100 or other device from which sets are downloaded can comprise a microphone and audio recording software whereby commentary, personal music, accompaniment, or other audio recordings can be recorded, stored, and interspersed between commercial or pre-recorded audio signals, much in the manner that a radio show host or "disc jockey" might alter or supplement music. Such downloads can be accessible from a variety of sources including Internet web sites and private personal computers.

Automatic Generation of DJ 200 Control Signals

In this section, we will describe the automatic and manual generation of control signals for the DJ 200 transducers. The control signals are generally made to correspond to

audio signals played on the units 100, although it is within the spirit of the present invention for such control signals to be made separate from audio signals, and to be displayed on the digital jewelry independently of audio signals played on the unit 100. Figure 20 is a time-amplitude trace of an audio signal automatically separated into beats. Beats 1180, 1182 and 1183 are denoted by vertical dashed and dotted lines and, as described below, are placed at locations on the basis of their rapid rise in low-frequency amplitude relative to the rest of the trace. As can be seen, the beats 1180 are generally of higher amplitude than the other beats 1182 and 1183, and represent the primary beats of a 4/4 time signature. The beat 1183 is of intermediate nature between the characteristics of the beats 1180 and 1182. It represents the third beat of the second measure. Overall, the audio signal thus displayed can be orally represented as ONE-two-Three-four-ONE-two-Three-four ("one" is heavily accented, and the "three" is more lightly accented), which is common in the 4/4 time signature.

Processing of this data can proceed via a number of different methods. Figure 21A is a block flow diagram of a neural network method of creating DJ 200 transducer control signals from an audio signal as shown in Figure 20. In a step 1200, audio data is received either at the unit 100 or the DJ 200. It should be noted that the creation of control signals from audio signals can, within the present invention, take place at either the unit 100 or the DJ 200, or even at a device or system not part of or connected to the unit 100 or DJ 200 (as will be described in more detail below). In an optional step 1202, the data is low pass filtered and/or decimated so that the amount of data is reduced for computational purposes. Furthermore, the data can be processed for automatic gain to normalize the data for recording volume differences. Furthermore, the automatic gain filtering can provide control signals of significant or comparable magnitude throughout the audio data.

In general, the creation of the audio signal depends on audio representing a period of time, which can be tens of milliseconds to tens of seconds, depending on the method. Thus, the audio data from the step 1202 is stored in a prior data array 1204 for use in subsequent processing and analysis. At the same time, the current average amplitude, computed over an interval of preferably less than 50 milliseconds, is computed in a step 1208. In broad outline, the analysis of the signal compares the current average amplitude against the amplitude history stored in the prior data analysis. In the embodiment of Figure 21A, the comparison takes places through neural network processing in a step 1206, preferably with a cascading time back propagation network which takes into account a slowly varying time signal (that is, the data in the prior data array changes only fractionally at each computation, with most of the data remaining the same). The use of prior steps of neural network processing in the

current step of neural network processing is indicated by the looped arrow in the step 1206. The output of the neural network is a determination whether the current time sample is a primary or a secondary beat. The neural network is trained on a large number of different music samples, wherein the training output is identified manually as to the presence of a beat.

The output of the neural network is then converted into a digital jewelry signal in a step 1210, in which the presence of a primary or secondary beat determines whether a particular light color, tactile response, etc., is activated. This conversion can be according to either fixed, predetermined rules, or can be determined by rules and algorithms that are externally specified. Such rules can be according to the aesthetics of the user, or can alternatively be determined by the specific characteristics of the transducer. For example, some transducers can have only a single channel or two or three channels. While light transducers will generally work well with high frequency signals, other transducers, such as tactile transducers, will want signals that are much more slowly varying. Thus, there can be algorithm parameters, specified for instance in configuration files that accompany DJ 200 transducers, that assist in the conversion of beats to transducer control signals that are appropriate for the specific transducer.

Figure 21B is a block flow diagram of a deterministic signal analysis method of creating DJ 200 transducer control signals from an audio signal as shown in Figure 20. The data is received in the step 1200. In this case, a running average over a time sufficient to remove high frequencies, and preferably less than 50 milliseconds, is performed in a step 1212. Alternatively, a low pass filter and/or data decimation as in the step 1202 can be performed.

In a step 1214, the system determines whether there has been a rise of X-fold in average amplitude over the last Y milliseconds, where X and Y are predetermined values. The value of X is preferably greater than two-fold and is even more preferably three-fold, while the value of Y is preferably less than 100 milliseconds and is even more preferably less than 50 milliseconds. This rise relates to the sharp rises in amplitude found in the signal at the onset of a beat, as shown in Figure 20 by the beat demarcations 1180, 1182, and 1183. If there has not been a rise meeting the criteria, the system returns to the step 1200 for more audio input.

If the signal does meet the criteria, it is checked to ensure that the rise in amplitude is not the "tail end" of a previously identified beat. For this, in a step 1216, the system determines whether there has been a previous beat in the past Z milliseconds, where Z is a predetermined value preferably less than 100 milliseconds, and even more preferably less

than 50 milliseconds. If there has been a recent beat, the system returns to the step 1200 for more audio input. If there has not been a recent beat, then a digital jewelry signal is used to activate a transducer. The level of transduction can be modified according to the current average amplitude which is determined in a step 1208 from, in this case, the running average computed in the step 1212.

The embodiment of Figure 21B provides transducer activation signals at each rapid rise in amplitude, with the activation signal modulated according to the strength of the amplitude. This will capture much of the superficial musical quality of the audio signal, but will not capture or express more fundamental patterns within the audio signal.

Figure 21C is a schematic flow diagram of a method to extract fundamental musical patterns from an audio signal to create DJ 200 control signals. In the step 1200, the audio data is received into a buffer for calculations. In a step 1220, a low pass filter is applied to remove high frequency signal. Such high frequency signals can alternatively be removed via decimation, running averages, and other means as set forth in the embodiments of Figures 21A and B. As in the embodiment of Figure 21B, beat onsets are extracted from the audio signal in the steps 1214 and 1216, and a current average amplitude is computed in a step 1208.

The amplitudes and times of the onsets of beats are placed into an array in a step 1222. From this array, a musical model is created in a step 1224. This model is based on the regularity of beats and beat emphasis – as seen in the amplitudes – that is independent of the beats and amplitudes in any one short section of music (corresponding, for instance, to a measure of music).

In general, music is organized into repeating patterns, as represented in a time signature such as 3/4, 4/4, 6/8 and the like. Within each time signature, there are primary and secondary beats. In general, the downbeat to a measure is the first beat, representing the beginning of the measure. The downbeat is generally the strongest beat within a measure, but in any given measure, another beat may be given more emphasis. Indeed, there will be high amplitude beats that may not be within the time signature whatsoever (such as an eighth note in ¾ time that is not on one of the beats). Thus, by correlating the beats to standard amplitude patterns, the output to the music model identifies the primary (down) beats, secondary beats (e.g. the third beat in 4/4 time) and the tertiary beats (e.g. the second and fourth beats in 4/4/ time).

Figure 21D is a schematic flow diagram of an algorithm to identify a music model, resulting in a time signature. In a step 1600, the minimum repeated time interval is

determined, using the array of beat amplitude and onset 1222. This is, over a period of time, the shortest interval for a quarter note equivalent is determined, wherein the time signature beat frequency (i.e. the note value of the denominator of the time signature, such as 8 in 6/8) is preferably limited to between 4 per second and one every two seconds, and even more preferably limited to between 3 per second and 1.25 per second. This is considered the beat time.

From the array of beat amplitudes and onsets 1222, the average and maximum amplitudes over a time period of preferably 3-10 seconds is computed in a step 1604. For the beginning of the audio signal, shorter periods of time can be used, though they will tend to give less reliable DJ 200 control signals. Indeed, in this embodiment, the initial times of an audio signal will tend to follow audio signal amplitude and changes in amplitude more than fundamental musical patterns until the patterns are elicited.

In a step 1606, the amplitude of a beat is compared with the maximum amplitude determined in the step 1604. If the beat is within a percentage threshold of the maximum amplitude, wherein the threshold is preferably 50% and more preferably 30% of the maximum amplitude, the beat is designated a primary beat in a step 1612. In a step 1608, the amplitude of non-primary beats is compared with the maximum amplitude determined in the step 1604. If the beat is within a percentage threshold of the maximum amplitude, wherein the threshold is preferably 75% and more preferably 50% of the maximum amplitude, and the beat is greater than a predetermined fraction of the average amplitude, wherein the fraction is preferably greater than 40% and even more preferably greater than 70% of the average beat amplitude, the beat is designated a secondary beat in a step 1614. The remaining beats are denoted tertiary beats in the step 1610.

In a step 1616, the sequence of the three types of beats is compared with that of established time signatures, such as 4/4, 3/4, 6/8, 2/4 and others, each with their own preferred sequence of primary, secondary and tertiary beats, in order to determine the best fit. This best fit is identified as the time signature in a step 1618.

Returning to Figure 21C, the channels of the DJ are pre-assigned to four different beats in a step 1225. Thus, if there are four channels, each channel is given a separate assignment. With a smaller number of channels, a single channel is assigned multiple beats. Some beats can also be unassigned, thus not being represented in a DJ 200 transducer output. Thus, a high jewelry signal, medium jewelry signal, low jewelry signal and an amplitude dependent signal are each assigned to a channel for DJ 200 transduction.

In a step 1226, a beat determined to be a primary/down beat is assigned to a high jewelry signal 1228. In a step 1230, a beat determined to be a secondary beat is assigned to a medium jewelry signal 1232. In a step 1234, a beat determined to be a tertiary beat is assigned to a low jewelry signal 1236. Beats which are then unassigned, and which will generally be beats that occur not within the music model of the step 1224 (e.g. rapid beats not falling on beats of the time signature) are then assigned in a step 1238 to an amplitude dependent (and not music model dependent) signal 1240.

It should be noted that the computations performed in the flow methods of Figures 21A-C may take time on the order of milliseconds, such that if the computations are made in real time during the playing of music, the activation of the transducers in the DJ 200 are "behind" in time relative to the audio playing of the corresponding music in the audio unit 100. This can be compensated for by carrying out the computations while the audio signal is still in buffers prior to being played in the unit 100, as is described above for numerous embodiments of the present invention. Thus, signals to the DJ 200 can then be made simultaneously with respect to the audio signal to which it corresponds.

It should be noted that many of the parameters described above can conveniently be affected by manual controls either on the DJ 200 or the unit 100 that transmits signals to the DJ 200. For example, if can be convenient for the user to be able to set, for a given DJ 200 response amplitude, the threshold audio amplitude level at which the output transducer (e.g. light transducer 240) responds, or to set the output transducer amplitude corresponding to a maximum audio amplitude, or to set the frequency bands for which different DJ 200 channels respond, or to set other similar parameters. The manual controls for such parameters can comprise dials, rocker switches, up/down button, voice or display menu choices, or other such controls as are convenient for users. Alternatively, these choices can be set on a computer or other user input device, for download onto the unit 100 or DJ 200.

A preferable means of setting the parameters is for the parameters to be stored in a configuration file that can be altered either on the unit 100, the DJ 200 or a computer, so that the same DJ 200 can take on different characteristics dependent on the configuration settings within the file. The configuration settings can then be optimized for a particular situation, or set to individual preference, and be traded or sold between friends or as commercial transactions, for instance over the Internet. For a most preferable use of these configuration files, each file with its set of configurations can be considered to represent a "mode" of operation, and multiple configuration files can be set on the DJ 200 or the unit 100, depending on where the automatic generation of control signals is performed. The user can

then select from the resident configuration files, appearing to the user as different modes, for use of his system, and can change the mode at will. This can be arranged as a series of choices on a voice or display menuing system, as a list toggled through by pressing a single button, or through other convenient user interfaces.

Manual Generation of DJ 200 Control Signals

In the description above, the use of filtering and digital modification of audio signals can be used to create control signals for DJ 200 transducers 240, 250, and 260. In addition, manual choreography of DJ 200 signals can be accomplished. For example, buttons or other interface features (e.g. areas on a touch-screen) on the unit 100 can correspond to different arrays of transducers, such as the LED arrays 290 and 292 of Figure 2A. While playing the audio signal, the user can press the buttons, where pressing of the buttons can correspond to a control signal for a transducer being ON, and otherwise the signal can be off. To aid in choreography where rapid changes in transducers are desired, the audio can be played at less than normal speed.

Figure 22A is a top-view diagram of an audio unit 100 user interface 1250, demonstrating the use of buttons to create DJ 200 control signals. The interface 1250 comprises a display screen (e.g. LCD or oLED), which can display information to the user, such as shown in Figures 18 A-B. Standard music control buttons 1254 for playing, stopping, pausing, and rewinding allow the user to control the audio signal musical output. Buttons 1252 further control aspects of the music output, such as volume control, musical tracks, and downloading and uploading of music. The number of buttons 1252 is conveniently three as shown, but can be more or less than three.

In addition, buttons are provided to allow the user to input DJ 200 control signals, comprising a record button 1256, a first channel button 1258, a second channel button 1260 and a third channel button 1262. The channel buttons 1258, 1260 and 1262 are prominent and accessible, since the user will want to easily depress the buttons. A record button 1256 allows the user to activate the channel buttons 1258, 1260 and 1262, and has a low profile (even below the nominal surface of the interface 1250) so that it is not accidentally activated. The record button can serve various purposes, including recording into a permanent storage file the sequence of DJ control signals relative to music being played, or controlling the DJ transducers in realtime, synchronously with music being played on the audio unit 100.

Pressing the buttons 1258, 1260 and 1262 create DJ control signals for the corresponding channels. The number of buttons is conveniently three as shown, but can also

be two or four or more buttons. If a telephone is being used as the unit 100, keys on the telephone keypad can alternatively be used. The channel buttons will generally be used with thumbs, and the buttons are spaced so that two of the buttons can be depressed with a single thumb, so that all three buttons can be activated with only two fingers. In is also convenient for the two secondary buttons 1260 and 1262 to be spaced more closely together, as it will be a preferred mode of operation that the secondary buttons be operated together from time to time.

To further aid in the choreography of the DJs 200, a separate "keyboard" with the number of keys related to the number of possible arrays can be used. The amplitude of the corresponding transducer signal can be modified either according to the pressure on the keys, according to the length of time that a key is depressed, or according to a foot pedal. Figure 22B is a top-view diagram of a hand-pad 1270 for creating DJ control signals. The hand-pad 1270 comprises a platform 1271, a primary transducer 1272, a secondary transducer 1274 and a tertiary transducer 1276. The platform 1271 has a generally flat top and bottom, and can conveniently be placed on a table, or held in the user's lap. The size of the platform is such that two hands are conveniently placed across it, being preferably more than 6 inches across, and even more preferably more than 9 inches across. The pressure transducers 1272, 1274 and 1276 respond to pressure by creating a control signal, with said control signal preferably capturing both the time and amplitude of the pressure applied to the corresponding transducer. The primary transducer 1272 creates a primary control signal, the secondary transducer 1274 creates a secondary control signal and the tertiary transducer 1276 creates a tertiary control signal. The sizes and placements of the transducers can be varied within the spirit of the present invention, but it is convenient for the primary transducer 1272 to be larger and somewhat separate from that of the other transducers 1274 and 1276. In one more method of user interaction, both hands can be rapidly and alternately used to make closely spaced control signals on the primary transducer 1272. In addition, it can be convenient on occasion for the user to activate both the secondary transducer 1274 and the tertiary transducer 1276 with different fingers on one hand, and thus these can be conveniently placed relatively near to one another. In general, while a single transducer will provide minimal function, it is preferable for there to be at least two transducers, and even more preferable that there be three transducers.

The control signals can be transferred to the audio unit 100 for playing and/or storage, or to the DJ 200 unit directly for playing, either wirelessly, or through wired communication. In addition, the hand-pad can also be configured to create percussive or other sounds, either

directly through the incorporation of hollow chambers in the manner of a drum, or preferably by the synthesis of audio waveform signals that can be played through the audio unit 100 (and other audio units 100 participating in a cluster 700), or directly through speakers within the hand-pad 1270 or attached to the hand-pad 1270 through wired or wireless communications. Such audible, percussive feedback can aid the user in the aesthetic creation of control signals.

It is within the spirit of the present invention for the hand-pad to take on various sizes and configurations. For instance, it is also convenient for the hand-pad 1270 to be configured for the use of index and middle fingers, being of dimensions as small as two by four inches or less. Such a hand-pad is highly portable, and can be battery powered.

Additionally, DJ 200 control signals can also be manually generated live, during broadcast at a party, for example, by a percussionist playing a set of digital drums. Figure 22C is a schematic block diagram of a set of drums used for creating DJ control signals. The set of drums comprises four percussive instruments 1280, 1282, 1284 and 1286, which can include snare drums, foot drums, cymbals, foot cymbals and other percussive musical instruments, such as might be found with a contemporary musical "band". Microphones 1290 are positioned so as to receive audio input primarily from instruments to which they are associated. One microphone can furthermore be associated with multiple instruments, as with the drums 1282 and 1284. The microphones 1290 are connected with a controller 1292 that takes the input and creates DJ control signals therefrom. For example, the drums 1282 and 1284 can be associated with the primary channel, the drum 1280 can be associated with the secondary channel, and the drum 1286 can be associated with the tertiary channel. The association of the microphone input with the channel can be determined in many ways. For example, the jack in the controller 1292 to which each microphone 1290 attaches can correspond to a given channel. Alternatively, the user can associate the jacks in the controller to different channels, with such control being manual through a control panel with buttons or touch control displays, or even through prearranged "sets". That is, a set is a pre-arranged configuration of associations of microphones to channels, and thus a set can be chosen with a single choice that instantiates a group of microphone-channel associations.

In general, the inputs from the microphones 1290 will be filtered in frequency and also to enhance audio contrast. For instance, control signals can be arranged to be the highest when the low-frequency envelope is rising the quickest (i.e. the beat or sound onset). The algorithms for conversion of audio signal to DJ control signal can be pre-configured in the controller 1292, or can be user selectable.

It should be noted that the methods and systems of Figures 22 A-C need to synchronize the control signals so generated with the audio files to which they correspond. This can be accomplished in many ways. For example, the first control signal can be understood to correspond to the first beat within the audio file. Alternatively, the audio unit 100 or other device that is playing the audio signal to which the control signal is to correspond can send a signal to the device that is creating the control signals indicating the onset of playing of the audio file. The control signal can then be related to the time from the onset of the audio file. In addition, with regards to this synchronization, the user manually inputting the control signals will always be listening to the music during the control signal input. If the device on which control signals are being input is the same as the device that is playing the music, a control signal input cam be easily related to the sound that is currently being played by the audio output – many such devices allow information to within less than a millisecond of what sample or time within the audio files is currently being output by the audio device. With the arrangement of the control signal input device being also an audio player, close calibration of the control signals and the audio output is easily accomplished.

DJ 200 Control Signal Files

The control signals can be in a variety of formats within the spirit of the present invention. Such formats include pairs of locations within the associated music file and the corresponding amplitudes of the various DJ channels, and pairs of locations and the amplitudes of those DJ channels which are different from before. The locations can be either time from the start of the song (e.g. in milliseconds) or in terms of sample number. If the location is given in terms of sample number, the sample rate of the music will generally also be provided, since the same song can be recorded at different sample rates, and the invariant in terms of location will generally be time from onset of the music.

Other formats include an amplitude stream, corresponding to each DJ channel, provided in a constant stream with a fixed sample rate, which may be equal to or different from that of the corresponding music file. This format can be stored, for example, as additional channels into the music file, such that one channel corresponds to monoaural sound, two channels correspond to stereo sound, three channels correspond to stereo sound and one channel of control signals, and additional channels correspond to stereo sound plus additional channels of DJ control signals. Another arrangement is to allow for only a small number of states of the transduction in the control signal, so that multiple channels of control signal can be multiplexed into a single transmitted channel for storage and transmission with

the audio signal. For example, if the audio is stored as a 16-bit signal, 3 channels of 5 bit DJ 200 control signal could be stored in a single channel along side the one or two audio channels normally used.

It should be appreciated that these different control signal storage formats are largely interchangeable. For instance, as described above, control signals can be stored as if they are additional audio channels within a music file, but then be extracted from the file for separate transfer (e.g. over the Internet), and then be reintegrated into an audio file at the destination location.

It should be appreciated that there are a number of means by which DJ 200 control signals can be generated, either automatically or manually, and can include the use of devices other than the unit 100 that can have sophisticated digital or analog filtering and modification hardware and software. The control signals so created can be stored in files that are associated with the music files (e.g. MP3) that the control signals are meant to accompany. To aid in their distribution, particularly in reference to limitations on the commercial and private distribution of the corresponding music files, the signal files will generally be separate from the music files, and transferable between units 100 either through inter-unit communication mediated by the inter-unit transmitter/receiver 110, or alternatively through computers or computer networks to which the unit 100 can be connected.

The audio signals and the DJ control signals should also be well synchronized during playback. Figure 23 is a schematic block flow diagram of the synchronized playback of an audio signal file with a DJ control signal file, using transmission of both audio and control signal information. For purposes of convenience in discussion, the audio signal file will be called a "song file" and the "control signal file" will be called a "dance file." In a step 1300, the user is provided a list of song files for display, preferably on the display 1170. In a step 1302, the user then selects a song from the display to play. In a step 1304, the dance files that are associated with the selected song file from the step 1302 are displayed for the user. These song files can be either locally resident on the unit 100, or can alternatively be present on other audio units 100 to which the audio unit 100 is connected, as in a cluster, or can alternatively be on the Internet, if the audio unit 100 is connected to the Internet. If there is a dance file that has been previously preferred in association with the song file, this file can be more prominently displayed than other associated dance files.

In a step 1306, the user selects the dance file to play along with the song file. This association is stored in a local database of song file/dance file associations in a step 1307, to be later used in a subsequent step 1304, should such an association not have been previously

made, or if the preferred association is different from the previously preferred association. If the dance file is not locally resident, it can be copied to the audio unit 100 to ensure that the dance file is available throughout the duration of the song file playback.

In a step 1308, a timer is initialized at the beginning of the song file playback. In the step 1310, the song file is played on the local unit 100, and is also streamed to the other units 100 within the cluster 700. The corresponding DJ control signal accompanies the streaming song, either multiplexed within the song file audio signal, on another streaming socket, or through other communications (e.g. a TCP socket) channels between the two units. In a step 1312, the time advances along with the playback of the music. In a step 1314, this timer information is used to obtain current control signals from the dance file – that is, the dance file is arranged so that at each moment, the status of the different transducer channels can be determined. The control signals to be streamed along with the song file information can be either the current status of each transducer, or alternatively, can only send changes from the current transducer state.

The matching of the files in the database of song file and dance file associations of the step 1307 can be performed both within a machine, but also over a local or wide area network. In such cases, the association can either be external to the file – that is, using the name of the file, that is available the normal system file routines – or can use information internal to one or both files. For example, the dance file can have stored within it a reference to the song to which it is associated, either as the name of the song file, the name and/or other characteristics of the song (such as the recording artist, year of publication, music publisher) or alternatively as a numerical or alphanumerical identifier associated with the song. Then, given a song file, the relationship of the dance file with the song file can be easily determined.

For ease in creating an association, it is convenient for the names of the song files and the associated dance files to have a relationship with one another that is easily understood by casual users. For example, given a song file with the name "oops.mp3", it is convenient for an associated dance file to share the same root (in this case "oops") with a different extension, creating for example the dance file name "oops.dnc". Because of the multiplicity of dance files that will often be associated with a particular song file, the root itself can be extended to allow for either a numerical or descriptive filename, which can be preferably done in conjunction with a known punctuation mark to separate the song file root from the dance file description, such as the file names "oops.david2.dnc" or "oops$wild.dnc". It is

preferable to use a punctuation mark that is allowed within a range of different operating systems.

Dance files can be stored on the Internet or other wide area network in a store for access by users who want dance files associated with a particular song file. In such case, if the storage is through the root of the filename, the user, requesting dance files corresponding to "oops.mps" would then be returned the names of related files such as "oops$wild.dnc". If the dance file internally carries the relationship with "oops.mps" as described above, either through the name or other characteristics, or alternatively, through a numerical or alphanumerical identifier, it is preferable to store the information in a database on the storage computer or unit 100, so that it is not necessary to open the file each time for perusal of the dance file information. Thus, if the music file has a substantially unique identifier associated with it internally, it is also useful for the dance file to also have the same identifier associated internally as well. In such case, the identifier is conveniently used to reference both files within a database.

In operation, a remote user would request a dance file for a particular song file by providing the name of the song file, along possibly with other information about the song file, which could include the name of the choreographer, the number of channels of DJ 200 transduction, the specific brand or type of DJ 200, or other information. The database would then return a listing of the various dance file that met the criteria requested. The remote user would then choose one or more of the files to download to the remote computer, and then the database would retrieve the dance files from storage and then transmits the dance file over the wide area network. On the remote computer or unit 100, the dance file would become associated with the corresponding song file through means such as naming the dance file appropriately or making an association between the song file and the dance file in a database or indexing file. Alternatively, the dance file can be integrated into the song file as mentioned elsewhere within this specification.

It can be useful to preview a dance file for its desirability or suitability. Since the dance files can be retrieved from a wide area network such as the Internet, it is convenient for such an emulator to operate on a computer that may not be portable or have the proper transmitter that allows communications with a DJ 200. In such case, it is preferable to have an emulator which places an image or drawing of a DJ 200 on the screen, and which is provided the name of a song file and a dance file, and which then plays the song file through the audio of the computer and displays appropriate images or drawings of transducers being activated within the emulator image or drawing. The characteristics of the DJ 200 being

emulated (e.g. colors of lights, frequency responses, levels of illumination, arrangement of lights, response to amplitude, etc.) can be simulated by a number of means. For example, the user can move slider controls, set checkboxes and radio boxes, enter numerical values, click-and-drag icons and use other standard user interface controls to make the DJ 200 operate as desired. Alternatively, manufacturers of DJ 200s can create configuration files (including, for example, bitmaps of photos of the actual DJ 200) that can be downloaded for this purpose (and which can also be used by prospective purchasers to view the "virtual" operation of the DJ 200 prior to purchase, for example, through an Internet merchant). The configuration files would contain the information necessary for the emulator to properly display the operation of the specific DJ.

Alternatively, as described above, the dance file information can be stored within the song file as, for example, another channel in place of an audio channel, or alternatively within MP3 header or other file information. In such case, the step 1307 would have the alternative function of looking through song files to find the song file with the particular desired embedded dance file within.

In addition to sending dance files from computers to units 100 or between units 100, the dance files can be streamed from unit 100 to unit 100 through the normal unit-to-unit communications, in the manners described above for audio communications. This is particularly convenient given that DJ 200 displays can be used to show group identification, and such displays can be more effective if the DJs for each user are nearly identical (which might not be the case if the users were using, for example, different dance files). The dance file control signal information can be transmitted in a variety of ways, including multiplexing the control signals into the same packets as the audio information as if it were a different audio channel, alternating packets of control signals with packets of the audio information, or broadcasting control signals on a different UDP socket as the audio. Alternatively, if the receiving unit has a copy of the dance file corresponding to the song file being transferred by unit-to-unit communication, the receiving unit can determine the current time being played, and to extract from the local dance file the control signals for the receiving unit DJ 200.

It should be known that most streaming protocols have relatively small data packets that are communicated, due to the fact that reception at the source is not guaranteed and it is not desirable to lose a large amount of information in any one stream. Thus, it is possible with smaller transmission buffers and higher data rates to send a single DJ control signal in each transmission. For example, with a buffer size of 600 bytes, and an audio rate of 22,050 Hz with two single byte channels, each transmission covers only about 12 milliseconds, and

any signal would therefore be at most 13 milliseconds from its correct time. Alternatively, each control signal can be accompanied by an offset in time from the beginning of the transmitted audio signal. Also, the time or packet number of each transmission buffer can be sent, as well as the time or packet number of the DJ audio signals, so that the audio unit 100 can compute the proper offset.

Stationary Transducers

DJs 200 that have been previously described are portable devices, usually associated with a particular user and unit 100. Figures 5A and 5B indicate the ways in which DJs 200 associated with multiple users can be controlled by a single unit 100.

It is also convenient for transducers to be non-portable and stationary. Consider, for example, a user who is at home listening to music. Instead of a DJ 200 worn by the user, the user can alternatively have a bank of lights or other transducers in fixed locations through the room that operate under the same or similar control signals as to which DJs respond. Such fixed transducers can operate at far higher power than portable DJs 200, and can each incorporate a large number of separate transducers.

Furthermore, in a party, concert or other large social gathering, the effects of portable DJs worn by guests can be supplemented by large transducers that are generally perceptible by most guests. For example, such transducers can include spark or smoke generators, strobe lights, laser painters, arrays of lights similar to Christmas light strings, or mechanical devices with visible (e.g. a flag waving device) or tactile effects (e.g. a machine that pounds the floor). In general, transducers for large gatherings will not communicate with a unit 100, but will be directed by a wide-area broadcast unit 360, as in Figure 5B.

Because of the large area over which such stationary transducers can operate, the communications between the unit 100 and the stationary transducers can be through wired rather than wireless transmission. Furthermore, there can be mixed communication, such as wireless transmission of control signals from a portable unit 100 to a stationary receiver, and thence wired transmission to one or multiple transducers.

Modular Configurations

In the embodiments above, the audio player 130 is directly integrated with the inter-unit and unit-to-DJ communications. This requires both a re-engineering of existing audio players (e.g. CD, MP3, MO and cassette players), and furthermore does not allow the communications functionality to be reused between players.

An alternative embodiment of the present invention is to place the communications functions external to the audio playing functions, and to adjustably connect the two via the audio output port of the audio player. Figure 12A is a schematic diagram of a modular audio unit 132. Audio player 131 is a conventional audio player (e.g. CD or MP3 player) without the functionality of the present invention. Analog audio output is sent via audio output port 136 through the cable 134 to the audio input port 138 of the modular audio unit 132. The modular audio unit 132 comprises the inter-unit transmitter/receiver 110 and the DJ transmitter 120, which can send and receive inter-unit and unit-to-DJ communications in a manner similar to an audio unit 100. A switch 144 chooses between audio signals from the audio player 131 and from the inter-unit transmitter/receiver 110 for output to the output audio port 142 to the earphone 901 via cable 146 (the earphone 901 can also be a wireless earphone, wherein the output port 142 can be a wireless transmitter, which can also be a DJ transmitter 120). A convenient configuration for the switch 144 is a three way switch. In an intermediate position, the unit 132 acts simply as a pass-through, in which output from the audio player 131 is conveyed directly to the earphone 901, and the transmitter/receiver functions of the unit 132 do not operate. In another position, the unit 132 operates as a receiver, and audio from the inter-unit transmitter/receiver 110 is conveyed to the earphone 901.

When the combined system operates as a broadcast unit 710, audio input from the audio unit 131 is directed to the inter-unit transmitter/receiver 110 for transmission to receive units 730, as well as for output to the earphone 901 (which can be direct to the earphone 901 through the switch, or indirectly through the inter-unit transmitter/receiver 110).

When the combined system operates as a conventional audio player, the switch directs audio signals from the input port 138 directly through to the output port 142. In this mode of operation, it can be arranged for the audio output to traverse the modular audio unit 132 without the unit being powered up. In case there is a transmission delay to the receive unit 730 such that audio played locally through the earphone 901 and audio played remotely on the receive unit 730 are not in synchrony, the system can incorporate a time delay in the output port 142 such that the local and remote audio output play with a common time delay, and are thus in synchrony.

When the combined system operates as a receiver unit 730, audio input from the input port 138 is ignored, and signals to the audio output port are delivered solely through the inter-unit transmitter/receiver 110.

It is convenient for the modular audio unit 132 to be able to operate independently of the associated audio player 131. In such a case, the unit 132 must have an independent energy store, such as one or more batteries, which can be rechargeable. In that case, the unit 132 has no audio signals locally to listen to through the earphone 901 or to transmit over the transmitter/receiver 110. However, the unit 132 can in that case receive external audio signals sent by other units 132 or units 100 for listening.

The audio player 131 can be placed in a backpack, purse, or other relatively inaccessible storage location, while the modular audio unit is, like a "remote control", accessible for interaction with other users.

Video

While the units 100 described above have comprised audio players 130, within the spirit of the present invention, such units can also comprise video or audio/visual players (both of which are referred to below as video players). Such video players would be used generally for different entertainment and educational purposes, not limited to films, television, industrial training and music videos. Such video enabled units can operate similarly to audio units, including the capability of sharing video signals, synchronously played, with nearby units through inter-unit communication, as well as the use of DJ's that can produce human-perceptible signals (such as light transduction for accompaniment of audio signals in music videos). It should be noted, however, that there is a larger bandwidth requirement for the inter-unit transmitter/receiver 110 for the communication of video signals as compared with audio signals. In the case of shared video, wire connections (e.g. FireWire) between two units can allow simultaneous viewing of a single video signal.

In addition, text, including language-selectable closed caption and video subtitling, can accompany such video, as well as chat or dubbing to allow the superposition of audio over the audio normally accompanying such video.

Music Distribution Using Audio Units

The music industry is suffering from reduced sales due to the advent of Internet-based music file sharing; in addition, the manufacturers of personal audio devices are bringing to market audio devices that can wirelessly transfer music files between the devices. Such sharing-enabled devices could significantly reduce the sales of music. Audio units of the present invention, however, can be used to provide new means of music distribution and thereby increase the sales of music.

Figure 25 is a schematic flow diagram indicating music sharing using audio devices, providing new means of distributing music to customers. Three entities are involved in the transactions – the DJ (operating a broadcast unit 710), the cluster member (operating a receive unit 730), and the music distributor, and their actions are tracked in separate columns. In this case, the term DJ is used to indicate the person operating a broadcast unit 710, and has no meaning with respect to a DJ unit 200. Indeed, the DJ unit 200 is a part of the system only inasmuch as it provides for heightened pleasure of the DJ and the member in enhancing their experience of the music. For the rest of this section, DJ will refer specifically to the person operating the broadcast unit 710.

In a first step 1340, the DJ registers with the distributor, who places information about the DJ into a database in a step 1342. Part of this information is a DJ identifier (the DJ ID), which is unique to the DJ, and which DJ ID is provided to the DJ as part of the registration process. This ID is stored in the unit 100 for later retrieval. The DJ at some later time broadcasts music of the type distributed by the distributor, in a step 1344. The broadcast of the music by the DJ can be adventitious (that is, without respect to the prior registration of the DJ with the distributor), or the distributor can provide the music to the DJ either free of charge, at a reduced charge, or free of charge for a limited period of time.

In a step 1346, the member becomes a part of the cluster 700 of which the DJ is the broadcaster broadcasting the distributor's music, and has thereby an opportunity to listen to the music. Along with the transfer of the audio signal of the music, in a step 1348, the DJ can send information about the song, which can include a numerical identifier of the music or album from which the music is derived. Furthermore, the DJ ID is provided to the member, and is associated with the music ID and stored in a database on the member unit 100 in a step 1350. In order to prevent this database from becoming too sizable, music IDs and DJ IDs can be purged from it on a regular basis (for example, IDs which are older than 60 or 120 days can be removed).

If the member requests purchase of the music from the distributor in a step 1352, in a step 1354, the distributor stores the member information, the music ID, and the DJ ID associated with the music (i.e. the person who introduced the member to the music). The distributor then completes the transaction with the member, providing a copy of the music in exchange for money, in a step 1356. As the member receives the music copy, he also becomes registered as a DJ as well in a step 1358. Thus, if the member now becomes the DJ of his own cluster, and introduces people to this music, he will also be known to the distributor as an introducer of the music.

In a step 1360, the distributor provides points to the DJ who introduced the member to the music and facilitated the sale of the music. In a step 1362, the DJ accumulates points related to the sale of the music to the member, as well as points related to the sale of other music to other members. These points can at that point or later be redeemed for money, discounted music, free music, gifts, access to restricted activities (e.g. seats at a concert) or other such real or virtual objects of value to the DJ.

In a step 1364, the DJ is optionally further linked to the music and member for whom he has received points. If this member introduces the music to yet other members, who are induced to buy the music from the distributor, the DJ is further awarded points in a step 1366, given that the "chain" of members introduced directly or indirectly to the music includes the original DJ.

This set of interactions does not decrease music sales as does file sharing, but rather increases sales of music, as the DJ has incentives to encourage others to buy the music, and the offering of the music by the DJ through his broadcasts introduces music to people who may not have already had the opportunity to hear the music.

Figure 31 contains tables of DJ, song and transaction information according to the methods of Figure 25. A USER table 1810 comprises information about the USER, which can include the name of the person (Alfred Newman), their nickname/handle ("WhatMeWorry"), their email address (AEN@mad.com), and the machine ID of their unit 100 (B1B25C0). This information is permanently stored in the audio unit 100. A second set of information relates to music that the USER has heard while in other clusters 700 that the USER liked, and which is indicated as the USER's "wish list". This set of information includes a unique ID associated with the song (or other music or audio signal), which is transmitted by the broadcast unit 710 of the cluster 700. This information can alternatively or additionally include other information about the music, such as an album name, an artist name, a track number, or other such information that can uniquely identify the music of interest.

Along with each song ID is a DJ identifier, indicating the unique ID associated with the DJ who introduced the desired music to the USER. Additionally or alternatively, the information can comprise the DJ's email address, personal nickname/handle, name, or other uniquely identifying information.

The Wish List can either be permanent, or it can be that each song entry is dated, and that after a predetermined amount of time, which can be set by the user, the songs that are still on the Wish List are removed. It is also convenient that songs that are purchased

according to the methods of the present invention, such as Figure 25, are also removed from the list automatically.

A DISTRIBUTOR table 1812 comprises information about purchases made by USERS with the DISTRIBUTOR. The table 1812 has numerous records keyed according to unique USER identifiers, which in this case is the MAC ID of the unit 100. A single record from the table is provided, of which there can be hundreds of thousands or millions of such records stored.

The record can include contact information about the USER, including name, email address, or other business related information such as credit card number. In addition, each record comprises a list of all of the songs known to have been purchased through the DISTRIBUTOR, as identified by a unique song ID. In addition, the DJ associated with the purchase of the given song by the USER is also noted. This information was previously transmitted from the USER table 1810, which includes the associated DJ identifier along with the song identifier, at the time of purchase of the song. This association allows the DISTRIBUTOR to compensate the DJ for his part in introducing the USER to the song.

It should also be noted that such an arrangement of information allows the compensation, if desired, of the individual who introduced the DJ to the song, prior to the DJ introducing the USER to the song. For example, when the user purchased the song with song ID 230871C40, points were credited with the DJ whose ID is 42897DD. Looking in the record for the DJ 42897DD, one can determine whether there is another individual (DJ) associated with the purchase of the song 230871C40 by the DJ. If so, that individual can also receive compensation for the purchase of the song by the USER.

Use of Internet Connections

It is within the teachings of the present invention to allow normal Internet connections of the audio unit 100 with non-mobile devices connected with the Internet. Figure 29A is a schematic block diagram of the connection of an Internet-enabled audio unit 100 with an Internet device through the Internet cloud 1708, using an Internet access point 1704. An Internet-enabled audio unit 1700, unit A, is wirelessly connected to an audio unit 100, denoted unit B, as members in a cluster 700. The dashed line connecting the two units A and B indicates that the connection is wireless, whereas the solid connecting lines indicate wired connections. The unit A is connected to a wireless access point 1704, such as an 802.11 access point, which is connected to an Internet device 1706 via wired connections through the Internet cloud 1708.

Figure 29B is a schematic block diagram of the connection of an Internet-enabled audio unit 1702 with an Internet device through the Internet cloud, with an audio unit 1702 directly connected to the Internet cloud 1708. In this case the audio unit 1702 is capable of directly connecting to the Internet cloud 1708, and thence to the Internet device 1706, through a wired connection. This could be through a high speed connection (such as a twisted wire Ethernet connection) or through a lower speed connection (e.g. a serial port connection, or a dial-up modem).

The connection of the unit 1700 or unit 1702 is illustrated in Figure 30, tables of ratings of audio unit 100 users. As described above, members of a cluster can decide whether or not to admit a new member to the cluster using a variety of automatic or manual methods. One method of determining the suitability of a user to become a member of the cluster 700 is to determine the user's ratings by members of other clusters to which the user has previously been a member. In this case, the Internet device 1706 is a computer hosting a database, which can be queried and to which information can be supplied by the unit A (either 1700 or 1702). On the Internet device 1706 are stored ratings of units 100, as indicated by the table 1802. The left hand column is the primary key of the database, and is a unique identifier associated with each unit 100. This ID can be a numerical MAC ID, associated with the hardware and software of each unit 100, a unique nickname or word handle (e.g. "Jen412smash") associated with each audio unit user, or other such unique identifier.

The second and third columns, indicated as numbers with dollar signs, are the total summed positive ratings (column two) and the negative ratings (column three) registered with each user by another member of a cluster 700 with which the user has been associated, and in which the user was operating the broadcast unit 710. This rating can, for example, reflect the perceived quality of music provided by the user. The fourth and fifth columns are the total, summed ratings of the user by other members of clusters 700 with which the user has been associated, in which the user was the operator of a receive unit 730. This rating can, for example, indicate the good spirits, friendliness, dress or other characteristics of the user as perceived by other members of the cluster. The sixth column indicates the largest cluster 700 for which the user has been the broadcaster. This is a good indicator of a broadcaster's popularity, since a poor or unpopular broadcaster would not be able to attract a large group of members for a cluster.

There are many other characteristics that can be stored in such a database, and can also include IDs of other members of groups with which the user has been associated (so that members can accept new members who have been associated with friends of those in the

cluster), specific music that the user has played (in order to determine musical compatibility), information on the individuals making each rating (in order to determine rating reliability), and gradations of ratings (rather than simply a positive or negative response).

The cluster members can access the ratings of the user requesting membership in the cluster 700 in order to determine their desirability and suitability. This would require a connection with the Internet device 1706 at the time that the user was requesting to join, and would preferably involve a wireless connection through an access point, as in Figure 29A. The information from the database on the device 1706 can either be displayed to the members of the cluster 700, or can be used by an automatic algorithm to determine whether the person can join.

The table 1800 represents the ratings of a cluster 700 of 5 total members (comprising a broadcaster with ID 12089AD, and four additional members with IDs E1239AC, F105AA3, B1B25C0, and ED5491B). The ratings are supplied by ED5491B (whose ID is preceded by a zero), and then specific ratings of each member are made. The DJ is indicated by a dollar sign preceding his ID. These ratings can be made by putting the nicknames/handles of the cluster members on a screen, and allowing the member to indicate positive or negative ratings by pressing one of two buttons. A plus in the first column indicates a positive response, and a minus sign indicates a negative response. These ratings can then be sent during either wired communications directly to the Internet device 1706 or via the access point 1704. It should be noted that the ratings, once made, can be stored on the unit 1700 or 1702 indefinitely, until connection with the Internet cloud 1708 can be made. As indicated by the arrow, the information for B1B25C0 can be added to the table 1802 – in this case, by incrementing the value in the fourth column (a positive rating for a user who is not the broadcaster).

Other applications of connections to Internet devices 1706 include exchanging (via uploading and downloading) dance files with distant individuals, and obtaining music via downloading , which can include transactions with distributors similar to that seen in Figure 25. Such connections also allow the integration of other connectivity, such as telephone and messaging capabilities, expanding the usefulness and attractiveness of audio units 100.

Many Embodiments Within the Spirit of the Present Invention

It should be apparent to one skilled in the art that the above-mentioned embodiments are merely illustrations of a few of the many possible specific embodiments of the present invention. For example, the elements of a unit 100, including the inter-unit transmitter/receiver 110 protocol and hardware, the DJ transmitter 120 and the audio player

130 can be chosen from a range of available technologies, and can be combined with user interface elements (keyboards, keypads, touch screens, and cursor buttons, without significantly affecting the operation of the unit 100. Furthermore, many different transducers can be combined into DJs 200, which can further comprise many decorative and functional pieces (e.g. belt clasps, functional watches, microphones, or wedding rings) within the spirit of the present invention. Indeed, the unit 100, itself, can comprise transducers 240, 250 or 260.

It should also be appreciated that communications protocols provide a nearly uncountable number of arrangements of communications links between units in a cluster, that the links can be of mixed software protocols (e.g. comprising both TCP and UDP protocols, and even non-IP protocols) over a variety of hardware formats, including DECT, Bluetooth, 802.11 a, b, and g, Ultra-Wideband, 3G/GPRS, and i-Beans, and that communications can include not only digital but also analog communications modes. Furthermore, communications between audio units and digital jewelry can further comprise analog and digital communications, and a variety of protocols (both customized as well as well-established IP protocols).

It is important, as well, to note that the inter-unit communication and the unit-to-DJ communication can operate and provide significant benefits independently of one another. For example, members listening to music together gain the benefits of music sharing, even without the use of DJs 200. Alternatively, an individual's appreciation of music and personal expression can be augmented through use of a DJ 200, even in the absence of music sharing. However, the combination of music sharing along with enhanced personal expression through a DJ 200 provides a synergistic benefit to all members sharing the music.

Numerous and varied other arrangements can be readily devised by those skilled in the art without departing from the spirit and scope of the invention. Moreover, all statements herein reciting principles, aspects and embodiments of the present invention, as well as specific examples thereof, are intended to encompass both structural and functional equivalents thereof. Additionally, it is intended that such equivalents include both currently known equivalents as well as equivalents developed in the future, i.e. any elements developed that perform the same function, regardless of structure.

In the specification hereof any element expressed as a means for performing a specified function is intended to encompass any way of performing that function. The invention as defined by such specification resides in the fact that the functionalities provided by the various recited means are combined and brought together in the manner which the

specification calls for. Applicant thus regards any means which can provide those functionalities as equivalent as those shown herein.

Claims

What is claimed is:

1. A method for sharing music from stored musical signals between a first user with a first music player device and at least one second user with at least one second music player device, comprising:

playing the musical signals for the first user on the first music player device while essentially simultaneously wirelessly transmitting the musical signals from the first music player device to the at least one second music player device; and

receiving the musical signals by the at least one second player device, wherein the musical signals can be played on the at least one second player device essentially simultaneously with the playing of the musical signals on the first music player device;

wherein the first and the at least one second users are mobile and maintain less than a predetermined distance.

2. The method of claim 1 wherein the predetermined distance is 40 feet.

3. The method of claim 1 wherein the predetermined distance is 100 feet.

4. The method of claim 1 wherein the predetermined distance is 500 feet.

5. The method of claim 1 wherein the transmitting and receiving involve the use of an Internet protocol.

6. The method of claim 5 wherein the Internet protocol uses a connection-less transport layer.

7. The method of claim 5 wherein the Internet protocol uses a link layer involving a protocol chosen from the group consisting of 802.11, Bluetooth, ultrawideband, HiperLAN and WDCT.

8. The method of claim 1 wherein the transmitting and receiving use analog radio frequency transmission.

9. The method of claim 1 wherein the transmitted music signals are digital.

10. The method of claim 9 wherein the transmitted digital music signals are compressed.

11. The method of claim 1 wherein the first music player device is carried by the first user.

12. The method of claim 1, additionally comprising displaying by the first user a humanly-perceptible light display that is roughly synchronized with the beat of the musical signals.

13. The method of claim 1, additionally comprising communicating an identifier of the at least one second user to the first music player device, and indicating the identifier to the first user on a display located on the first music player device.

14. The method of claim 1, additionally comprising transferring a substantially unique music identifier of the musical signal wherein the personal identifier is stored on the at least one second music player device in association with the musical identifier.

15. The method of claim 14 wherein the user of the at least one second music player device retrieves the musical signal from a remote system by using the music identifier.

16. The method of claim 15 wherein the remote system is located on the Internet.

17. The method of claim 15, additionally comprising a substantially unique personal identifier of the first user wherein the personal identifier is transferred to the at least one second music player device and stored in association with the music identifier of the music transmitted by the first music player device to the at least one second music player device.

18. The method of claim 1, additionally comprising monitoring the strength of the musical signal received by the at least one second music player device and communicating with the at least one second user the strength of the musical signal.

19. A system of music sharing for a plurality of users, comprising:

a first sharing device and at least one second sharing device, each comprising a musical signal store, a musical signal transmitter, a musical signal receiver, and a musical signal player;

a broadcast user operating the first sharing device; and

at least one member user operating the at least one second sharing device, wherein the broadcast user plays the musical signal for his own enjoyment on the first sharing device and simultaneously transmits the musical signal to the receiver of the at least one second sharing device of the at least one member user, on which the musical signal is played for the at least one member user, wherein the broadcast user and the at least one member user hear the musical signal substantially simultaneously.

20. The system of claim 19, additionally comprising a manual switch on at least one second sharing device, wherein actuation of the switch causes the first sharing device to cease transmitting the musical signal and for the second sharing device to begin transmitting musical signal from the musical signal store of the second sharing device, wherein the broadcast user and the at least one member user hear the musical signal from the second sharing device substantially simultaneously.

21. The system of claim 19 wherein the first sharing device and the at least one second sharing device additionally comprise a substantially unique identifier, wherein the first sharing device communicates its substantially unique identifier to the at least one second sharing device and the at least one second sharing device communicates its substantially unique identifier to the first sharing device.

22. The system of claim 21, wherein the identifier includes a least one piece of information about the user of the sharing device taken from the group consisting of the name of the user, the nickname of the user, a numerical identifier, an alphanumerical identifier, the musical preferences of the user, the age of the user, and the educational institution that the user attends.

23. The system of claim 22 wherein the first sharing device and the at least one second sharing device additionally comprise a visual display on which the substantially unique identifiers which have been communicated to the respective device are displayed.

24. The system of claim 19 additionally comprising at least one requesting member user operating an at least one requesting sharing device, wherein the at least one requesting member user requests to receive transmission of the musical signal from the first sharing device.

25. The system of claim 24 wherein the at least one requesting sharing device is associated with a substantially unique identifier which is communicated by the requesting sharing device to the first sharing device and wherein the first sharing device additionally comprises a display on which the identifier is displayed.

26. The system of claim 25 wherein the broadcast user elects whether to transmit the musical signal from the first sharing device to the requesting sharing device.

27. The system of claim 25 wherein the broadcast user and the at least one member user elect through voting whether to transmit the musical signal from the first sharing device, wherein a majority vote is determinative.

28. The system of claim 25 wherein the broadcast user and the at least one member user elect through voting whether to transmit the musical signal from the first sharing device, wherein a unanimous vote is required for transmitting.

29. The system of claim 25 wherein the identifier of the requesting sharing device is compared with identifiers stored in the first device, such that the match of the identifier of the requesting sharing device with a stored identifier allows the first sharing device to transmit the musical signal to the requesting sharing device.

30. The system of claim 29 wherein the stored identifiers are retrieved from the Internet.

31. The system of claim 19 wherein the first sharing device additionally comprises a visual display.

32. The system of claim 31 wherein the visual display indicates the number of member users receiving the musical signal.

33. The system of claim 31 wherein the visual display can display whether the musical being played by the sharing device comes from the musical signal store of the same device, or from the musical signal store another sharing device.

34. The system of claim 19 wherein the first sharing device is mobile.

35. The system of claim 19 wherein the first sharing device is carried by the user.

36. The system of claim 19 wherein the first sharing device and at least one sharing device maintain less than a predetermined distance.

37. The system of claim 36 wherein the predetermined distance is 40 feet.

38. The system of claim 36 wherein the predetermined distance is 100 feet.

39. The system of claim 36 wherein the predetermined distance is 500 feet.

40. The system of claim 19, additionally comprising a substantially unique user identifier associated with the broadcast user and a substantially unique musical signal identifier associated with each musical signal, wherein the user identifier and the musical signal identifier are transmitted from the first sharing device to the at least one second sharing device and stored in association with one another on the at least one second sharing device.

41. A wireless communications system for sharing audio entertainment between a first mobile device and a second mobile device in the presence of a non-participating third mobile device, comprising:

an announcement signal transmitted by the first mobile device for which the second mobile device and the third mobile device are receptive;

a response signal transmitted by the second mobile device in response to the announcement signal for which the first mobile device is receptive and for which the third mobile device is not receptive;

an identifier signal transmitted by the first mobile device to the second mobile device in response to the response signal, and which is not receptive to the third mobile device; and

a broadcast signal comprising audio entertainment that is transmitted by the first mobile device, and which is receptive by the second mobile device on the basis of the reception of the identifier signal.

42. The system of claim 41 wherein the announcement signal is transmitted via an Internet protocol.

43. The system of claim 42 wherein the Internet protocol is a connection-oriented transport protocol.

44. The system of claim 43 wherein the connection-oriented protocol is TCP.

45. The system of claim 42 wherein the Internet protocol further comprises a link layer protocol selected from the group consisting of 802.11, Bluetooth, ultrawideband, HiperLAN and WDCT.

46. The system of claim 41 wherein the broadcast signal is transmitted via an Internet protocol.

47. The system of claim 46 wherein the Internet protocol is a connectionless transport protocol.

48. The system of claim 47 wherein the connectionless transport protocol is UDP.

49. The system of claim 46 wherein the Internet protocol further comprises a link layer protocol selected from the group consisting of 802.11, Bluetooth, ultrawideband, HiperLAN and WDCT.

50. The system of claim 41 wherein the broadcast signal is an analog radio frequency signal.

51. The system of claim 41 wherein the broadcast signal is transmitted via telephony.

52. The system of claim 41 wherein the broadcast signal is compressed.

53. The system of claim 52 wherein the compression is MP3.

54. The system of claim 41 wherein the broadcast signal is encrypted.

55. The system of claim 41 wherein the identifier signal is an IP socket address.

56. The system of claim 41 wherein the identifier signal is a decryption code.

57. The system of claim 41 wherein the identifier signal is a digital password.

58. The system of claim 41 wherein the announcement signal and the broadcast signal use different IP transport protocols.

59. The system of claim 41 additionally comprising a transfer signal that is transmitted by the first mobile device upon halting the transmission of its broadcast signal, and which transfer signal is receptive by the second mobile device, such that reception by the second mobile device causes the device to transmit a second broadcast signal comprising audio entertainment for which the first unit is receptive.

60 The system of claim 41 wherein the broadcast signal is transmitted via peer-to-peer communications.

61 The system of claim 60, additionally comprising a switch, wherein actuation of the switch causes the device to toggle between peer-to-peer communications and infrastructure-based communications.

62 The system of claim 41 wherein the broadcast signal is transmitted via infrastructure-based communications.

63. An audio entertainment device, comprising:

a signal store that stores an audio entertainment signal;

a transmitter that can transmit the stored audio entertainment signal;

a receiver that can receive the transmitted audio entertainment signal from a transmitter of another such device; and

a player that can play audio entertainment from a member selected from the group of stored audio entertainment signals or audio entertainment signals transmitted from the transmitter of another such device.

64. The device of claim 63 wherein the transmitter transmits to the receiver of the same device in order to play through the player of that device.

65. The device of claim 63 wherein the player can play audio entertainment from signals obtained from either the receiver or from the audio store.

66. The device of claim 63 additionally comprising a visual display that is perceptive to the user of the device, indicating information about the audio entertainment signal.

67. The device of claim 66 wherein a third device can receive the transmitted audio signal, wherein all such devices transmitting and receiving the same audio entertainment signal constitute a sharing group.

68. The device of claim 67 additionally comprising a substantially unique identifier that identifies the user of the device, wherein at least one device of the sharing group displays the identifier of the user of another device within the sharing group.

69. The device of claim 68 wherein the identifier includes a least one piece of information about the user taken from the group consisting of the name of the user, the nickname of the user, the musical preferences of the user, a numerical identifier, an alphanumerical identifier, the age of the user, or the educational institution that the user attends.

70. The device of claim 68 wherein the identifier is interspersed within the transmitted audio entertainment signal.

71. The device of claim 66 wherein the visual display can display the number of devices in the sharing group.

72. The device of claim 66 wherein the visual display can display whether the audio entertainment being played by a device comes from the audio store of the device of which the display is a part, or from another device.

73. The device of claim 63 wherein the second device provides information to the second user about the strength of the audio signal received by the receiver.

74. The device of claim 73 additionally comprising a visual display, wherein the signal strength information is provided in the visual display

75. The device of claim 73 wherein the signal strength information is provided secondary audio signals played substantially simultaneously with the audio entertainment signal through the player of the device.

76. The device of claim 63 wherein the signal store stores a compressed audio signal.

77. The device of claim 63 wherein the compressed audio signal is transmitted through the transmitter.

78. The device of claim 63 wherein the compressed audio signal is decompressed prior to being transmitted through the transmitter.

79. The device of claim 63 wherein the identifier is interspersed within the transmitted audio signal.

80. The device of claim 63 additionally comprising a delay circuit.

81. The device of claim 80 wherein the playing of the audio entertainment signal is delayed in each device so that the audio entertainment signal in all devices playing from the same stored audio entertainment signal is played substantially simultaneously.

82. The device of claim 63 additionally comprising a manual control with at least three states, wherein in the first state the player plays a stored audio entertainment signal and does not transmit the stored entertainment signal, wherein in the second state the player plays a stored audio entertainment signal and substantially simultaneously transmits the same audio entertainment signal, and wherein in the third state the player plays an audio entertainment signal received by the receiver.

83. The device of claim 63 wherein the ambient sound admitted by the earphone can be manually controlled.

84. The device of claim 63 additionally comprising a microphone wherein the device can transmit vocal sounds from the user of the device.

85. The device of claim 63 additionally comprising a manual control such that actuation transmits a message.

86. The device of claim 63 wherein the device can connect with the telephone network.

87. The device of claim 63 wherein the device is mobile and can be carried by a user.

88. The device of claim 63 additionally comprising a control signal transmitter that transmits control signals that controls a humanly-perceptible display in a wearable accessory.

89. The device of claim 88 wherein the control signal transmitter is the same as the transmitter that transmits stored audio entertainment signals.

90. The device of claim 88, additionally comprising a store of control signals.

91. The device of claim 88, additionally comprising a generator wherein the generator coverts stored audio entertainment signals into control signals.

92. An audio entertainment device comprising:

a wireless transmitter for the transmission of audio entertainment signals;

a wireless receiver for the reception of the transmitted audio entertainment signals from a transmitter of audio entertainment signals;

a first manually-separable connector for electrically connecting with an audio player, wherein the connector allows transfer of audio entertainment signals from the player to the device;

a second connector for connecting with a speaker; and

a control to manually switch between at least three states, wherein in the first state the speaker plays audio entertainment signals from the audio player and the transmitter does not transmit the audio entertainment signals, wherein in the second state the speaker plays audio entertainment signals from the audio player and the transmitter essentially simultaneously transmits the audio entertainment signals, and wherein in the third state the speaker plays audio entertainment signals received by the receiver.

93. The device of claim 92 wherein the first connector connects to the audio player through the audio output jack.

94. The device of claim 92 wherein the speaker is chosen from the group consisting of earphone and headphone.

95. The device of claim 92 wherein the device additionally comprises a control signal generator that converts audio entertainment signals into control signals that control a humanly-perceptible display in a wearable accessory.

96. The device of claim 92 additionally comprising a second transmitter that transmits control signals that control a humanly-perceptible display in a wearable accessory.

97. The device of claim 92 additionally comprising a display that indicates the state of the device.

98. The device of claim 92 wherein the receiver is receptive of the audio entertainment signals transmitted by the transmitter.

99. The device of claim 92 wherein the receiver is operational when the first connector is manually separated from the audio player.

100. A system for the sharing of stored music between a first user and a second user, comprising:

a first device for playing music to the first user, comprising a store of musical signals, a first controller that prepares musical signals from the first store for transmission and playing, a first player that takes musical signals from the first controller and plays the signals for the first user, and a transmitter, that is capable of taking the musical signals from the controller and transmitting the musical signals via wireless broadcast; and

a second device for playing music to the second user, comprising a receiver receptive of the transmissions from the transmitter of the first device, a second controller that prepares musical signals from the receiver for playing, and a second player that takes musical signal from the second controller and plays the signals for the second user;

wherein the first user and the second user hear the musical signals at substantially the same time.

101. The system of claim 100 wherein the store of musical signals comprises compressed digital audio files.

102. The system of claim 100 wherein the compressed digital audio files are in formats taken from the group consisting of mp3, wav, cda, wma, and au.

103. The system of claim 100 wherein the preparation of the musical signals by the first controller comprises decompression of compressed binary files.

104. The system of claim 100 wherein the transmitter transmits via an Internet protocol.

105. The system of claim 104 wherein the Internet protocol has a link level selected from the group consisting of 802.11, Bluetooth, ultrawideband, HiperLAN and WDCT.

106. The system of claim 100 wherein the first device is a telephone.

107. The system of claim 100 wherein the first device transmits compressed musical signals and the second controller decompresses the compressed musical signals received by the receiver.

108. The system of claim 100 wherein the transmitter additionally transmits an annunciator signal indicating the presence of a transmitter of musical signals.

109. The system of claim 100 wherein the second device receiver becomes receptive of transmissions from the first device by obtaining an identifier from the first device.

110. The system of claim 109 wherein the first device additionally comprises an associated visual transducer and the second device additionally comprises a visual receiver, wherein the visual transducer outputs a visual identifier signal that is received by the visual receiver.

111. The system of claim 109 wherein the first device additionally comprises a first electrical connector and the second device additionally comprises a second electrical connector, wherein the identifier is communicated via an electrical signal passed from the first electrical connector to the second electrical connector.

112. The system of claim 111 wherein the communication of the identifier signal occurs via an inductive coupling between the first connector and the second connector.

113. The system of claim 109 wherein the second device additionally comprises a microphone, wherein the identifier is communicated in an audio signal that is produced by the first player and which is received by the microphone.

114. The system of claim 109 wherein the identifier is transmitted by the transmitter and received by the receiver.

115. The system of claim 100 wherein the second device receiver becomes receptive of transmissions from the first device by transmitting an identifier to the first device.

116. The system of claim 115 wherein the second device additionally comprises a visual transducer and the first device additionally comprises a visual receiver, wherein the visual transducer outputs a visual identifier signal that is received by the visual receiver.

117. The system of claim 115 wherein the first device additionally comprises a first electrical connector and the second device additionally comprises a second electrical connector, wherein the identifier is communicated via an electrical signal passed from the second electrical connector to the first electrical connector.

118. The system of claim 117 wherein the communication of the identifier signal occurs via an inductive coupling between the first connector and the second connector.

119. The system of claim 115 wherein the first device additionally comprises a microphone, wherein the identifier is communicated in an audio signal that is produced by the second player and which is received by the microphone.

120. The system of claim 115 wherein the identifier is transmitted by the transmitter and received by the receiver.

121. The system of claim 100 wherein the first controller additionally generates control signals related to the musical signals.

122. The system of claim 121 wherein the store additionally comprises stored control signals that are used by the first controller.

123. The system of claim 121 wherein the controller generates control signals from the musical signals.

124. The system of claim 123 wherein the controller performs frequency analysis of the musical signals in order to generate the control signals.

125. The system of claim 121 wherein the first device additionally transmits the control signals to the second device using the first transmitter.

126. The system of claim 125 wherein the control signals and the musical signals transmitted by the first device are interspersed within a single digital stream.

127. The system of claim 125 wherein the control signals and the musical signals are transmitted in two digital streams.

128. The system of claim 121 wherein the first device additionally comprises a second transmitter that transmits control signals.

129. The system of claim 128 wherein the second transmitter uses an IP protocol.

130. The system of claim 129 wherein the Internet protocol uses a link layer selected from the group consisting of 802.11, Bluetooth, and ultrawideband.

131. The system of claim 121 wherein the second transmitter uses analog radio frequency transmissions.

132. The system of claim 100, additionally comprising a wearable accessory, comprising a control signal receiver that is receptive of control signals transmitted by the first device, an accessory controller that generates transducer signals from the received control signals, and a transducer that generates humanly-perceptive music accompaniment signals from the transducer signals.

133. The system of claim 132 wherein the music accompaniment signals are visual signals.

134. The system of claim 133 wherein the visual signals are generated by a light emitting diode

135. The system of claim 132 wherein the music accompaniment signals are tactile signals.

136. The system of claim 132 wherein the accessory controller performs frequency analysis on the received control signals.

137. The system of claim 132 wherein the wearable accessory additionally comprises a switch, wherein actuation of the switch causes the accessory to be receptive to an entraining signal from the first device, such that the accessory will accept signals from the first device.

138. An earphone for listening to audio entertainment allowing for the controlled reception of ambient sound by a user, comprising

a speaker that is oriented towards the user's ear;

an enclosure that reduces the amount of ambient noise perceptive to the user; and

a manually-adjustable characteristic of the enclosure that adjusts the amount of ambient sound perceptive to the user.

139. The earphone of claim 138 wherein the manually-adjustable characteristic is the size of apertures in the enclosure that admit ambient sound.

140. The earphone of claim 138 wherein the manually-adjustable characteristic is the number of apertures in the enclosure that admit ambient sound.

141. The earphone of claim 138 wherein the manually-adjustable characteristic is the surface area of the enclosure.

142. The earphone of claim 138 wherein the manually-adjustable characteristic is the thickness of the enclosure.

143. The earphone of claim 138, additionally comprising a spacer that attaches to the earphone and which increases the distance between the speaker and the user's ear.

144. The earphone of claim 138, additionally comprising a secondary enclosure that encloses the enclosure and the speaker, and which is detachable from the earphone.

145. A mobile device for the transmission of audio entertainment signals comprising

an audio signal store for the storage of the audio entertainment signals;

an audio signal player for the playing of the audio entertainment signals;

an wireless transmitter for the transmission of the audio entertainment signals; and

a transmitter control to manually switch between two states consisting of the operation and the non-operation of the audio transmitter.

146. The device of claim 145 wherein the transmitter is connected to the audio jack of the audio signal player.

147. The device of claim 146 wherein the audio signal that is transmitted via the audio jack is subsequently divided between the transmitter and the headphone.

148. The device of claim 146 wherein the transmitter additionally comprises a second jack to which the earphone plug can connect, such that the insertion of the earphone plug into the second jack allows the audio signal to be heard in the earphone

149. A mobile device for the reception of digital audio entertainment signals comprising:

an audio signal store for the storage of the digital audio entertainment signals;

an audio receiver for the reception of external digital audio entertainment signals from a mobile audio signal transmitter located within a predetermined distance of the audio receiver;

a receiver control with at least a first state and a second state; and

an audio signal player that plays digital audio entertainment signals from the audio signal store when the receiver control is in the first state, and that plays digital audio entertainment signals from the audio receiver when the receiver control is in the second state.

150. The device of claim 149, additionally comprising a jack to which a listening device chosen from the group consisting of earphone and headphone can be connected.

151. A method for the shared enjoyment of music from stored musical signals between a first user with a first music player device and at least one second user with at least one second music player device, comprising:

playing the musical signals for the first user on the first music player device while essentially simultaneously wirelessly transmitting synchronization signals from the first music player device to the at least one second music player device; and

receiving the synchronization signals by the at least one second player device, wherein the synchronization signals allow the musical signals on the at least one second player device to be played essentially simultaneously with the playing of the musical signals on the first music player device;

wherein the first and the at least one second users are mobile.

152. The method of claim 151 wherein the same musical signals are stored on the first music player device and the at least one second music player device.

153. The method of claim 152 wherein the musical signals are stored as compressed audio files.

154. The method of claim 151 wherein the synchronization signals comprise the current sample number being played on the first music player device.

155. The method of claim 154 wherein the synchronization signals further comprise the sample rate of the musical signal being played on the first music player device.

156. The method of claim 151 wherein the synchronization signals comprise the current time relative to the beginning of the musical signal being played on the first music player device.

157. The method of claim 151, additionally comprising transferring the musical signals from the first music player device to the at least one second music player device.

158. A wireless communications system for sharing audio entertainment between a first mobile device and a second mobile device, comprising:

a broadcast identifier transmitted by the first mobile device to the second mobile device;

a personal identifier transmitted by the second mobile device to the first mobile device; and

a broadcast signal comprising audio entertainment that is transmitted by the first mobile device of which the second device is receptive;

wherein the first mobile device and the second mobile device have displays which can display the identifier signal that they receive and the second mobile device can play the audio entertainment from the broadcast signal that it receives.

159. The system of claim 158 wherein the identifier signals are transmitted via an Internet protocol.

160. The system of claim 159 wherein the Internet protocol is a connection-oriented transport protocol.

161. The system of claim 159 wherein the Internet protocol further comprises a link layer protocol selected from the group consisting of 802.11, Bluetooth, ultrawideband, HiperLAN and WDCT.

162. The system of claim 158 wherein the broadcast signal is transmitted via an Internet protocol.

163. The system of claim 162 wherein the Internet protocol is a connectionless transport protocol.

164. The system of claim 162 wherein the Internet protocol further comprises a link layer protocol selected from the group consisting of 802.11, Bluetooth, ultrawideband, HiperLAN and WDCT.

165. The system of claim 158 wherein the broadcast signal is an analog radio frequency signal.

166. The system of claim 158 wherein the broadcast signal is transmitted via telephony.

167. The system of claim 158 wherein the broadcast signal is compressed.

168. The system of claim 158 wherein the broadcast signal is encrypted.

169. The system of claim 158, additionally comprising a transfer signal that is transmitted by the first mobile device upon halting the transmission of its broadcast signal, and which transfer signal is receptive by the second mobile device, such that reception by the second mobile device causes the device to transmit a second broadcast signal comprising audio entertainment for which the first unit is receptive.

170. The system of claim 158 wherein the broadcast signal is transmitted via peer-to-peer communications.

171. The system of claim 170 additionally comprising a switch, wherein actuation of the switch causes the device to toggle between peer-to-peer communications and infrastructure-based communications.

172. The system of claim 158 wherein the broadcast signal is transmitted via infrastructure-based communications.

173. The system of claim 158 wherein the personal identifier includes a least one piece of information about the user of the second mobile device taken from the group consisting of the name of the user, the nickname of the user, a numerical identifier, an alphanumerical identifier, the musical preferences of the user, the age of the user, and the educational institution that the user attends.

174. The system of claim 158 additionally comprising a third mobile device whose user requests to receive the broadcast signal and additionally comprising a requester identifier signal transmitted by the third mobile device to the first mobile device.

175. The system of claim 174 wherein the requester identifier signal is additionally transmitted by the third mobile device to the second mobile device.

176. The system of claim 174 wherein the requester identifier signal is additionally transmitted by the first mobile device to the second mobile device.

177. The system of claim 174 wherein the requester identifier is displayed on the display of the first mobile device.

178. The system of claim 174 wherein the third mobile device receives the broadcast signal without the affirmative action of the first mobile device.

179. The system of claim 174 wherein the third mobile device must receive the affirmative action of the first mobile device in order to receive the broadcast signal.

180. The system of claim 174 wherein the third mobile device must receive the affirmative action of the first mobile device and all other mobile devices receptive of the broadcast signal.

181. The system of claim 180 wherein the affirmative action is a unanimous vote.

182. The system of claim 180 wherein the affirmative action is a majority vote.

183. The system of claim 158, additionally comprising a substantially unique musical identifier related to the audio entertainment, wherein the musical identifier is transmitted by the first mobile device to the second mobile device in association with the broadcast signal, and the musical identifier is stored on the second mobile device in association with the broadcast identifier.

184. A system for identifying a first device that introduces a music selection to a second device, comprising:

a mobile music transmitter operated by the first device;

a mobile music receiver operated by the second device;

a music signal comprising the music selection transmitted by the transmitter and received by the receiver;

an individual musical identifier that is associated with the music selection; and

an individual transmitter identifier that identifies the transmitter, wherein the transmitter identifier and the individual music identifier are stored in association with each other in the receiver.

185. A method for enhancing enjoyment of a musical selection, comprising:

obtaining control signals related to the musical selection;

transmitting the control signals wirelessly;

receiving the control signals; and

converting the control signals to a humanly-perceptible form.

186. The method of claim 185 wherein the distance over which the control signals are transmitted has a minimum range of at least 1 foot.

187. The method of claim 185 wherein the distance over which the control signals are transmitted has a maximum range of less than 500 feet.

188. The method of claim 185 wherein the humanly-perceptible form is a visually-perceptible form.

189. The method of claim 185 wherein the humanly-perceptible form is a tactilely-perceptible form.

190. The method of claim 185 wherein the transmitting and receiving use radio frequency communications.

191. The method of claim 185 wherein the radio frequency communications use an Internet protocol.

192. The method of claim 185 wherein the Internet protocol involves a link layer chosen from the group consisting of 802.11, Bluetooth, ultrawideband, HiperLAN and WDCT.

193. The method of claim 185 wherein obtaining additionally comprises extracting the musical beat from the musical selection.

194. The method of claim 185 wherein obtaining additionally comprises retrieving the control signals from a store of control signals.

195. The method of claim 185 wherein the control signals are divided into multiple channels.

196. The method of claim 195 wherein the step of converting comprises a step of demultiplexing the control signals.

197. The method of claim 185 wherein the steps of transmitting and receiving are performed using means conveyed by a single person.

198. The method of claim 197 wherein the user is mobile.

199. The method of claim 185 wherein the step of transmitting is performed by a stationary user and the step of receiving is performed by a mobile user.

200. The method of claim 185 wherein the musical selection is provided by musicians in concert playing essentially simultaneously with obtaining control signals.

201. The method of claim 185 wherein the control signals are transmitted with a substantially unique identifier.

202. The method of claim 185 wherein the control signals control the appearance of a wearable device.

203. A method for generating and storing control signals corresponding to musical signals, comprising:

playing musical signals for a user;

receiving manual input signals from the user that are produced substantially in synchrony with the music;

generating control signals from the input signals; and

storing the control signals so that they can be retrieved with the musical signals.

204. The method of claim 203 wherein receiving comprises transducing tactile input from the user.

205. The method of claim 204 wherein the tactile input is produced using a user appendage selected from the group consisting of hand, finger and foot.

206. The method of claim 203 wherein receiving comprises transducing sound output produced by the user.

207. The method of claim 203 wherein storing comprises interleaving control signals within the musical signals.

208. The method of claim 203 wherein storing additionally comprises creating a digital file.

209. The method of claim 203 wherein storing additionally comprises associating the digital file with the musical signals.

210. A wearable personal accessory, comprising:

an input transducer taken from the group consisting of a microphone and an accelerometer, said transducer generating a time-varying input transduction signal;

a controller which accepts the input transduction signal, and which generates an output transducer signal whose signal varies in amplitude with time;

an output transducer receptive of the output transducer signal, and which provides a humanly-perceptible signal; and

an energy source for powering the input transducer, controller and output transducer.

211. The accessory of claim 210 wherein the accelerometer responds to the movements of the user in one axis of movement.

212. The accessory of claim 210 wherein the accelerometer responds to the movements of the user in more than one axis of movement.

213. The accessory of claim 210 wherein the controller performs automatic gain control on the signal from the input transducer.

214. The accessory of claim 210 wherein the output transducer is a light transducer.

215. The accessory of claim 214 wherein the output transducer comprises a light emitting diode.

216. The accessory of claim 214 wherein the output transducer comprises a liquid crystal display.

217. The accessory of claim 214 wherein the output transducer comprises an organic light emitting diode.

218. The accessory of claim 210 wherein the output transducer is a tactile transducer.

219. The accessory of claim 210 wherein the output transducer comprises a first and a second transduction element.

220. The accessory of claim 219 wherein the first transduction element responds preferentially to a first band of frequencies of the input transduction signal and the second transduction element responds preferentially to a second band of frequencies of the input transduction signal.

221. The accessory of claim 219 wherein the first transduction element responds preferentially to a first band of amplitudes of the input transduction signal and the second transduction element responds preferentially to a second band of amplitudes of the input transduction signal.

222. A wearable personal accessory controlled via wireless communications, comprising:

a wireless communications receiver, receptive of an external control signal;

a controller that accepts the external control signal and that generates a time-varying visual output transducer signal;

a visual output transducer that is receptive of the output transducer signal, and which provides a humanly-perceptible visual signal; and

an energy store for powering the receiver, controller and output transducer;

wherein the visual output transducer generates visually-perceptive output.

223. The accessory of claim 222 wherein the external control signal is a time-varying amplitude signal.

224. The accessory of claim 223 wherein the amplitude signal is analog.

225. The accessory of claim 223 wherein the amplitude signal is digital.

226. The accessory of claim 223 wherein the amplitude signal is substantially a musical signal.

227. The accessory of claim 226 wherein the controller performs frequency filtering of the external control signal.

228. The accessory of claim 226 wherein the controller determines the primary beat of the music.

229. The accessory of claim 228 wherein the primary beat of the music is determined according to the positions in the music where the amplitude rises at least a predetermined factor over less than a predetermined time.

230. The accessory of claim 222 wherein the external control signal is a digital signal of pairs comprising positions within the music and the amplitudes to which the output transducers are set at such times.

231. The accessory of claim 230 wherein the positions are chosen from the group of characteristics consisting of samples and times.

232. The accessory of claim 222 wherein the receiver is an analog radio frequency receiver.

233. The accessory of claim 232 wherein the analog radio frequency communications are at a frequency set by a crystal oscillator, wherein the crystal oscillator can be manually exchanged in order to change the frequency of operation.

234. The accessory of claim 222 wherein the receiver receives Internet protocol signals.

235. The accessory of claim 234 wherein the Internet protocol further comprises a link layer selected from the group consisting of Bluetooth, 802.11 and ultrawideband.

236. The accessory of claim 222 wherein the visual output transducer comprises a first light transducer and a second light transducer.

237. The accessory of claim 236 wherein the external control signal comprises a multiplexed first signal and second signal, and wherein the first signal relates to the first light transducer and the second signal relates to a second light transducer.

238. The accessory of claim 237 wherein the controller comprises at least two ports, wherein the first port relates to the first light generating element and the second port related to the second light generating element.

239. The accessory of claim 236 wherein the first light transducer responds preferentially to a first range of frequency in the musical signal, and the second light transducer responds preferentially to a second range of frequency in the musical signal.

240. The accessory of claim 236 wherein the first light transducer responds preferentially to a first range of amplitude in the musical signal, and the second light transducer responds preferentially to a second range of amplitude in the musical signal.

241. The accessory of claim 222 wherein the visual output transducer comprises a light emitting diode.

242. The accessory of claim 222 wherein the visual output transducer comprises an organic light emitting diode.

243. The accessory of claim 222 wherein the visual output transducer comprises a liquid crystal display.

244. The accessory of claim 222 additionally comprising a tactile transducer responsive to the external control signal.

245. The accessory of claim 222 additionally comprising a sound transducer responsive to the external control signal.

246. The accessory of claim 222 wherein the external control signal additional comprises a digital identifier.

247. The accessory of claim 246 wherein the visual output transducer additionally provides an identifier signal that corresponds to the digital identifier, wherein such identifier signal is higher frequency than is humanly perceptible, and wherein the identifier signal is output substantially simultaneously with the humanly-perceptible visual signal.

248. The accessory of claim 222 additionally comprising a manually-separable electrical connection, wherein the connection connects the receiver to the controller.

249. The accessory of claim 222 additionally comprising a manually-separable electrical connection, wherein the connection connects the controller to the output transducer.

250. The accessory of claim 222 wherein the accessory functions as an accoutrement selected from the group consisting of a ring, a pendant, a hair clip, an earring, a necklace, a brooch, a belt buckle, watches, and a bracelet.

251. The accessory of claim 222 wherein the accessory is incorporated into an article of clothing selected from the group consisting of a backpack, a wallet, a purse, a hat and a shoe.

252. The accessory of claim 222 additionally comprising a manual control whose operation sets either a first state and a second state and wherein in the first state the accessory responds exclusively to external control signals comprising a substantially unique identifier.

253. The accessory of claim 252 wherein the identifier to which the accessory responds in the first state is received by the receiver when the accessory is set to the second state.

254. The accessory of claim 253 additionally comprising a third state which is manually set, wherein in the third state, the accessory responds to external control signals in the absence of a substantially unique identifier.

255. A device for converting user tactile responses to stored music into a stored control signal, comprising:

a player that plays stored music audible to the user;

a manually-operated transducer that outputs an electrical signal, wherein the transducer is actuated by the user in response to the music;

a controller that receives the electrical signal and outputs a control signal; and

a store that receives the control signal and stores it.

256. The device of claim 255 wherein the control signal comprises a list of positions, wherein the positions are chosen from the set of characteristics consisting of samples and time relative to the beginning of the stored music.

257. The device of claim 255 wherein the control signal further comprises a list of amplitudes corresponding to the positions.

258. The device of claim 257 additionally comprising a second manually-operated transducer that outputs a second electrical signal, wherein the controller additionally receives the second electrical signal.

259. The device of claim 255 wherein the transducer comprises a musical drum operated by the user whose audible output is received by a microphone that generates the electrical signal received by the controller.

260. The device of claim 255 wherein the controller performs frequency filtering on the electrical signal.

261. The device of claim 255 wherein the transducer is a tactile transducer that is operated by an appendage of the user selected from the group consisting of finger, hand, and foot.

262. The device of claim 255 wherein the store is a file system and the control signal is stored as control file and the stored music is stored as a music file, wherein the control file and the music file are stored in association with each other.

263. The device of claim 255 wherein the control signals are stored in an interleaved fashion with the stored music.

264. The device of claim 71 additionally comprising a visual display on which a visual representation of the control signals is displayed.

265. A music player that wirelessly transmits control signals related to the music, wherein the control signals control a wearable electronic accessory, comprising:

a store of music signal files;

a controller that reads a musical signal file from the store and generates audio signals and that further generates the control signals;

a transducer that converts the audio signals into sound audible to the user; and

a wireless transmitter that transmits the control signal to the wearable electronic accessory.

266. The player of claim 265 wherein the music signal files comprise compressed audio files.

267. The player of claim 265 wherein the controller performs frequency filtering of the external control signal.

268. The player of claim 265 wherein the controller determines the primary beat of the music.

269. The player of claim 268 wherein the primary beat of the music is determined according to the positions in the music where the amplitude rises at least a predetermined factor over less than a predetermined time.

270. The player of claim 265 wherein the control signals are pairs comprising positions within the music and the amplitudes to which the output transducers are set at such times.

271. The player of claim 265 wherein the positions are chosen from the set of characteristics consisting of samples and times.

272. The player of claim 265 wherein the transmitter is an analog radio frequency transmitter.

273. The player of claim 265 wherein the transmitter transmits according to an Internet protocol.

274. The player of claim 273 wherein the Internet protocol has a link layer selected from the group consisting of 802.11, Bluetooth, ultrawideband, HiperLAN and WDCT.

275. The player of claim 265 additionally comprising a visual transducer that generates visually-preceptive output in response to the control signal, wherein the visual transducer receives the control signal from the controller.

276. The player of claim 275 wherein the visual transducer is a light emitting diode.

277. The player of claim 275 wherein the visual transducer is a liquid crystal display.

278. The player of claim 265 additionally comprising a store of configuration files wherein the configuration files include parameters read by the controller that affect the generation of control signals from the musical signals.

279. A music player that wirelessly transmits control signals related to the music, wherein the control signals control a wearable electronic accessory, comprising:

a store of music signal files;

a second store of control signal files associated with the music signal files;

a controller that reads a musical signal file from the store and generates audio signals and that further reads an associated control signal file;

a transducer that converts the audio signals into sound audible to the user, and

a wireless transmitter that transmits the control signals from the associated control signal file to the wearable electronic accessory.

280. The player of claim 279 wherein the music signal files comprise compressed audio files.

281. The player of claim 279 wherein the transmitter is an analog radio frequency transmitter.

282. The player of claim 279 wherein the transmitter transmits according to an Internet protocol.

283. The player of claim 282 wherein the Internet protocol has a link layer selected from the group consisting of 802.11, Bluetooth, ultrawideband, HiperLAN and WDCT.

284. The player of claim 279 wherein the store of control information comprises at least one binary data file that relates to a single music signal file.

285. The player of claim 279 wherein the store of control information is interleaved with a related music signal file in binary data format.

286. The player of claim 279 additionally comprising a visual transducer that generates visually-preceptive output in response to the electrical control signal, wherein the visual transducer receives the electrical control signal from the controller.

287. A system for exhibition of music enjoyment, comprising:

a source of music signals;

a controller that generates control signals from the music signals;

a transmitter of the control signals, wherein the transmission of the control signals is synchronized with the playing of the music signals;

a receiver of the control signals; and

a transducer that responds to the control signals.

288. The system of claim 287 wherein the source of music signals are stored binary music files.

289. The system of claim 288 wherein the stored binary data files are compressed audio files.

290. The system of claim 288 wherein the music signals are stored on a medium selected from the group consisting of compact disk, digital video disk, hard disk drive, removable disks, flash memory, magneto-optical disk, and audio tape.

291. The system of claim 287 wherein the source of music signals are musicians playing live music.

292. The system of claim 287 wherein the receiver and transducer are worn by a user.

293. The system of claim 287 wherein the receiver and the transducer are physically separated from the transmitter by more than 5 feet.

294. The system of claim 287 wherein the transducer is a visual transducer that generates humanly-perceptible visual signals.

295. The system of claim 294 wherein the visual transducer is a light emitting diode

296. The system of claim 294 wherein the visual transducer is a liquid crystal display

297. The system of claim 287 wherein the transducer is a tactile transducer.

298. The system of claim 287 wherein the minimum range of the transmitter for which the receiver is receptive is at least 1 foot.

299. The system of claim 287 wherein the maximum range of the transmitter for which the receiver is receptive is less than 500 feet.

300. The system of claim 287 wherein the controller determines the primary beat of the music.

301. The system of claim 287 wherein the controller performs frequency filtering.

302. The system of claim 287, additionally comprising a second controller that is electrically connected to the receiver, and which generates transducer control signals from the control signals received by the receiver, wherein the transducer control signals directly control the transducer.

303. The system of claim 302 wherein the control signals are time-varying amplitude signals and wherein the second controller performs frequency filtering.

304. The system of claim 302 additionally comprising a switch that is electrically connected to the second controller such that actuation of the switch places the second controller into entraining mode, wherein the receiver becomes receptive of entraining signals from the transmitter.

305. The system of claim 304 wherein a receiver that is entrained will receive control signals other than from the transmitter in the absence of signals from the transmitter.

306. The system of claim 304 wherein additional actuation of the switch will allow reception of control signals by the receiver from a transmitter to which it is not entrained.

307. The system of claim 287 wherein the transmitter and receiver communicate via an Internet protocol.

308. The system of claim 287 wherein the transmitter and receiver communicate via digital radio frequency communications.

309. The system of claim 287 wherein the transmitter and receiver communicate via analog radio frequency communications.

310. The system of claim 309 wherein the analog radio frequency communications are at a frequency set by a crystal oscillator and wherein the crystal oscillator can be manually replaced.

311. The system of claim 287 additionally comprising a transferror for the transmission of music signals.

312. The system of claim 311 wherein the transmitter and the transferror are the same.

313. The system of claim 311 wherein the transmitter and the transferror are different, and communicate via different communications protocols.

314. The system of claim 287 additionally comprising a second receiver for the reception of music signals, wherein the music signals that are received are played by the music player.

315. A method for transferring a wearable-accessory control file stored on a first device to a second device in which an associated music file is stored, comprising:

storing on the first device the name of the music file in conjunction with the control file with which it is associated;

requesting by the second device of the first device for a control file stored in conjunction with the name of the music file; and

transferring the control file from the first device to the second device, wherein the control file is stored on the second device in conjunction with the name of the associated music file.

316. The method of claim 315 additionally comprising the step of extracting, wherein the control file is stored on the first device interspersed with the music file as a combined file, and the control file is extracted from the music file prior to being transferred to the second device.

317. The method of claim 315 wherein the requesting is performed over a wide area network.

318. The method of claim 315 wherein the name of the control file comprises the root of the name of the associated music file appended with a predetermined extension that is different from that of the music file.

319. The method of claim 315 wherein the control file and the associated music file both comprise the same substantially unique identifier.

320. A device for transmitting control signals to a wearable accessory receptive of such control signals, comprising:

a manually-separable input connector for connecting to an output port of an audio player, wherein audio signals are conveyed from the audio player to the device across the connector;

a controller for generating control signals from the audio signals; and

a transmitter for transmitting the control signals.

321. The device of claim 320 additionally comprising a second connector, wherein a portion of the audio signal conveyed to the device is conveyed to the second connector, and wherein an earphone can connect to the second connector.

322. The player of claim 320 wherein the controller performs frequency filtering of the audio signal.

323. The player of claim 320 wherein the controller determines the primary beat of the audio signal.

324. The player of claim 320 additionally comprising a visual transducer that generates visually-preceptive output in response to the control signals, wherein the visual transducer receives the control signals from the controller.

325. The system of claim 320 wherein the transmitter transmits via an Internet protocol.

326. The system of claim 320 wherein the transmitter transmits via digital radio frequency communications.

327. The system of claim 320 wherein the transmitter transmits via analog radio frequency communications.

328. The system of claim 327 wherein the analog radio frequency communications are at a frequency set by a crystal oscillator and wherein the crystal oscillator can be manually replaced.

329. The system of claim 320 wherein the transmitter additionally transmits the audio signal to wireless earphones.

1/34

Fig. 1

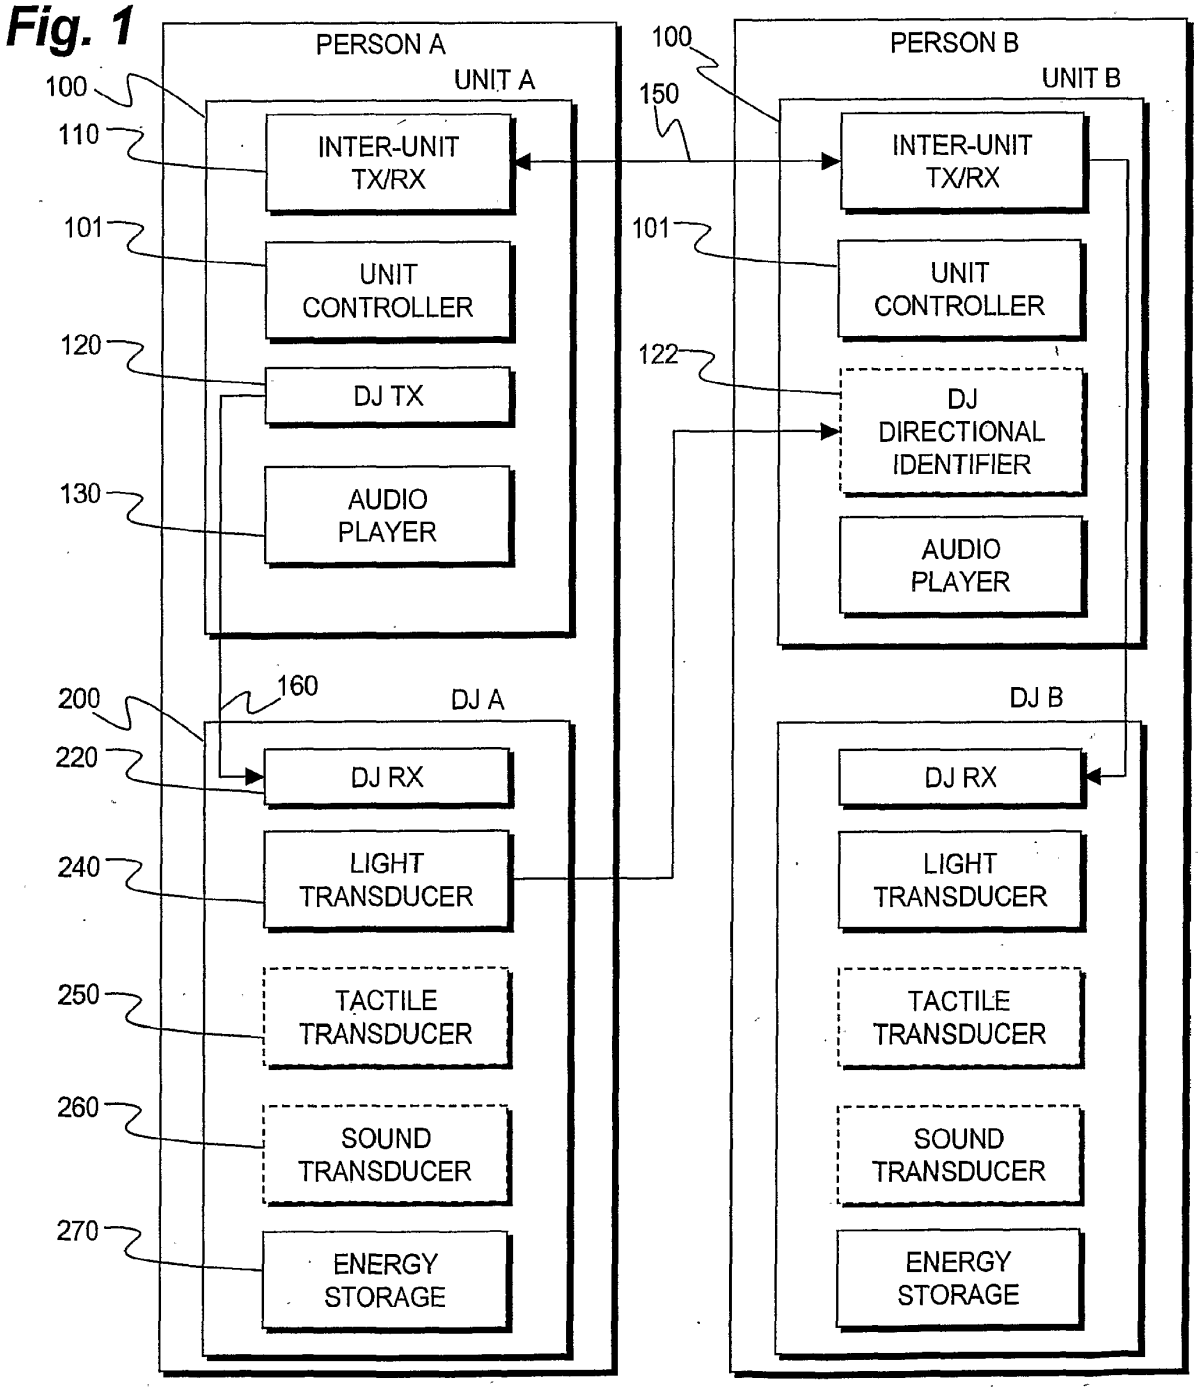

PERSON A
UNIT A
100
110
INTER-UNIT TX/RX
101
UNIT CONTROLLER
120
DJ TX
130
AUDIO PLAYER
160
200
DJ A
220
DJ RX
240
LIGHT TRANSDUCER
250
TACTILE TRANSDUCER
260
SOUND TRANSDUCER
270
ENERGY STORAGE
PERSON B
UNIT B
100
150
INTER-UNIT TX/RX
101
UNIT CONTROLLER
122
DJ DIRECTIONAL IDENTIFIER
AUDIO PLAYER
DJ B
DJ RX
LIGHT TRANSDUCER
TACTILE TRANSDUCER
SOUND TRANSDUCER
ENERGY STORAGE

Fig. 2A

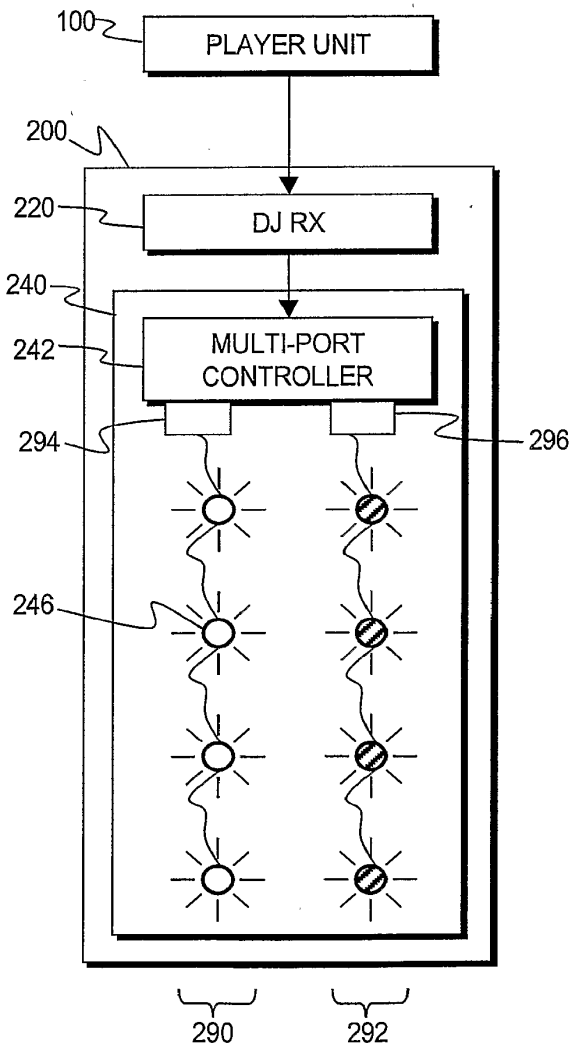

100
PLAYER UNIT
200
220
DJ RX
240
242
MULTI-PORT CONTROLLER
294
296
246
290
292

Fig. 2B

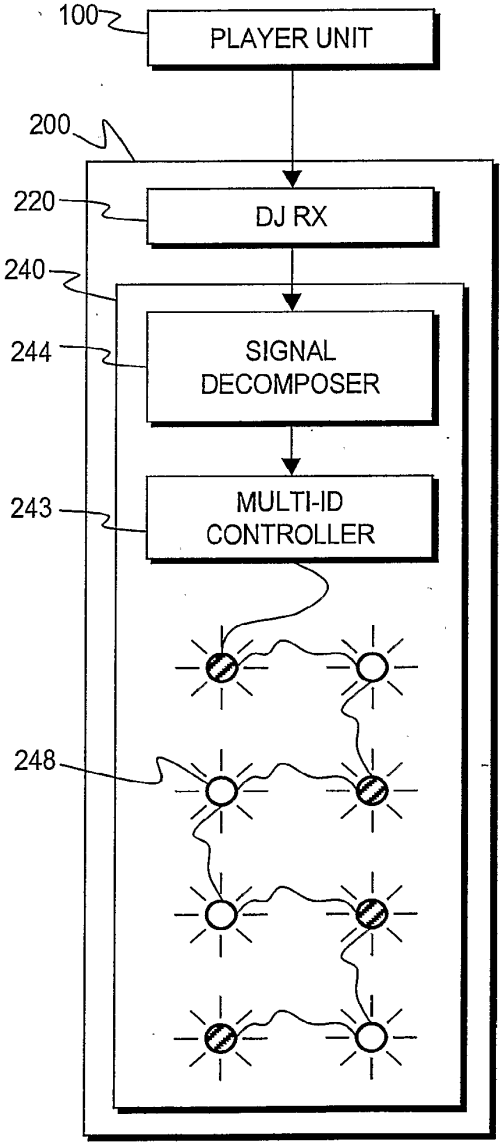

100
PLAYER UNIT
200
220
DJ RX
240
244
SIGNAL DECOMPOSER
243
MULTI-ID CONTROLLER
248

Fig. 3A

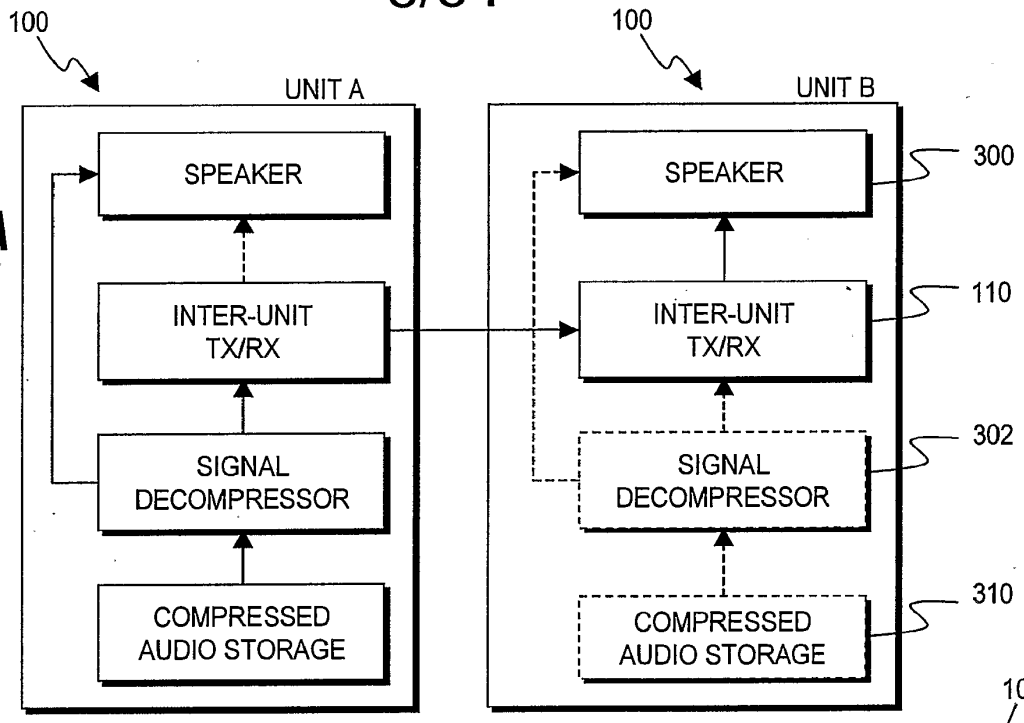

100
UNIT A
SPEAKER
INTER-UNIT TX/RX
SIGNAL DECOMPRESSOR
COMPRESSED AUDIO STORAGE
100
UNIT B
SPEAKER
300
INTER-UNIT TX/RX
110
SIGNAL DECOMPRESSOR
302
COMPRESSED AUDIO STORAGE
310

Fig. 3B

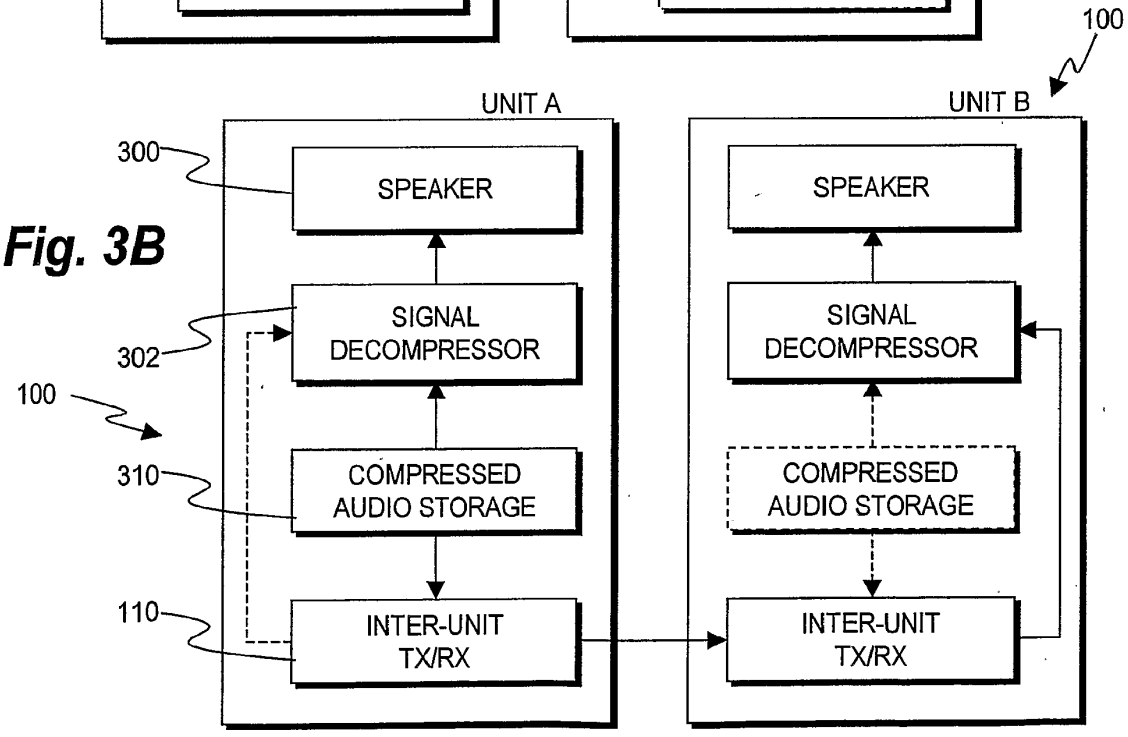

100
UNIT A
UNIT B
300
SPEAKER
SPEAKER
302
SIGNAL DECOMPRESSOR
SIGNAL DECOMPRESSOR
100
310
COMPRESSED AUDIO STORAGE
COMPRESSED AUDIO STORAGE
110
INTER-UNIT TX/RX
INTER-UNIT TX/RX

Fig. 3C

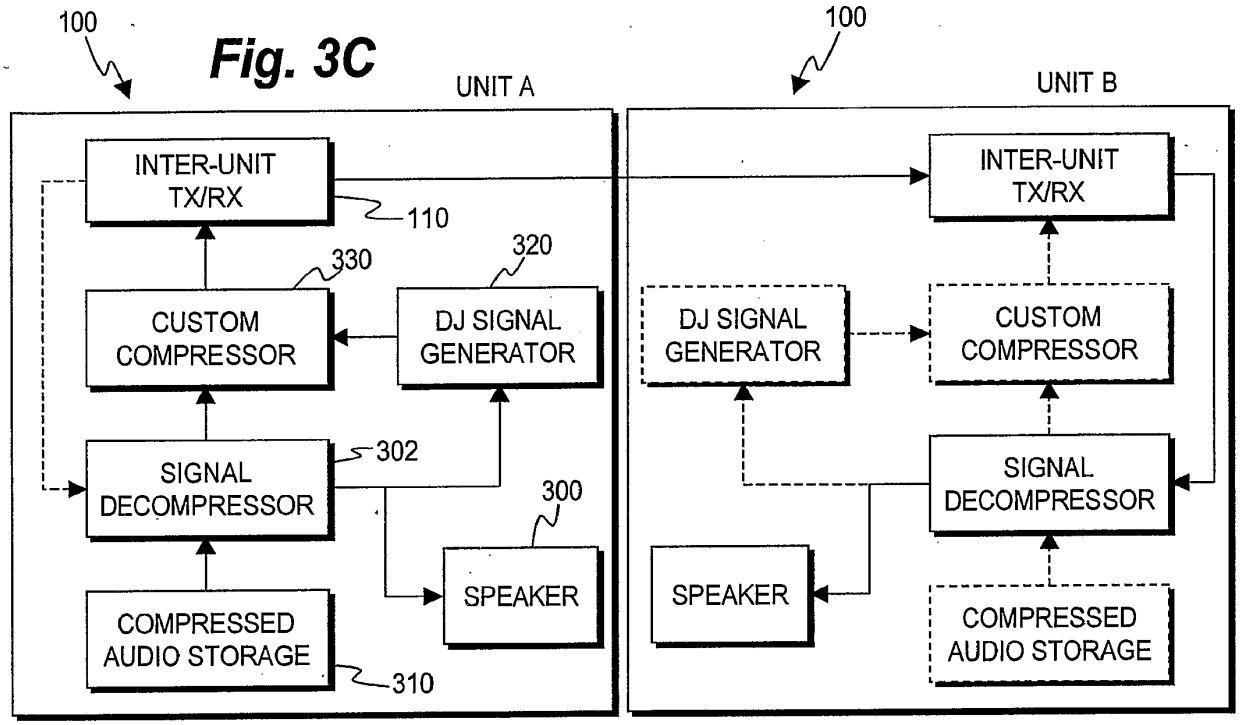

100
UNIT A
INTER-UNIT TX/RX
110
330
320
CUSTOM COMPRESSOR
DJ SIGNAL GENERATOR
SIGNAL DECOMPRESSOR
302
300
SPEAKER
COMPRESSED AUDIO STORAGE
310
100
UNIT B
INTER-UNIT TX/RX
DJ SIGNAL GENERATOR
CUSTOM COMPRESSOR
SIGNAL DECOMPRESSOR
SPEAKER
COMPRESSED AUDIO STORAGE

Fig. 4

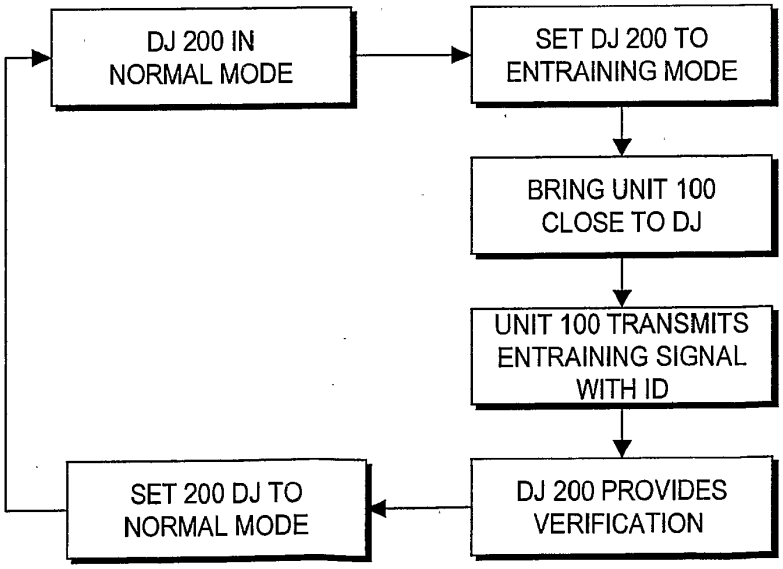

DJ 200 IN NORMAL MODE
SET DJ 200 TO ENTRAINING MODE
BRING UNIT 100 CLOSE TO DJ
UNIT 100 TRANSMITS ENTRAINING SIGNAL WITH ID
DJ 200 PROVIDES VERIFICATION
SET 200 DJ TO NORMAL MODE

Fig. 5A

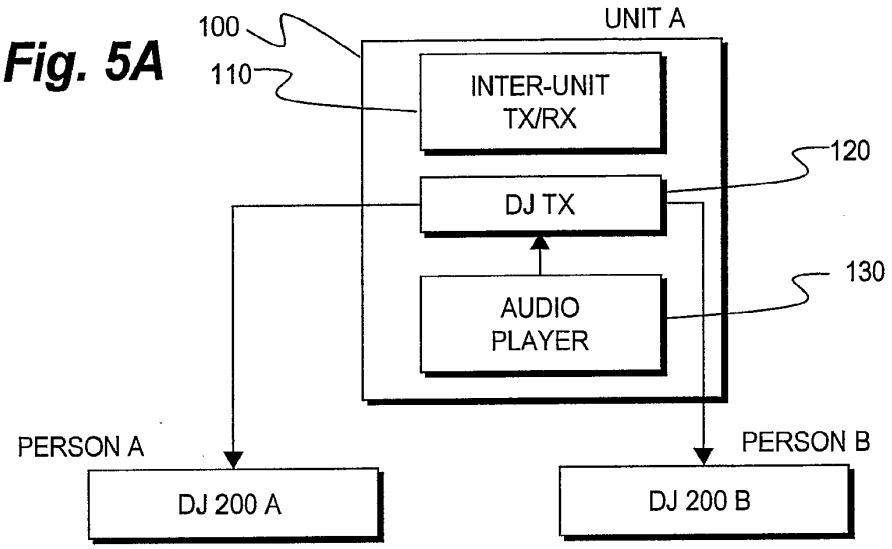

100
UNIT A
110
INTER-UNIT TX/RX
120
DJ TX
130
AUDIO PLAYER
PERSON A
PERSON B
DJ 200 A
DJ 200 B

Fig. 5B

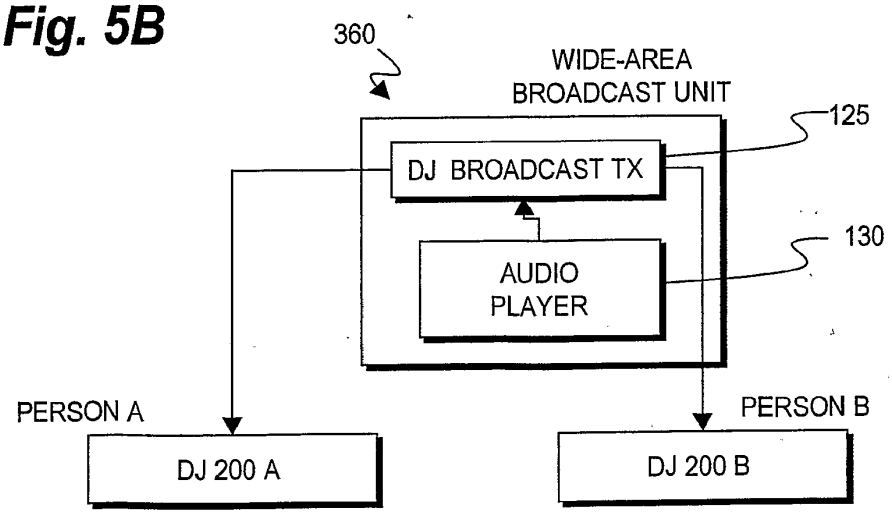

360
WIDE-AREA BROADCAST UNIT
125
DJ BROADCAST TX
130
AUDIO PLAYER
PERSON A
PERSON B
DJ 200 A
DJ 200 B

Fig. 6

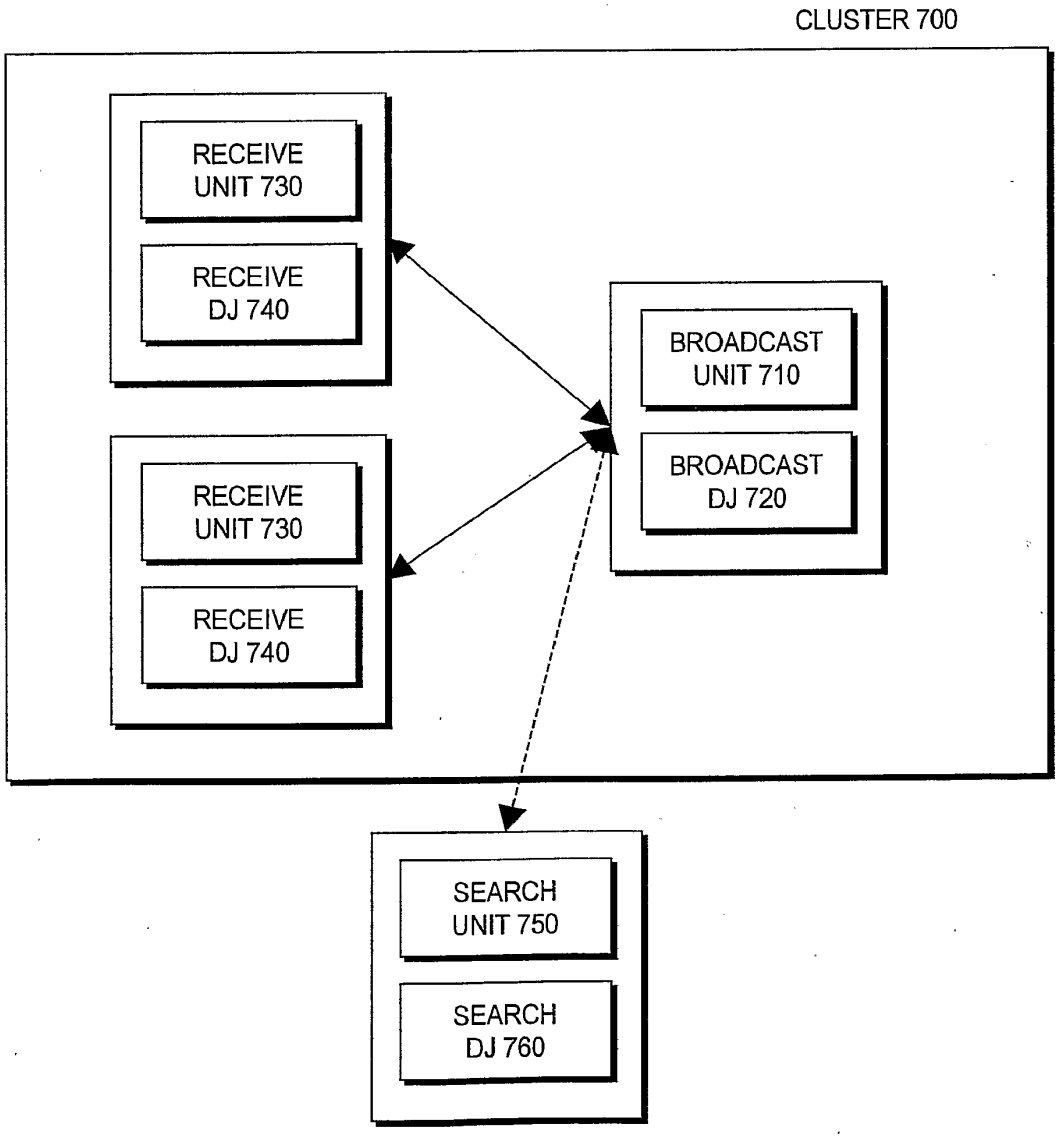

CLUSTER 700
RECEIVE UNIT 730
RECEIVE DJ 740
RECEIVE UNIT 730
RECEIVE DJ 740
BROADCAST UNIT 710
BROADCAST DJ 720
SEARCH UNIT 750
SEARCH DJ 760

Fig. 7

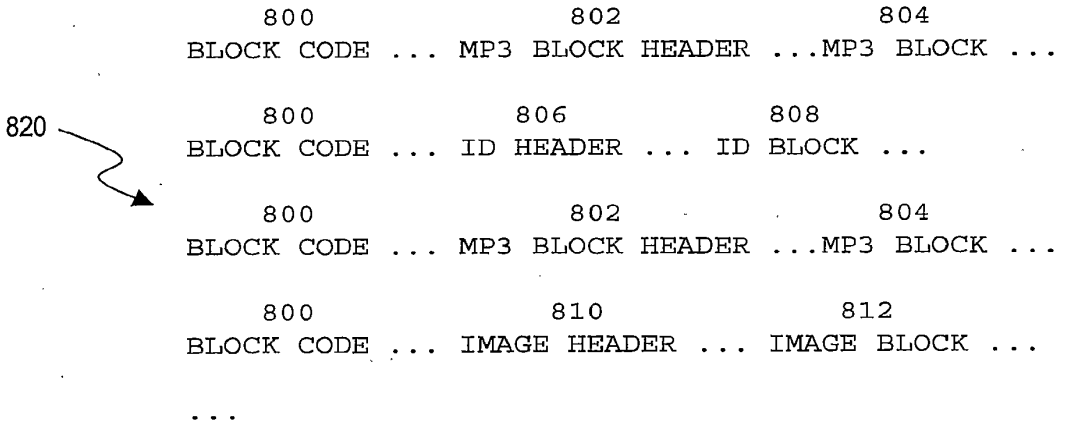

800
802
804
BLOCK CODE ... MP3 BLOCK HEADER ...MP3 BLOCK ...
820
800
806
808
BLOCK CODE ... ID HEADER ... ID BLOCK ...
800
802
804
BLOCK CODE ... MP3 BLOCK HEADER ...MP3 BLOCK ...
800
810
812
BLOCK CODE ... IMAGE HEADER ... IMAGE BLOCK ...
...

Fig. 36

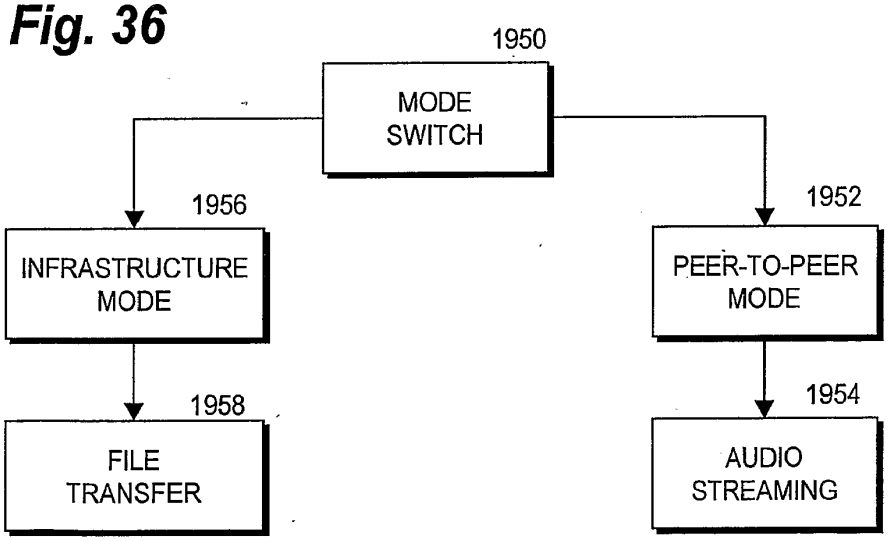

1950
MODE SWITCH
1956
INFRASTRUCTURE MODE
1952
PEER-TO-PEER MODE
1958
FILE TRANSFER
1954
AUDIO STREAMING

Fig. 8A

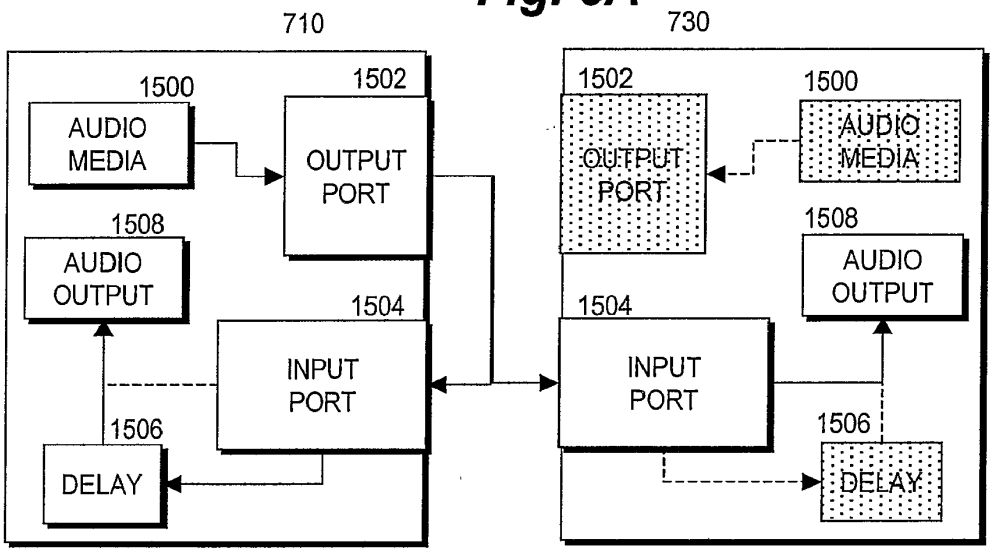

710
730
1500
AUDIO MEDIA
1502
OUTPUT PORT
1508
AUDIO OUTPUT
1504
INPUT PORT
1506
DELAY
1502
OUTPUT PORT
1500
AUDIO MEDIA
1508
AUDIO OUTPUT
1504
INPUT PORT
1506
DELAY

Fig. 8B

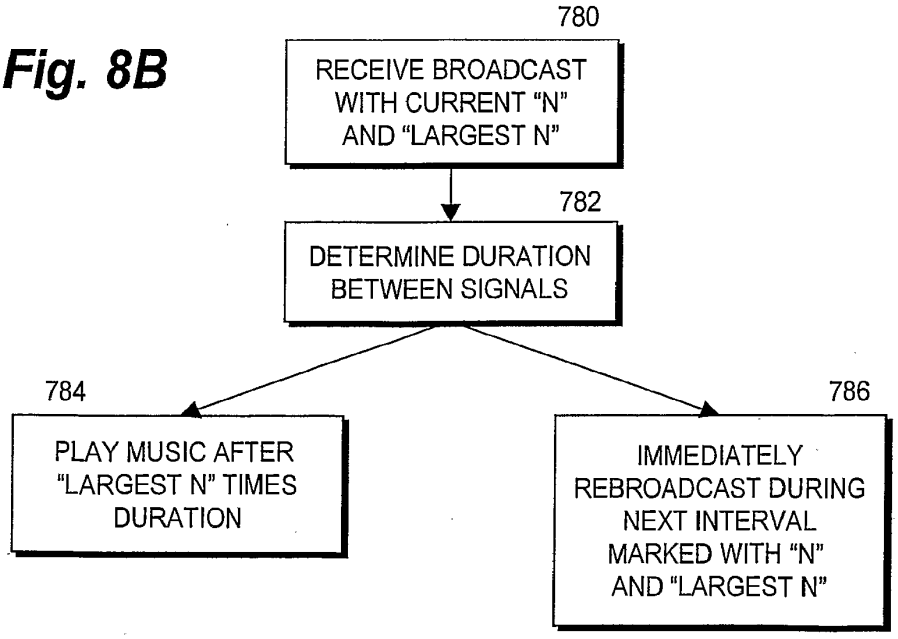

780
RECEIVE BROADCAST WITH CURRENT "N" AND "LARGEST N"
782
DETERMINE DURATION BETWEEN SIGNALS
784
PLAY MUSIC AFTER "LARGEST N" TIMES DURATION
786
IMMEDIATELY REBROADCAST DURING NEXT INTERVAL MARKED WITH "N" AND "LARGEST N"

Fig. 9A

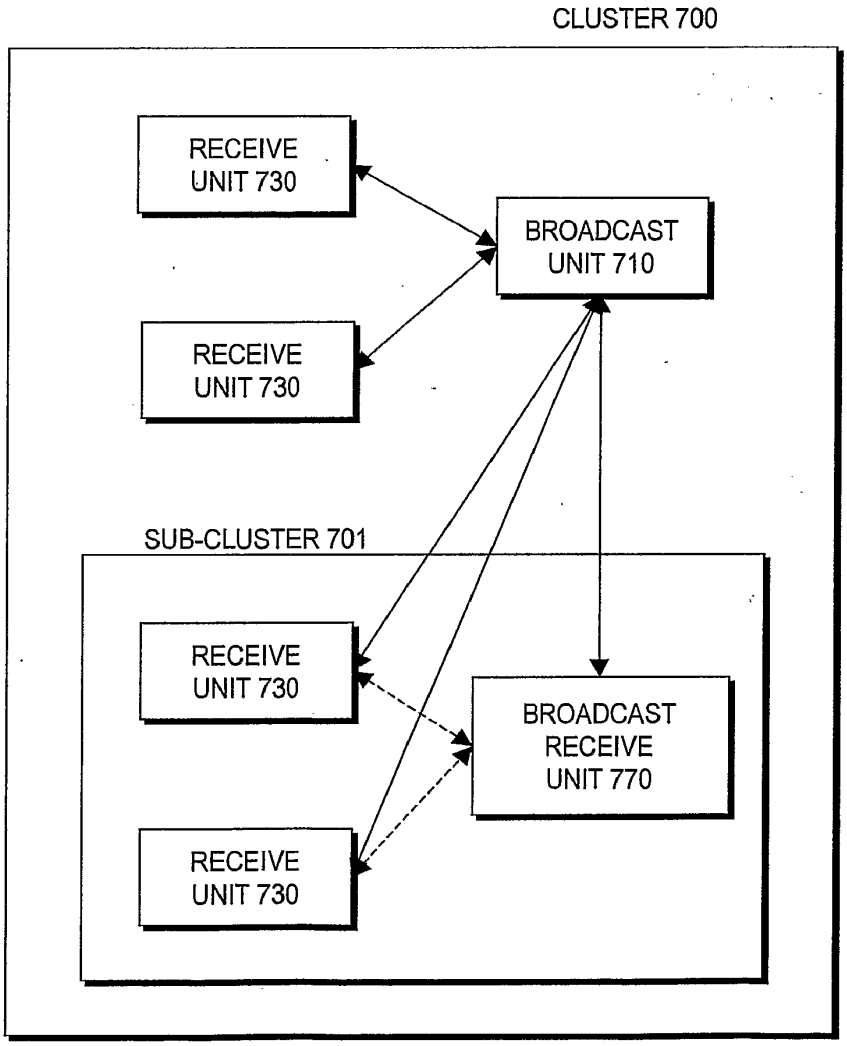

CLUSTER 700
RECEIVE UNIT 730
BROADCAST UNIT 710
RECEIVE UNIT 730
SUB-CLUSTER 701
RECEIVE UNIT 730
BROADCAST RECEIVE UNIT 770
RECEIVE UNIT 730

Fig. 9B

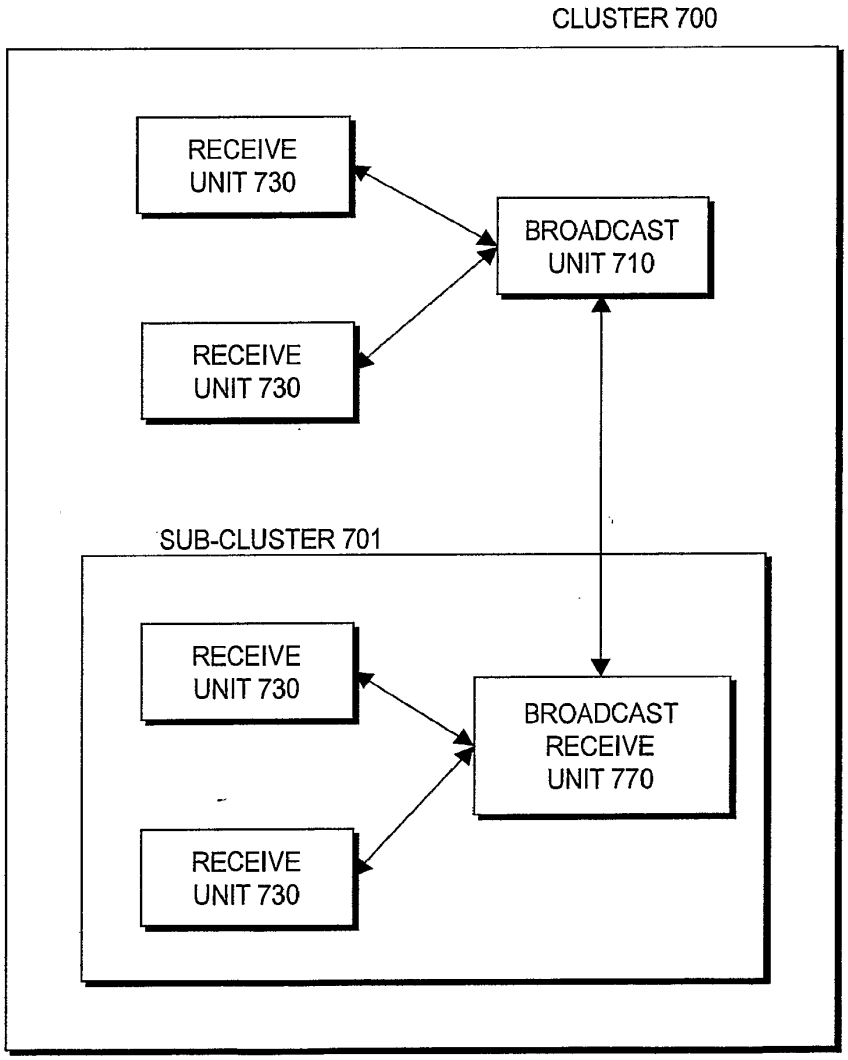

CLUSTER 700
RECEIVE UNIT 730
BROADCAST UNIT 710
RECEIVE UNIT 730
SUB-CLUSTER 701
RECEIVE UNIT 730
BROADCAST RECEIVE UNIT 770
RECEIVE UNIT 730

Fig. 10

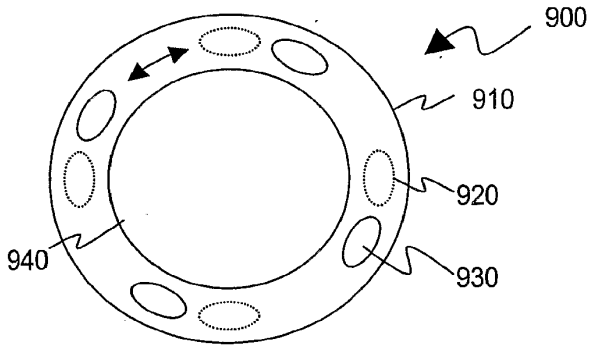


Fig. 11A

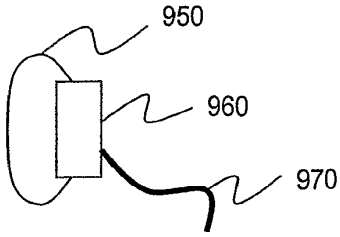


Fig. 11B

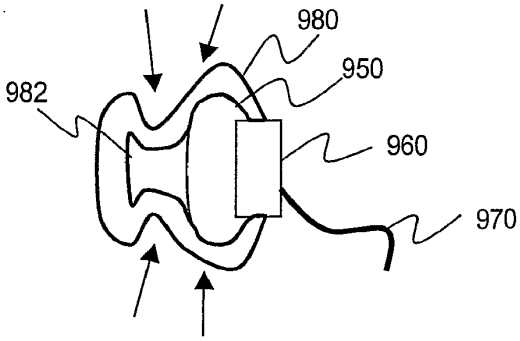


Fig. 12A

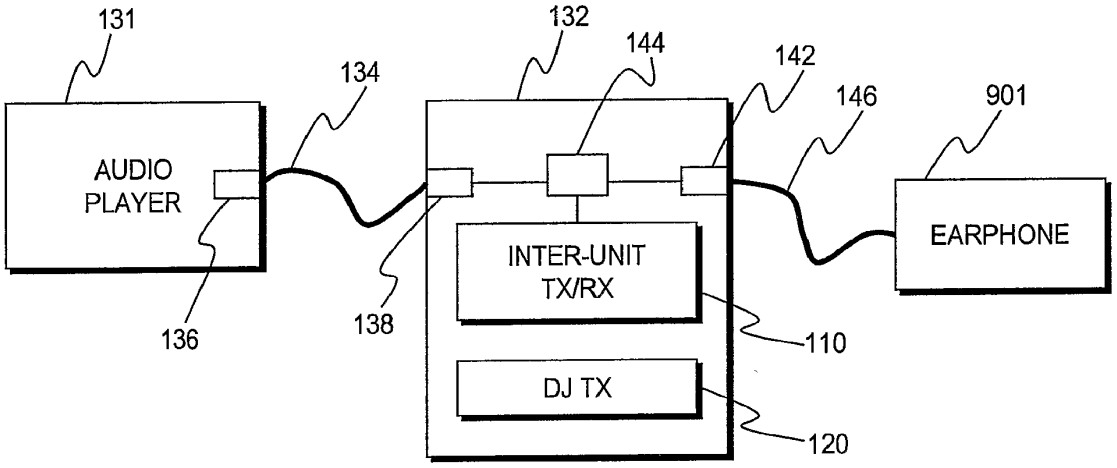


Fig. 12B

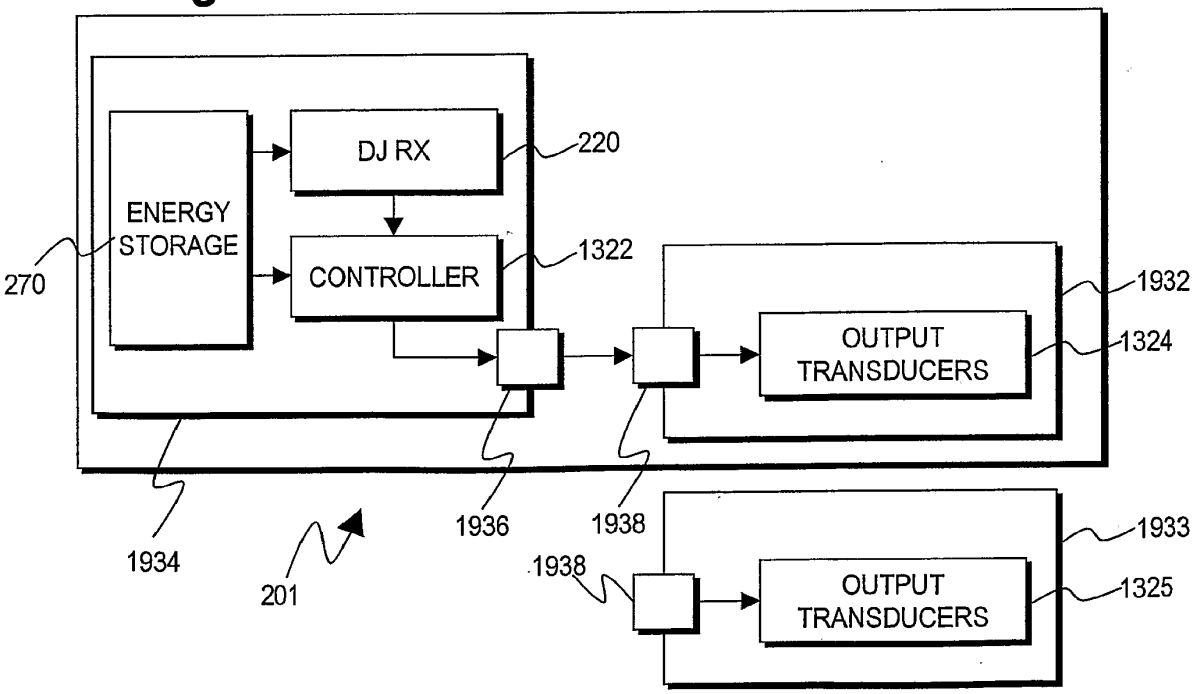
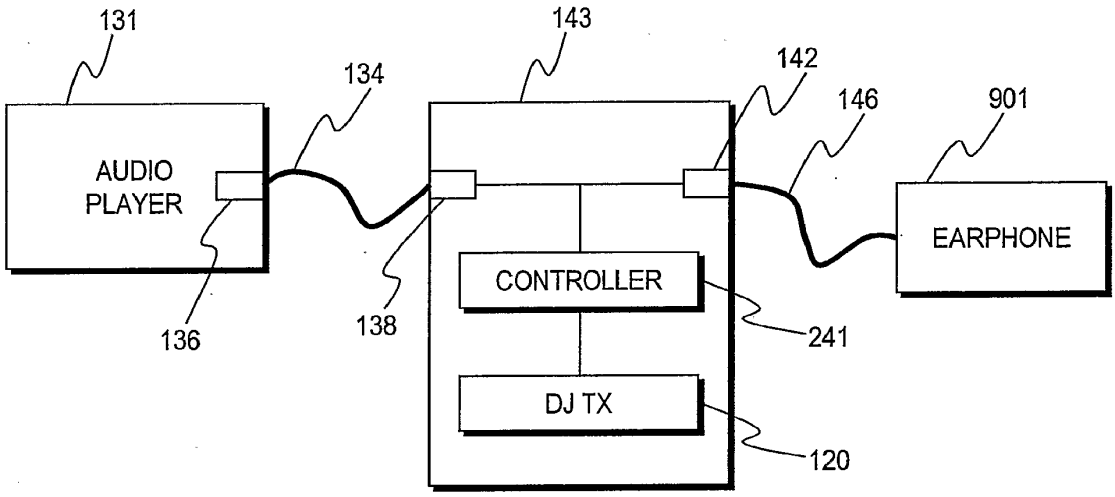


Fig. 12C


131
134
143
142
146
901
AUDIO PLAYER
EARPHONE
CONTROLLER
DJ TX
136
138
241
120

BROADCAST UNIT 710 SEARCH UNIT 750

Fig. 13A

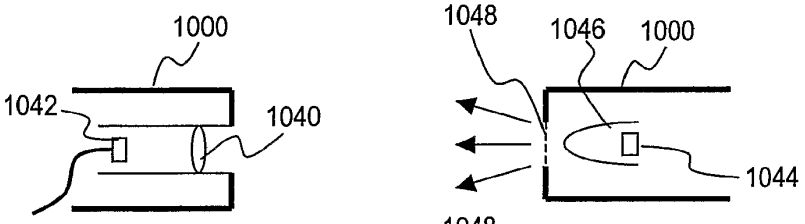


Fig. 13B

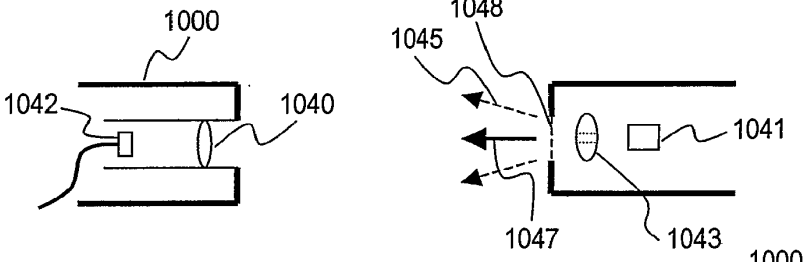


Fig. 13C

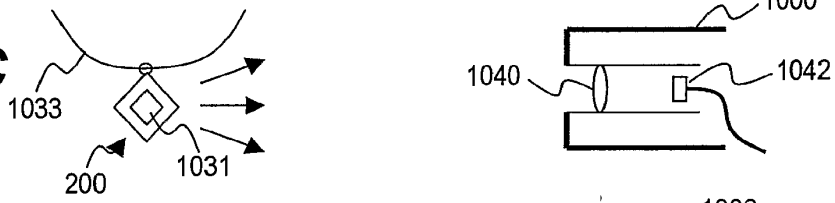


Fig. 13D

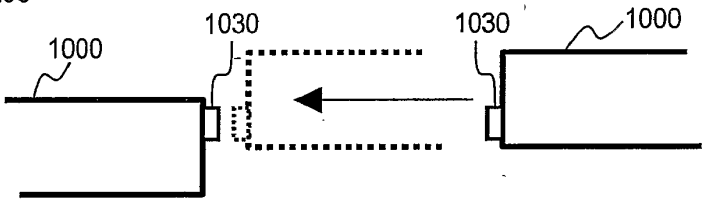


Fig. 13E

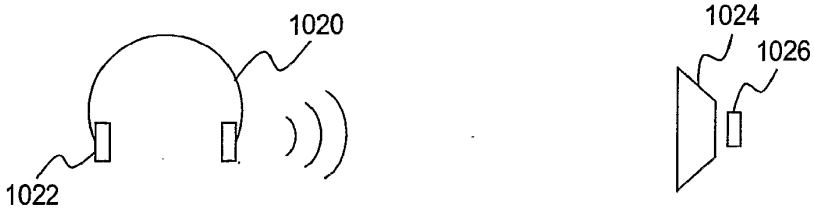


Fig. 13F



Fig. 14A

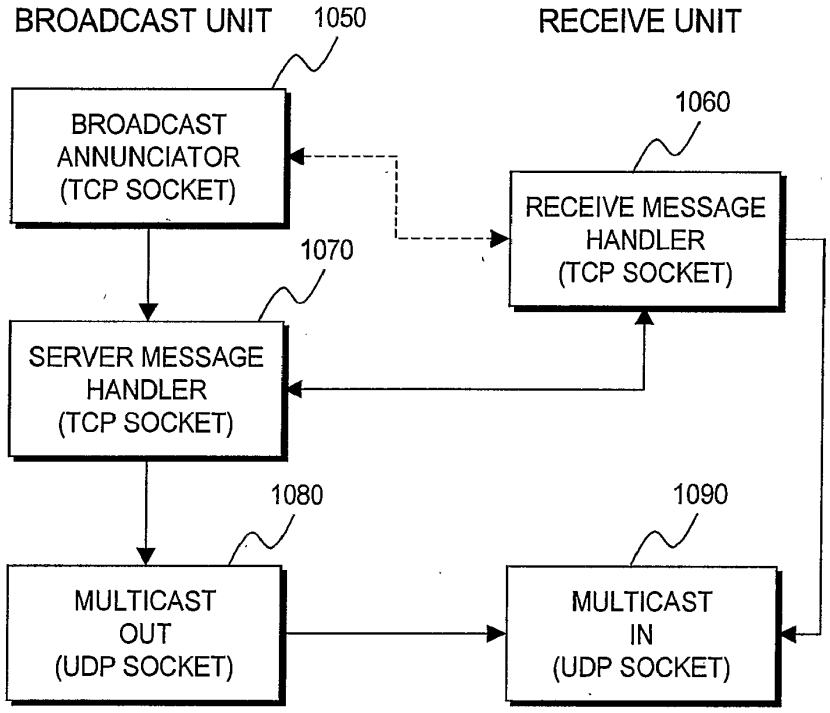

BROADCAST UNIT
1050
RECEIVE UNIT
BROADCAST ANNUNCIATOR (TCP SOCKET)
1060
RECEIVE MESSAGE HANDLER (TCP SOCKET)
1070
SERVER MESSAGE HANDLER (TCP SOCKET)
1080
1090
MULTICAST OUT (UDP SOCKET)
MULTICAST IN (UDP SOCKET)

Fig. 14B

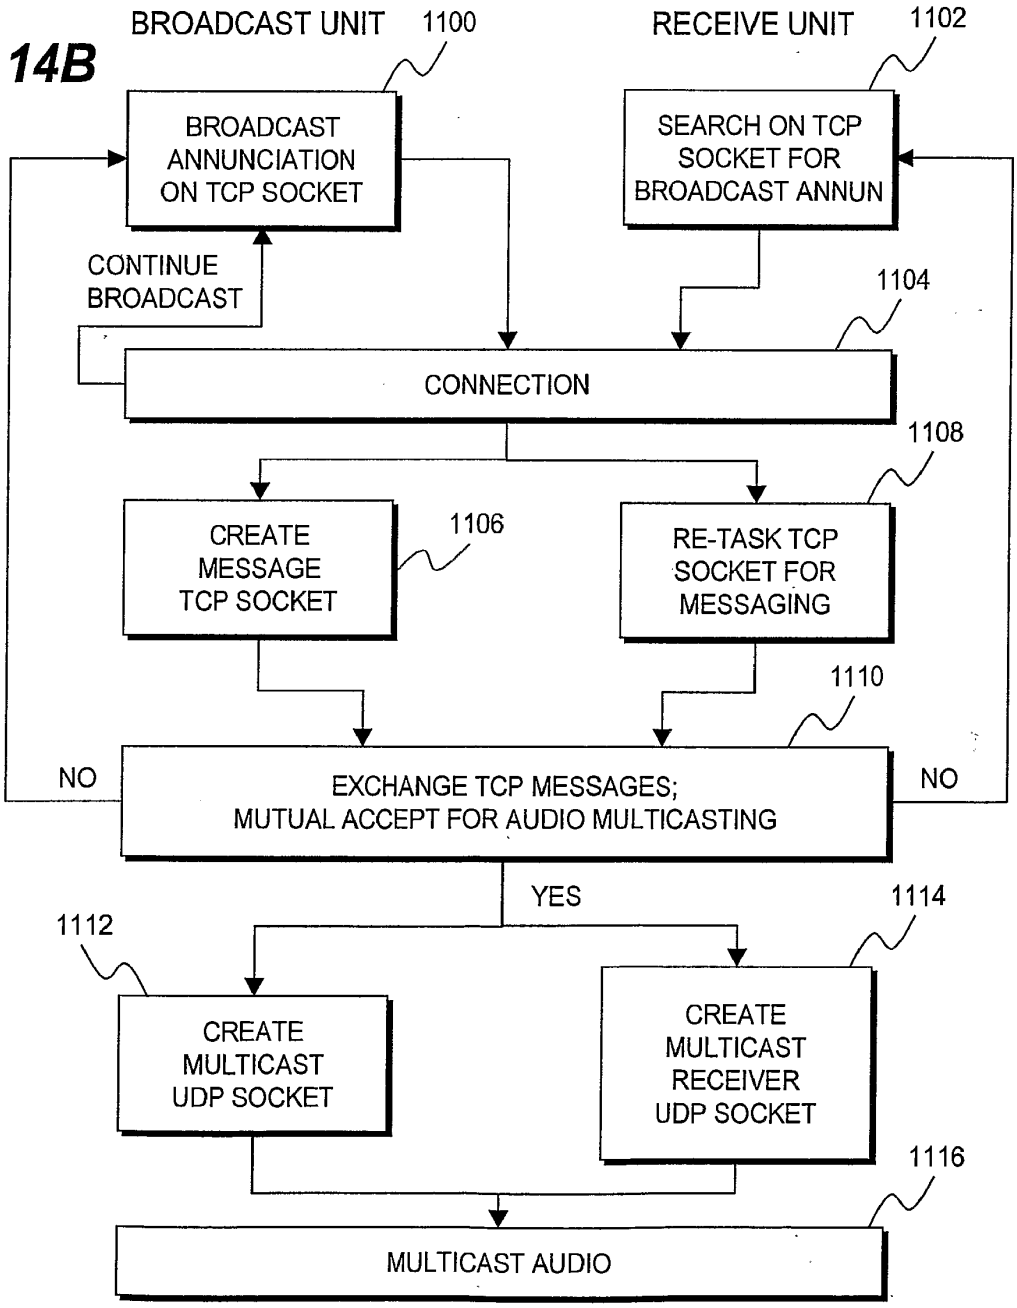

BROADCAST UNIT
1100
RECEIVE UNIT
1102
BROADCAST ANNUNCIATION ON TCP SOCKET
SEARCH ON TCP SOCKET FOR BROADCAST ANNUN
CONTINUE BROADCAST
1104
CONNECTION
1108
CREATE MESSAGE TCP SOCKET
1106
RE-TASK TCP SOCKET FOR MESSAGING
1110
NO
EXCHANGE TCP MESSAGES; MUTUAL ACCEPT FOR AUDIO MULTICASTING
NO
YES
1112
1114
CREATE MULTICAST UDP SOCKET
CREATE MULTICAST RECEIVER UDP SOCKET
1116
MULTICAST AUDIO

Fig. 15

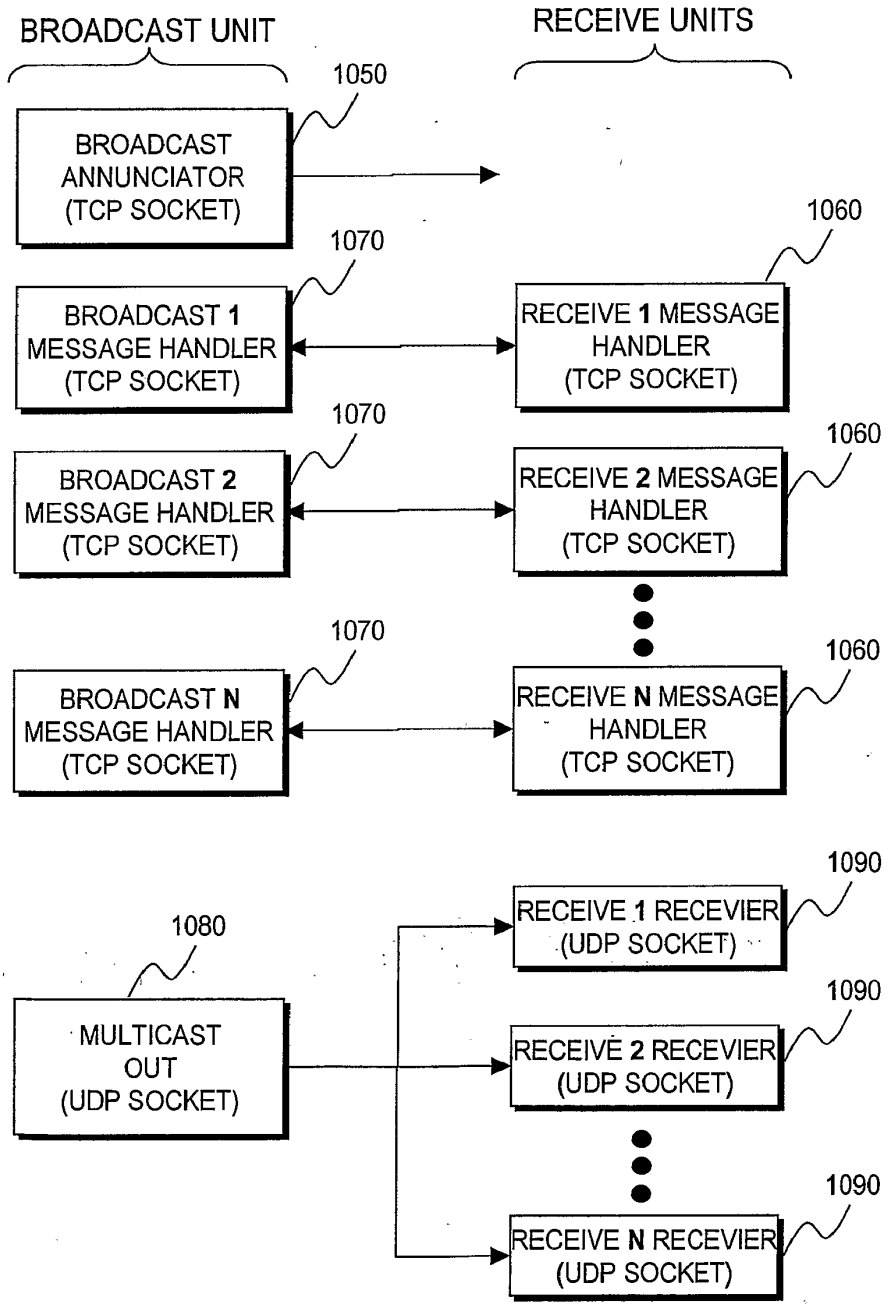

BROADCAST UNIT
RECEIVE UNITS
1050
BROADCAST ANNUNCIATOR (TCP SOCKET)
1060
1070
BROADCAST 1 MESSAGE HANDLER (TCP SOCKET)
RECEIVE 1 MESSAGE HANDLER (TCP SOCKET)
1070
1060
BROADCAST 2 MESSAGE HANDLER (TCP SOCKET)
RECEIVE 2 MESSAGE HANDLER (TCP SOCKET)
1070
1060
BROADCAST N MESSAGE HANDLER (TCP SOCKET)
RECEIVE N MESSAGE HANDLER (TCP SOCKET)
1090
1080
RECEIVE 1 RECEVIER (UDP SOCKET)
1090
MULTICAST OUT (UDP SOCKET)
RECEIVE 2 RECEVIER (UDP SOCKET)
1090
RECEIVE N RECEVIER (UDP SOCKET)

Fig. 16

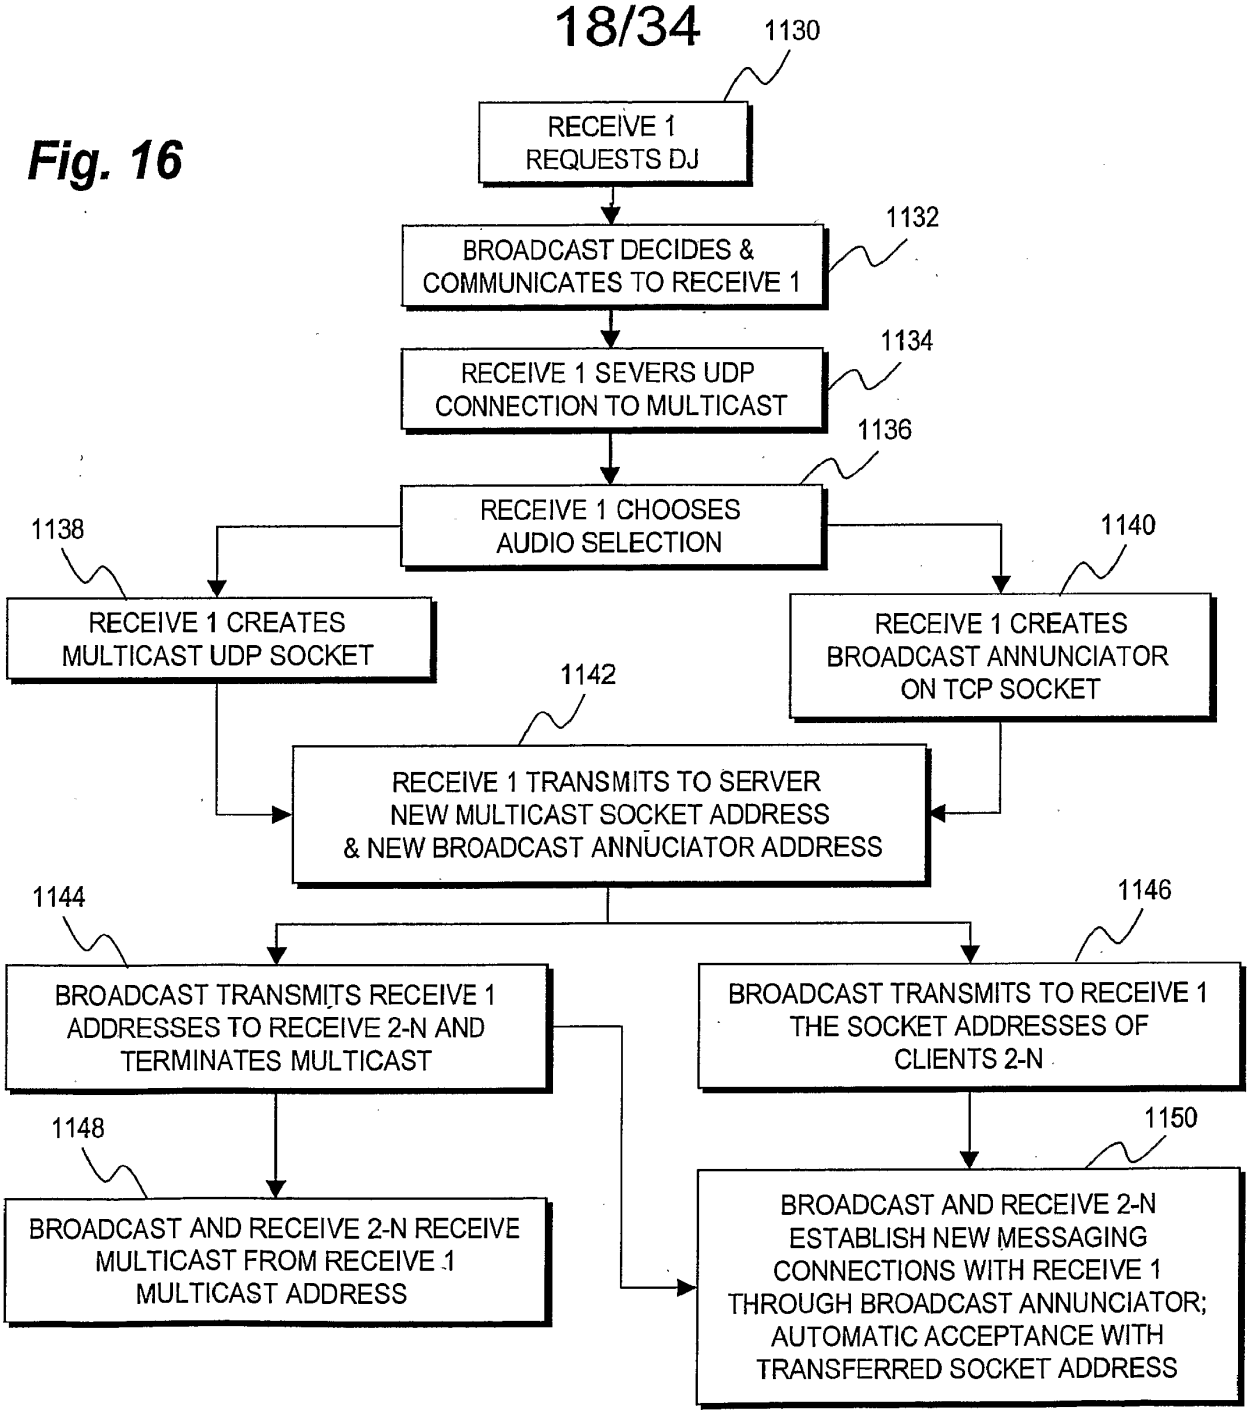

1130
RECEIVE 1 REQUESTS DJ
1132
BROADCAST DECIDES & COMMUNICATES TO RECEIVE 1
1134
RECEIVE 1 SEVERS UDP CONNECTION TO MULTICAST
1136
RECEIVE 1 CHOOSES AUDIO SELECTION
1138
RECEIVE 1 CREATES MULTICAST UDP SOCKET
1140
RECEIVE 1 CREATES BROADCAST ANNUNCIATOR ON TCP SOCKET
1142
RECEIVE 1 TRANSMITS TO SERVER NEW MULTICAST SOCKET ADDRESS & NEW BROADCAST ANNUCIATOR ADDRESS
1144
BROADCAST TRANSMITS RECEIVE 1 ADDRESSES TO RECEIVE 2-N AND TERMINATES MULTICAST
1146
BROADCAST TRANSMITS TO RECEIVE 1 THE SOCKET ADDRESSES OF CLIENTS 2-N
1148
BROADCAST AND RECEIVE 2-N RECEIVE MULTICAST FROM RECEIVE 1 MULTICAST ADDRESS
1150
BROADCAST AND RECEIVE 2-N ESTABLISH NEW MESSAGING CONNECTIONS WITH RECEIVE 1 THROUGH BROADCAST ANNUNCIATOR; AUTOMATIC ACCEPTANCE WITH TRANSFERRED SOCKET ADDRESS

Fig. 17

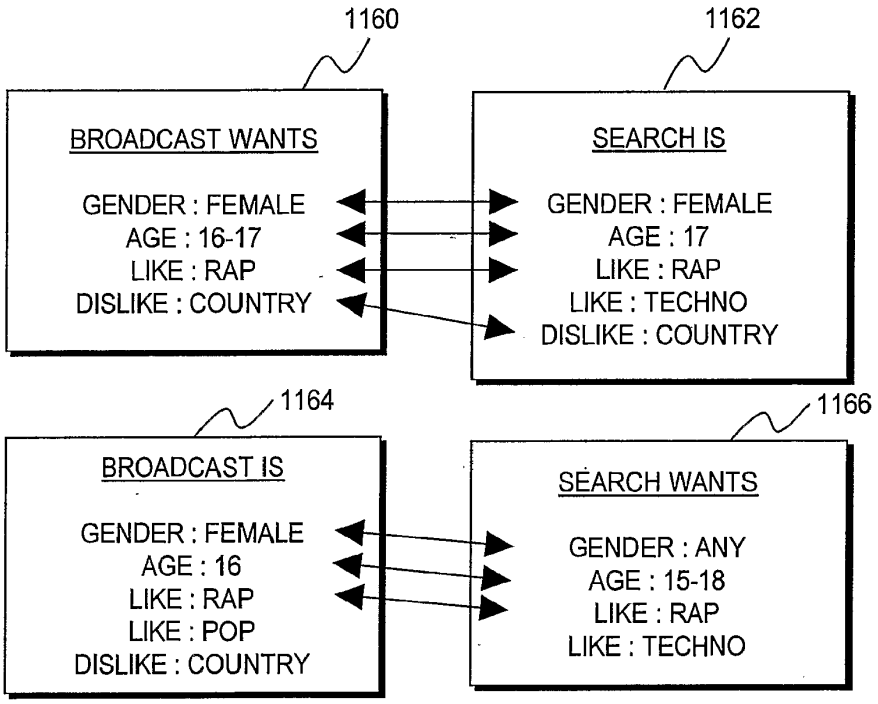


Fig. 18A

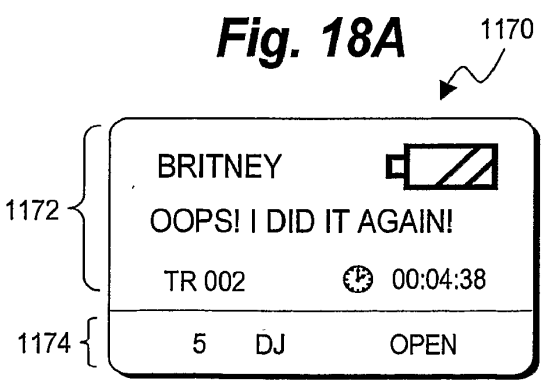


Fig. 18B

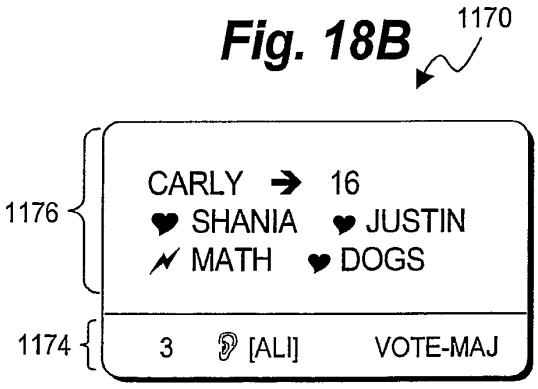


Fig. 19

RULE	METHOD
BROADCASTER	ACCEPTED WHEN BROADCASTER INDICATES "YES". REJECTED WHEN BROADCASTER INDICATES "NO"
MAJORITY	ACCEPTED WHEN (N/2)+1 INDICATE "YES". REJECTED WHEN (N/2) INDICATE "NO".
UNANIMOUS	ACCEPTED WHEN N MEMBERS INDICATE "YES". REJECTED WHEN 1 MEMBER INDICATES "NO".
TIMED MAJORITY	MACHINE INDICATES ONSET OF VOTING, AND COUNTS DOWN TO VOTING TERMINATION. ACCEPTED IF (Q/2)+1 INDICATE "YES" BEFORE TERMINATION; OTHERWISE "NO".
SYNCHRONIZED MAJORITY	MACHINE COUNTS DOWN TO VOTE. ACCEPTED IF (Q/2)+1 INDICATE "YES" WITHIN PREDETERMINED TIME OF COUNTDOWN. OTHERWISE "NO".

Fig. 20

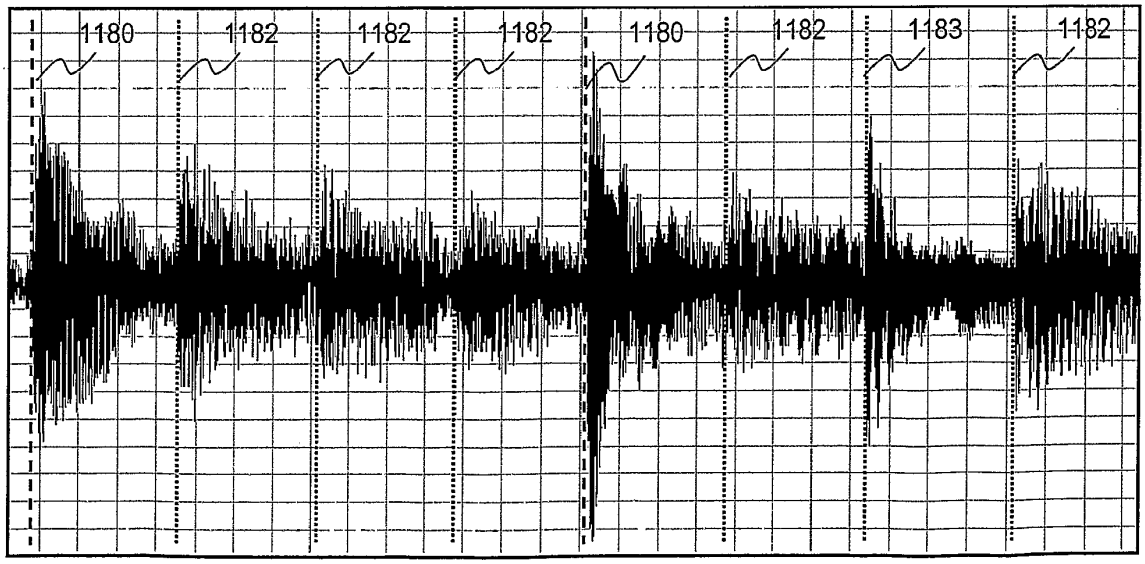


Fig. 21A

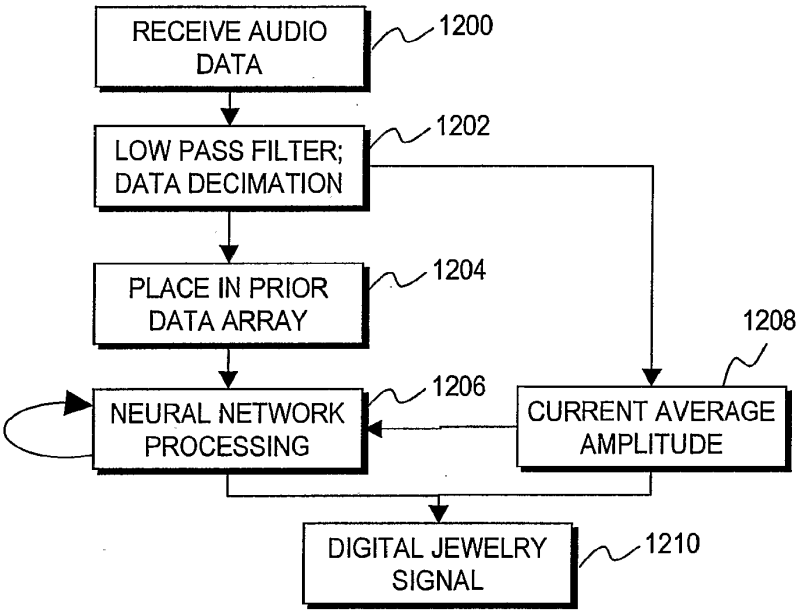


Fig. 21B

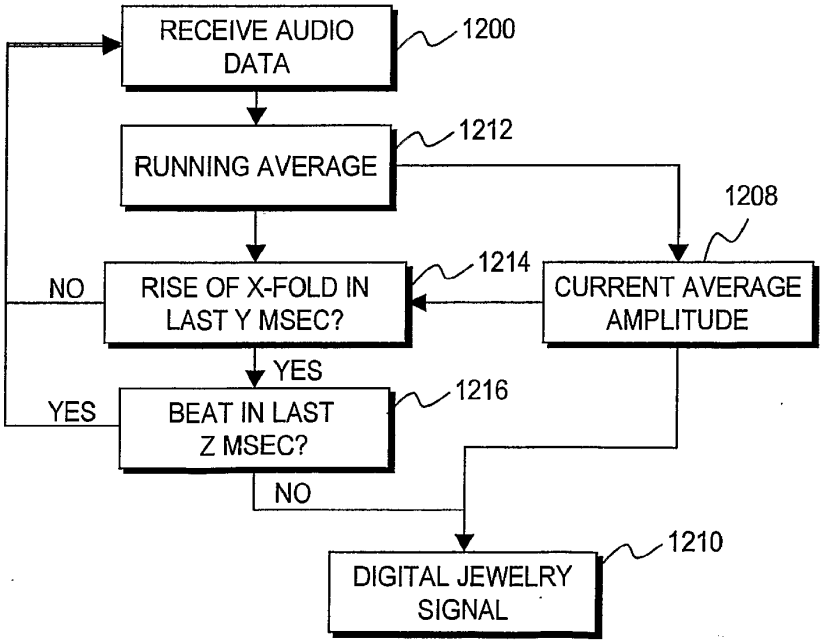


Fig. 21C

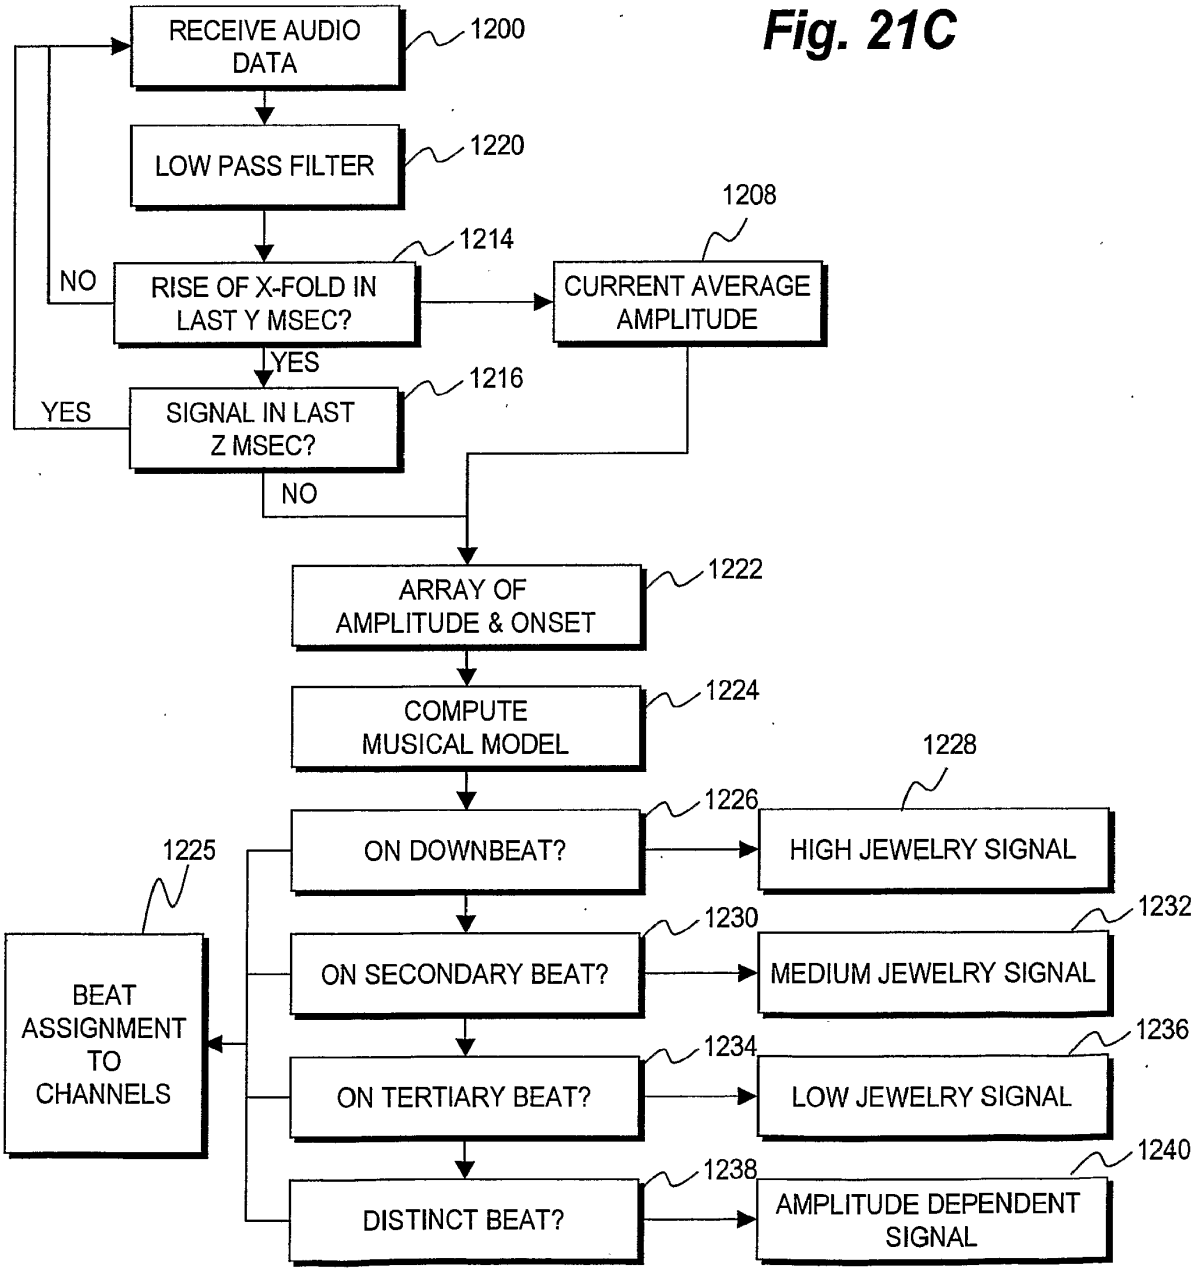
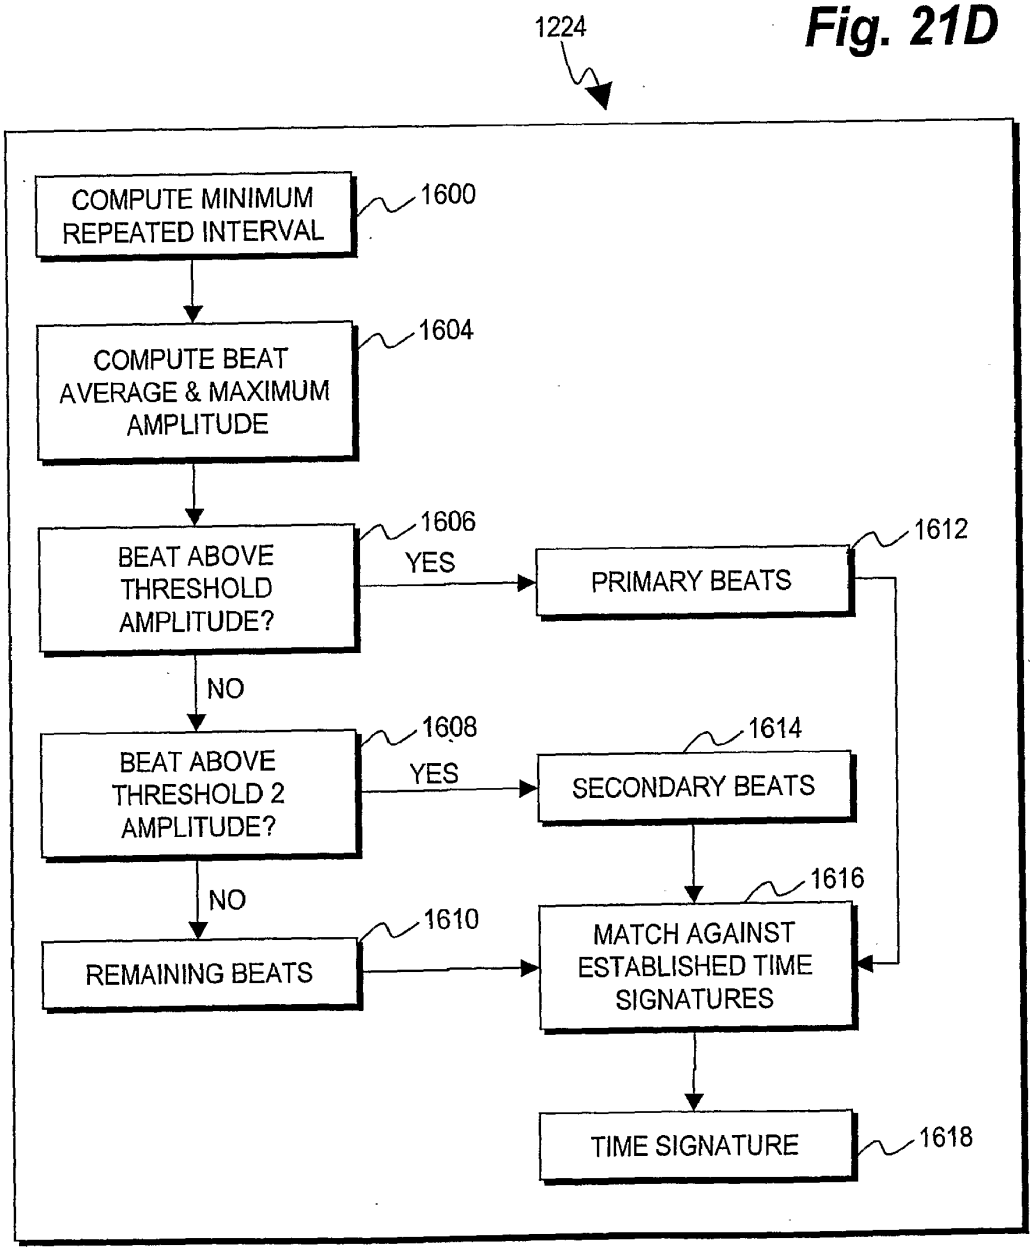

RECEIVE AUDIO DATA
1200
LOW PASS FILTER
1220
1208
1214
NO
RISE OF X-FOLD IN LAST Y MSEC?
CURRENT AVERAGE AMPLITUDE
YES
1216
YES
SIGNAL IN LAST Z MSEC?
NO
ARRAY OF AMPLITUDE & ONSET
1222
COMPUTE MUSICAL MODEL
1224
1228
1226
ON DOWNBEAT?
HIGH JEWELRY SIGNAL
1225
1230
1232
ON SECONDARY BEAT?
MEDIUM JEWELRY SIGNAL
BEAT ASSIGNMENT TO CHANNELS
1234
1236
ON TERTIARY BEAT?
LOW JEWELRY SIGNAL
1238
1240
DISTINCT BEAT?
AMPLITUDE DEPENDENT SIGNAL

Fig. 21D


1224
COMPUTE MINIMUM REPEATED INTERVAL
1600
COMPUTE BEAT AVERAGE & MAXIMUM AMPLITUDE
1604
BEAT ABOVE THRESHOLD AMPLITUDE?
1606
YES
PRIMARY BEATS
1612
NO
BEAT ABOVE THRESHOLD 2 AMPLITUDE?
1608
YES
SECONDARY BEATS
1614
NO
REMAINING BEATS
1610
MATCH AGAINST ESTABLISHED TIME SIGNATURES
1616
TIME SIGNATURE
1618

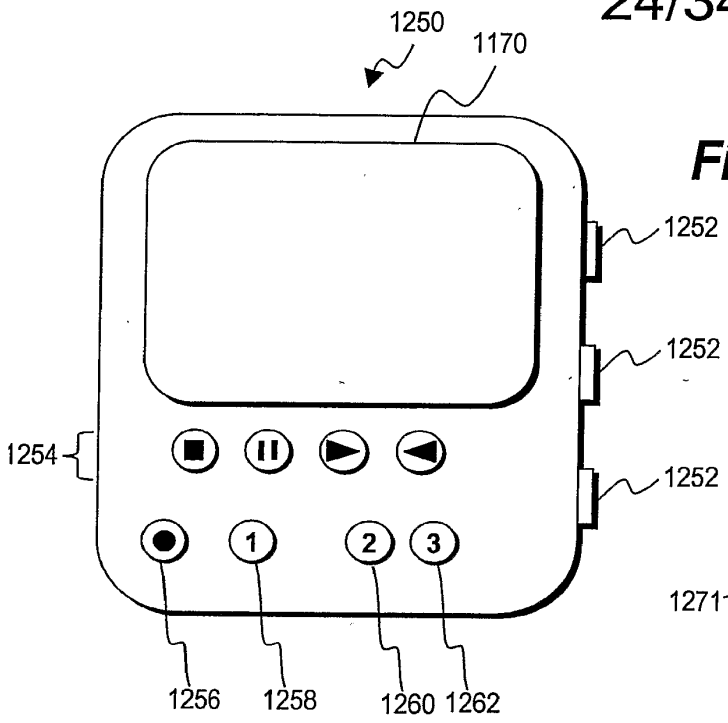


Fig. 22A

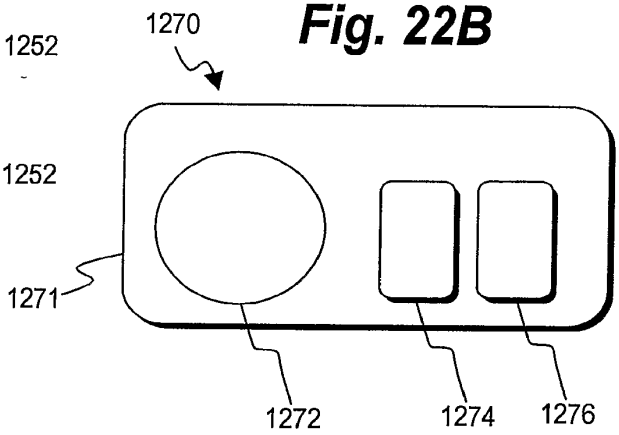


Fig. 22B

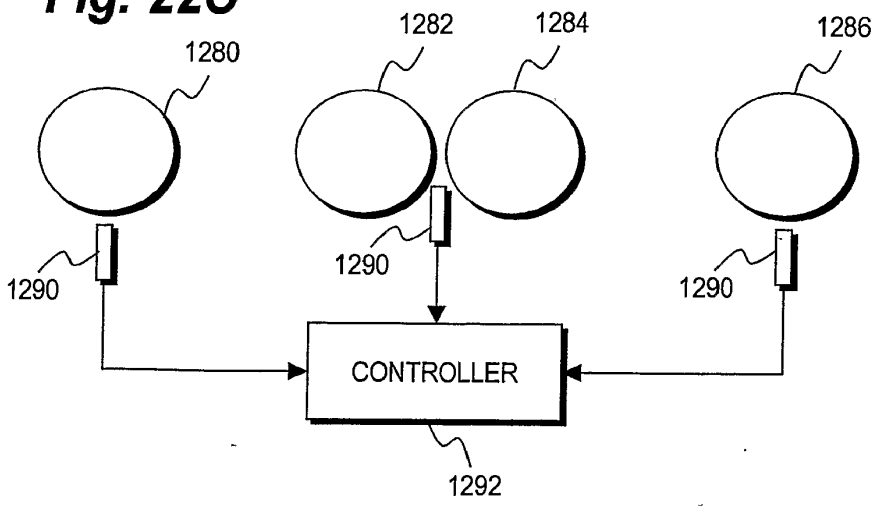


Fig. 22C

Fig. 23

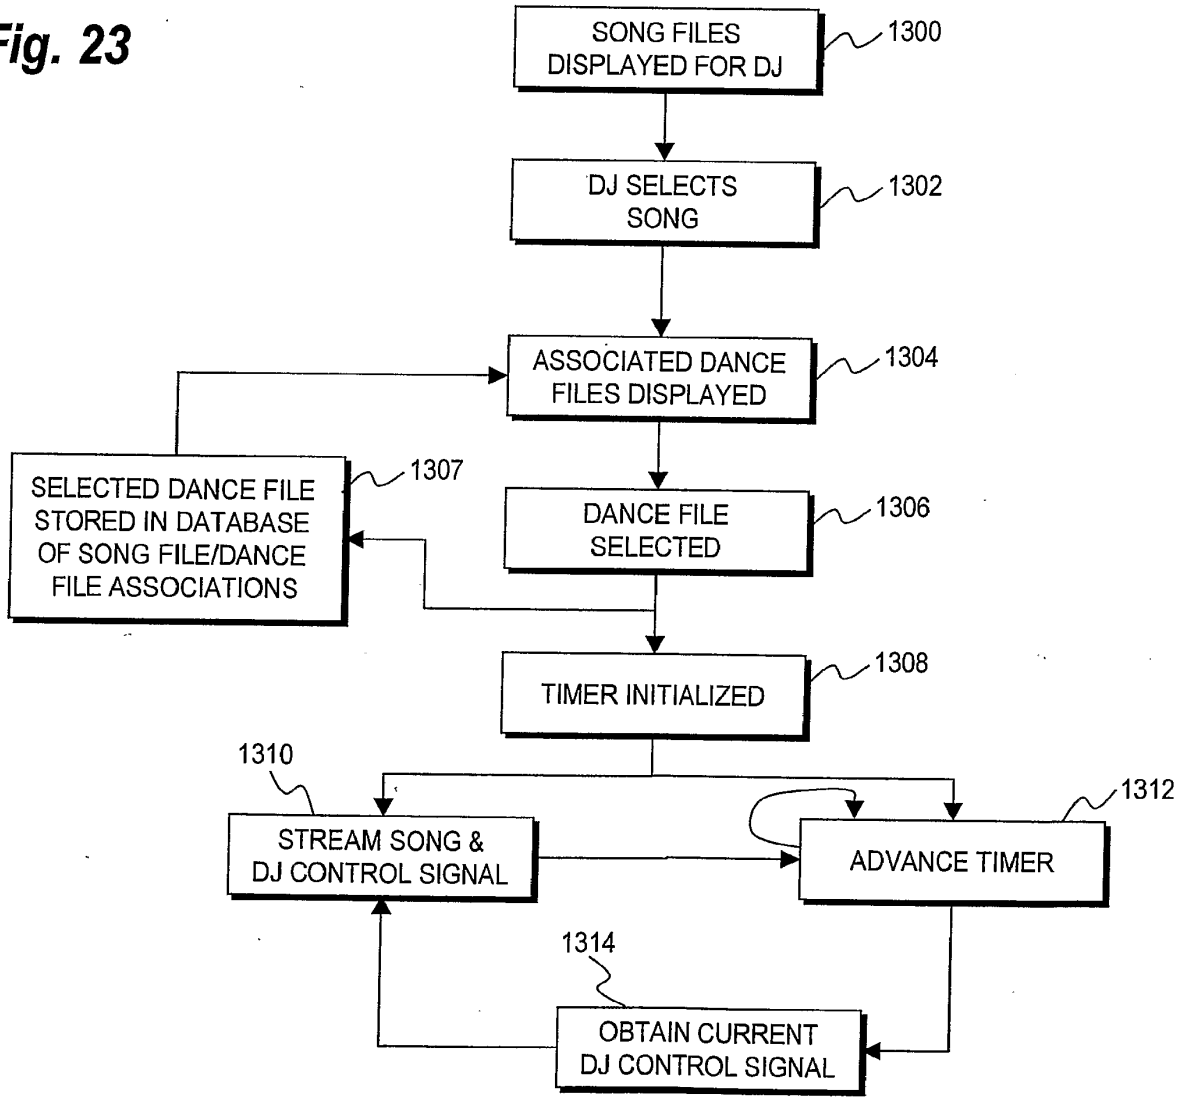

SONG FILES DISPLAYED FOR DJ
1300
DJ SELECTS SONG
1302
ASSOCIATED DANCE FILES DISPLAYED
1304
SELECTED DANCE FILE STORED IN DATABASE OF SONG FILE/DANCE FILE ASSOCIATIONS
1307
DANCE FILE SELECTED
1306
TIMER INITIALIZED
1308
1310
STREAM SONG & DJ CONTROL SIGNAL
ADVANCE TIMER
1312
1314
OBTAIN CURRENT DJ CONTROL SIGNAL

Fig. 24

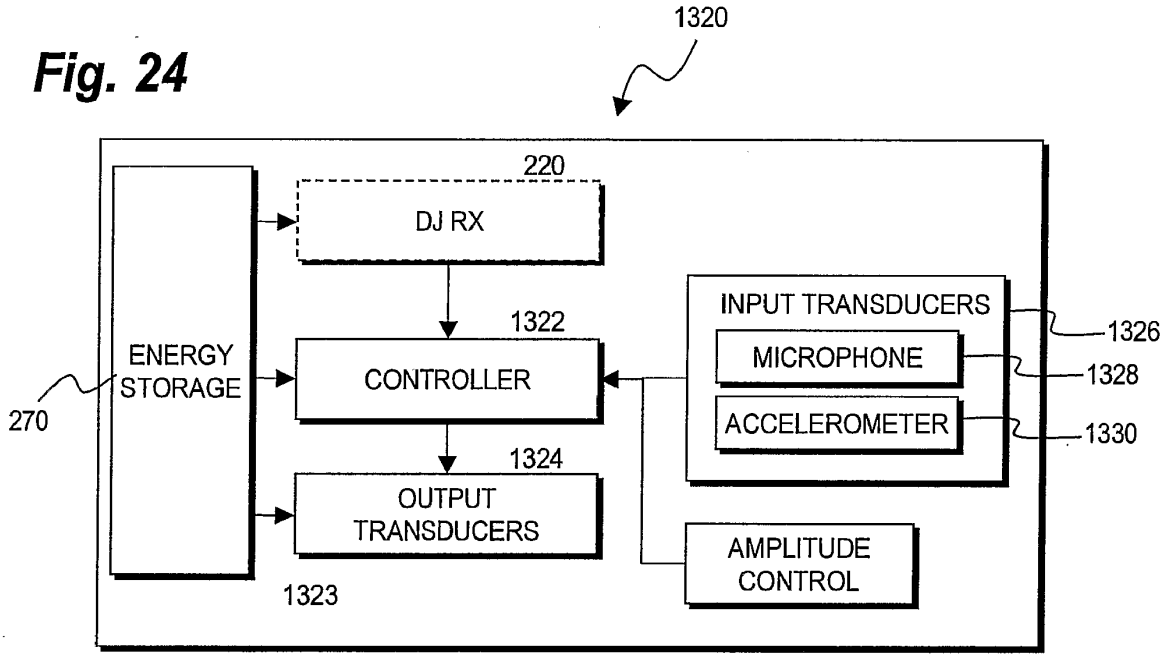

1320
220
DJ RX
1322
ENERGY STORAGE
CONTROLLER
270
1324
OUTPUT TRANSDUCERS
1323
INPUT TRANSDUCERS
MICROPHONE
ACCELEROMETER
1326
1328
1330
AMPLITUDE CONTROL

Fig. 25

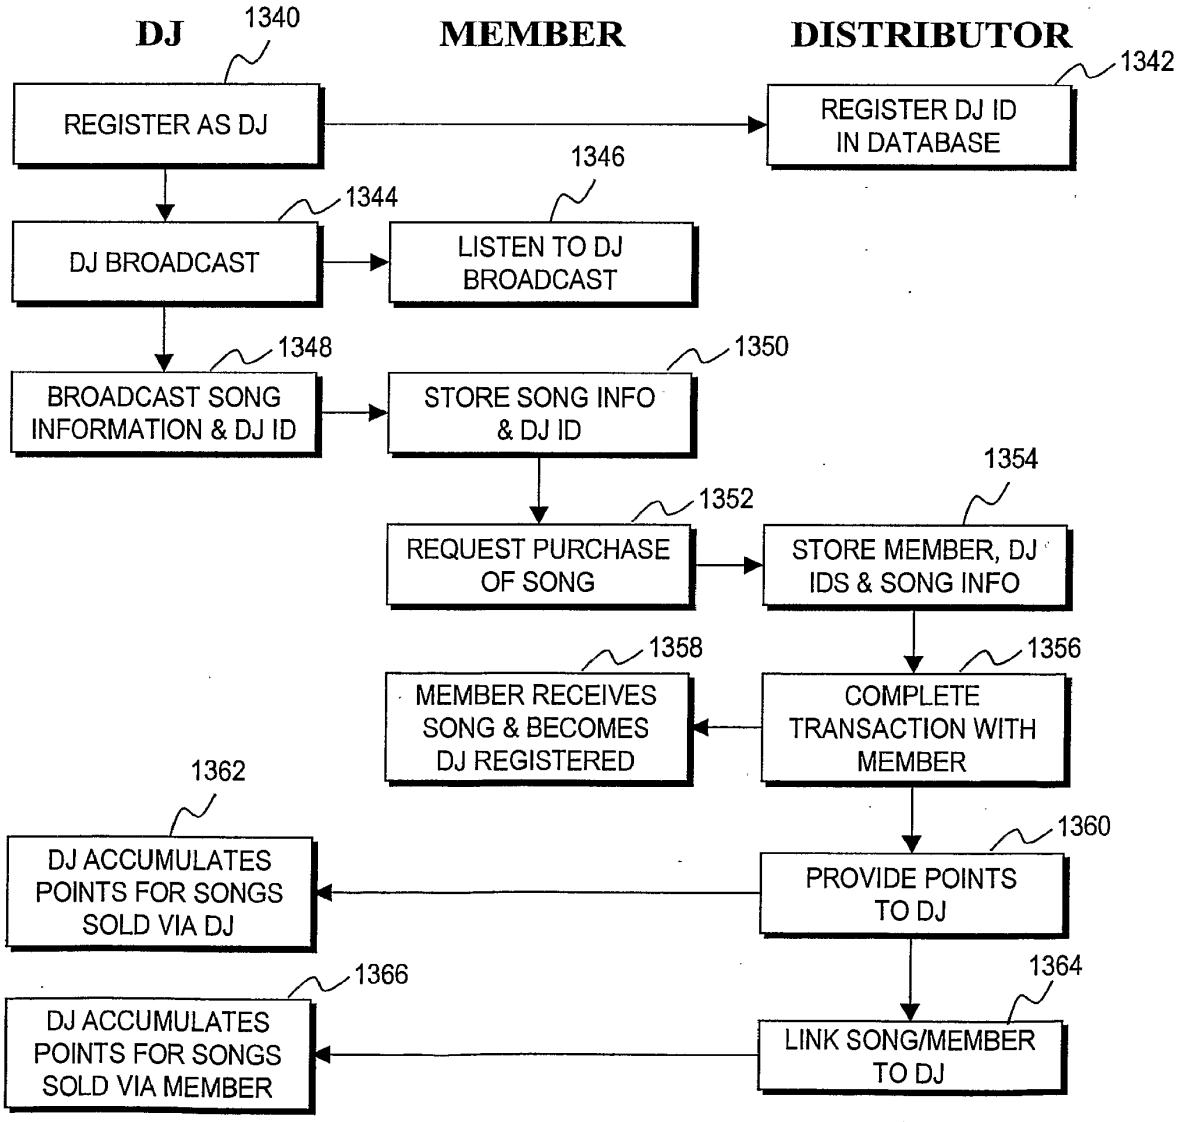

DJ
MEMBER
DISTRIBUTOR
1340
1342
REGISTER AS DJ
REGISTER DJ ID IN DATABASE
1346
1344
DJ BROADCAST
LISTEN TO DJ BROADCAST
1348
1350
BROADCAST SONG INFORMATION & DJ ID
STORE SONG INFO & DJ ID
1354
1352
REQUEST PURCHASE OF SONG
STORE MEMBER, DJ IDS & SONG INFO
1358
1356
MEMBER RECEIVES SONG & BECOMES DJ REGISTERED
COMPLETE TRANSACTION WITH MEMBER
1362
1360
DJ ACCUMULATES POINTS FOR SONGS SOLD VIA DJ
PROVIDE POINTS TO DJ
1366
1364
DJ ACCUMULATES POINTS FOR SONGS SOLD VIA MEMBER
LINK SONG/MEMBER TO DJ

Fig. 26

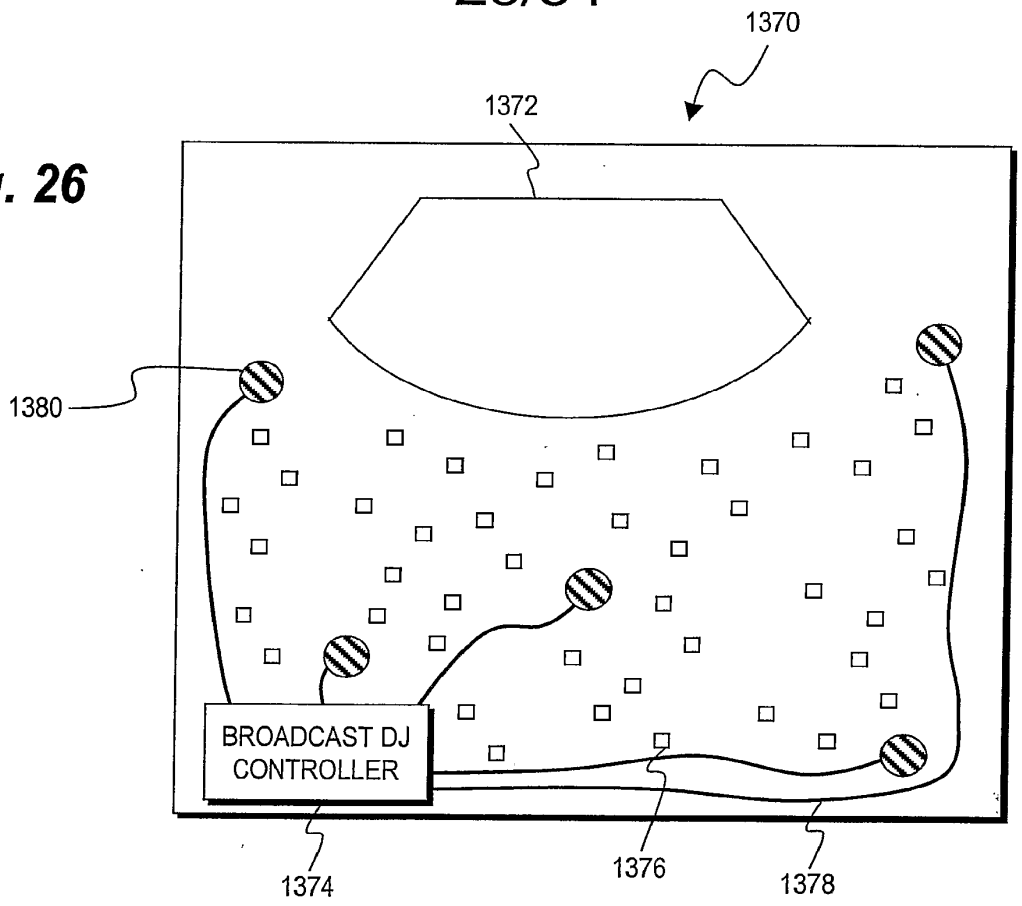

1370
1372
1380
BROADCAST DJ CONTROLLER
1374
1376
1378

Fig. 35

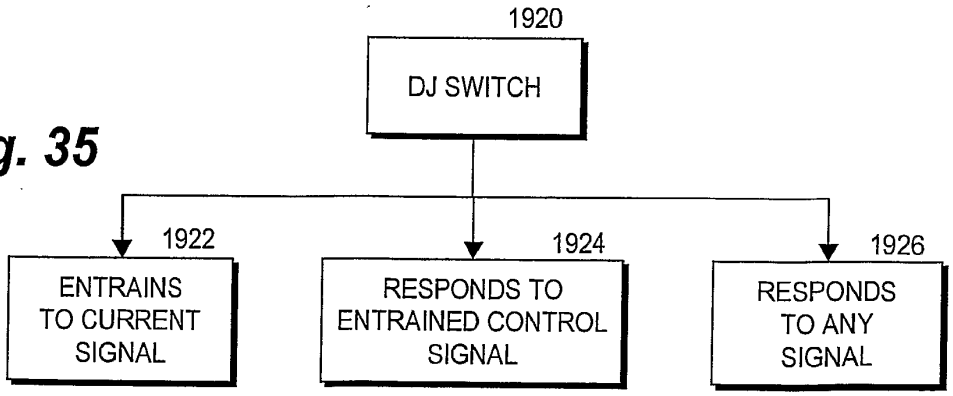

1920
DJ SWITCH
1922
ENTRAINS TO CURRENT SIGNAL
1924
RESPONDS TO ENTRAINED CONTROL SIGNAL
1926
RESPONDS TO ANY SIGNAL

Fig. 27

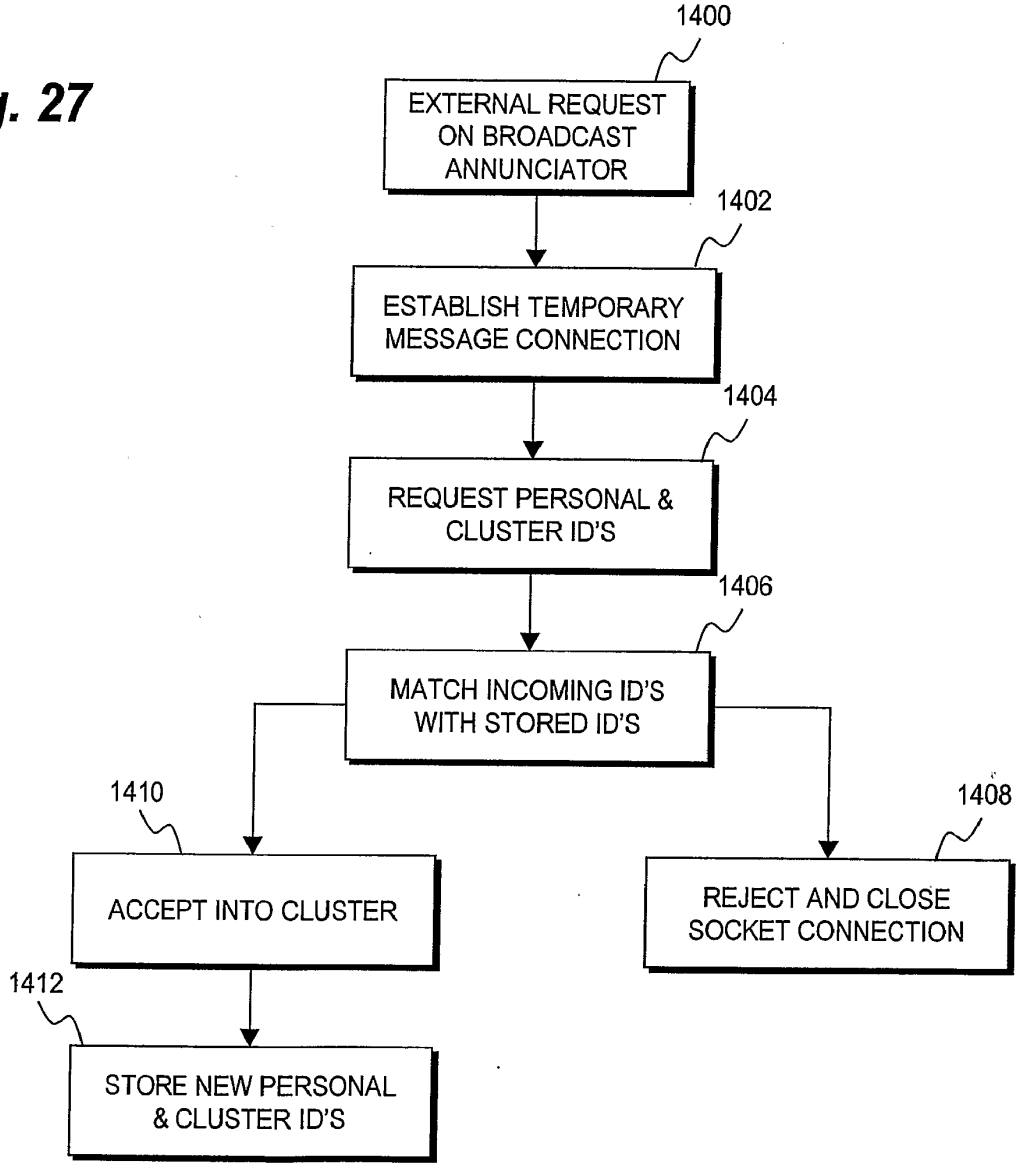

1400
EXTERNAL REQUEST ON BROADCAST ANNUNCIATOR
1402
ESTABLISH TEMPORARY MESSAGE CONNECTION
1404
REQUEST PERSONAL & CLUSTER ID'S
1406
MATCH INCOMING ID'S WITH STORED ID'S
1410
ACCEPT INTO CLUSTER
1408
REJECT AND CLOSE SOCKET CONNECTION
1412
STORE NEW PERSONAL & CLUSTER ID'S

Fig. 28

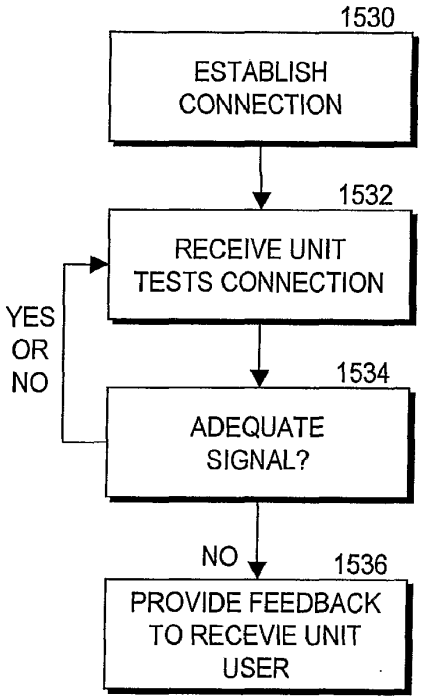

1530
ESTABLISH CONNECTION
1532
RECEIVE UNIT TESTS CONNECTION
YES OR NO
1534
ADEQUATE SIGNAL?
NO
1536
PROVIDE FEEDBACK TO RECEVIE UNIT USER

Fig. 29A

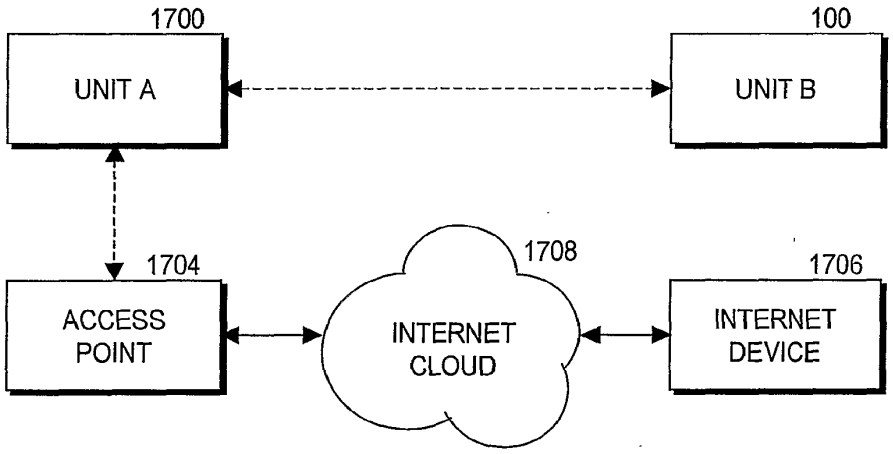


Fig. 29B

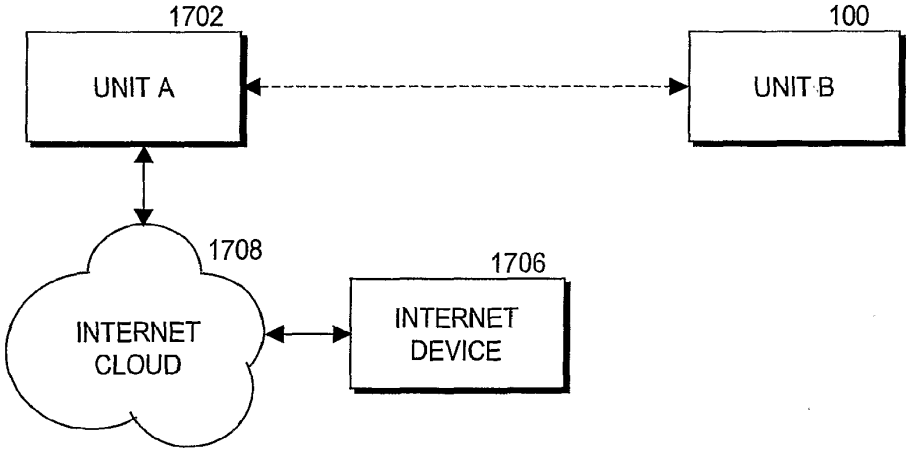


Fig. 30

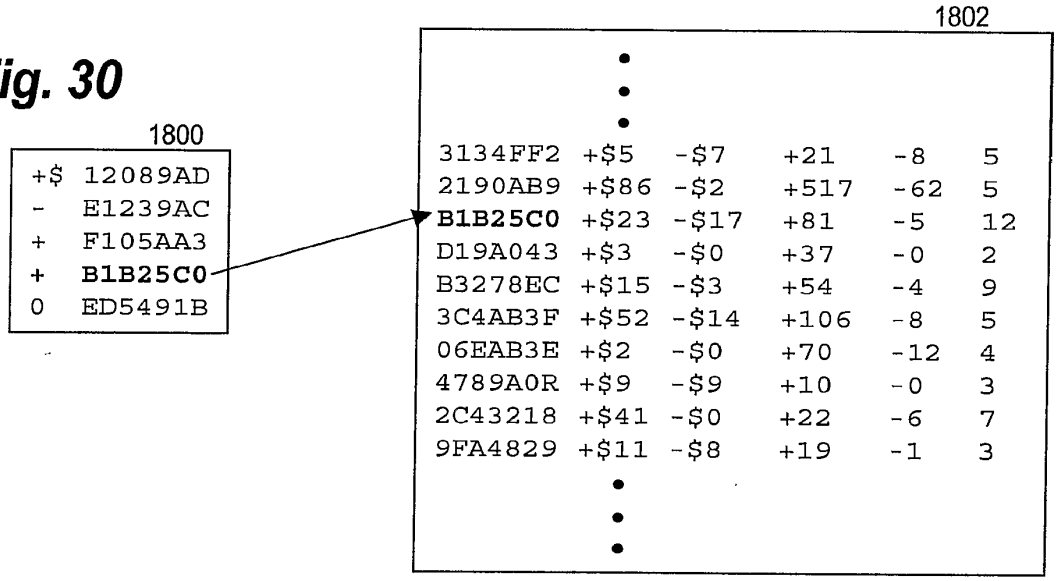

1800
+$ 12089AD
- E1239AC
+ F105AA3
+ B1B25C0
0 ED5491B
1802
3134FF2 +$5 -$7 +21 -8 5
2190AB9 +$86 -$2 +517 -62 5
B1B25C0 +$23 -$17 +81 -5 12
D19A043 +$3 -$0 +37 -0 2
B3278EC +$15 -$3 +54 -4 9
3C4AB3F +$52 -$14 +106 -8 5
06EAB3E +$2 -$0 +70 -12 4
4789A0R +$9 -$9 +10 -0 3
2C43218 +$41 -$0 +22 -6 7
9FA4829 +$11 -$8 +19 -1 3

Fig. 31

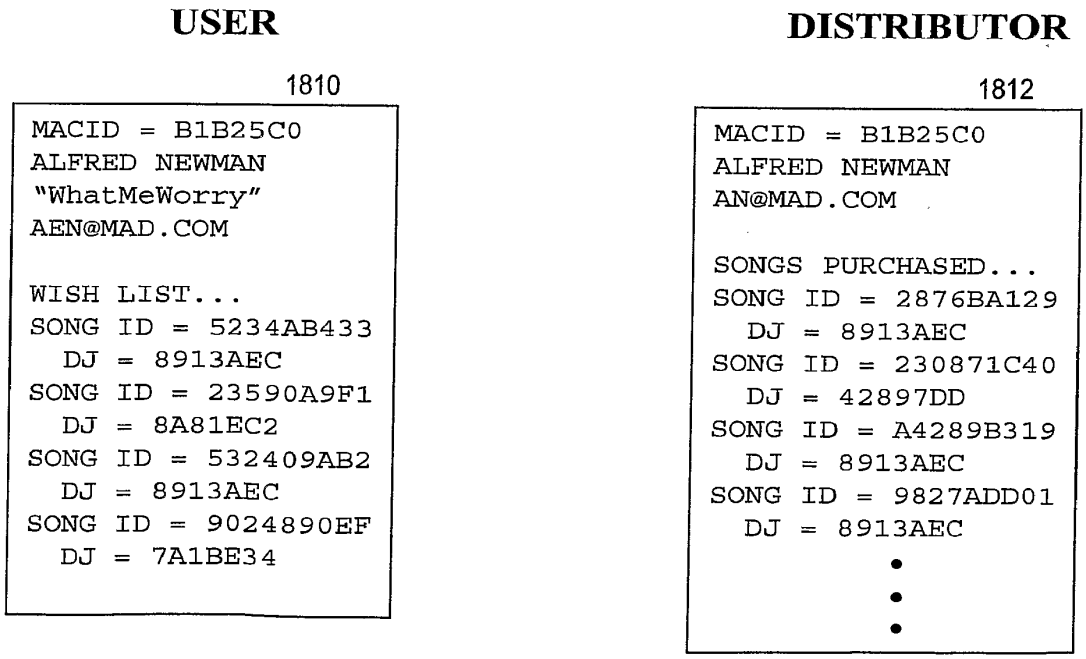

USER
DISTRIBUTOR
1810
MACID = B1B25C0
ALFRED NEWMAN
"WhatMeWorry"
AEN@MAD.COM
WISH LIST...
SONG ID = 5234AB433
DJ = 8913AEC
SONG ID = 23590A9F1
DJ = 8A81EC2
SONG ID = 532409AB2
DJ = 8913AEC
SONG ID = 9024890EF
DJ = 7A1BE34
1812
MACID = B1B25C0
ALFRED NEWMAN
AN@MAD.COM
SONGS PURCHASED...
SONG ID = 2876BA129
DJ = 8913AEC
SONG ID = 230871C40
DJ = 42897DD
SONG ID = A4289B319
DJ = 8913AEC
SONG ID = 9827ADD01
DJ = 8913AEC

Fig. 32A

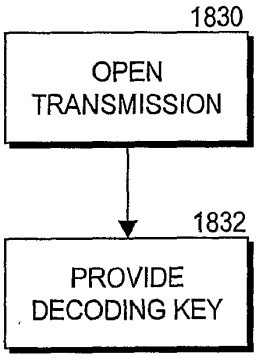


Fig. 32B

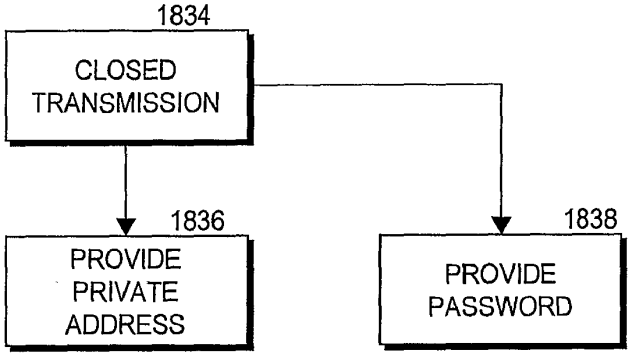


Fig. 33

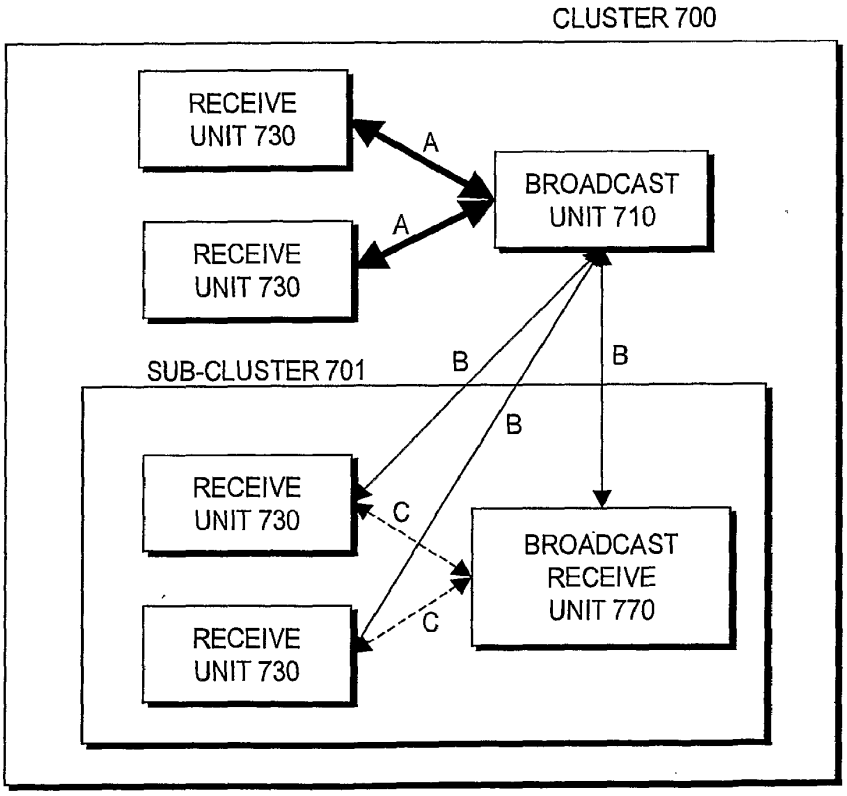


Fig. 34A

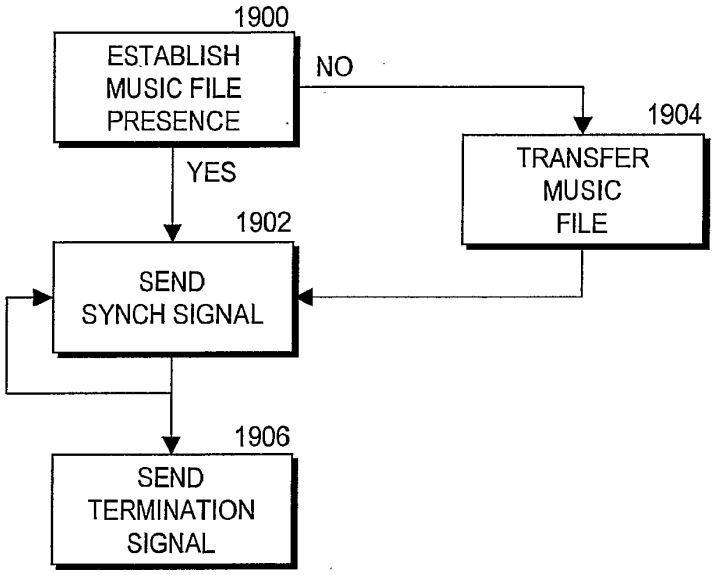
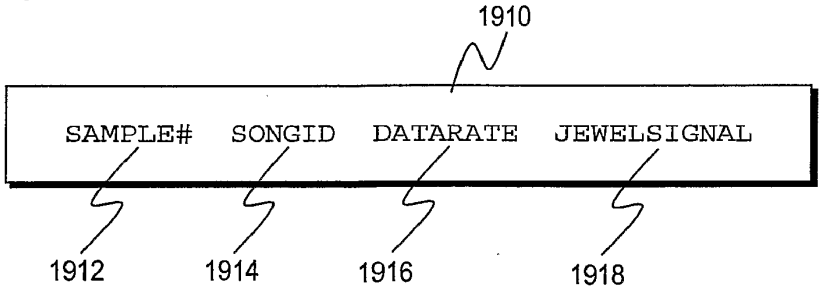


Fig. 34B



PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	13864248
Filing Date	2013-04-17
First Named Inventor	Nicholas A. J. Millington
Art Unit	2442
Examiner Name	Oleg Survillo
Attorney Docket Number	SNS-71301.US.10 (13-0401)

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	5239458		1993-08-24	SUZUKI; Takashi	
	2	5299266		1994-03-29	LUMSDEN; John L.	
	3	5751819		1998-05-12	DORROUGH; Michael L.	
	4	5808662		1998-09-15	KINNEY; Daniel Reid et al.	
	5	6108686		2000-08-22	WILLIAMS Jr.; ,Henry R.	
	6	6255961		2001-07-03	VAN; Ryzin John M. et al.	
	7	6349339		2002-02-19	WILLIAMS; Henry R.	
	8	7113999		2006-09-26	PESTONI; Florian et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2442
	Examiner Name	Oleg Survillo
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	9	7218708		2007-05-15	BEREZOWSKI; Andrew G. et al.	
	10	7236739		2007-06-26	CHANG; Chin-Chi	
	11	7293060		2007-11-06	KOMSI; Asko	
	12	7483958		2009-01-27	ELABBADY; Tarek Z. et al.	
	13	7571014		2009-08-04	LAMBOURNE; Robert A. et al.	
	14	7599685		2009-10-06	GOLDBERG; David et al.	
	15	7657224		2010-02-02	GOLDBERG; David et al.	
	16	7668990		2010-02-23	KRZYZANOWSKI; Paul et al.	
	17	7669219		2010-02-23	SCOTT; Samuel T.III.	
	18	7702279		2010-04-20	KO; Steve et al.	
	19	7835689		2010-11-16	GOLDBERG; David et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2442
	Examiner Name	Oleg Survillo
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	20	7865137		2011-01-04	GOLDBERG; David et al.	
	21	7885622		2011-02-08	KRAMPF; Steven S. et al.	
	22	7916877		2011-03-29	GOLDBERG; David et al.	
	23	7917082		2011-03-29	GOLDBERG; David et al.	
	24	8023663		2011-09-20	GOLDBERG; David	
	25	8028038		2011-09-27	WEEL; Martin	
	26	8028323		2011-09-27	WEEL; Martin	
	27	8112032		2012-02-07	KO; Steve et al.	
	28	8214873		2012-07-03	WEEL; Martin	
	29	8230099		2012-07-24	WEEL; Martin	
	30	8315555		2012-11-20	KO; Steve et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2442
	Examiner Name	Oleg Survillo
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20020003548		2002-01-10	KRUSCHE; Arnd et al.	
	2	20020072816		2002-06-13	SHDEMA; Yoav et al.	
	3	20020109710		2002-08-15	HOLTZ; Alex et al.	
	4	20020124097		2002-09-05	ISELY; Larson J. et al.	
	5	20020188762		2002-12-12	TOMASSETTI; Stephen Robert et al.	
	6	20030020763		2003-01-30	MAYER; Theodore et al.	
	7	20030235304		2003-12-25	EVANS; Glenn F. et al.	
	8	20040008852		2004-01-15	ALSO; Masaru et al.	
	9	20040010727		2004-01-15	FUJINAMI; Yasushi	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2442
	Examiner Name	Oleg Survillo
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	10	20040015252		2004-01-22	AISO; Masaru et al.	
	11	20040131192		2004-07-08	METCALF; Randall B.	
	12	20050047605		2005-03-03	LEE; Noel et al.	
	13	20050125357		2005-06-09	SAADAT; Abbas S. et al.	
	14	20050181348		2005-08-18	CAREY; Tadhg M. et al.	
	15	20050195205		2005-09-08	ABRAMS; Thomas A.Jr	
	16	20050289224		2005-12-29	DESLIPPE; Mark H. et al.	
	17	20060095516		2006-05-04	WIJERATNE; Viranga L.	
	18	20060143236		2006-06-29	WU; Michael	
	19	20070048713		2007-03-01	PLASTINA; Daniel et al.	
	20	20070142022		2007-06-21	MADONNA; Robert P. et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2442
	Examiner Name	Oleg Survillo
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	21	20070169115		2007-07-19	KO; Steve et al.	
	22	20070299778		2007-12-27	HAVESON; Ryan A. et al.	
	23	20080091771		2008-04-17	ALLEN; Jeffrey L. et al.	
	24	20080144861		2008-06-19	MELANSON; John L. et al.	
	25	20090031336		2009-01-29	CHAVEZ; Timothy R. et al.	
	26	20090222115		2009-09-03	MALCOLM; David et al.	
	27	20100049835		2010-02-25	KO; Steve et al.	
	28	20100228740		2010-09-09	CANNISTRARO; Alan et al.	
	29	20110066943		2011-03-17	BRILLON; Francois et al.	
	30	20120030366		2012-02-02	COLLART; Todd et al.	
	31	20120060046		2012-03-08	MILLINGTON; Nicholas A. J.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2442
	Examiner Name	Oleg Survillo
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	32	20120129446		2012-05-24	KO; Steve et al.	
	33	20130080599		2013-03-28	KO; Steve et al.	
	34	20130191454		2013-07-25	OLIVER; Michael R. et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	2005013047	WO	A2	2005-02-10	SONOS INC [US], et al.		☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	"INTERNATIONAL BUREAU,"Search Report" issued in connection with International Patent application No. PCT/ US2013/046372, mailed August 26, 2013, 3 pages."	☐
	2	"INTERNATIONAL BUREAU,"Written opinion" issued in connection with International Patent application No. PCT/ US2013/046372, mailed August 26, 2013, 4 pages."	☐
	3	"International Search Report, issued by the International Searching Authority in connection with PCT Application No. PCT/US2013/046386, on September 30, 2013, 3 pages."	☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2442
	Examiner Name	Oleg Survillo
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	4	M. NILSSON., "ID3 Tag Version 2", March 26, 1998, 28 Pages.	☐
	5	"Non-final Office Action in connection with Application No. 13/619,237, mailed April 10, 2013", United States Patent and Trademark Office, 10 pages.	☐
	6	"Non-Final Office Action mailed 25 November 2013 for U.S. Application No. 13/533105 filed 26 June 2012."	☐
	7	"Notice of Allowability in Application No. 13/619,237", September 6, 2013, 4 pages.	☐
	8	"Polycom Conference Composer manual: copyright 2001".	☐
	9	"Written Opinion of the International Searching Authority. issued by the International Searching Authority in connection with PCT Application No. PCT/US2013/046386, on September 30, 2013, 6 pages."	☐
	10	"Yamaha DME 32 manual: copyright 2001".	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2442
	Examiner Name	Oleg Survillo
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Robert J. Irvine III/	Date (YYYY-MM-DD)	2014-02-27
Name/Print	Robert J. Irvine III	Registration Number	41865

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt	
EFS ID:	18327250
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 (13-0401)
Receipt Date:	04-MAR-2014
Filing Date:	17-APR-2013
Time Stamp:	14:41:25
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Foreign Reference	WO2005013047A2.pdf	9269608 953769d47f04a9db684f7db18a01a31e0aa9bbb1	no	143
Warnings:					
Information:					

2	Non Patent Literature	ISR-PCTUS2013-046372-82613.pdf	128842 f665a6c13841b22a8c7fd50160dbdfd6f69e00ee	no	3
Warnings:					
Information:					
3	Non Patent Literature	WrittenOpinionPCTUS2013-046372-82613.pdf	166042 aed54d2da7fafd5488d9e48565bde415cddf01fc	no	4
Warnings:					
Information:					
4	Non Patent Literature	ISR-PCT-US2013-046386-93013.pdf	128924 c360063c4cfcb334e572ef3064f7682218404801	no	3
Warnings:					
Information:					
5	Non Patent Literature	Nilsson_1998.pdf	1104759 6ee8bd721afb9a6b25c0a0043a35bc92a3ebe796	no	28
Warnings:					
Information:					
6	Non Patent Literature	13-619237-Non-FinalOAmailed41013.pdf	329715 29e4c915f6e0cdc22f63f1be82b2b404ac6906ac	no	10
Warnings:					
Information:					
7	Non Patent Literature	13-533105-Non-FinalOA112513.pdf	559192 fce2751da5ceb20902ab46325e70f3c2e3b9e0ad	no	19
Warnings:					
Information:					
8	Non Patent Literature	13-619237-NoticeAllowability9613.pdf	163025 33f07f9b3ef4c7485c6c8541c5c0893e12b8cf4f	no	4
Warnings:					
Information:					
9	Non Patent Literature	PolycomConfManual_2001.pdf	2282175 03c0aed8dbcb96508286ab4466d6a49e748fde05	no	29
Warnings:					
Information:					
10	Non Patent Literature	WrittenOpinionPCT-US2013-046386-93013.pdf	232837 6aa1639ae8799b1400f1fda0fec57d72587c8eaa	no	6
Warnings:					
Information:					

11	Non Patent Literature	YAMAHA_DME_32_OWNERS_MANUAL.pdf	22132154 e919a88a90129a3511de896ee616d56e4eb83127	no	296
Warnings:					
Information:					
12	Information Disclosure Statement (IDS) Form (SB08)	SNS-71301US10-Supplemental_IDS.pdf	709206 27c898a8721ac419dce7157fb68a586ffc70dadf	no	10
Warnings:					
Information:					
		Total Files Size (in bytes):	37206479		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
10 February 2005 (10.02.2005)

PCT

(10) International Publication Number
WO 2005/013047 A2

(51) International Patent Classification[7]: G06F

(21) International Application Number:
PCT/US2004/023102

(22) International Filing Date: 2 July 2004 (02.07.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/490,768 28 July 2003 (28.07.2003) US
10/816,217 1 April 2004 (01.04.2004) US

(71) Applicant *(for all designated States except US)*: **SONOS, INC.** [US/US]; 506 Chapala, Santa Barbara, CA 93101 (US).

(72) Inventor; and
(75) Inventor/Applicant *(for US only)*: **MILLINGTON, Nicholas, A., J.** [US/US]; 1221 Salsipuedes Street, Santa Barbara, CA 93103 (US).

(74) Agent: **JORDAN, Richard, A.**; P.O. Box 81363, Wellesley Hills, MA 02481-0004 (US).

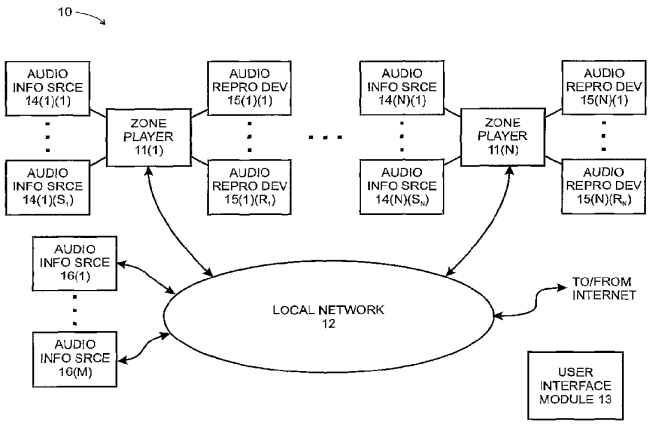
(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES



(57) Abstract: A system is described for maintaining synchrony of operations among a plurality of devices hat have independent clocking arrangements. The system includes a task distribution device that distributes tasks to a synchrony group comprising a plurality of devices that are to perform the tasks distributed by the task distribution device in synchrony. The task distribution device distributes each task to the members of the synchrony group over a network. Each task is associated with a time stamp that indicates a time, relative to a clock maintained by the task distribution device, at which the members of the synchrony group are to execute the task. Each member of the synchrony group periodically obtains from the task distribution device an indication of the current time indicated by its clock, determines a time differential between the task distribution device's clock and its respective clock and determines therefrom a time at which, according to its respective clock, the time stamp indicates that it is to execute the task.

WO 2005/013047 A2

System and Method For Synchronizing Operations Among A Plurality Of Independently Clocked Digital Data Processing Devices

Field Of The Invention

The present invention relates generally to the field of digital data processing devices, and more particularly to systems and methods for synchronizing operations among a plurality of independently-clocked digital data processing devices. The invention is embodied in a system for synchronizing operations among a plurality of devices, in relation to information that is provided by a common source. One embodiment of the invention enables synchronizing of audio playback as among two or more audio playback devices that receive audio information from a common information source, or channel.

More generally, the invention relates to the field of arrangements that synchronize output generated by a number of output generators, including audio output, video output, combinations of audio and video, as well as other types of output as will be appreciated by those skilled in the art, provided by a common channel. Generally, the invention will find utility in connection with any type of information for which synchrony among independently-clocked devices is desired.

Background Of The Invention

There are a number of circumstances under which it is desirable to maintain synchrony of operations among a plurality of independently-clocked digital data processing devices in relation to, for example, information that is provided thereto by a common source. For example, systems are being developed in which one audio information source can distribute audio information in digital form to a number of audio playback devices for playback. The audio playback devices receive the digital information and convert it to analog form for playback. The audio playback devices may be located in the same room or they may be distributed in different rooms in a residence such as a house or an apartment, in different offices in an office building, or the like. For example, in a system installed in a residence, one audio playback device may be located in a living room, while another audio playback device is be located in a kitchen, and yet other audio playback devices may be located in various bedrooms of a house. In such an arrangement, the audio information that is distributed to various audio playback devices may relate to the same audio program, or the information may relate to different audio programs. If the audio information source provides audio information relating to the same audio program to two or more audio playback devices at the same time, the audio playback devices will generally contemporaneously play the same program. For

example, if the audio information source provides audio information to audio playback devices located in the living room and kitchen in a house at the same time, they will generally contemporaneously play the same program.

One problem that can arise is to ensure that, if two or more audio playback devices are contemporaneously attempting to play back the same audio program, they do so simultaneously. Small differences in the audio playback devices' start times and/or playback speeds can be perceived by a listener as an echo effect, and larger differences can be very annoying. Differences can arise because for a number of reasons, including delays in the transfer of audio information over the network. Such delays can differ as among the various audio playback devices for a variety of reasons, including where they are connected into the network, message traffic and other reasons as will be apparent to those skilled in the art.

Another problem arises from the following. When an audio playback device converts the digital audio information from digital to analog form, it does so using a clock that provides timing information. Generally, the audio playback devices that are being developed have independent clocks, and, if they are not clocking at precisely the same rate, the audio playback provided by the various devices can get out of synchronization.

SUMMARY OF THE INVENTION

The invention provides a new and improved system and method for synchronizing operations among a number of digital data processing devices that are regulated by independent clocking devices. Generally, the invention will find utility in connection with any type of information for which synchrony among devices connected to a network is desired. The invention is described in connection with a plurality of audio playback devices that receive digital audio information that is to be played back in synchrony, but it will be appreciated that the invention can find usefulness in connection with any kind of information for which coordination among devices that have independent clocking devices would find utility.

In brief summary, the invention provides, in one aspect, a system for maintaining synchrony of operations among a plurality of devices that have independent clocking arrangements. The system includes a task distribution device that distributes tasks to a synchrony group comprising a plurality of devices that are to perform the tasks distributed by the task distribution device in synchrony. The task distribution device distributes each task to the members of the synchrony group over a network. Each task is associated with a time stamp that indicates a time, relative to a clock maintained by the task distribution device, at which the members of the synchrony group are to execute the task. Each

member of the synchrony group periodically obtains from the task distribution device an indication of the current time indicated by its clock, determines a time differential between the task distribution device's clock and its respective clock and determines therefrom a time at which, according to its respective clock, the time stamp indicates that it is to execute the task.

In one embodiment, the tasks that are distributed include audio information for an audio track that is to be played by all of the devices comprising the synchrony group synchronously. The audio track is divided into a series of frames, each of which is associated with a time stamp indicating the time, relative to the clock maintained by an audio information channel device, which, in that embodiment, serves as the task distribution device, at which the members of the synchrony group are to play the respective frame. Each member of the synchrony group, using a very accurate protocol, periodically obtains the time indicated by the audio information channel device, and determines a differential between the time as indicated by its local clock and the audio information channel device's clock. The member uses the differential and the time as indicated by the time stamp to determine the time, relative to its local clock, at which it is to play the respective frame. The members of the synchrony group do this for all of the frames, and accordingly are able to play the frames in synchrony.

BRIEF DESCRIPTION OF THE DRAWINGS

This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

FIG. 1 schematically depicts an illustrative networked audio system, constructed in accordance with the invention;

FIG. 2 schematically depicts a functional block diagram of a synchrony group utilizing a plurality of zone players formed within the networked audio system depicted in FIG. 1;

FIG. 2A schematically depicts two synchrony groups, illustrating how a member of one synchrony group can provide audio information to the members of another synchrony group;

FIG. 3 depicts an functional block diagram of a zone player for use in the networked audio system depicted in FIG. 1; and

FIG. 4 is useful in understanding an digital audio information framing methodology useful in the network audio system depicted in FIG. 1.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

FIG. 1 depicts an illustrative network audio system 10 constructed in accordance with the invention. With reference to FIG. 1, the network audio system 10 includes a plurality of zone players 11(1) through 11(N) (generally identified by reference numeral 11(n)) interconnected by a local network 12, all of which operate under control of one or more user interface modules generally identified by reference numeral 13. One or more of the zone players 11(n) may also be connected to one or more audio information sources, which will generally be identified herein by reference numeral 14(n)(s), and/or one or more audio reproduction devices, which will generally be identified by reference numeral 15(n)(r). In the reference numeral 14(n)(s), index "n" refers to the index "n" of the zone player 11(n) to which the audio information source is connected, and the index "s" ($s=1,...,S_n$) refers to the "s-th" audio information source connected to that "n-th" zone player 11(n). Thus, if, for example, a zone player 11(n) is connected to four audio information sources 14(n)(1) through 14(n)(4), the audio information sources may be generally identified by reference numeral 14(n)(s), with $S_n=4$. It will be appreciated that the number of audio information sources S_n may vary as among the various zone players 11(n), and some zone players may not have any audio information sources connected thereto. Similarly, in the reference numeral 15(n)(r), index "n" refers to the index "n" of the zone player 11(n) to which the audio reproduction device is connected, and the index "r" ($r=1,...,R_n$) refers to the "r-th" audio information source connected to that "n-th" zone player 11(n). In addition to the audio information sources 14(n)(s), the network audio system 10 may include one or more audio information sources 16(1) through 16(M) connected through appropriate network interface devices (not separately shown) to the local network 12. Furthermore, the local network may include one or more network interface devices (also not separately shown) that are configured to connect the local network 12 to other networks, including a wide area network such as the Internet, the public switched telephony network (PSTN) or other networks as will be apparent to those skilled in the art, over which connections to audio information sources may be established.

The zone players 11(n) associated with system 10 may be distributed throughout an establishment such as residence, an office complex, a hotel, a conference hall, an amphitheater or auditorium, or other types of establishments as will be apparent to those skilled in the art or the like. For example, if the zone players 11(n) and their associated audio information source(s) and/or audio reproduction device(s) are distributed throughout a residence, one, such as zone player 11(1) and its associated audio information source(s) and audio reproduction device(s) may be located in a living room, another may be located in a kitchen, another may be located in a dining room, and yet others

may be located in respective bedrooms, to selectively provide entertainment in those rooms. On the other hand, if the zone players 11(n) and their associated audio information source(s) and/or audio reproduction device(s) are distributed throughout an office complex, one may, for example, be provided in each office to selectively provide entertainment to the employees in the respective offices. Similarly, if the zone players 11(n) and associated audio information source(s) and/or audio reproduction device(s) are used in a hotel, they may be distributed throughout the rooms to provide entertainment to the guests. Similar arrangements may be used with zone players 11(n) and associated audio information source(s) and/or audio reproduction device(s) used in an amphitheater or auditorium. Other arrangements in other types of environments will be apparent to those skilled in the art. In each case, the zone players 11(n) can be used to selectively provide entertainment in the respective locations, as will be described below.

The audio information sources 14(n)(s) and 16(m) may be any of a number of types of conventional sources of audio information, including, for example, compact disc ("CD") players, AM and/or FM radio receivers, analog or digital tape cassette players, analog record turntables and the like. In addition, the audio information sources 14(n)(s) and 16(m) may comprise digital audio files stored locally on, for example, personal computers (PCs), personal digital assistants (PDAs), or similar devices capable of storing digital information in volatile or non-volatile form. As noted above, the local network 12 may also have an interface (not shown) to a wide area network, over which the network audio system 10 can obtain audio information. Moreover, one or more of the audio information sources 14(n)(s) may also comprise an interface to a wide area network such as the Internet, the public switched telephony network (PSTN) or any other source of audio information. In addition, one or more of the audio information sources 14(n)(s) and 16(m) may comprise interfaces to radio services delivered over, for example, satellite. Audio information obtained over the wide area network may comprise, for example, streaming digital audio information such as Internet radio, digital audio files stored on servers, and other types of audio information and sources as will be appreciated by those skilled in the art. Other arrangements and other types of audio information sources will be apparent to those skilled in the art.

Generally, the audio information sources 14(n)(s) and 16(m) provide audio information associated with audio programs to the zone players for playback. A zone player that receives audio information from an audio information source 14(n)(s) that is connected thereto can provide playback and/or forward the audio information, along with playback timing information, over the local network 12 to other zone players for playback. Similarly, each audio information source 16(m)

that is not directly connected to a zone player can transmit audio information over the network 12 to any zone player 11(n) for playback. In addition, as will be explained in detail below, the respective zone player 11(n) can transmit the audio information that it receives either from an audio information source 14(n)(s) connected thereto, or from an audio information source 16(m), to selected ones of the other zone players 11(n'), 11(n''),... (n not equal to n', n'',...) for playback by those other zone players. The other zone players 11(n'), 11(n''),... to which the zone player 11(n) transmits the audio information for playback may be selected by a user using the user interface module 13. In that operation, the zone player 11(n) will transmit the audio information to the selected zone players 11(n'), 11(n''),... over the network 12. As will be described below in greater detail, the zone players 11(n), 11(n'), 11(n''),... operate such that the zone players 11(n'), 11(n''),... synchronize their playback of the audio program with the playback by the zone player 11(n), so that the zone players 11(n), 11(n'), 11(n'') provide the same audio program at the same time.

Users, using user interface module 13, may also enable different groupings or sets of zone players to provide audio playback of different audio programs synchronously. For example, a user, using a user interface module 13, may enable zone players 11(1) and 11(2) to play one audio program, audio information for which may be provided by, for example, one audio information source 14(1)(1). The same or a different user may, using the same or a different user interface module 13, enable zone players 11(4) and 11(5) to contemporaneously play another audio program, audio information for which may be provided by a second audio information source, such as audio information source 14(5)(2). Further, a user may enable zone player 11(3) to contemporaneously play yet another audio program, audio information for which may be provided by yet another audio information source, such as audio information source 16(1). As yet another possibility, a user may contemporaneously enable zone player 11(1) to provide audio information from an audio information source connected thereto, such as audio information source 14(1)(2), to another zone player, such as zone player 11(6) for playback.

In the following, the term "synchrony group" will be used to refer to a set of one or more zone players that are to play the same audio program synchronously. Thus, in the above example, zone players 11(1) and 11(2) comprise one synchrony group, zone player 11(3) comprises a second synchrony group, zone players 11(4) and 11(5) comprise a third synchrony group, and zone player 11(6) comprises yet a fourth synchrony group. Thus, while zone players 11(1) and 11(2) are playing the same audio program, they will play the audio program synchronously. Similarly, while zone players 11(4) and 11(5) are playing the same audio program, they will play the audio program

synchronously. On the other hand, zone players that are playing different audio programs may do so with unrelated timings. That is, for example, the timing with which zone players 11(1) and 11(2) play their audio program may have no relationship to the timing with which zone player 11(3), zone players 11(4) and 11(5), and zone player 11(6) play their audio programs. It will be appreciated that, since "synchrony group" is used to refer to sets of zone players that are playing the same audio program synchronously, zone player 11(1) will not be part of zone player 11(6)'s synchrony group, even though zone player 11(1) is providing the audio information for the audio program to zone player 11(6).

In the network audio system 10, the synchrony groups are not fixed. Users can enable them to be established and modified dynamically. Continuing with the above example, a user may enable the zone player 11(1) to begin providing playback of the audio program provided thereto by audio information source 14(1)(1), and subsequently enable zone player 11(2) to join the synchrony group. Similarly, a user may enable the zone player 11(5) to begin providing playback of the audio program provided thereto by audio information source 14(5)(2), and subsequently enable zone player 11(4) to join that synchrony group. In addition, a user may enable a zone player to leave a synchrony group and possibly join another synchrony group. For example, a user may enable the zone player 11(2) to leave the synchrony group with zone player 11(1), and join the synchrony group with zone player 11(6). As another possibility, the user may enable the zone player 11(1) to leave the synchrony group with zone player 11(2) and join the synchrony group with zone player 11(6). In connection with the last possibility, the zone player 11(1) can continue providing audio information from the audio information source 14(1)(1) to the zone player 11(2) for playback thereby.

A user, using the user interface module 13, can enable a zone player 11(n) that is currently not a member of a synchrony group to join a synchrony group, after which it will be enabled to play the audio program that is currently being played by that synchrony group. Similarly, a user, also using the user interface module 13, can enable a zone player 11(n) that is currently a member of one synchrony group, to disengage from that synchrony group and join another synchrony group, after which that zone player will be playing the audio program associated with the other synchrony group. For example, if a zone player 11(6) is currently not a member of any synchrony group, it, under control of the user interface module 13, can become a member of a synchrony group, after which it will play the audio program being played by the other members of the synchrony group, in synchrony with the other members of the synchrony group. In becoming a member of the synchrony group, zone player 11(6) can notify the zone player that is the master device for the synchrony group

that it wishes to become a member of its synchrony group, after which that zone player will also transmit audio information associated with the audio program, as well as timing information, to the zone player 11(6). As the zone player 11(6) receives the audio information and the timing information from the master device, it will play the audio information with the timing indicated by the timing information, which will enable the zone player 11(6) to play the audio program in synchrony with the other zone player(s) in the synchrony group.

Similarly, if a user, using the user interface module 13, enables a zone player 11(n) associated with a synchrony group to disengage from that synchrony group, and if the zone player 11(n) is not the master device of the synchrony group, the zone player 11(n) can notify the master device, after which the master device can terminate transmission of the audio information and timing information to the zone player 11(n). If the user also enables the zone player 11(n) to begin playing another audio program using audio information from an audio information source 14(n)(s) connected thereto, it will acquire the audio information from the audio information source 14(n)(s) and initiate playback thereof. If the user enables another zone player 11(n') to join the synchrony group associated with zone player 11(n), operations in connection therewith can proceed as described immediately above.

As yet another possibility, if a user, using the user interface module 13, enables a zone player 11(n) associated with a synchrony group to disengage from that synchrony group and join another synchrony group, and if the zone player is not the master device of the synchrony group from which it is disengaging, the zone player 11(n) can notify the master device of the synchrony group from which it is disengaging, after which that zone player will terminate transmission of audio information and timing information to the zone player 11(n) that is disengaging. Contemporaneously, the zone player 11(n) can notify the master device of the synchrony group that it (that is, zone player 11(n)) is joining, after which the master device can begin transmission of audio information and timing information to that zone player 11(n). The zone player 11(n) can thereafter begin playback of the audio program defined by the audio information, in accordance with the timing information so that the zone player 11(n) will play the audio program in synchrony with the master device.

As yet another possibility, a user, using the user interface module 13, may enable a zone player 11(n) that is not associated with a synchrony group, to begin playing an audio program using audio information provided to it by an audio information source 14(n)(s) connected thereto. In that case, the user, also using the user interface module 13 or a user interface device that is specific to the audio information source 14(n)(s), can enable the audio information source 14(n)(s) to provide audio information to the zone player 11(n). After the zone player 11(n) has begun playback, or

contemporaneously therewith, the user, using the user interface module 13, can enable other zone players 11(n'), 11(n''),... to join zone player 11(n)'s synchrony group and enable that zone player 11(n) to transmit audio information and timing information thereto as described above, to facilitate synchronous playback of the audio program by the other zone players 11(n'), 11(n'')....

A user can use the user interface module 13 to control other aspects of the network audio system 10, including but not limited to the selection of the audio information source 14(n)(s) that a particular zone player 11(n) is to utilize, the volume of the audio playback, and so forth. In addition, a user may use the user interface module 13 to turn audio information source(s) 14(n)(s) on and off and to enable them to provide audio information to the respective zone players 11(n).

Operations performed by the various devices associated with a synchrony group will be described in connection with FIG. 2, which schematically depicts a functional block diagram of a synchrony group in the network audio system 10 described above in connection with FIG. 1. With reference to FIG. 2, a synchrony group 20 includes a master device 21 and zero or more slave devices 22(1) through 22 (G) (generally identified by reference numeral 22(g)), all of which synchronously play an audio program provided by an audio information channel device 23. Each of the master device 21, slave devices 22(g) and audio information channel device 23 utilizes a zone player 11(n) depicted in FIG. 1, although it will be clear from the description below that a zone player may be utilized both for the audio information channel device for the synchrony group 20, and the master device 21 or a slave device 22(g) of the synchrony group 20. As will be described below in more detail, the audio information channel device 23 obtains the audio information for the audio program from an audio information source, adds playback timing information, and transmits the combined audio and playback timing information to the master device 21 and slave devices 22(g) over the network 12 for playback. The playback timing information that is provided with the audio information, together with clock timing information provided by the audio information channel device 23 to the various devices 21 and 22(g) as will be described below, enables the master device 21 and slave devices 22(g) of the synchrony group 20 to play the audio information simultaneously.

The master device 21 and the slave devices 22(g) receive the audio and playback timing information, as well as the clock timing information, that are provided by the audio information channel device 23, and play back the audio program defined by the audio information. The master device 21 is also the member of the synchrony group 20 that communicates with the user interface module 13 and that controls the operations of the slave devices 22(g) in the synchrony group 20. In addition, the master device 21 controls the operations of the audio information channel device 23 that

provides the audio and playback timing information for the synchrony group 20. Generally, the initial master device 21 for the synchrony group will be the first zone player 11(n) that a user wishes to play an audio program. However, as will be described below, the zone player 11(n) that operates as the master device 21 can be migrated from one zone player 11(n) to another zone player 11(n'), which preferably will be a zone player that is currently operating as a slave device 22(g) in the synchrony group.

In addition, under certain circumstances, as will be described below, the zone player 11(n) that operates as the audio information channel device 23 can be migrated from one zone player to another zone player, which also will preferably will be a zone player that is currently operating as a member of the synchrony group 20. It will be appreciated that the zone player that operates as the master device 21 can be migrated to another zone player independently of the migration of the audio information channel device 23. For example, if one zone player 11(n) is operating as both the master device 21 and the audio information channel device 23 for a synchrony group 20, the master device 21 can be migrated to another zone player 11(n') while the zone player 11(n) is still operating as the audio information channel device 23. Similarly, if one zone player 11(n) is operating as both the master device 21 and the audio information channel device 23 for a synchrony group 20, the audio information channel device 23 can be migrated to another zone player 11(n') while the zone player 11(n) is still operating as the master device 21. In addition, if one zone player 11(n) is operating as both the master device 21 and the audio information channel device 23 for a synchrony group 20, the master device 21 can be migrated to another zone player 11(n') and the audio information channel device can be migrated to a third zone player 11(n").

The master device 21 receives control information from the user interface module 13 for controlling the synchrony group 20 and provides status information indicating the operational status of the synchrony group to the user interface module 13. Generally, the control information from the user interface module 13 enables the master device 21 to, in turn, enable the audio information channel device 23 to provide audio and playback timing information to the synchrony group to enable the devices 21 and 22(g) that are members of the synchrony group 20 to play the audio program synchronously. In addition, the control information from the user interface module 13 enables the master device 21 to, in turn, enable other zone players to join the synchrony group as slave devices 22(g) and to enable slave devices 22(g) to disengage from the synchrony group. Control information from the user interface module 13 can also enable the zone player 11(n) that is currently operating as the master device 21 to disengage from the synchrony group, but prior to

doing so that zone player will enable the master device 21 to transfer from that zone player 11(n) to another zone player11(n'), preferably to a zone player 11(n') that is currently a slave device 22(g) in the synchrony group 20. The control information from the user interface module 13 can also enable the master device 21 to adjust its playback volume and to enable individual ones of the various slave devices 22(g) to adjust their playback volumes. In addition, the control information from the user interface module 13 can enable the synchrony group 20 to terminate playing of a current track of the audio program and skip to the next track, and to re-order tracks in a play list of tracks defining the audio program that is to be played by the synchrony group 20.

The status information that the master device 21 may provide to the user interface module 13 can include such information as a name or other identifier for the track of the audio work that is currently being played, the names or other identifiers for upcoming tracks, the identifier of the zone player 11(n) that is currently operating as the master device 21, and identifiers of the zone players that are currently operating as slave devices 22(g). In one embodiment, the user interface module 13 includes a display (not separately shown) that can display the status information to the user.

It will be appreciated that the zone player 11(n) that is operating as the audio information channel device 23 for one synchrony group may also comprise the master device 21 or any of the slave devices 22(g) in another synchrony group. This may occur if, for example, the audio information source that is to provide the audio information that is to be played by the one synchrony group is connected to a zone player also being utilized as the master device or a slave device for the other synchrony group. This will be schematically illustrated below in connection with FIG. 2A. Since, as noted above, the zone player 11(n) that is operating as the audio information channel device 23 for the synchrony group 20 may also be operating as a master device 21 or slave device 22(g) for another synchrony group, it can also be connected to one or more audio reproduction devices 15(n)(r), although that is not depicted in FIG. 2. Since the master device 21 and slave devices 22(g) are all to provide playback of the audio program, they will be connected to respective audio reproduction devices 15(n)(r). Furthermore, it will be appreciated that one or more of the zone players 11(n) that operate as the master device 21 and slave devices 22(g) in synchrony group 20 may also operate as an audio information channel device for that synchrony group or for another synchrony group and so they may be connected to one or more audio information sources 14(n)(s), although that is also not depicted in FIG. 2. In addition, it will be appreciated that a zone player 11(n) can also operate as a audio information channel device 23 for multiple synchrony groups.

If the audio information channel device 23 does not utilize the same zone player as the master device 21, the master device 21 controls the audio information channel device by exchanging control information over the network 12 with the audio information channel device 23. The control information is represented in FIG. 2 by the arrow labeled CHAN_DEV_CTRL_INFO. The control information that the master device 21 provides to the audio information channel device 23 will generally depend on the nature of the audio information source that is to provide the audio information for the audio program that is to be played and the operation to be enabled by the control information. If, for example, the audio information source is a conventional compact disc, tape, or record player, broadcast radio receiver, or the like, which is connected to a zone player 11(n), the master device 21 may merely enable the zone player serving as the audio information channel device 23 to receive the audio information for the program from the audio information source. It will be appreciated that, if the audio information is not in digital form, the audio information channel device 23 will convert it to digital form and provide the digitized audio information, along with the playback timing information, to the master device 21 and slave devices 22(g).

On the other hand, if the audio information source is, for example, a digital data storage device, such as may be on a personal computer or similar device, the master device 21 can provide a play list to the audio information channel device 23 that identifies one or more files containing the audio information for the audio program. In that case, the audio information channel device 23 can retrieve the files from the digital data storage device and provide them, along with the playback timing information, to the master device 21 and the slave devices 22(g). It will be appreciated that, in this case, the audio information source may be directly connected to the audio information channel device 23, as, for example, an audio information source 14(n)(s), or it may comprise an audio information source 16(m) connected to the network 12. As a further alternative, if the audio information source is a source available over the wide area network, the master device 21 can provide a play list comprising a list of web addresses identifying the files containing the audio information for the audio program that is to be played, and in that connection the audio information channel device 23 can initiate a retrieval of the files over the wide area network. As yet another alternative, if the audio information source is a source of streaming audio received over the wide area network, the master device 21 can provide a network address from which the streaming audio can be received. Other arrangements by which the master device 21 can control the audio information channel device 23 will be apparent to those skilled in the art.

The master device 21 can also provide control information to the synchrony group's audio information channel device 23 to enable a migration from one zone player 11(n) to another zone player 11(n'). This may occur if, for example, the audio information source is one of audio information sources 16 or a source accessible over the wide area network via the network 12. The master device 21 can enable migration of the audio information channel device 23 for several reasons, including, for example, to reduce the loading of the zone player 11(n), to improve latency of message transmission in the network 12, and other reasons as will be appreciated by those skilled in the art.

As noted above, the audio information channel device 23 provides audio and playback timing information for the synchrony group to enable the master device 21 and slave devices 22(g) to play the audio program synchronously. Details of the audio and playback timing information will be described in detail below in connection with FIGS. 3 and 4, but, in brief, the audio information channel device 23 transmits the audio and playback timing information in messages over the network 12 using a multi-cast message transmission methodology. In that methodology, the audio information channel device 23 will transmit the audio and playback timing information in a series of messages, with each message being received by all of the zone players 11(n) comprising the synchrony group 20, that is, by the master device 21 and the slave devices 22(g). Each of the messages includes a multi-cast address, which the master device 21 and slave devices 22(g) will monitor and, when they detect a message with that address, they will receive and use the contents of the message. The audio and playback timing information is represented in FIG. 2 by the arrow labeled "AUD+PBTIME_INF0," which has a single tail, representing a source for the information at the audio information channel device 23, and multiple arrowheads representing the destinations of the information, with one arrowhead extending to the master device 21 and other arrowheads extending to each of the slave devices 22(g) in the synchrony group 20. The audio information channel device 23 may make use of any convenient multi-cast message transmission methodology in transmitting the audio and playback timing information to the synchrony group 20. As will be described in detail in connection with FIG. 4, the audio and playback timing information is in the form of a series of frames, with each frame having a time stamp. The time stamp indicates a time, relative to the time indicated by a clock maintained by the audio information channel device 23, at which the frame is to be played. Depending on the size or sizes of the messages used in the selected multi-cast message transmission methodology and the size or sizes of the frames, a message may contain one frame, or multiple frames, or, alternatively, a frame may extend across several messages.

The audio information channel device 23 also provides clock time information to the master device 21 and each of the slave devices 22(g) individually over network 12 using a highly accurate clock time information transmission methodology. The distribution of the clock time information is represented in FIG. 2 by the arrows labeled "AICD_CLK_INF (M)" (in the case of the clock time information provided to the master device 21) and "AICD_CLK_INF (S_1)" through "AICD_CLK_INF (S_G)" (in the case of audio information channel device clock information provided to the slave devices 22(g)). In one embodiment, the master device 21 and slave devices 22(g) make use of the well-known SNTP (Simple Network Time Protocol) to obtain current clock time information from the audio information channel device 23. The SNTP makes use of a unicast message transfer methodology, in which one device, such as the audio information channel device 23, provides clock time information to a specific other device, such as the master device 21 or a slave device 22(g), using the other device's network, or unicast, address. Each of the master device 21 and slave devices 22(g) will periodically initiate SNTP transactions with the audio information channel device 23 to obtain the clock time information from the audio information channel device 23. As will be described below in more detail, the master device 21 and each slave device 22(g) make use of the clock time information to determine the time differential between the time indicated by the audio information channel device's clock and the time indicated by its respective clock, and use that time differential value, along with the playback time information associated with the audio information and the respective device's local time as indicated by its clock to determine when the various frames are to be played. This enables the master device 21 and the slave devices 22(g) in the synchrony group 20 to play the respective frames simultaneously.

As noted above, the control information provided by the user to the master device 21 through the user interface module 13 can also enable the master device 21 to, in turn, enable another zone player 11(n') to join the synchrony group as a new slave device 22(g). In that operation, the user interface module 13 will provide control information, including the identification of the zone player 11(n') that is to join the synchrony group to the master device 21. After it receives the identification of the zone player 11(n') that is to join the synchrony group, the master device 21 will exchange control information, which is represented in FIG. 2 by the arrows labeled SLV_DEV_CTRL_INF (S_1) through SLV_DEV_CTRL_INF (S_G) group slave control information, over the network 12 with the zone player 11(n') that is identified in the control information from the user interface module 13. The control information that the master device 21 provides to the new zone player 11(n') includes the network address of the zone player 11(n) that is operating as the audio information channel

device 23 for the synchrony group, as well as the multi-cast address that the audio information channel device 23 is using to broadcast the audio and playback timing information over the network. The zone player that is to operate as the new slave device 22(g') uses the multi-cast address to begin receiving the multi-cast messages that contain the audio information for the audio program being played by the synchrony group.

It will be appreciated that, if the zone player 11(n) that is operating as the master device 21 for the synchrony group 20 is also operating the audio information channel device 23, and if there are no slave devices 22(g) in the synchrony group 20, the audio information channel device 23 may not be transmitting audio and playback timing information over the network. In that case, if the new slave device 22(g') is the first slave device in the synchrony group, the zone player 11(n) that is operating as both the master device 21 and audio information channel device 23, can begin transmitting the audio and playback timing information over the network 12 when the slave device 22(g') is added to the synchrony group 20. The zone player 11(n) can maintain a count of the number of slave devices 22(g) in the synchrony group 20 as they join and disengage, and, if the number drops to zero, it can stop transmitting the audio and playback timing information over the network 12 to reduce the message traffic over the network 12.

The new slave device 22(g') added to the synchrony group 20 uses the network address of the audio information channel device 23 for several purposes. In particular, the new slave device 22(g') will, like the master device 21 (assuming the zone player 11(n) operating as the master device 21 is not also the audio information channel device 23), engage in SNTP transactions with the audio information channel device 23 to obtain the clock timing information from the audio information channel device 23. In addition, the new slave device 22(g') can notify the audio information channel device 23 that it is a new slave device 22(g') for the synchrony group 20 and provide the audio information channel device 23 with its network address. As will be described below, in one embodiment, particularly in connection with audio information obtained from a source, such as a digital data storage device, which can provide audio information at a rate that is faster than the rate at which it will be played, the audio information channel device 23 will buffer audio and timing information and broadcast it over the network 12 to the synchrony group 20 generally at a rate at which it is provided by the source. Accordingly, when a new slave device 22(g') joins the synchrony group 20, the playback timing information may indicate that the audio information that is currently being broadcast by the audio information channel device 23 using the multi-cast methodology is to be played back some time in the future. To reduce the delay with which the new slave device 22(g')

will begin playback, the audio information channel device 23 can also retransmit previously transmitted audio and timing information that it had buffered to the new slave device 22(g') using the unicast network address of the slave device 22(g').

The master device 21 can also use the slave device control information exchanged with the slave devices 22(g) for other purposes. For example, the master device 21 can use the slave device control information to initiate a migration of the master from its zone player 11(n) to another zone player 11(n'). This may occur for any of a number of reasons, including, for example, that the master device 21 is terminating playback by it of the audio program and is leaving the synchrony group 20, but one or more of the other devices in the synchrony group is to continue playing the audio program. The master device 21 may also want to initiate a migration if it is overloaded, which can occur if, for example, the zone player 11(n) that is the master device 21 for its synchrony group is also operating as an audio information channel device 23 for another synchrony group.

The user can also use the user interface module 13 to adjust playback volume by the individual zone players 11(n) comprising the synchrony group. In that operation, the user interface module 13 provides information identifying the particular device whose volume is to be adjusted, and the level at which the volume is to be set to the master device 21. If the device whose volume is to be adjusted is the master device 21, the master device 21 can adjust its volume according to the information that it receives from the user interface module 13. On the other hand, if the device whose volume is to be adjusted is a slave device 22(g), the master device 21 can provide group slave control information to the respective slave device 22(g), to enable it to adjust its volume.

The user can also use the user interface module 13 to enable a synchrony group 20 to cancel playing of the track in an audio program that is currently being played, and to proceed immediately to the next track. This may occur, for example, if the tracks for the program is in the form of a series of digital audio information files, and the user wishes to cancel playback of the track that is defined by one of the files. In that case, when the master device 21 receives the command to cancel playback of the current track, it will provide channel device control information to the audio information channel device 23 so indicating. In response, the audio information channel device 23 inserts control information into the audio and playback timing information, which will be referred to as a "resynchronize" command. In addition, the audio information channel device 23 will begin transmitting audio information for the next track, with timing information to enable it to be played immediately. The resynchronize command can also enable playback of a track to be cancelled before it has been played. Details of these operations will be described below.

As noted above, there may be multiple synchrony groups in the network audio system 10, and further that, for example, a zone player 11(n) may operate both as a master device 21 or a slave device 22(g) in one synchrony group, and as the audio information channel device 23 providing audio and playback timing information and clock timing information for another synchrony group. An illustrative arrangement of this will be described in connection with FIG. 2A. With reference to FIG. 2A, that FIG. depicts elements of two synchrony groups, identified by reference numerals 20(1) and 20(2), respectively. For clarity, FIG. 2A does not show a number of elements, the presence of which would be evident from FIGS. 1 and 2 as described above. For example, FIG. 2A does not depict the audio information sources from which audio information is obtained for the synchrony groups or the audio reproduction devices that are used to produce sound for the master and slave devices, which are depicted in both FIGS. 1 and 2. In addition, FIG. 2A does not depict arrows that represent control information provided by the respective master devices to the slave devices in the respective synchrony groups, or to the audio information channel devices that provide audio and timing information for the respective synchrony groups, which are depicted in FIG. 2. In addition, FIG. 2A does not depict the arrows that represent the clock timing information provided by the audio information channel devices to the respective members of the respective synchrony groups, which are also depicted in FIG. 2. As will be noted below, however, FIG. 2A does depict arrows representing the audio and playback timing information provided by the respective audio information channel devices for the respective synchrony groups 20(1), 20(2), to the master and slave devices comprising the respective synchrony groups 20(1), 20(2).

Each synchrony group 20(1), 20(2) comprises elements of a number of zone players. A functional block diagram of a zone player will be described below in connection with FIG. 3. Synchrony group 20(1) includes a master device 21(1) and "K" slave devices 22(1)(1) through 22(K)(1) (the index "1" in reference numeral 21(1) and the last index in reference numeral 22(1)(1) through 21(K)(1) corresponds to the index of the synchrony group 20(1) to which they belong) utilize zone players 11(1) and 11(K+1) respectively. Similarly, synchrony group 20(2) includes a master device 21(2) and "L" slave devices 22(1)(2) through 22(L)(2) that utilize zone players 11(K+2) through 11(K+L+2). In the illustrative arrangement depicted in FIG. 2A, both synchrony groups 20(1) and 20(2) are controlled by the user interface module 13, which can provide control information to, and receive status information from, the master devices 21(1) and 21(2) independently. It will be appreciated that separate user interface modules may be provided to

provide control information to, and receive status information from, the respective master devices 21(1), 21(2).

As noted above, the slave device 22(1)(2) in synchrony group 20(2) utilizes zone player 11(K+3). In the illustrative arrangement depicted in FIG. 2A, the audio information channel device 23(1) that provides audio and playback timing information to the master and slave devices 21(1), 22(1)(1),..., 22(K)(1) of synchrony group 20(1) also utilizes zone player 11(K+3). As noted above, this may occur if, for example, the audio information source that is to provide audio information to be played by the synchrony group 20(1) is connected to the zone player 11(K+3). Thus, when the master device 21(1) of synchrony group 20(1) exchanges channel device control information with the audio information channel device 23(1), it is effectively exchanging channel device control information with the zone player 11(K+3). Similarly, when the master and slave devices 21(1), 22(1)(1),..., 22(K)(1) of synchrony group 20(1) receive audio and playback timing information, as well as clock timing information, from the audio information channel device 23(1), they are effectively receiving the information from the zone player 11(K+3). FIG. 2A depicts a multi-headed arrow representing audio and playback timing information transmitted by the zone player 11(K+3), as audio information channel device 23(1), to the master and slave devices 21(1), 22(1)(1),...,11(K)(1) comprising synchrony group 20(1).

On the other hand, in the illustrative arrangement depicted in FIG. 2A, the synchrony group 20(2) utilizes a zone player 11(K+L+3) as its audio information channel device 23(2). As with synchrony group 20(1), when the master device 21(2) of synchrony group 20(2) exchanges channel device control information with the audio information channel device 23(2), it is effectively exchanging channel device control information with the zone player 11(K+L+3). Similarly, when the master and slave devices 21(2), 22(1)(2),..., 22(L)(2) of synchrony group 20(2) receive audio and playback timing information, as well as clock timing information, from the audio information channel device 23(2), they are effectively receiving the information from the zone player 11(K+L+3). FIG. 2A depicts a multi-headed arrow representing audio and playback timing information transmitted by the zone player 11(K+3) as audio information channel device 23(2) to the master and slave devices 21(2), 22(1)(2),...,22(L)(2) comprising synchrony group 20(2).

In the illustrative arrangement depicted in FIG. 2A, zone player 11(K+L+3), which is the audio information channel device 23(2) for synchrony group 20(2), is not shown as being either a master or a slave device in another synchrony group. However, it will be appreciated that zone player 11(K+L+3) could also be utilized as the master device or a slave device for another synchrony

group. Indeed, it will be appreciated that the zone player that is utilized as the audio information channel device for synchrony group 20(2) may also be a zone player that is utilized as the master device 21(1) or a slave device 22(1)(1),..., 22(K)(1) in the synchrony group 20(1).

A zone player 11(n) that is utilized as a member of one synchrony group may also be utilized as the audio information channel device for another synchrony group if the audio information source that is to supply the audio information that is to be played by the other synchrony group is connected to that zone player 11(n). A zone player 11(n) may also be utilized as the audio information channel device for the other synchrony group if, for example, the audio information source is an audio information source 16(m) (FIG. 1) that is connected to the network 12 or an audio information source that is available over a wide area network such as the Internet. The latter may occur if, for example, the zone player 11(n) has sufficient processing power to operate as the audio information channel device and it is in an optimal location in the network 12, relative to the zone players comprising the other synchrony group (that is the synchrony group for which it is operating as audio information channel device) for providing the audio and playback timing information to the members of the other synchrony group. Other circumstances under which the zone player 11(n) that is utilized as a member of one synchrony group may also be utilized as the audio information channel device for another synchrony group will be apparent to those skilled in the art.

As was noted above, the master device 21 for a synchrony group 20 may be migrated from one zone player 11(n) to another zone player 11(n'). As was further noted above, the audio information channel device 23 for a synchrony group 20 may be migrated from one zone player 11(n) to another zone player 11(n'). It will be appreciated that the latter may occur if, for example, the audio information source that provides the audio program for the synchrony group is not connected to the zone player 11(n) that is operating as the audio information channel device 23, but instead is one of the audio information sources 16(m) connected to the network 12 or a source available over a wide area network such as the Internet. Operations performed during a migration of an audio information channel device 23 from one zone player 11(n) to another zone player 11(n') will generally depend on the nature of the audio information that is being channeled by the audio information channel device 23. For example, if the audio information source provides streaming audio, the zone player 11(n) that is currently operating as the audio information channel device 23 for the synchrony group 20, can provide the following information to the other zone player 11(n') that is to become the audio information channel device 23 for the synchrony group 20:

(a) the identification of the source of the streaming audio information,

(b) the time stamp associated with the frame that the zone player 11(n) is currently forming, and

(c) the identifications of the zone players that are operating as the master device 21 and slave devices 22(g) comprising the synchrony group 20.

After the zone player 11(n') receives the information from the zone player 11(n), it will begin receiving the streaming audio from the streaming audio information source identified by the zone player 11(n), assemble the streaming audio information into frames, associate each frame with a time stamp, and transmit the resulting audio and playback timing information over the network 12. The zone player 11(n') will perform these operations in the same manner as described above, except that, instead of using the time indicated by its digital to analog converter clock 34 directly in generating the time stamps for the frames, the initial time stamp will be related to the value of the time stamp that is provided by the zone player 11(n) (reference item (b) above), with the rate at which the time stamps are incremented corresponding to the rate at which its (that is, the zone player 11(n')'s) clock increments. In addition, the zone player 11(n') will notify the zone players that are operating as the master device 21 and slave devices 22(g) of the synchrony group 20 that it is the new audio information channel device 23 for the synchrony group 20, and provide the multi-cast address that it will be using to multi-cast the audio and playback timing information, as well as its unicast network address. After the members of the synchrony group 20 receive the notification from the zone player 11(n') indicating that it is the new audio information channel device 23 for the synchrony group 20, they will receive the audio and playback timing information from the zone player 11(n') instead of the zone player 11(n), using the multi-cast address provided by the zone player 11(n'). In addition, they can utilize the zone player 11(n')'s unicast network address to obtain current time information therefrom. It will be appreciated that the zone player 11(n') will determine its current time in relation to the time stamp that is provided by the zone player 11(n) (reference item (b) above) or the current time information that it received from the zone player 11(n) using the SNTP protocol as described above.

Generally similar operations can be performed in connection with migrating the audio information channel device from one zone player 11(n) to another zone player 11(n') if the audio information is from one or more audio information files, such as may be the case if the audio information comprises MP3 or WAV files that are available from sources such as sources 16(m) connected to the network 12 or over from sources available over a wide area network such as the Internet, except for differences to accommodate the fact that the audio information is in files. In that

case, the zone player 11(n) that is currently operating as the audio information channel device 23 for the synchrony group 20 can provide the following information to the zone player 11(n') that is to become the audio information channel device 23 for the synchrony group 20:

(d) a list of the audio information files containing the audio information that is to be played;

(e) the identification of the file for which the zone player 11(n) is currently providing audio and playback timing information, along with the offset into the file for which the current item of audio and playback timing information is being generated and the time stamp that the zone player 11(n) is associating with that frame, and

(f) the identifications of the zone players that comprise the master device 21 and slave devices 22(g) comprising the synchrony group 20.

After the zone player 11(n') receives the information from the zone player 11(n), it will begin retrieving audio information from the file identified in item (e), starting at the identified offset. In addition, the zone player 11(n') can assemble the retrieved audio information into frames, associate each frame with a time stamp and transmit the resulting audio and playback timing information over the network 12. The zone player 11(n') will perform these operations in the same manner as described above, except that, instead of using the time indicated by its digital to analog converter clock 34 directly in generating the time stamps for the frames, the value of the initial time stamp will be related to the time stamp that is provided by the zone player 11(n) (reference item (e) above), with the rate at which the time stamps are incremented corresponding to the rate at which its (that is, the zone player 11(n')'s) clock increments. In addition, the zone player 11(n') will notify the zone players that are operating as the master device 21 and slave devices 22(g) of the synchrony group 20 that it is the new audio information channel device 23 for the synchrony group 20, and provide the multi-cast address that it will be using to multi-cast the audio and playback timing information, as well as its unicast network address. After the members of the synchrony group 20 receive the notification from the zone player 11(n') indicating that it is the new audio information channel device 23 for the synchrony group 20, they will receive the audio and playback timing information from the zone player 11(n') instead of the zone player 11(n), using the multi-cast address provided by the zone player 11(n'). In addition, they can utilize the zone player 11(n')'s unicast network address to obtain current time information therefrom. It will be appreciated that the zone player 11(n') will determine its current time in relation to the time stamp that is provided by the zone player 11(n) (reference item (b) above) or the current time information that it received from the zone player 11(n) using the SNTP

protocol as described above. The zone player 11(n') will process successive audio information files in the list that it receives from the zone player 11(n) (reference item (d)).

Operations performed by the zone players 11(n) and 11(n') in connection with migration of the audio information channel device 23 for other types of audio information will be apparent to those skilled in the art. In any case, preferably, the zone player 11(n) will continue operating as an audio information channel device 23 for the synchrony group 20 for at least a brief time after it notifies the zone player 11(n') that it is to become audio information channel device for the synchrony group, so that the zone player 11(n') will have time to notify the zone players in the synchrony group 20 that it is the new audio information channel device 23 for the synchrony group.

Before proceeding further in describing operations performed by the network audio system 10, it would be helpful to provide a detailed description of a zone player 11(n) constructed in accordance with the invention. FIG. 3 depicts a functional block diagram of a zone player 11(n) constructed in accordance with the invention. All of the zone players in the network audio system 10 may have similar construction. With reference to FIG. 3, the zone player 11(n) includes an audio information source interface 30, an audio information buffer 31, a playback scheduler 32, a digital to analog converter 33, an audio amplifier 35, an audio reproduction device interface 36, a network communications manager 40, and a network interface 41, all of which operate under the control of a control module 42. The zone player 11(n) also has a device clock 43 that provides timing signals that control the general operations of the zone player 11(n). In addition, the zone player 11(n) includes a user interface module interface 44 that can receive control signals from the user interface module 13 (FIGS. 1 and 2) for controlling operations of the zone player 11(n), and provide status information to the user interface module 13.

Generally, the audio information buffer 31 buffers audio information, in digital form, along with playback timing information. If the zone player 11(n) is operating as the audio information channel device 23 (FIG. 2) for a synchrony group 20, the information that is buffered in the audio information buffer 31 will include the audio and playback timing information that will be provided to the devices 21 and 22(g) in the synchrony group 20. If the zone player 11(n) is operating as the master device 21 or a slave device 22(g) for a synchrony group, the information that is buffered in the audio information buffer 31 will include the audio and playback timing information that the zone player 11(n) is to play.

The audio information buffer 31 can receive audio and playback timing information from two sources, namely, the audio information source interface 30 and the network communications

manager 40. In particular, if the zone player 11(n) is operating as the audio information channel device 23 for a synchrony group 20, and if the audio information source is a source 14(n)(s) connected to the zone player 11(n), the audio information buffer 31 will receive and buffer audio and playback timing information from the audio information source interface 30. On the other hand, if the zone player 11(n) is operating as the audio information channel device 23 for a synchrony group 20, and if the audio information source is a source 16(m) connected to the network 12, or a source available over the wide area network, the audio information buffer 31 will receive and buffer audio and playback timing information from the network communications manager 40. It will be appreciated that, if the zone player 11(n) is not a member of the synchrony group, the zone player 11(n) will not play this buffered audio and playback timing information.

On yet another hand, if the zone player 11(n) is operating as the master device 21 or a slave device 22(g) in a synchrony group, and if the zone player 11(n) is not also the audio information channel device 23 providing audio and playback timing information for the synchrony group 20, the audio information buffer 31 will receive and buffer audio and playback timing information from the network communications manager 40.

The audio information source interface 30 connects to the audio information source(s) 14(n)(s) associated with the zone player 11(n). While the zone player 11(n) is operating as audio information channel device 23 for a synchrony group 20, and if the audio information is to be provided by a source 14(n)(s) connected to the zone player 11(n), the audio information source interface 30 will selectively receive audio information from one of the audio information source(s) 14(n)(s) to which the zone player is connected and store the audio information in the audio information buffer 21. If the audio information from the selected audio information source 14(n)(s) is in analog form, the audio information source interface 30 will convert it to digital form. The selection of the audio information source 14(n)(s) from which the audio information source interface 30 receives audio information is under control of the control module 42, which, in turn, receives control information from the user interface module through the user interface module interface 44. The audio information source interface 30 adds playback timing information to the digital audio information and buffers the combined audio and playback timing information in the audio information buffer 21.

More specifically, as noted above, the audio information source interface 30 receives audio information from an audio information source 14(n)(s), converts it to digital form if necessary, and buffers it along with playback timing information in the audio information buffer 21. In addition,

the audio information source interface 30 will also provide formatting and scheduling information for the digital audio information, whether as received from the selected audio information source 14(n)(s) or as converted from an analog audio information source. As will be made clear below, the formatting and scheduling information will control not only playback by the zone player 11(n) itself, but will also enable other zone players 11(n'), 11(n''),... that may be in a synchrony group for which the zone player 11(n) is the master device, to play the audio program associated with the audio information in synchrony with the zone player 11(n).

In one particular embodiment, the audio information source interface 30 divides the audio information associated with an audio work into a series of frames, with each frame comprising digital audio information for a predetermined period of time. As used herein, an audio track may comprise any unit of audio information that is to be played without interruption. On the other hand, an audio program may comprise a series of one or more audio tracks that are to be played in succession. It will be appreciated that the tracks comprising the audio program may also be played without interruption, or alternatively playback between tracks may be interrupted by a selected time interval. FIG. 4 schematically depicts an illustrative framing strategy used in connection with one embodiment of the invention for a digital audio stream comprising an audio work. More specifically, FIG. 4 depicts a framed digital audio stream 50 comprising a sequence of frames 51(1) through 51(F) (generally identified by reference numeral 51(f)). Each frame 51(f), in turn, comprises a series of audio samples 52(f)(1) through 52(f)(S) (generally identified by reference numeral 52(f)(s)) of the audio track. Preferably all of the frames will have the same number "S" of audio samples, although it will be appreciated from the following that that is primarily for convenience. On the other hand, it will be appreciated that, the number of audio samples may differ from "S"; this may particularly be the case if the frame 51(f) contains the last audio samples for the digital audio stream for a particular audio work. In that case, the last frame 51(F) will preferably contain samples 52(F)(1) through 52(F)(x), where "x" is less than "S." Generally, it is desirable that the number of samples be consistent among all frames 51(f), and in that case padding, which will not be played, can be added to the last frame 51(F).

Associated with each frame 51(f) is a header 55(f) that includes a number of fields for storing other information that is useful in controlling playback of the audio samples in the respective frame 51(f). In particular, the header 55(f) associated with a frame 51(f) includes a frame sequence number field 56, an encoding type field 57, a sampling rate information field 58, a time stamp field 60, an end of track flag 61, and a length flag field 62. The header 55(f) may also include fields (not shown)

for storing other information that is useful in controlling playback. Generally, the frame sequence number field 56 receives a sequence number "f" that identifies the relative position of the frame 51(f) in the sequence of frames 51(1)..51(f)...51(F) containing the digital audio stream 50. The encoding type field 57 receives a value that identifies the type of encoding and/or compression that has been used in generating the digital audio stream. Conventional encoding or compression schemes include, for example, the well-known MP3 and WAV encoding and/or compression schemes, although it will be appreciated that other schemes may be provided for as well. The sampling rate information field 58 receives sampling rate information that indicates the sampling rate for the audio samples 52(f)(s). As will be apparent to those skilled in the art, the sampling rate determines the rate at which the zone player 11(n) is to play the audio samples 52(f)(s) in the frame, and, as will be described below, determines the period of the digital to analog converter clock 34.

The condition of the end of work flag 61 indicates whether the frame 51(f) contains the last digital audio samples for the audio track associated with the framed digital audio work 50. If the frame 51(f) does not contain the audio samples that are associated with the end of the digital audio stream 50 for a respective audio work, the end of work flag will be clear. On the other hand, if the frame 51(f) does contain the audio samples that are associated with the end of the digital audio stream 50 for a respective audio work, the end of work flag 61 will be set. In addition, since the number of valid audio samples 52(F)(s) in the frame 51(F), that is, the samples that are not padding, may be less than "S," the default number of audio samples in a frame 51(f), the length flag field 62 will contain a value that identifies the number of audio samples in 52(F)(s) in the last frame 51(F) of the audio work 50. If, as noted above, the frames have a consistent number "S" of samples, the samples 52(F)(x+1) through 52(F)(S) will contain padding, which will not be played.

The time stamp field 60 stores a time stamp that identifies the time at which the zone player 11(n) is to play the respective frame. More specifically, for each frame of a framed digital audio stream 50 that is buffered in the audio information buffer 21, the audio information source interface 30, using timing information from the digital to analog converter clock 34, will determine a time at which the zone player 11(n) is to play the respective frame, and stores a time stamp identifying the playback time in the time stamp field 60. The time stamp associated with each frame will later be used by the playback scheduler 32 to determine when the portion of the digital audio stream stored in the frame is to be coupled to the digital to analog converter 33 to initiate play back. It will be appreciated that the time stamps that are associated with frames in sequential frames 51(1), 51(2),...,51(F), will be such that they will be played back in order, and without an interruption

between the sequential frames comprising the digital audio stream 50. It will further be appreciated that, after a time stamp has been determined for the first frame, stored in frame 51(1), of a digital audio stream 50, the audio information source interface 30 can determine time stamps for the subsequent frame 51(2), 51(3),....,51(F) in relation to the number of samples "S" in the respective frames and the sample rate. The time stamps will also preferably be such that frames will be played back after some slight time delay after they have been buffered in the audio information buffer 21; the purpose for the time delay will be made clear below.

Returning to FIG. 3, in addition to dividing the digital audio information into frames, the audio information source interface 30 also aggregates and/or divides the frames 51(f) as necessary into packets, each of which will be of a length that would fit into a message for transmission over the network, and associates each packet with a packet sequence number. For example, if a packet will accommodate multiple frames 51(f), 51(f+1),... 51(f+y-1), it will aggregate them into a packet and associate them with a packet number, for example p(x). If the entire frames 51(f) and 51(f+y-1) was accommodated in packet p(x), where "x" is the sequence number, which will occur if the size of a packet is an exact multiple of the frame size, the next packet, p(x+1) will begin with frame 51(f+y) and will include frames 51(f+y),..., 51(f+2y-1). Subsequent packets p(x+2),.... will be formed in a similar manner. On the other hand, if the packet length will not accommodate an exact multiple of the frame size, the last frame in the packet will be continued at the beginning of the next packet.

If the audio information source interface 30 is aware of track boundaries, which may be the case if the tracks are divided into files, the packets will reflect the track boundaries, that is, the packets will not contain frames from two tracks. Thus, if the last frames associated with a track are insufficient to fill a packet, the packet will contain padding from the last frame associated with the track to the end of the packet, and the next packet will begin with the first frames associated with the next track.

In one embodiment, the audio information source interface 30 stores the packets in the audio information buffer 31 in a ring buffer. As is conventional, a ring buffer includes a series of storage locations in the buffer. Each entry will be sufficient to store one packet. Four pointers are used in connection with the ring buffer, a first pointer pointing to the beginning of the ring buffer, a second pointer pointing to the end of the ring buffer, an third "write" pointer pointing to the entry into which a packet will be written and a fourth "read" pointer pointing to the entry from which packet will be read for use in playback. When a packet is read from the ring buffer for playback, it will be read

from the entry pointed to by the read pointer. After the packet has been read, the read pointer will be advanced. If the read pointer points beyond the end of the ring buffer, as indicated by the end pointer, it will be reset to point to the entry pointed to by the beginning pointer, and the operations can be repeated.

On the other hand, when the audio information source interface 30 stores a packet in the ring buffer, first determine whether the entry pointed to by the write pointer points to the same entry as the entry pointed to by the read pointer. If the write pointer points to the same entry as the entry pointed to by the read pointer, the entry contains at least a portion of a packet that has not yet been read for playback, and the audio information source interface 30 will delay storage of the packet until the entire packet has been read and the read pointer advanced. After the read pointer has been advanced, the audio information source interface 30 can store the packet in the entry pointed to by the write pointer. After the packet has been stored, the audio information source interface 30 will advance the write pointer. If the write pointer points beyond the end of the ring buffer, as indicated by the end pointer, it will be reset to point to the entry pointed to by the beginning pointer, and the operations can be repeated.

As noted above, the zone player 11(n) can operate both as an audio information channel device 23 and a member of the synchrony group 20 of which it is a member. In that case, the audio information buffer 31 can contain one ring buffer. On the other hand, the zone player 11(n) can operate as an audio information channel device 23 for one synchrony group 20(1) (FIG. 2A) and a member of another synchrony group 20(2). In that case, the audio information buffer 31 would maintain two ring buffers, one for the audio and timing information associated with synchrony group 20(1), and the other for the audio and timing information associated with synchrony group 20(2). It will be appreciated that, in the latter case, the zone player 11(n) will only use the audio and timing information that is associated with synchrony group 20(2) for playback.

The playback scheduler 32 schedules playback of the audio information that is buffered in the audio information buffer 31 that is to be played by the zone player 11(n). Accordingly, under control of the playback scheduler 32, the digital audio information that is buffered in the audio information buffer 21 that is to be played by the zone player 11(n) is transferred to the digital to analog converter 33 for playback. As noted above, if the zone player 11(n) is operating as an audio information channel device 23 for a synchrony group 20 for which it is not a member, the playback scheduler 32 will not schedule the digital audio information that is to be played by that synchrony group 20 for playback. The playback scheduler 32 only schedules the digital audio information, if

any, that is buffered in the audio information buffer 31 that is associated with a synchrony group for which the zone player 11(n) is a member, whether as master device 21 or a slave device 22(g).

Essentially, the playback scheduler 32 makes use of the read pointer associated with the circular buffer that contains the audio and playback timing information that is to be played by the zone player 11(n). The playback scheduler 32 retrieves the packet information from the entry of the ring buffer pointed to by the read pointer, and then advances the ring pointer as described above. The playback scheduler 32 determines the boundaries of the frames in the packet and uses the time stamps in the time stamp fields 60 associated with the respective frame 51(f), along with timing information provided by the zone player 11(n)'s digital to analog converter clock 34, to determine when the respective frame is to be transferred to the digital to analog converter 33. Generally, when the time stamp associated with a buffered digital audio information frame corresponds to the current time as indicated by the digital to analog converter clock 34, the playback scheduler 32 will enable the respective frame to be transferred to the digital to analog converter 33.

The digital to analog converter 33, also under control of the digital to analog converter clock 34, converts the buffered digital audio information to analog form, and provides the analog audio information to the audio amplifier 35 for amplification. The amplified analog information, in turn, is provided to the audio reproduction devices 15(n)(r) through the audio reproduction device interface 36. The audio reproduction devices 15(n)(r) transform the analog audio information signal to sound thereby to provide the audio program to a listener. The amount by which the audio amplifier 35 amplifies the analog signal is controlled by the control module 42, in response to volume control information provided by the user through the user interface module 13.

The network communications manager 40 controls network communications over the network 12, and the network interface 41 transmits and receives message packets over the network 12. The network communications manager 40 generates and receives messages to facilitate the transfer of the various types of information described above in connection with FIG. 2, including the channel device control information, slave device control information, audio and playback timing information and the audio information channel device's clock timing information. In connection with the channel device control information and the slave device control information, the network communications manager 40 will generate messages for transfer over the network 12 in response to control information from the control module 42. Similarly, when the network communications manager 40 receives messages containing channel device control information and slave device

control information, the network communications manager will provide the information to the control module 42 for processing.

With regards to the audio information channel device's clock timing information, as noted above, the master device 21 and slave devices 22(g) of the synchrony group 20 obtain the clock timing information from the audio information channel device 23 using the well-known SNTP. If the zone player 11(n) is operating as the audio information channel device 23 for a synchrony group, during the SNTP operation, it will provide its current time, particularly a current time as indicated by its digital to analog converter clock 34. On the other hand, if the zone player 11(n) is operating as the master device 21 or slave device 22(g) of a synchrony group 20, it will receive the clock timing information from the audio information channel device 23. After the respective device 21, 22(g) has obtained the audio information channel device's clock timing information, it will generate a differential time value ΔT representing the difference between the time T indicated by its digital to analog converter clock 34 and the current time information from the audio information channel device 23. The differential time value will be used to update the time stamps for the frames of the digital audio stream 50 (FIG. 4) that are received from the audio information channel device.

With regards to the audio and playback timing information, operations performed by the network communications manager 40 will depend on whether

(i) the audio and playback timing information has been buffered in the audio information buffer 31 for transmission, as audio information channel device 23, over the network 12 to the master device 21 and/or slave devices 22(g) of a synchrony group, or

(ii) the audio and playback timing information has been received from the network 12 to be played by the zone player 11(n) as either the master device 21 for a synchrony group or a slave device in a synchrony group.

It will be appreciated that the network communications manager 40 may be engaged in both (i) and (ii) contemporaneously, since the zone player 11(n) may operate both as the audio information channel device 23(1) for a synchrony group 20(1) (reference FIG. 2A) of which it is not a member, and a member of another synchrony group 20(2) for which another zone player 11(n') is the audio information channel device 23(2). With reference to item (i) above, after a packet that is to be transmitted has been buffered in the respective ring buffer, the network communications manager 40 retrieves the packet, packages it into a message and enables the network interface 41 to transmit the message over the network 12. If the control module 42 receives control information from the user interface module 13 (if the master device 21 is also the audio information channel device 23 for

the synchrony group 20) or from the master device (if the master device 21 is not the audio information channel device 23 for the synchrony group 20) that would require the transmission of a "resynchronize" command as described above, the control module 42 of the audio information channel device 23 enables the network communications manager 40 to insert the command into a message containing the audio and playback timing information. Details of the operations performed in connection with the "resynchronize" command will be described below. As noted above, the "resynchronize" command is used if the user enables a synchrony group to terminate the playback of a track that is currently being played, or cancel playback of a track whose playback has not begun.

On the other hand, with reference to item (ii) above, if network interface 41 receives a message containing a packet containing frames of audio and playback timing information that the zone player 11(n) is to play either as a master device 21 or a slave device for a synchrony group 20, the network interface 41 provides the audio and playback timing information to the network communications manager 40. The network communications manager 40 will determine whether the packet contains a resynchronize command and, if so, notify the control module 42, which will enable operations to be performed as described below. In any case, the network communications manager 40 will normally buffer the various frames comprising the audio and playback timing information in the audio information buffer 31, and in that operation will generally operate as described above in connection with the audio information source interface 30. Before buffering them, however, the network communications manager 40 will update their time stamps using the time differential value described above. It will be appreciated that the network communications manager 40 will perform similar operations whether the messages that contain the packets were multi-cast messages or unicast messages as described above

The updating of the time stamps by the master device 21 and the slave devices 22(g) in the synchrony group 20 will ensure that they all play the audio information synchronously. In particular, after the network communications manager 40 has received a frame 51(f) from the network interface 41, it will also obtain, from the digital to analog converter clock 34, the zone player 11(n)'s current time as indicated by its digital to analog converter clock 34. The network communications manager 40 will determine a time differential value that is the difference between the slave device's current clock time, as indicated by its digital to analog converter 34, and the audio information channel device's time as indicated by the audio information channel device's clock timing information. Accordingly, if the master or slave device's current time has a value T_S and the audio information channel device's current time, as indicated by the clock timing information, has a value T_C, the time

differential value $\Delta T=T_S-T_C$. If the current time of the master or slave device in the synchrony group 20, as indicated by its digital to analog converter clock 34, is ahead of the audio information channel device's clock time as indicated by the clock timing information received during the SNTP operation, the time differential value will have a positive value. On the other hand, if the master or slave device's current time is behind the audio information channel device's clock time, the time differential value ΔT will have a negative value. If the zone player 11(n) obtains clock timing information from the audio information channel device 23 periodically while it is a member of the synchrony group 20, the network communications manager 40 can generate an updated value for the time differential value ΔT when it receives the clock timing information from the audio information channel device 23, and will subsequently use the updated time differential value.

The network communications manager 40 uses the time differential value ΔT that it generates from the audio information channel device timing information and zone player 11(n)'s current time to update the time stamps that will be associated with the digital audio information frames that the zone player 11(n) receives from the audio information channel device. For each digital audio information frame that is received from the audio information channel device, instead of storing the time stamp that is associated with the frame as received in the message in the audio information buffer 21, the network communications manager 40 will store the updated time stamp with the digital audio information frame. The updated time stamp is generated in a manner so that, when the zone player 11(n), as a member of the synchrony group plays back the digital audio information frame, it will do so in synchrony with other devices in the synchrony group.

More specifically, after the zone player 11(n)'s network interface 41 receives a message containing a packet that, in turn, contains one or more frames 51(f), it will provide the packet to the network communications manager 40. For each frame 51(f) in the packet that the network communications manager 40 receives from the network interface 41, the network communications manager 40 will add the time differential value ΔT to the frame's time stamp, to generate the updated time stamp for the frame 51(f), and store the frame 51(f), along with the header 55(f) with the updated time stamp in the audio information buffer 31. Thus, for example, if a frame's time stamp has a time value T_F, the network communications manager 40 will generate an updated time stamp T^U_F having a time value $T^U_F=T_F+\Delta T$. Since time value T^U_F according to the slave device's digital to analog converter clock 34 is simultaneous to the time value T_F according to the audio information channel device's digital to analog converter clock 34, the zone player 11(n) device will play the digital audio information frame at the time determined by the audio information channel device 23.

Since all of the members of the synchrony group 20 will perform the same operations, generating the updated time stamps T^U_F for the various frames 51(f) in relation to their respective differential time values, all of the zone players 11(n) in the synchrony group 20 will play them synchronously. The network communications manager 40 will generate updated time stamps T^U_F for all of the time stamps 60 in the packet, and then store the packet in the audio information buffer 31.

It will be appreciated that, before storing a packet in the audio information buffer 21, the network communications manager 40 can compare the updated time stamps T^U_F associated with the frames in the packet to the slave device's current time as indicated by its digital to analog converter clock 34. If the network communications manager 40 determines the time indicated by the updated time stamps of frames 51(f) in the packet are earlier than the zone player's current time, it can discard the packet instead of storing it in the audio information buffer 21, since the zone player 11(n) will not play them. That is, if the updated time stamp has a time value T^U_F that identifies a time that is earlier than the zone player's current time T_S as indicated by the zone player's digital to analog converter clock 34, the network communications manager 40 can discard the packet.

If the zone player 11(n) is operating as the master device 21 of a synchrony group 20, when the user, through the user interface module 13, notifies the zone player 11(n) that another zone player 11(n') is to join the synchrony group 20 as a slave device 22(g), the control module 42 of the zone player 11(n) enables the network communications manager 40 to engage in an exchange of messages, described above in connection with FIG. 2, to enable the other zone player 11(n') to join the synchrony group 20 as a slave device. As noted above, during the message exchange, the messages generated by the network communications manager 40 of the zone player 11(n) will provide the network communications manager of the zone player 11(n') that is to join the synchrony group 20 with information such as the multi-cast address being used by the audio information channel device 23 that is providing the audio program to the synchrony group 20, as well as a unicast network address for the audio information channel device 23. After receiving that information, the network communications manager and network interface of the zone player 11(n') that is to join the synchrony group 20 can begin receiving the multi-cast messages containing the audio program for the synchrony group, engage in SNTP transactions with the audio information channel device 23 to obtain the latter's current time, and also enable the audio information channel device 23 to send the zone player 11(n') frames 51(f) that it had previously broadcast using the unicast message transmission methodology as described above.

On the other hand, if the network communications manager 40 and network interface 41 of the zone player 11(n) receive a message over the network 12 indicating that it is to become a slave device 22(g) of a synchrony group for which another zone player 11(n') is the master device, the network communications manager 40 for zone player 11(n) will provide a notification to the control module 42 of zone player 11(n). Thereafter, the control module 42 of zone player 11(n) can enable the network communications manager 40 of zone player 11(n) to perform the operations described above to enable it to join the synchrony group 20.

As noted above, the user, using user interface module 13, can enable the synchrony group to terminate playback of a track of an audio program that is currently being played. After playback of a track that is currently being played has been terminated, playback will continue in a conventional manner with the next track that has been buffered in the audio information buffer 31. It will be appreciated that that could be the next track that is on the original play list, or a previous track. In addition, the user can enable the synchrony group 20 to cancel playback of a track that it has not yet begun to play, but for which buffering of packets has begun in the synchrony group 20. Both of these operations make use of the "resynchronize" command that the master device 21 of the synchrony group 20 can enable the audio information channel device 23 to include in the multi-cast message stream that it transmits to the synchrony group 20. Generally, in response to receipt of the resynchronize command, the members of the synchrony group 20 flush the ring buffer containing the packets that they are to play in the future. In addition, if the members of the synchrony group provide separate buffers for their respective digital to analog converters 33, the members will also flush those buffers as well. After the audio information channel device transmits a packet containing the resynchronize command:

(i) in the case of the use of the resynchronize command to terminate playing of a track that is currently being played, the audio information channel device 23 will begin multi-casting packets for the next track, to begin play immediately, and will continue through the play list in the manner described above; and

(ii) in the case of the use of the resynchronize command to cancel play of a track for which buffering has begun, but which is to be played in the future, the audio information channel device 23 will begin multi-casting packets for the track after the track that has been cancelled, to be played beginning at the time the cancelled track was to begin play, and will also continue through the play list in the manner described above.

It will be appreciated that,

(a) in the first case (item (i) directly above), the resynchronize command can enable the read pointer to be set to point to the entry in the circular buffer into which the first packet for the next track will be written, which will correspond to the entry to which the write pointer points, but

(b) in the second case (item (ii) directly above), the resynchronize command can enable the write pointer for the circular buffer to be set to point to the entry that contains the first packet for the track whose play is being cancelled.

It will further be appreciated that, if a track is cancelled for which buffering has not begun, the resynchronize command will generally not be necessary, since the audio information channel device 23 for the synchrony group 20 need only delete that track from the play list.

Operations performed in connection with use of the resynchronize command to cancel playback of a track that is currently being played will be described in connection with Packet Sequence A below, and operations performed in connection with the use of the resynchronize command to cancel playback of a track that is has not yet begun to play, but for which buffering of packets has begun, will be described in connection with Packet Sequence B below.

Packet Sequence A

(A1.0) [packet 57]

(A1.1 [continuation of frame 99]

(A1.2) [frame 100, time = 0:00:01, type = mp3 audio]

(A1.3) [frame 101, time = 0:00:02, type = mp3 audio]

(A1.4) [frame 102, time = 0:00:03, type = mp3 audio]

(A2.0) [packet 58]

(A2.1) [continuation of frame 102]

(A2.2) [frame 103, time = 0:00:04, type = mp3 audio]

(A2.3) [frame 104, time = 0:00:05, type = mp3 audio]

(A2.4) [frame 105, time = 0:00:06, type = mp3 audio]

(A3.0) [packet 59]

(A3.1) [continuation of frame 105]

(A3.2) [frame 106, time = 0:00:07, type = mp3 audio]

(A3.3) [frame 107, time = 0:00:08, type = mp3 audio]

(A3.4) [frame 108, time = 0:00:09, type = mp3 audio]

(A4.0) [packet 60]

(A4.1) [continuation of frame 108]

(A4.2) [frame 109, time = 0:00:10, type = mp3 audio]

(A4.3) [Resynchronize command]

(A4.4) [Padding, if necessary]

(A5.0) [packet 61]

(A5.1) [frame 1, time = 0:00:07, type = mp3 audio]

(A5.2) [frame 2, time = 0:00:08, type = mp3 audio]

(A5.3) [frame 3, time = 0:00:09, type = mp3 audio]

(A5.4) [frame 4, time = 0.00.10, type = mp3 audio]

(A6.0) [packet 62]

(A6.1) [continuation of frame 4]

(A6.2) [frame 5, time = 0:00:11, type = mp3 audio]

(A6.3) [frame 6, time = 0:00:12, type = mp3 audio]

(A6.4) [frame 7, time = 0:00:13, type = mp3 audio]

Packet Sequence A comprises a sequence of six packets, identified by packet 57 through packet 62, that the audio information channel device 23 multi-casts in respective messages to the members of a synchrony group 20. It will be appreciated that the series of messages that the audio information channel device 23 may multi-cast to the synchrony group 20 may include a messages prior to the packet 57, and may also include messages after packet 62. Each packet comprises a packet header, which is symbolized by lines (A1.0), (A2.0),...(A6.0) in Packet Sequence A, and will generally also include information associated with at least a portion of a frame. In the packets represented in Packet Sequence A, each packet includes information associated with a plurality of frames. Depending on the lengths of the packets, each packet may contain information associated with a portion of a frame, an entire frame, or multiple frames. In the illustration represented by Packet Sequence A, it is assumed that each packet may contain information associated with multiple frames. In addition, it is assumed that a packet does not necessarily contain information associated with an integral number of frames; in that case, a packet may contain information associated with

a portion of a frame, and the next packet will contain the information associated with the rest of a frame.

The frames and associated header playback timing information contained in the various packets are symbolized by lines (A1.1), (A1.2),...,(A1.4), (A2.1),....(A6.4) of Packet Sequence A. Thus, for example, line (A1.2) of packet 57 represents the one-hundredth frame, that is, frame 51(100) (reference FIG. 4), of the track whose audio information is being transmitted in the sequence of packets that includes packet 57. The frame 51(100) is to be played at an illustrative time, according to the audio information channel device's digital to analog converter clock, of "time = 0:00:01," and the frame is encoded and/or compressed using the well-known MP3 encoding and compression methodology. In that case, the legend"time = 0:00:01" represents the time stamp that would be included in field 60 (FIG. 4) of the header associated with the frame 50(100) as multi-cast by the audio information channel device for the synchrony group. It will be appreciated that the playback time and encoding/compression methodology will be referred in the header 55(100) that is associated with the frame 51(100). It will also be appreciated that the header may also contain additional information as described above.

Similarly, line (A1.3) of packet 57 represents the one-hundred and first frame, that is, frame 51(101), of the track whose audio information is being transmitted in the sequence of packets that includes packet 57. The frame 51(101) is to be played at an illustrative time, according to the audio information channel device's digital to analog converter clock, of "0:00:02," and the frame is also encoded and/or compressed using the MP3 encoding and compression methodology. Line (A1.4) of packet 57 represents similar information, although it will be appreciated that, depending on the length of packet 57, the line may not represent the information for an entire frame 51(102) and/or its associated header. If the length of packet 57 is not sufficient to accommodate the information for the entire frame 51(102) and/or associated header, the information will continue in packet 58, as represented by line (A2.1) in Packet Sequence A. Similarly, if the length of packet 56 was not sufficient to contain the information for an entire frame 51(99) preceding frame 51(100), packet 57 (lines (A1.0) through 1.4) may contain any information from frame 51(99) that packet 56 was unable to accommodate.

As noted above, when the master device 21 or a slave device 22(g) in the synchrony group 20 receives the packet 57, its respective network communications manager 40 will update the time stamps associated with the various frames 51(f) as described above before buffering the respective frames in the respective audio information buffer 31.

Packets 58 and 59 contain information that is organized along the lines described above in connection with packet 57.

Packet 60 also contains, as represented by lines (A4.1) and (A4.2), information that is organized along the lines of the information represented by lines (Ax.1) and (Ax.2) ("x" equals an integer) described above in connection with packets 57 through 59. On the other hand, packet 60 contains a resynchronize command, as represented by line (A4.3). Packet 60 also may contain padding, as represented by line 4.4, following the resynchronize command. As noted above, the master device 21 of a synchrony group 20 will enable the audio information channel device 23 that is providing audio information to the synchrony group 20 to multi-cast a message containing the resynchronize command when it receives notification from the user interface module 13 that the user wishes to cancel playback of a track that is currently being played. In the example depicted in Packet Sequence A, as will be described below, the audio information channel device 23 receives notification from the master device 21 that the user wishes to cancel playback of a track at a time corresponding to "time = 0:00:07" according to its digital to analog converter clock 34, and, in line (A4.3) of packet 60 it will provide the resynchronize command, followed by padding, if necessary.

As will be apparent from examining lines (A3.1) through (A3.4) of packet 59 and lines (A4.1) and (A4.2) of packet 60, although the audio information channel device 23 has received the notification from the synchrony group's master device 21 to multi-cast the resynchronize command at a time corresponding to "time = 0:00:07" according to the clock time indicated by its digital to analog converter clock 34, it (that is, the audio information channel device 23) has already multi-cast messages containing frames that are to be played at that time and subsequently. That is, the audio information channel device 23 has, multi-cast in packet 59, frames 51(106) through 51(108) that contain time stamps "time = 0:00:07," "time = 0:00:08" and "time = 0:00:09," respectively, and, in packet 60, in addition to the continuation of frame 51(108), frame 51(109) that contains time stamp "time = 0:00:10." (It will be appreciated that the times indicated by the illustrative time stamps are for illustration purposes only, and that in an actual embodiment the time stamps may have different values and differentials.)

As noted above, the audio information channel device 23 multi-casts a message containing a packet that, in turn, contains the resynchronize command when it receives the notification from the master device 21 to do so. In the example depicted in Packet Sequence A, the packet will be multi-cast when the audio information channel device's digital to analog converter clock time corresponds to "0:00:07." Subsequently, two things happen. In one, aspect, when the master device 21 and slave

devices 22(g) receive the packet that contains the resynchronize command, they will stop playing the audio program that they are playing.

In addition, the audio information channel device 23 will begin transmitting frames containing audio information for the next track, including therewith time stamps immediately following the digital to analog converter clock time at which the packet including the resynchronize command was transmitted. Accordingly, and with further reference to Packet Sequence A, the audio information channel device 23 will multi-cast a message containing packet 61. As indicated above, packet 61 contains, as represented in lines (A5.1) through (A5.3), frames 51(1) through 51(3), which are the first three frames of the next track of the audio program that is to be played. It is also compressed and encoded using the MP3 encoding and compression scheme, and it is accompanied by time stamps "time = 0:00:07," "time = 0:00:08" and "time = 0:00:10." As noted above, the time stamp "time = 0:00:07" corresponds to the clock time at which the audio information channel device 23 multi-casts the resynchronize command, and, when the master device 21 and slave devices 22(g) receive these frames, they would be expected to begin playing them very shortly, if not immediately after the audio information channel device 23 multi-casts the message containing the packet that, in turn, contains the resynchronize command. Packet 61 also includes at least a portion of the next frame, that is, frame 51(4), for that track. In addition, Packet Sequence A depicted above further includes a subsequent packet, namely, packet 62, that contains any necessary continuation of frame 51(4), as well as three subsequent frames. If any additional packets are required for the track, as well as for subsequent tracks, they can be multi-cast in a similar manner.

As further noted above, the resynchronize command can also be used to cancel playing of one or more tracks for which playback has begun. This will be illustrated in connection with Packet Sequence B:

Packet Sequence B

(B1.0) [packet 157]

(B1.1) [continuation of frame 99]

(B1.2) [frame 100, time = 0:00:01, type = mp3 audio]

(B1.3) [frame 101, time = 0:00:02, type = mp3 audio]

(B1.4) [frame 102, time = 0:00:03, type = mp3 audio]

(B2.0) [packet 158]

(B2.1) [continuation of frame 102]

(B2.2) [frame 103, time = 0:00:04, type = mp3 audio]

(B2.3) [frame 104, time = 0:00:05, type = mp3 audio]

(B2.4) [frame 105, time = 0:00:06, type = mp3 audio]

(B3.0) [packet 159]

(B3.1) [continuation of frame 105]

(B3.2) [frame 106, time = 0:00:07, type = mp3 audio]

(B3.3) [track boundary notification]

(B3.4) [Padding, if necessary]

(B4.0) [packet 160]

(B4.1) [frame 1, time = 0:00:08, type = mp3 audio]

(B4.2) [frame 2, time = 0:00:09, type = mp3 audio]

(B4.3) [frame 3, time = 0:00:10, type = mp3 audio]

(B5.0) [packet 161]

(B5.1) [continuation of frame 3]

(B5.2) [frame 4, time = 0:00:11, type = mp3 audio]

(B5.3) [Resynchronize, after packet 159]

(B5.4) [Padding, if necessary]

(B6.0) [packet 162]

(B6.1) [frame 1, time = 0:00:08, type = mp3 audio]

(B6.2) [frame 2, time = 0:00:09, type = mp3 audio]

(B6.3) [frame 3, time = 0:00:10, type = mp3 audio]

(B6.4) [frame 4, time = 0:00:11, type = mp3 audio]

(B7.0) [packet 163]

(B7.1) [continuation of frame 4]

(B7.2) [frame 5, time = 0:00:12, type = mp3 audio]

(B7.3) [frame 6, time = 0:00:13, type = mp3 audio]

(B7.4) [frame 7, time = 0:00:14, type = mp3 audio]

Packet Sequence B comprises a series of seven packets, identified by packet 157 through 163 , that the audio information channel device 23 multi-casts to the members of a synchrony group 20. As with Packet Sequence A, it will be appreciated that the series of packets that the audio information channel device 23 may multi-cast to the synchrony group 20 may include packets prior to the packet 157, and may also include packets after packet 162. Each packet comprises a packet header, which is symbolized by lines (B1.0), (B2.0),...(B7.0) in Packet Sequence B. As in Packet Sequence A, each packet will also generally include information associated with at least a portion of a frame 51(f)along with its associated frame 55(f). As in the packets represented in Packet Sequence A, each packet includes information associated with a plurality of frames. Depending on the lengths of the packets, each packet may contain information associated with a portion of a frame, an entire frame, or multiple frames. Further, as with Packet Sequence A, it is assumed that each packet may contain information associated with multiple frames. In addition, it is assumed that a packet does not necessarily contain information associated with an integral number of frames; in that case, a packet may contain information associated with a portion of a frame, and the next packet will contain the information associated with the rest of a frame.

The structures of the packets represented by Packet Sequence B are similar to those described above in connection with Packet Sequence A, and will not be repeated here. Generally, Packet Sequence B illustratively contains a sequence of packets that represent at least portions of three tracks that may have been selected from, for example, a play list. In particular, packets 157 through 159 represent frames from a portion of one track, packets 160 and 161 represent frames from a second track and packets 162 and 163 represent frames from a third track. The play list indicated that the first, second and third tracks were to be played in that order. With particular reference to Packet Sequence B, it should be noted that line (B3.3) indicates that packet 159 includes an indication that that packet contains the last frame for the track, and line (B3.4) provides for padding to the end of the packet. The first frame of the next track begins in packet 160.

In connection with the use of the resynchronize command to cancel playback of a track, at least a portion of which the audio information channel device 23 has multi-cast to the members of the synchrony group, packet 161, in line (B5.3) represents a resynchronize command that indicates that resynchronization is to occur after packet 159, that is, immediately after the packet that contains the last frame of the first of the three tracks represented by the packets in Packet Sequence B. It should be noted that the resynchronize command is in the packet 161, while the resynchronization is to occur at packet 160, that is, the synchrony group is to not play the track starting with packet

160, but instead is to begin playing the track frames for which begin with the next packet, that is, packet 162. As with Packet Sequence A, in Packet Sequence B the audio information channel device 23, in packet 162 and 163, multi-casts frames whose time stamps indicate that they are to be played when the frames that were multi-cast in packets 160 and 161 were to be played. By use of the resynchronize command and specifying a packet in this manner, the audio information channel device can cancel playback of a track for which playback has not yet begun.

It will be appreciated that the resynchronize command is generally not necessary for cancelling play back of a track that the audio information channel device 23 has not started multi-casting to the synchrony group 20, since the audio information channel device 23 itself can re-order the play list to accommodate the cancellation.

The invention provides a number of advantages. In particular, the invention provides a network audio system in which a number of devices share information can reproduce audio information synchronously, notwithstanding the fact that packets, which may contain digital audio information, transmitted over the network to the various zone players connected thereto may have differing delays and the zone players operate with independent clocks. Moreover, although the invention has been described in connection with audio information, it will be appreciated that the invention will find utility in connection with any type of isochronous information for which synchrony among devices is desired. The system is such that synchrony groups are created and destroyed dynamically, and in such a manner as to avoid requiring a dedicated device as the master device.

It will be appreciated that a number of changes and modifications may be made to the network audio system 10 as described above. For example, although the invention has been described as providing that the audio information channel device 23 provides digital audio information to the members synchrony group 20 that has been encoded using particular types of encoding and compression methodologies, it will be appreciated that the audio information channel device 23 can provide digital audio information to various members of the synchrony group 20 that have been encoded and compressed using different types of encoding and compression methodologies, and, moreover, for which different sampling rates have been used. For example, the audio information channel device 23 may provide digital audio information to the master device 21 and slave devices 22(1) through 22(g_1) using the MP3 methodology at a specified sampling rate, the digital audio information for the same program to slave devices 22(g_1+1) through 22(g_2) using the WAV methodology at one specified sampling rate, and to slave devices 22(g_2+1) through 22(G)

using the WAV methodology at another specified sampling rate. In that case, the audio information channel device 23 can specify the particular encoding and compression methodology that has been used in the encoding type field 57 associated with each frame and the sampling rate in the sampling rate field 58. Moreover, since the encoding and compression type and sampling rate are specified for each frame, the encoding and compression type and sampling rate can be changed from frame to frame. The audio information channel device 23 may use different multi-cast addresses for the different encoding and compression types and sampling rates, but it will be appreciated that that would not be required.

It will be appreciated that two advantages of providing that the encoding and compression methodology and the sampling rate is provided on a frame-by-frame basis, instead of on, for example, a track-by-track basis, is that that would facilitate a slave device joining the synchrony group 20 at a frame mid-track, without requiring, for example, the master device 21 or the audio information channel device 23 to notify it of the encoding and compression methodology or the sampling rate.

Another modification is that, instead of the network communications manager 40 of a member of a synchrony group 20 generating the updated time stamp T^u_F for a digital audio information frame by adding the time differential value ΔT to the time stamp T_F associated with a frame, the network communications manager 40 may instead generate the updated time stamp T^U_F by subtracting the differential time value ΔT from the member's current time T_S as indicated by the member's digital to analog converter clock 34 at the time at which the digital audio information is received. It will be appreciated, however, that there may be variable time delays in processing of messages by the slave device's network communications manager 40, and so it may be preferable to generate the time differential value ΔT using the time stamp T_F provided by the audio information channel device 23.

In addition, instead of the network communications manager 40 of a member of a synchrony group generating an updated time stamp to reflect the difference between the times indicated by the member's digital to analog converter clock and the audio information channel device's digital to analog converter clock, the network communications manager 40 can generate the time differential value ΔT and provide it to the member's playback scheduler 32. In that case, the member's network communications manager 40 can store each digital audio information frame along with the time stamp T_F as received from the master device in the audio information buffer 21. The playback scheduler 32 can utilize the time differential value ΔT, and the time stamps T_F associated with the

digital audio information frames, to determine when the respective digital audio information frames are to be played. In determining when a digital audio information frame is to be played, the playback scheduler can add the time differential value to the time stamp T_F associated with the digital audio frame, and enable the digital audio frame to be coupled to the digital to analog converter 33 when the time indicated by the sum corresponds to the current time as indicated by the slave device's digital to analog converter clock 34. Alternatively, when the member's digital to analog converter clock 34 updates its current time T_S, the playback scheduler can generate an updated current time T'_S by subtracting the differential time value ΔT from the current time T_S, and using the updated current time T'_S to determine when to play a digital audio information frame.

As described above, the members of a synchrony group 20 periodically obtain the audio information channel device's current time value and uses the current time value that it receives from the audio information channel device to periodically update the time differential value ΔT that it uses in updating the time stamps associated with the various frames. It will be appreciated that, if the digital to analog converter clock(s) associated with the member(s) of a synchrony group 20 are ensured to have the same rate as the digital to analog converter clock, a member need only obtain the current time value from the audio information channel device once, at the beginning of playback.

As another alternative, if the zone players are provided with digital to analog converter clock 34 whose time and rate can be set by an element such as the network communications manager 40, when a zone player 11(n) is operating as a member of a synchrony group 20, its network communications manager 40 can use the various types of timing information that it receives from the audio information channel device 23, including the current time information and the playback timing information indicated by the time stamps that are associated with the various frames 51(f) comprising the audio and playback timing information that it receives, to adjust the synchrony group member's digital to analog converter clock's time value and/or the clock rate that it uses for playback. If the clock's time value is to be adjusted, when the synchrony group member's network communications manager 40 initially receives the current time information from the audio information channel device 23 for the synchrony group 20, the network communications manager 40 can set the synchrony group member's digital to analog converter clock 34 to the current time value as indicated by the audio information channel device's current time information. The network communications manager 40 can set the clock 34 to the current time value indicated by the audio information channel device's current time information once, or periodically as it receives the current time information.

Alternatively or in addition, the synchrony group member's network communications manager 40 can use one or both of the current time information and/or the playback timing information in the time stamps associated with the respective frames 51(f) to adjust the clock rate of the clock 34 that it uses for playback. For example, when the synchrony group member's network communications manager 40 receives a frame 51(f_X) having a time stamp having a time value T_{f_X}, it can generate the updated time value $T^U{}_{f_X} = T_{f_X} + \Delta T$ as described above, and store the frame with the time stamp with the updated time value in the audio information buffer 30. In addition, since both the number of samples in a frame and the sampling rate, which determines the rate at which the frame is to be played, are known to the network communications manager 40, it can use that information, along with the updated time value $T^U{}_{F_X}$ that is to be used for frame 51(f_X) to generate an expected updated time value $T^E{}_{f_{X+1}}$ that is expected for the updated time stamp of the next frame 51(f_{X+1}). After the synchrony group member's network communications manager 40 receives the next frame 51(f_{X+1}), it can generate the updated time value $T^U{}_{f_{X+1}}$ and compare that value to the expected updated time value $T^E{}_{f_{X+1}}$. If the two time values do not correspond, or if the difference between them is above a selected threshold level, the clock that is used by the audio information channel device 23 to generate the time stamps is advancing at a different rate than the synchrony group member's digital to analog converter clock 34, and so the network communications manager 40 can adjust the rate of the digital to analog converter clock 34 to approach that of the clock used by the audio information channel device 23 so that the differential time value ΔT is constant. On the other hand, if the two time values do correspond, then the time differential value ΔT is constant, or the difference is below a threshold level, and the network communications manager 40 need not change the clock rate of the digital to analog converter clock 34. It will be appreciated that, if the clock rate is to be adjusted, the rate adjustment can be fixed, or it can vary based on, for example, the difference between the updated time value $T^U{}_{f_{X+1}}$ and the expected updated time value $T^E{}_{f_{X+1}}$.

It will also be appreciated that, if no rate adjustment is performed for one frame 51(f_{X+1}), the synchrony group member's network communications manager 40 can generate an expected updated

time value $T^{E}_{f_{X+2}}$ that is expected for the updated time stamp of the next frame 51(f_{X+2}) using the updated time value $T^{U}_{F_X}$ determined for frame 51(f_X), along with the number of samples in a frame and the sampling rate, and compare the expected updated time value $T^{E}_{f_{X+2}}$ to the updated time value $T^{U}_{f_{X+2}}$ that it generates when it receives frame 51(f_{X+2}). At that point, if the network communications manager 41 determines that two time values do not correspond, or if the difference between them is above a selected threshold level, it can adjust the rate of the digital to analog converter clock 34. Similar operations can be performed if no rate adjustment is performed for several successive frames 51(f_{X+1}), 51(f_{X+2}),.... This will accommodate the possibility that the rate differential between the clock 34 and the clock used by the audio information channel device 23 in generating the time stamps have rates that differ by an amount sufficiently small that it cannot be detected using time stamps of two or more successive frames.

Instead or in addition to adjusting the clock rate as described above, the synchrony group member's network communications manager 40 can perform similar operations in connection with adjusting the clock rate in connection with the current time information that it receives from the audio information channel device 23.

Furthermore, although the network audio system 10 has been described such that the master device 21 of a synchrony group 20 can, in response to control information provided thereto by a user through the user interface module 13, provide a notification to a zone player 11(n) that it is to become a member of its synchrony group 20 as a slave device 22(g), it will be appreciated that the user interface module 13 can provide the notification directly to the zone player 11(n) that is to become a member of the synchrony group 20. In that case, the zone player 11(n) can notify the master device 21 that it is to become a slave device 22(g) in the synchrony group 20, after which the master device 21 can provide information regarding the synchrony group 20, including the multi-cast and unicast addresses of the audio information channel device and other information as described above.

Similarly, although the network audio system 10 has been described such that the master device 21 of a synchrony group 20 can, in response to control information provided thereto by a user through the user interface module 13, provide a command to a slave device 22(g) to enable the slave device 22(g) to adjust its volume, it will be appreciated that the user interface module 13 can provide

control information directly to the slave device 22(g) to enable the slave device 22(g) to adjust its volume.

In addition, although the network audio system 10 has been described such that each frames 51(f) is associated with a frame sequence number (reference field 56, FIG. 4), it will be appreciated that, if the packets described above in connection with Packet Sequence A and Packet Sequence B are provided with packet sequence numbers, the frame sequence numbers need not be provided, since the packet sequence numbers can suffice for defining the frame sequencing.

Furthermore, although the network audio system 10 has been described such that the zone players 11(n) are provided with an audio amplifier 35 for amplifying the analog signal provided by the respective digital to analog converters 33, it will be appreciated that a zone player may be provided that does not itself include an audio amplifier. In that case, the analog signal may be coupled to an external amplifier for amplification as necessary before being provided to the audio reproduction device(s) 15(n)(r). It will be appreciated that a single zone player 11(n) may be provided with multiple audio amplifiers and audio reproduction device interfaces, and, if necessary, multiple digital to analog converters 33, to provide audio programs for corresponding numbers of synchrony groups.

Similarly, although the zone players 11(n) have been described such that they may be connected to one or more audio information sources, it will be appreciated that an audio information source may form part of and be integrated into a zone player 11(n). For example, a zone player may include a compact disk player, cassette tape player, broadcast radio receiver, or the like, that has been integrated into it. In addition, as noted above, an individual zone player 11(n) may be connected to multiple audio information sources and may contemporaneously operate as the audio information channel device 23 for multiple synchrony groups .

In addition, although FIG. 1 shows the network audio system 10 as including one user interface module 13, it will be appreciated that the system 10 may include a plurality of user interface modules. Each user interface module be useful for controlling all of the zone players as described above, or alternatively one or more of the user interface modules may be useful for controlling selected subsets of the zone players.

Moreover, it will be appreciated that, although the invention has been described in connection with audio information, it will be appreciated that the invention will find utility in connection with any type of information for which synchrony among devices connected to a network is desired.

As noted above, while a zone player 11(n) is operating as audio information channel device 23 for a synchrony group 20, when the zone player 11(n)'s audio information source interface 30 or network communications manager 40 stores digital audio information frames based on audio information from an audio information source 14(n)(s) in the audio information buffer 31, it will provide time stamps for the respective frames to schedule them for playback after some time delay after they have been buffered in the audio information buffer 31. The delay is provided so that, for other zone players 11(n'), 11(n''),... that are operating as members of a synchrony group, there will be sufficient time for the audio and playback timing information to be transferred over the network 12 to those other zone players 11(n'), 11(n''),... so that it can be processed and played by them at the appropriate time as described above. The time period that is selected for the time delay may be fixed or variable, and in either case may be based on a number of factors. If the time period selected for the time delay is fixed, it may be based on, for example, factors such as an estimate of the maximum latency in the network 12, the estimated maximum loading of the various components comprising the zone players 11(n), and other estimates as will be appreciated by those skilled in the art.

The time delay may be the same for audio information from all types of audio information sources, and may be constant over the entire period that the synchrony group 20 is playing an audio work. Alternatively, different time delays may be utilized based on various criteria. For example, if the audio information is to be played independently of information associated with other types of media, the time delay may be selected to be relatively long, on the order of a significant fraction of a second, or longer. On the other hand, if the audio information is to be played contemporaneously with, for example, video information, which may be supplied by, for example, a video disk, video tape cassette, over cable, satellite, or broadcast television, which may not be buffered or which may be displayed independently of the network audio system 10, it may be undesirable to provide for such a lengthy delay, since the time delay of the audio playback, in relation to the video display, may be noticeable. In that case, the zone player 11(n) may provide for a much shorter time delay. In one embodiment, the time delay provided for audio information to be played concurrently with video information is selected to be generally on the order of fifty milliseconds, which would barely, if at all, be perceptible to someone viewing the video. Other desirable time delays for information from other types of sources will be apparent to those skilled in the art.

As yet a further possibility, the zone player 11(n), when operating as an audio information channel device 23 for a synchrony group 20, can dynamically determine the time delay based on a number of conditions in the network audio system 10, including, for example, the message transfer

latency in network 12, the loading of microprocessors or other components that are used in the various zone players 11(n'), 11(n''),... that may comprise a synchrony group 20, as well as other factors. For example, if the audio information channel device 23 determines that the latency in the network 12 has increased beyond a selected threshold, the audio information channel device 23 can adjust the delay to increase the likelihood that the members of the synchrony group 20 will be able to receive the packets and process the frames so that they will be able to play them at the appropriate times. Similarly, if the audio information channel device 23 is notified that a member of the synchrony group 20 to which it provides audio information requires additional time to receive and process the frames that it transmits, the audio information channel device 23 can adjust the delay accordingly. It will be appreciated that, to reduce or minimize possible discontinuities in the audio playback by the members of the synchrony group, the audio information channel device 23 can, instead of adjusting the time delay during a particular audio track, adjust the time delay between tracks, during silent periods of a track or otherwise as will be appreciated by those skilled in the art. In addition, the audio information channel device 23 can use conventional audio compression methodologies to facilitate a speeding up and/or slowing down of playback of an audio track while it is in the process of providing additional time delay. Generally, the members of the synchrony group 20 can provide notifications to the audio information channel device 23 if they determine that they will need an additional time delay, and the audio information channel device 23 can adjust the time delay in accordance with the notifications from the members of the synchrony group 20.

It will be appreciated that a system in accordance with the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof, any portion of which may be controlled by a suitable program. Any program may in whole or in part comprise part of or be stored on the system in a conventional manner, or it may in whole or in part be provided in to the system over a network or other mechanism for transferring information in a conventional manner. In addition, it will be appreciated that the system may be operated and/or otherwise controlled by means of information provided by an operator using operator input elements (not shown) which may be connected directly to the system or which may transfer the information to the system over a network or other mechanism for transferring information in a conventional manner.

The foregoing description has been limited to a specific embodiment of this invention. It will be apparent, however, that various variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention. It is the object of the appended

claims to cover these and such other variations and modifications as come within the true spirit and scope of the invention.

What is claimed as new and desired to be secured by Letters Patent is:

CLAIMS

1. A system comprising plurality of devices, one of the devices operating as a task source device and at least one other device operating as a member of a synchrony group,

A. the task source device being configured to distribute a series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to a clock maintained by the task source device, at which the devices comprising the synchrony group are to execute the respective task;

B each member device being configured to:

(i) periodically obtain from the task source device an indication of a current time value indicated by the task source device's clock, and

(ii) determine, from the time stamp associated with each respective task and a time differential value representing a difference between the current time value indicated by the task source device's clock, and a current time value indicated by its respective clock, a time, relative to its respective clock, at which it is to execute the task.

2. A system as defined in claim 1 in which the synchrony group comprises a plurality of member devices.

3. A system as defined in claim 2 in which each device comprising a member of the synchrony group is further configured to execute each task that it receives from the task source device at the determined time, thereby to facilitate execution by them of respective tasks in the series in respective timing relationships relative to the time indicated by the task source device's clock.

4. A system as defined in claim 3 in which the member devices are configured to execute respective tasks in synchrony.

5. A system as defined in claim 2 in which one of the member devices operates as a master device for the synchrony group, any other member devices comprising respective slave devices, the master device being configured to perform at least one type of synchrony group management operation in connection with the member devices comprising the synchrony group.

6. A system as defined in claim 5 further including a user interface module configured to control the master device's performance of said at least one type of synchrony group management operation, the master device being further configured to provide status information relating to the status of the synchrony group to the user interface module.

7. A system as defined in claim 6 in which the status information includes identifications of the devices comprising the synchrony group.

8. A system as defined in claim 6 in which the status information includes an identification of the task currently being executed.

9. A system as defined in claim 5, the system comprising at least one additional device, in which, in one type of synchrony group management operation, the master device is configured to enable the at least one additional device to join the synchrony group as a slave device.

10. A system as defined in claim 9 in which the task source device is configured to distribute tasks to the member devices using a selected multi-cast transmission methodology, the member devices being configured to buffer the tasks until they are to be executed, and further in which, when the at least one additional device joins the synchrony group as a slave device, the task source device is enabled to transmit at least one previously distributed task to the slave device using a selected unicast transmission methodology.

11. A system as defined in claim 5 in which, in at least one type of synchrony group management operation, the master device is configured to enable the task source device to join the synchrony group as a slave device, the task source device continuing to operate as the task source device for the synchrony group.

12. A system as defined in claim 5 in which, in at least one type of synchrony group management operation, the master device is configured to enable a slave device to disengage from the synchrony group, the disengaged slave device thereafter not being a member device in the synchrony group.

13. A system as defined in claim 5 in which, in at least one type of synchrony group management operation, the master device is configured to control the series of tasks to be distributed by the task source device.

14. A system as defined in claim 5 in which, in at least one type of synchrony group management operation, the master device is further configured to control execution by the member devices of the tasks that have been distributed by the task source device.

15. A system as defined in claim 14 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, and further in which in at least one type of synchrony group management operation, the master device is configured to enable the member devices to terminate execution of a task sequence that is currently being executed.

16. A system as defined in claim 15 in which, in the synchrony group management operation enabling the member devices to terminate execution of a task sequence that is currently being executed, the master device enables the task source device to distribute a command to enable the

member devices to terminate execution of the task sequence, the member devices being configured to, after receiving the command, terminate execution of the task sequence.

17. A system as defined in claim 14 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, and further in which at least one type of synchrony group management operation, the master device is configured to enable the member devices to cancel execution of a task sequence that has been previously distributed, but for which execution has not begun.

18. A system as defined in claim 17 in which, in the synchrony group management operation enabling the member devices to cancel execution of a task sequence that has been previously distributed but for which execution has not begun, the master device enables the task source device to distribute a command to enable the member devices to cancel execution of the task sequence, the member devices being configured to, after receiving the command, cancel execution of the task sequence.

19. A system as defined in claim 5 in which, in at least one type of synchrony group management operation, the member device operating as the master device is configured to enable the master device to migrate from one member device to another member device in the synchrony group.

20. A system as defined in claim 5 in which, in at least one type of synchrony group management operation, the master device is configured to enable the task source device to migrate from one device to another device in the system.

21. A system as defined in claim 20 in which

A. the device operating as the task source device is configured to, after being enabled to migrate the task source device to said other device, provide migration information to the other device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices,

B. the other device being configured to, after receiving the migration information,

   i. distribute the series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the task source device, at which the devices comprising the synchrony group are to execute the respective task, and

   ii. notify the members of the synchrony group that it is to thereafter operate as the task source device.

22. A system as defined in claim 21 in which the information respecting the tasks that are to be distributed is a source of streaming information, the other device being configured to, after receiving the migration information, assemble the streaming information into tasks and associate each task with a said time stamp.

23. A system as defined in claim 21 in which the information respecting the tasks that are to be distributed is in a file identified in the migration information, the other device being configured to obtain information from the identified file, assemble the information into tasks and associate each task with a said time stamp.

24.A system as defined in claim 23 in which the information respecting the tasks that are to be distributed begins at an offset, identified in the migration information, into the identified file, the other device being configured to assemble the information from a position in the file associated with the identified offset.

25. A system as defined in claim 23 in which the information respecting the tasks that are to be distributed is in a series of files identified in the migration information, the other device being configured to obtain information from successive files in the series assemble the information into tasks and associate each task with as said time stamp.

26. A system as defined in claim 1 in which the task source device is further configured to distribute the tasks using a selected multi-cast message transmission methodology.

27. A system as defined in claim 1 in which the task source device is configured to provide, associated with at least one task, a time stamp indicating a task execution time later than the time currently indicated by the clock maintained by the task source device.

28. A system as defined in claim 27 in which the task source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be executed by the member devices at the time indicated by the time stamp.

29. A system as defined in claim 28 in which the task source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be distributed to the member devices for execution at the time indicated by the time stamp

30. A system as defined in claim 27 in which task source device is configured to obtain information associated with the tasks from at least two types of information sources, the task source device being further configured to select, for each task the delay interval between the time indicated by the time

stamp and the current time based on the information source type that is associated with the respective task.

31. A system as defined in claim 1 in which at least one member device is further configured to adjust its clock rate in relation to a clock rate value maintained by the task source device's clock.

32. A system as defined in claim 31 in which the at least one member device is configured to adjust its clock rate in relation to time stamps associated with respective ones of the tasks received from the task source device.

33. A system as defined in claim 1 in which at least one other device operates as a task source device configured to distribute tasks to a second synchrony group, the device operating as the task source device for the first synchrony group also operating as a member device of a second synchrony group.

34. A system as defined in claim 1 in which at least one member device comprises:

A. a task receiving module configured to receive the series of tasks,

B. a current time retrieval module configured to obtain, from the task source device, the current time value as indicated by the task source's clock;

C. an execution time determination module configured to determine, from the time stamp associated with each respective task and a time differential value representing a difference between the current time value obtained by the current time retrieval module and a current time value indicated by a clock maintained by the device, a time, relative to the member device's respective clock at which the task is to be executed; and

D. a task execution module configured to execute each respective task at the time determined by the execution time determination module.

35. A system as defined in claim 34 further including a control module for controlling execution of commands received by said interface module.

36. A system as defined in claim 35 in which the at least one member device further comprises a buffer configured to buffer tasks that the task receiving module receives from the task source device until they are to be executed, the task execution module being configured to execute the buffered tasks at the times determined by the execution time determination module.

37. A system as defined in claim 36 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the interface module being further configured to receive task execution control commands for controlling task execution, the control module being configured to, in response to receipt of at least one task execution control command for controlling task execution, enable the task execution module to terminate execution of a task

sequence currently being executed, being configured to not execute subsequent tasks in the task sequence that have been buffered.

38. A system as defined in claim 37 in which the control module is configured to enable the task execution module to resume execution of at least one task associated with a subsequent task sequence at a time determined by the execution time determination module.

39. A system as defined in claim 36 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the interface module being further configured to receive task execution control commands for controlling task execution, the control module being configured to, in response to receipt of at least one task execution control command to cancel execution of a task sequence for which at least one task has been buffered but for which execution has not begun, enable the task execution module to not execute tasks from that task sequence that have been buffered.

40. A system as defined in claim 39 in which the control module is configured to enable the task execution module to resume execution of at least one task associated with a task sequence subsequent to the task sequence for which at least one task has been buffered but for which execution has not begun.

41. A system as defined in claim 34 in which the member device includes:

A. a migration information receiving module configured to receive migration information from the task source device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group, and

B. a migration control module configure to, after the migration information has been received.,

   i. distribute the series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the task source device, at which the devices comprising the synchrony group are to execute the respective task, and

   ii. notify the members of the synchrony group that it is to thereafter operate as the task source device.

42. A system as defined in claim 41 in which the information respecting the tasks that are to be distributed is a source of streaming information, the migration control module being configured to, after the migration information has been received, assembling the streaming information into tasks and associate each task with a said time stamp.

43. A system as defined in claim 41 in which the information respecting the tasks that are to be distributed is in a file identified in the migration information, the migration control module being configured to obtain information from the identified file, assemble the information into tasks and associate each task with a said time stamp.

44.A system as defined in claim 43 in which the information respecting the tasks that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control module being configured to assemble the information from a position in the file associated with the identified offset.

45. A system as defined in claim 43 in which the information respecting the tasks that are to be distributed is in a series of files identified in the migration information, the migration control module being configured to obtain information from successive files in the series assemble the information into tasks and associate each task with as said time stamp.

46. A system as defined in claim 34 in which the member device further includes a clock rate adjustment module configured to adjust the member device's clock rate in relation to a clock rate value maintained by the task source device's clock.

47. A system as defined in claim 46 in which the clock rate adjustment module is configured to adjust the clock rate in relation to time stamps associated with respective ones of the tasks received from the task source device.

48. A system as defined in claim 1 in which the task source device comprises:

A. a task acquisition module configured to obtain respective ones of the tasks;

B. a task execution time determination module configured to determine a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission module configured to transmit the series of tasks to said at least one other device; and

D. an execution control command generation module configured to, in response to a predetermined event, enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to modify the execution sequence of the tasks transmitted thereto.

49. A system as defined in claim 48 in which the task transmission module is configured to utilize a selected multi-cast transmission methodology.

50. A system as defined in claim 48 in which the predetermined event comprises input indicia provided by an operator.

51. A system as defined in claim 50 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to control execution of at least one task in relation to the time the predetermined event occurs.

52. A system as defined in claim 51 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to terminate execution of a task that it is currently executing and/or to not execute at least one task that is to be executed subsequent to the predetermined event.

53. A system as defined in claim 51 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation module being configured to, in response to said predetermined event,

(i) enable the task transmission module to insert into the series of tasks transmitted by the task transmission module, a command enabling the at least one other device to terminate execution of the currently executing task and/or to not execute at least one subsequent task, if any, in the task sequence including the currently executing task; and

(ii) enable

(a) the task execution time determination module to provide that time stamps associated with tasks of task sequences subsequent to the task sequence currently being executed will reflect the termination and/or non-execution of tasks from the task sequence of the task that is currently being executed, and

(b) the task transmission module to transmit the tasks of the subsequent task sequences to the at least one other device.

54. A system as defined in claim 53 in which the task transmission module is configured to insert a command to enable the at least one other device to not execute all tasks in the task sequence subsequent to the terminated and/or non-executed task.

55. A system as defined in claim 48 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert

into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to control execution of a task in relation to the task's position in the series of tasks.

56. A system as defined in claim 55 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to cancel execution of a task in relation to the task's position in the series of tasks.

57. A system as defined in claim 55 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted task in relation to the task's position in the series of tasks.

58. A system as defined in claim 57 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted task, in relation to the task's position in the series of tasks, for which execution has not begun.

59. A system as defined in claim 55 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation module being configured to, in response to said predetermined event,

(i) enable the task transmission module to insert into the series of tasks transmitted by the task transmission module, a command enabling the at least one other device to cancel execution of a task sequence for which execution has not begun; and

(ii) enable

  (a) the task execution time determination module to provide that time stamps associated with tasks of task sequences subsequent to the task sequence whose execution was cancelled will reflect the cancellation, and

  (b) the task transmission module to transmit the tasks of the subsequent task sequences to the at least one other device.

60 A system as defined in claim 1 in which the task source device comprises:

A. a task acquisition module configured to obtain respective ones of the tasks;

B. a task execution time determination module configured to determine a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission module configured to transmit the series of tasks to said at least one other device; and

D. a task source device migration control module configured to, in response to a predetermined event, provide migration information to another device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group.

61. A system as defined in claim 60 in which the task source device is configured to provide, associated with at least one task, a time stamp indicating a task execution time later than the time currently indicated by the clock maintained by the task source device.

62. A system as defined in claim 61 in which the task source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be executed by the member devices at the time indicated by the time stamp.

63. A system as defined in claim 62 in which the task source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be distributed to the member devices for execution at the time indicated by the time stamp

64. A system as defined in claim 62 in which task source device is configured to obtain information associated with the tasks from at least two types of information sources, the task source device being further configured to select, for each task the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective task.

65. A device for executing a series of tasks provided by a task source at times specified by the task source in relation to a clock maintained by the task source, the device comprising

A. an interface module configured to receive the series of tasks, each task being associated with a time stamp, each time stamp indicating a time value relative to a time indicated a clock maintained by the task source, at which the device is to execute the respective task;

B. a current time retrieval module configured to obtain, from the task source, a current time value as indicated by the task source's clock;

C. an execution time determination module configured to determine, from the time stamp associated with each respective task and a time differential value representing a difference between the current time value obtained by the current time retrieval module and a current time value indicated by a clock maintained by the device, a time, relative to the device's respective clock, at which the task is to be executed; and

D. a task execution module configured to execute each respective task at the time determined by the execution time determination module.

66. A device as defined in claim 65 further including a control module for controlling execution of commands received by said interface module.

67. A device as defined in claim 65 further comprising a buffer configured to buffer tasks that the interface module has received until they are to be executed, the task execution module being configured to execute the buffered tasks at the times determined by the execution time determination module.

68. A device as defined in claim 67 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the interface module being further configured to receive task execution control commands for controlling task execution, the control module being configured to, in response to receipt of at least one task execution control command for controlling task execution, enable the task execution module to terminate execution of a task sequence currently being executed, being configured to not execute subsequent tasks in the task sequence that have been buffered.

69. A device as defined in claim 68 in which the control module is configured to enable the task execution module to resume execution of at least one task associated with a subsequent task sequence at a time determined by the execution time determination module.

70. A device as defined in claim 69 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the interface module being further configured to receive task execution control commands for controlling task execution, the control module being configured to, in response to receipt of at least one task execution control command to cancel execution of a task sequence for which at least one task has been buffered but for which execution has not begun, enable the task execution module to not execute tasks from that task sequence that have been buffered.

71. A device as defined in claim 70 in which the control module is configured to enable the task execution module to resume execution of at least one task associated with a task sequence subsequent to the task sequence for which at least one task has been buffered but for which execution has not begun.

72. A device as defined in claim 66, the control module being configured to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive tasks for execution by the task execution module.

73. A device as defined in claim 72 in which the control module is configured to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive tasks currently being transmitted using a selected multi-cast transmission methodology, and further to retrieve at least one previously distributed task using a selected unicast transmission methodology.

74. A device as defined in claim 65 in which the control module is configured to, in receipt of a command to disengage from a synchrony group, disable the interface module from receiving tasks and the task execution module from executing previously received tasks that have not yet been executed.

75. A device as defined in claim 65 further comprising a user interface module interface module configured to receive control information from a user interface module, the control module being configured to enable selected operations to be performed in response to the control information.

76. A device as defined in claim 75 in which, in response to control information to enable another device to become a member of the device's synchrony group, the control module enables the interface module to transmit a command to the other device to enable the other device to become a member of the synchrony group.

77. A device as defined in claim 75 in which, in response to control information to enable another device that is a member of the device's synchrony group to disengage from the synchrony group, the control module enables the interface module to transmit a command to the other device to enable the other device to disengage from the synchrony group.

78. A device as defined in claim 75 in which, in response to control information to enable the task source to begin operating, the control module performs selected operations to enable the task source to begin operating.

79. A device as defined in claim 78, the device also operating as the task source, the device including a task generator module configured to, under control of the control module, obtain information from which the tasks are to be generated and generate the tasks.

80. A device as defined in claim 79 in which the interface module is further configured to transmit the tasks to at least one other device.

81. A device as defined in claim 79 in which the device is not the task source, the control module being configured to enable the interface module to transmit a command to enable the task source to operate.

82. A device as defined in claim 81 in which tasks are divided into a series of task sequences, and, in response to receipt of control information enabling execution of a sequence of tasks currently being executed to be terminated , the control module is configured to enable the interface module to transmit a command to enable the task source to insert a resynchronize command in the series of tasks associated with the task sequence to enable termination of execution of the task sequence.

83. A device as defined in claim 81 in which tasks are divided into a series of task sequences, and, in response to receipt of control information enabling execution of a sequence of tasks at least a portion of which have been received, but which are not being executed, the control module is configured to enable the interface module to transmit a command to enable the task source to insert a resynchronize command in the series of commands to disable execution of the task sequence.

84. A device as defined in claim 75 in which the control module is configured to enable the user interface module interface module to transmit status information to the user interface module.

85. A device as defined in claim 84 in which the status information includes the execution status of at least one task.

86. A device as defined in claim 65 further including:

A. a migration information receiving module configured to receive migration information from the task source device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group, and

B. a migration control module configured to, after the migration information has been received.,

   i. distribute the series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the task source device, at which the devices comprising the synchrony group are to execute the respective task, and

ii. notify the members of the synchrony group that it is to thereafter operate as the task source device.

87. A device as defined in claim 86 in which the information respecting the tasks that are to be distributed is a source of streaming information, the migration control module being configured to, after the migration information has been received, assembling the streaming information into tasks and associate each task with a said time stamp.

88. A device as defined in claim 86 in which the information respecting the tasks that are to be distributed is in a file identified in the migration information, the migration control module being configured to obtain information from the identified file, assemble the information into tasks and associate each task with a said time stamp.

89. A device as defined in claim 88 in which the information respecting the tasks that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control module being configured to assemble the information from a position in the file associated with the identified offset.

90. A device as defined in claim 89 in which the information respecting the tasks that are to be distributed is in a series of files identified in the migration information, the migration control module being configured to obtain information from successive files in the series assemble the information into tasks and associate each task with as said time stamp.

91. A device as defined in claim 65 further including a clock rate adjustment module configured to adjust the member device's clock rate in relation to a clock rate value maintained by the task source device's clock.

92. A task source device for distributing a series of tasks for execution by at least one other device at respective execution times, the task source device comprising:

A. a task acquisition module configured to obtain respective ones of the tasks;

B. a task execution time determination module configured to determine a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission module configured to transmit the series of tasks to said at least one other device; and

D. an execution control command generation module configured to, in response to a predetermined event, enable the task transmission module to insert into the series of tasks

transmitted by the task transmission module a command to enable the at least one other device to modify the execution sequence of the tasks transmitted thereto.

93. A task source device as defined in claim 92 in which the task transmission module is configured to utilize a selected multi-cast transmission methodology.

94. A task source device as defined in claim 92 in which the predetermined event comprises input indicia provided by an operator.

95. A task source device as defined in claim 94 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to control execution of at least one task in relation to the time the predetermined event occurs.

96. A task source device as defined in claim 95 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to terminate execution of a task that it is currently executing and/or to not execute at least one task that is to be executed subsequent to the predetermined event.

97. A task source device as defined in claim 95 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation module being configured to, in response to said predetermined event,

(i) enable the task transmission module to insert into the series of tasks transmitted by the task transmission module, a command enabling the at least one other device to terminate execution of the currently executing task and/or to not execute at least one subsequent task, if any, in the task sequence including the currently executing task; and

(ii) enable

(a) the task execution time determination module to provide that time stamps associated with tasks of task sequences subsequent to the task sequence currently being executed will reflect the termination and/or non-execution of tasks from the task sequence of the task that is currently being executed, and

(b) the task transmission module to transmit the tasks of the subsequent task sequences to the at least one other device.

98. A task source device as defined in claim 97 in which the task transmission module is configured to insert a command to enable the at least one other device to not execute all tasks in the task sequence subsequent to the terminated and/or non-executed task.

99. A task source device as defined in claim 92 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to control execution of a task in relation to the task's position in the series of tasks.

100. A task source device as defined in claim 99 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to cancel execution of a task in relation to the task's position in the series of tasks.

101. A task source device as defined in claim 99 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted task in relation to the task's position in the series of tasks.

102. A task source device as defined in claim 101 in which, in response to said predetermined event, the execution control command generation module is configured to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted task, in relation to the task's position in the series of tasks, for which execution has not begun.

103. A task source device as defined in claim 99 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation module being configured to, in response to said predetermined event,

(i) enable the task transmission module to insert into the series of tasks transmitted by the task transmission module, a command enabling the at least one other device to cancel execution of a task sequence for which execution has not begun; and

(ii) enable

(a) the task execution time determination module to provide that time stamps associated with tasks of task sequences subsequent to the task sequence whose execution was cancelled will reflect the cancellation, and

(b) the task transmission module to transmit the tasks of the subsequent task sequences to the at least one other device.

104. A task source device for distributing a series of tasks for execution by at least one other device at respective execution times, the task source device comprising:

A. a task acquisition module configured to obtain respective ones of the tasks;

B. a task execution time determination module configured to determine a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time, the task execution time determination module being configured to provide, associated with at least one task, a time stamp indicating a task execution time later than the time currently indicated by a clock maintained by the task source device; and

C. a task transmission module configured to transmit the series of tasks to said at least one other device.

105. A system as defined in claim 104 in which the task source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be executed by the member devices at the time indicated by the time stamp.

106. A system as defined in claim 105 in which the task source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be distributed to the member devices for execution at the time indicated by the time stamp

107. A system as defined in claim 104 in which task source device is configured to obtain information associated with the tasks from at least two types of information sources, the task source device being further configured to select, for each task the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective task.

108. A task source device for distributing a series of tasks for execution by at least one other device, the task source device comprising:

A. a task acquisition module configured to obtain respective ones of the tasks;

B. a task execution time determination module configured to determine a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission module configured to transmit the series of tasks to said at least one other device; and

D. a task source device migration control module configured to, in response to a predetermined event, provide migration information to another device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group.

109. A method of operating a system comprising plurality of devices, one of the devices operating as a task source device and at least one other device operating as a member of a synchrony group,

A. the task source device being enabled to distribute a series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to a clock maintained by the task source device, at which the devices comprising the synchrony group are to execute the respective task;

B each member device being enabled to:

(i) periodically obtain from the task source device an indication of a current time value indicated by the task source device's clock, and

(ii) determine, from the time stamp associated with each respective task and a time differential value representing a difference between the current time value indicated by the task source device's clock, and a current time value indicated by its respective clock, a time, relative to its respective clock, at which it is to execute the task.

110. A method as defined in claim 109 in which the synchrony group comprises a plurality of member devices.

111. A method as defined in claim 110 in which each device comprising a member of the synchrony group is further enabled to execute each task that it receives from the task source device at the determined time, thereby to facilitate execution by them of respective tasks in the series in respective timing relationships relative to the time indicated by the task source device's clock.

112. A method as defined in claim 111 in which the member devices are enabled to execute respective tasks in synchrony.

113. A method as defined in claim 110 in which one of the member devices operates as a master device for the synchrony group, any other member devices comprising respective slave devices, the master device being enabled to perform at least one type of synchrony group management operation in connection with the member devices comprising the synchrony group.

114. A method as defined in claim 113 further including a user interface module enabled to control the master device's performance of said at least one type of synchrony group management operation, the master device being further enabled to provide status information relating to the status of the synchrony group to the user interface module.

115. A method as defined in claim 114 in which the status information includes identifications of the devices comprising the synchrony group.

116. A method as defined in claim 114 in which the status information includes an identification of the task currently being executed.

117. A method as defined in claim 113, the method comprising at least one additional device, in which, in one type of synchrony group management operation, the master device is enabled to enable the at least one additional device to join the synchrony group as a slave device.

118. A method as defined in claim 117 in which the task source device is enabled to distribute tasks to the member devices using a selected multi-cast transmission methodology, the member devices being enabled to buffer the tasks until they are to be executed, and further in which, when the at least one additional device joins the synchrony group as a slave device, the task source device is enabled to transmit at least one previously distributed task to the slave device using a selected unicast transmission methodology.

119. A method as defined in claim 113 in which, in at least one type of synchrony group management operation, the master device is enabled to enable the task source device to join the synchrony group as a slave device, the task source device continuing to operate as the task source device for the synchrony group.

120. A method as defined in claim 113 in which, in at least one type of synchrony group management operation, the master device is enabled to enable a slave device to disengage from the synchrony group, the disengaged slave device thereafter not being a member device in the synchrony group.

121. A method as defined in claim 113 in which, in at least one type of synchrony group management operation, the master device is enabled to control the series of tasks to be distributed by the task source device.

122. A method as defined in claim 113 in which, in at least one type of synchrony group management operation, the master device is further enabled to control execution by the member devices of the tasks that have been distributed by the task source device.

123. A method as defined in claim 122 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, and further in which in at least one type of synchrony group management operation, the master device is enabled to enable the member devices to terminate execution of a task sequence that is currently being executed.

124. A method as defined in claim 123 in which, in the synchrony group management operation enabling the member devices to terminate execution of a task sequence that is currently being executed, the master device enables the task source device to distribute a command to enable the member devices to terminate execution of the task sequence, the member devices being enabled to, after receiving the command, terminating execution of the task sequence.

125. A method as defined in claim 122 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, and further in which at least one type of synchrony group management operation, the master device is enabled to enable the member devices to cancel execution of a task sequence that has been previously distributed, but for which execution has not begun.

126. A method as defined in claim 125 in which, in the synchrony group management operation enabling the member devices to cancel execution of a task sequence that has been previously distributed but for which execution has not begun, the master device enables the task source device to distribute a command to enable the member devices to cancel execution of the task sequence, the member devices being enabled to, after receiving the command, cancel execution of the task sequence.

127. A method as defined in claim 113 in which, in at least one type of synchrony group management operation, the member device operating as the master device is enabled to enable the master device to migrate from one member device to another member device in the synchrony group.

128. A method as defined in claim 113 in which, in at least one type of synchrony group management operation, the master device is enabled to perform the step of enabling the task source device to migrate from one device to another device in the system.

129. A method as defined in claim 128 in which

A. the device operating as the task source device is enable to perform the step of, after being enabled to migrate the task source device to said other device, providing migration

information to the other device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices,

B. the other device is enabled to, after receiving the migration information, perform the steps of

i. distributing the series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the task source device, at which the devices comprising the synchrony group are to execute the respective task, and

ii. notifying the members of the synchrony group that it is to thereafter operate as the task source device.

130. A method as defined in claim 129 in which the information respecting the tasks that are to be distributed is a source of streaming information, the other device being enabled to, after receiving the migration information, perform the steps of assembling the streaming information into tasks and associating each task with a said time stamp.

131. A method as defined in claim 129 in which the information respecting the tasks that are to be distributed is in a file identified in the migration information, the other device being enabled to perform the steps of obtaining information from the identified file, assembling the information into tasks and associating each task with a said time stamp.

132.A method as defined in claim 131 in which the information respecting the tasks that are to be distributed begins at an offset, identified in the migration information, into the identified file, the other device being enabled to perform the step of assembling the information from a position in the file associated with the identified offset.

133. A method as defined in claim 131 in which the information respecting the tasks that are to be distributed is in a series of files identified in the migration information, the other device being enabled to perform the step of obtaining information from successive files in the series assemble the information into tasks and associate each task with as said time stamp.

134. A method as defined in claim 109 in which the task source device is further enabled to distribute the tasks using a selected multi-cast message transmission methodology.

135. A method as defined in claim 109 in which the task source device is enabled to perform the step of providing, provide, associated with at least one task, a time stamp indicating a task execution time later than the time currently indicated by the clock maintained by the task source device.

136. A method as defined in claim 135 in which the task source device is enabled to perform the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be executed by the member devices at the time indicated by the time stamp.

137. A method as defined in claim 136 in which the task source device is enabled to perform the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be distributed to the member devices for execution at the time indicated by the time stamp

138. A method as defined in claim 135 in which task source device is enabled to perform the steps of obtaining information associated with the tasks from at least two types of information sources, and selecting, for each task, the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective task.

139. A method as defined in claim 109 in which at least one member device is further enabled to perform the step of adjusting its clock rate in relation to a clock rate value maintained by the task source device's clock.

140. A method as defined in claim 139 in which the at least one member device is enabled to perform the step of adjusting its clock rate in relation to time stamps associated with respective ones of the tasks received from the task source device.

141. A method as defined in claim 109 in which at least one other device operates as a task source device enabled to distribute tasks to a second synchrony group, the device operating as the task source device for the first synchrony group also operating as a member device of a second synchrony group.

142. A method as defined in claim 109 including, within a member device,

A. a task receiving step of receiving the series of tasks,

B. a current time retrieval step of obtaining, from the task source device, the current time value as indicated by the task source's clock;

C. an execution time determination step of determining, from the time stamp associated with each respective task and a time differential value representing a difference between the current time value obtained during the current time retrieval step and a current time value indicated by a clock maintained by the device, a time, relative to the member device's respective clock at which the task is to be executed; and

D. a task execution step of executing each respective task at the time determined during the execution time determination step.

143. A method as defined in claim 142 further including a control step for controlling execution of commands received during the task receiving step

144. A method as defined in claim 143 in which the at least one member device further comprises a buffer configured to buffer tasks received from the task source device during the task receiving step until they are to be executed, the task execution step including the step of executing the buffered tasks at the times determined during the execution time determination step.

145. A method as defined in claim 144 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the task receiving step being further enabled to receive task execution control commands for controlling task execution, the control step including the step of, in response to receipt of at least one task execution control command for controlling task execution, controlling the task execution step to terminate execution of a task sequence currently being executed, and further to not execute subsequent tasks in the task sequence that have been buffered.

146. A method as defined in claim 145 in which the control step includes the step of enabling the task execution step to resume execution of at least one task associated with a subsequent task sequence at a time determined during the execution time determination step.

147. A method as defined in claim 144 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the task receiving step including the step of receiving task execution control commands for controlling task execution, the control step including the step of, in response to receipt of at least one task execution control command to cancel execution of a task sequence for which at least one task has been buffered but for which execution has not begun, providing that tasks from that task sequence that have been buffered not be executed.

148. A method as defined in claim 147 in which the control step includes the step of enabling execution to resume of at least one task associated with a task sequence subsequent to the task sequence for which at least one task has been buffered but for which execution has not begun.

149. A method as defined in claim 142 in which the member device performs:

A. migration information receiving step of receiving migration information from the task source device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group, and

B. a migration control step of, after the migration information has been received.,

i. distributing the series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the task source device, at which the devices comprising the synchrony group are to execute the respective task, and

ii. notifying the members of the synchrony group that it is to thereafter operate as the task source device.

150. A method as defined in claim 149 in which the information respecting the tasks that are to be distributed is a source of streaming information, the migration control step includes the steps of, after the migration information has been received, assembling the streaming information into tasks and associating each task with a said time stamp.

151. A method as defined in claim 149 in which the information respecting the tasks that are to be distributed is in a file identified in the migration information, the migration control step including the steps of obtaining information from the identified file, assembling the information into tasks and associating each task with a said time stamp.

152. A method as defined in claim 151 in which the information respecting the tasks that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control step including the step of assembling the information from a position in the file associated with the identified offset.

153. A method as defined in claim 151 in which the information respecting the tasks that are to be distributed is in a series of files identified in the migration information, the migration control step including the steps of obtaining information from successive files in the series assemble the information into tasks and associating each task with as said time stamp.

154. A method as defined in claim 142 in which the member device is further enabled to perform a clock rate adjustment step including the step of adjusting the member device's clock rate in relation to a clock rate value maintained by the task source device's clock.

155. A method as defined in claim 154 in which the clock rate adjustment step includes the step of determining adjust the clock rate in relation to time stamps associated with respective ones of the tasks received from the task source device.

156. A method as defined in claim 109 including, within a task source device,

A. a task acquisition step of obtaining respective ones of the tasks;

B. a task execution time determination step of determining a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission step of transmitting the series of tasks to said at least one other device; and

D. an execution control command generation step of, in response to a predetermined event and during the task transmission step, the insertion into the series of tasks transmitted by the task transmission module of a command to enable the at least one other device to modify the execution sequence of the tasks transmitted thereto.

157. A method as defined in claim 156 in which the task transmission step includes the step of utilizing a selected multi-cast transmission methodology.

158. A method as defined in claim 156 in which the predetermined event includes the step of receiving input indicia provided by an operator.

159. A method as defined in claim 158 in which the execution control command generation step includes the step of enabling, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to control execution of at least one task in relation to the time the predetermined event occurs.

160. A method as defined in claim 159 in which, the execution control command generation step includes the step of enabling, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to terminate execution of a task that it is currently executing and/or to not execute at least one task that is to be executed subsequent to the predetermined event.

161. A method as defined in claim 159 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation step including the steps of, in response to said predetermined event,

(i) enabling during the task transmission step the insertion, into the series of tasks, of a command enabling the at least one other device to terminate execution of the currently executing task and/or to not execute at least one subsequent task, if any, in the task sequence including the currently executing task; and

(ii) enabling

(a) during the task execution time determination step, that time stamps associated with tasks of task sequences subsequent to the task sequence currently being executed will

reflect the termination and/or non-execution of tasks from the task sequence of the task that is currently being executed, and

(b) during the task transmission step, the transmission of the tasks of the subsequent task sequences to the at least one other device.

162. A method as defined in claim 161 in which the task transmission step includes the step of inserting a command into the task sequence to enable the at least one other device to not execute all tasks in the task sequence subsequent to the terminated and/or non-executed task.

163. A method as defined in claim 159 in which the execution control command generation step include the step of enabling, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to control execution of a task in relation to the task's position in the series of tasks.

164. A method as defined in claim 163 in which the execution control command generation step includes the step of, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to cancel execution of a task in relation to the task's position in the series of tasks.

165. A method as defined in claim 163 in which the execution control command generation step includes the step of enabling, in response to said predetermined event and during the task transmission step, the inserting into the series of tasks of a command to enable the at least one other device to cancel execution of a previously-transmitted task in relation to the task's position in the series of tasks.

166. A method as defined in claim 165 in which, the execution control command generation step includes the step of, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to cancel execution of a previously-transmitted task, in relation to the task's position in the series of tasks, for which execution has not begun.

167. A method as defined in claim 163 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation step including the steps of, in response to said predetermined event,

(i) enabling, during the task transmission step, the insertion, into the series of tasks, of a command enabling the at least one other device to cancel execution of a task sequence for which execution has not begun; and

(ii) enabling

(a) during the task execution time determination step, that time stamps associated with tasks of task sequences subsequent to the task sequence whose execution was cancelled will reflect the cancellation, and

(b) during the task transmission step, the transmission of the tasks of the subsequent task sequences to the at least one other device.

168. A method as defined in claim 109 in which the task source device is enabled to perform

A. a task acquisition step of obtaining respective ones of the tasks;

B. a task execution time determination step of determining a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission step of transmitting the series of tasks to said at least one other device; and

D. a task source device migration control step of, in response to a predetermined event, providing migration information to another device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group.

169. A method as defined in claim 109 in which the task source device is enabled to perform the step of providing, associated with at least one task, a time stamp indicating a task execution time later than the time currently indicated by the clock maintained by the task source device.

170. A method as defined in claim 169 in which the task source device is enabled to perform the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be executed by the member devices at the time indicated by the time stamp.

171. A method as defined in claim 170 in which the task source device is enabled to perform the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be distributed to the member devices for execution at the time indicated by the time stamp

172. A method as defined in claim 169 in which task source device is enabled to perform the step of obtaining information associated with the tasks from at least two types of information sources, and selecting for each task the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective task.

173. A method of operating a device for executing a series of tasks provided by a task source at times specified by the task source in relation to a clock maintained by the task source, the method including

A. a task receiving step of receiving the series of tasks, each task being associated with a time stamp, each time stamp indicating a time value relative to a time indicated a clock maintained by the task source, at which the device is to execute the respective task;

B. a current time retrieval module of obtaining, from the task source, a current time value as indicated by the task source's clock;

C. an execution time determination step of determining, from the time stamp associated with each respective task and a time differential value representing a difference between the current time value obtained during the current time retrieval step and a current time value indicated by a clock maintained by the device, a time, relative to the device's respective clock, at which the task is to be executed; and

D. a task execution module step of executing each respective task at the time determined during the execution time determination step.

174. A method as defined in claim 173 further including a control step for controlling execution of commands received by said interface module.

175. A method as defined in claim 174 further comprising a buffer configured to buffer tasks that have been received until they are to be executed, the task execution step including the step of executing the buffered tasks at the times determined during the execution time determination step.

176. A method as defined in claim 175 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the interface module being further enabled to receive task execution control commands for controlling task execution, the control module being enabled to, in response to receipt of at least one task execution control command for controlling task execution, enable the task execution module to terminate execution of a task sequence currently being executed, being enabled to not execute subsequent tasks in the task sequence that have been buffered.

177. A method as defined in claim 176 in which the control step includes the step of enabling task execution to resume execution with at least one task associated with a subsequent task sequence at a time determined during the execution time determination step.

178. A method as defined in claim 175 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the task receiving step including

the step of receiving task execution control commands for controlling task execution, the control step, in response to receipt of at least one task execution control command to cancel execution of a task sequence for which at least one task has been buffered but for which execution has not begun, including the step of providing that task execution not including execution of tasks from that task sequence that have been buffered.

179. A method as defined in claim 178 in which the control step includes the step of enabling task execution to resume execution of at least one task associated with a task sequence subsequent to the task sequence for which at least one task has been buffered but for which execution has not begun.

180. A method as defined in claim 173, the control step including the step of, in response to receipt of a command to become a member of a synchrony group, receiving tasks during the task receiving step for execution during the task execution step.

181. A method as defined in claim 180 in which the control step includes the step of, in response to receipt of a command to become a member of a synchrony group, enable the task receiving step to include the step of receiving tasks currently being transmitted using a selected multi-cast transmission methodology, and further to include the step of retrieving at least one previously distributed task using a selected unicast transmission methodology.

182. A method as defined in claim 173 in which the control step includes the steps of, in receipt of a command to disengage from a synchrony group, disabling the receipt of tasks and the execution of previously-received tasks that have not yet been executed.

183. A method as defined in claim 173 further comprising a user interface module control information receiving step of receiving control information from a user interface module, the control step including the step of enabling selected operations to be performed in response to the control information.

184 A method as defined in claim 183 in which, in response to control information to enable another device to become a member of the device's synchrony group, the control step includes the step of transmitting a command to the other device to enable the other device to become a member of the synchrony group.

185. A method as defined in claim 183 in which, in response to control information to enable another device that is a member of the device's synchrony group to disengage from the synchrony group, the control module enables the interface module to transmit a command to the other device to enable the other device to disengage from the synchrony group.

186. A method as defined in claim 185 in which, in response to control information to enable the task source to begin operating, the control step includes the step of performing selected operations to enable the task source to begin operating.

187. A method as defined in claim 186, the device also operating as the task source, the method including a task generation steps of, under control of the control step, obtaining information from which the tasks are to be generated and generating the tasks.

188. A method as defined in claim 187 further including the step of transmitting the tasks to at least one other device.

189. A method as defined in claim 187 in which the device is not the task source, the control step including the step of enabling the transmission of a command to enable the task source to operate.

190. A method as defined in claim 189 in which tasks are divided into a series of task sequences, and, in response to receipt of control information enabling execution of a sequence of tasks currently being executed to be terminated , the control step includes the step of transmitting a command to enable the task source to insert a resynchronize command in the series of tasks associated with the task sequence to enable termination of execution of the task sequence.

191. A method as defined in claim 189 in which tasks are divided into a series of task sequences, and, in response to receipt of control information enabling execution of a sequence of tasks at least a portion of which have been received, but which are not being executed, the control step includes the step of transmitting a command to enable the task source to insert a resynchronize command in the series of commands to disable execution of the task sequence.

192. A method as defined in claim 174 in which the control step includes the step of transmitting status information to the user interface module.

193. A method as defined in claim 192 in which the status information includes the execution status of at least one task.

194. A method as defined in claim 173 in which the member device performs:

A. migration information receiving step of receiving migration information from the task source device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group, and

B. a migration control step of, after the migration information has been received.,

   i. distributing the series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to the timing information that it receives

from the device operating as the task source device, at which the devices comprising the synchrony group are to execute the respective task, and

ii. notifying the members of the synchrony group that it is to thereafter operate as the task source device.

195. A method as defined in claim 194 in which the information respecting the tasks that are to be distributed is a source of streaming information, the migration control step includes the steps of, after the migration information has been received, assembling the streaming information into tasks and associating each task with a said time stamp.

196. A method as defined in claim 194 in which the information respecting the tasks that are to be distributed is in a file identified in the migration information, the migration control step including the steps of obtaining information from the identified file, assembling the information into tasks and associating each task with a said time stamp.

197. A method as defined in claim 196 in which the information respecting the tasks that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control step including the step of assembling the information from a position in the file associated with the identified offset.

198. A method as defined in claim 196 in which the information respecting the tasks that are to be distributed is in a series of files identified in the migration information, the migration control step including the steps of obtaining information from successive files in the series assemble the information into tasks and associating each task with as said time stamp.

199. A method as defined in claim 173 in which the member device is further enabled to perform a clock rate adjustment step including the step of adjusting the member device's clock rate in relation to a clock rate value maintained by the task source device's clock.

200. A method as defined in claim 199 in which the clock rate adjustment step includes the step of determining adjust the clock rate in relation to time stamps associated with respective ones of the tasks received from the task source device.

201. A method of operating a device for distributing a series of tasks for execution by at least one other device at respective execution times, the method comprising:

A. a task acquisition step of obtaining respective ones of the tasks;

B. a task execution time determination step of determining a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission step of transmitting the series of tasks to said at least one other device; and

D. an execution control command generation step of, in response to a predetermined event and during the task transmission step, the insertion into the series of tasks transmitted by the task transmission module of a command to enable the at least one other device to modify the execution sequence of the tasks transmitted thereto.

202. A method as defined in claim 201 in which the task transmission step includes the step of utilizing a selected multi-cast transmission methodology.

203. A method as defined in claim 201 in which the predetermined event includes the step of receiving input indicia provided by an operator.

204. A method as defined in claim 203 in which the execution control command generation step includes the step of enabling, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to control execution of at least one task in relation to the time the predetermined event occurs.

205. A method as defined in claim 204 in which, the execution control command generation step includes the step of enabling, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to terminate execution of a task that it is currently executing and/or to not execute at least one task that is to be executed subsequent to the predetermined event.

206. A method as defined in claim 203 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation step including the steps of, in response to said predetermined event,

(i) enabling during the task transmission step the insertion, into the series of tasks, of a command enabling the at least one other device to terminate execution of the currently executing task and/or to not execute at least one subsequent task, if any, in the task sequence including the currently executing task; and

(ii) enabling

(a) during the task execution time determination step, that time stamps associated with tasks of task sequences subsequent to the task sequence currently being executed will reflect the termination and/or non-execution of tasks from the task sequence of the task that is currently being executed, and

(b) during the task transmission step, the transmission of the tasks of the subsequent task sequences to the at least one other device.

207. A method as defined in claim 206 in which the task transmission step includes the step of inserting a command into the task sequence to enable the at least one other device to not execute all tasks in the task sequence subsequent to the terminated and/or non-executed task.

208. A method as defined in claim 201 in which the execution control command generation step include the step of enabling, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to control execution of a task in relation to the task's position in the series of tasks.

209. A method as defined in claim 208 in which the execution control command generation step includes the step of, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to cancel execution of a task in relation to the task's position in the series of tasks.

210. A method as defined in claim 208 in which the execution control command generation step includes the step of enabling, in response to said predetermined event and during the task transmission step, the inserting into the series of tasks of a command to enable the at least one other device to cancel execution of a previously-transmitted task in relation to the task's position in the series of tasks.

211. A method as defined in claim 210 in which, the execution control command generation step includes the step of, in response to said predetermined event and during the task transmission step, the insertion into the series of tasks of a command to enable the at least one other device to cancel execution of a previously-transmitted task, in relation to the task's position in the series of tasks, for which execution has not begun.

212. A method as defined in claim 208 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation step including the steps of, in response to said predetermined event,

(i) enabling, during the task transmission step, the insertion, into the series of tasks, of a command enabling the at least one other device to cancel execution of a task sequence for which execution has not begun; and

(ii) enabling

(a) during the task execution time determination step, that time stamps associated with tasks of task sequences subsequent to the task sequence whose execution was cancelled will reflect the cancellation, and

(b) during the task transmission step, the transmission of the tasks of the subsequent task sequences to the at least one other device.

213. A method of operating a device for distributing a series of tasks for execution by at least one other device at respective execution times, the method comprising:

A. a task acquisition step of obtaining respective ones of the tasks;

B. a task execution time determination step of determining a time at which each respective task is to be executed, and to associating the respective task with a time stamp indicating its execution time, the task execution time determination step including the step of providing, associated with at least one task, a time stamp indicating a task execution time later than the time currently indicated by a clock maintained by the task source device; and

C. a task transmission step of transmitting the series of tasks to said at least one other device.

214. A method as defined in claim 213 in which the task execution time determination step includes the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the device can be executed by at least one other device at the time indicated by the time stamp.

215. A method as defined in claim 214 in which the task execution time determination step includes the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the device can be distributed to other devices over a network for execution by the other devices at the time indicated by the time stamp.

216. A method as defined in claim 213 in which task acquisition step includes the step of obtaining information associated with the tasks from at least two types of information sources, the task execution time determination step including the step of selecting, for each task the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective task.

217. A method of operating a device for distributing a series of tasks for execution by at least one other device, the method comprising:

A. a task acquisition module configured to obtain respective ones of the tasks;

B. a task execution time determination module configured to determine a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission module configured to transmit the series of tasks to said at least one other device; and

D. a task source device migration control module configured to, in response to a predetermined event, provide migration information to another device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group.

218. A computer program product for use in connection with a computer to provide a device for executing a series of tasks provided by a task source at times specified by the task source in relation to a clock maintained by the task source, the computer program product comprising a computer-readable medium having encoded thereon:

A. an interface module configured to enable the computer to receive the series of tasks, each task being associated with a time stamp, each time stamp indicating a time value relative to a time indicated a clock maintained by the task source, at which the device is to execute the respective task;

B. a current time retrieval module configured to enable the computer to obtain, from the task source, a current time value as indicated by the task source's clock;

C. an execution time determination module configured to enable the computer to determine, from the time stamp associated with each respective task and a time differential value representing a difference between the current time value obtained by the current time retrieval module and a current time value indicated by a clock maintained by the device, a time, relative to the device's respective clock, at which the task is to be executed; and

D. a task execution module configured to enable the computer to execute each respective task at the time determined by the execution time determination module.

219. A computer program product as defined in claim 218 further including a control module for enabling said computer to control execution of commands received by the interface module.

220. A computer program product as defined in claim 219 further comprising a buffer module configured to enable the computer to buffer tasks received by the interface module until they are to

be executed, the task execution module being configured to enable the computer to execute the buffered tasks at the times determined by the execution time determination module.

221. A computer program product as defined in claim 220 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the interface module being further configured to enable the computer to receive task execution control commands for controlling task execution, the control module being configured to enable the computer to, in response to receipt of at least one task execution control command for controlling task execution, enable the task execution module to terminate execution of a task sequence currently being executed, being configured to enable the computer to not execute subsequent tasks in the task sequence that have been buffered.

222. A computer program product as defined in claim 220 in which the control module is configured to enable the computer to enable the task execution module to resume execution of at least one task associated with a subsequent task sequence at a time determined by the execution time determination module.

223. A computer program product as defined in claim 219 in which the series of tasks includes a series of task sequences with each task sequence including a subset of the series of tasks, the interface module being further configured to enable the computer to receive task execution control commands for controlling task execution, the control module being configured to enable the computer to, in response to receipt of at least one task execution control command to cancel execution of a task sequence for which at least one task has been buffered but for which execution has not begun, enable the task execution module to not execute tasks from that task sequence that have been buffered.

224. A computer program product as defined in claim 223 in which the control module is configured to enable the computer to enable the task execution module to resume execution of at least one task associated with a task sequence subsequent to the task sequence for which at least one task has been buffered but for which execution has not begun.

225. A computer program product as defined in claim 219, the control module being configured to enable the computer to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive tasks for execution by the task execution module.

226. A computer program product as defined in claim 225 in which the control module is configured to enable the computer to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive tasks currently being transmitted using a selected

multi-cast transmission methodology, and further to retrieve at least one previously distributed task using a selected unicast transmission methodology.

227. A computer program product as defined in claim 219 in which the control module is configured to enable the computer to, in receipt of a command to disengage from a synchrony group, disable the interface module from receiving tasks and the task execution module from executing previously received tasks that have not yet been executed.

228. A computer program product as defined in claim 219 further comprising a user interface module interface module configured to receive control information from a user interface module, the control module being configured to enable the computer to enable selected operations to be performed in response to the control information.

229. A computer program product as defined in claim 228 in which, in response to control information to enable another device to become a member of the device's synchrony group, the control module enables the interface module to transmit a command to the other device to enable the other device to become a member of the synchrony group.

230. A computer program product as defined in claim 228 in which, in response to control information to enable another device that is a member of the device's synchrony group to disengage from the synchrony group, the control module enables the interface module to transmit a command to the other device to enable the other device to disengage from the synchrony group.

231. A computer program product as defined in claim 228 in which, in response to control information to enable the task source to begin operating, the control module performs selected operations to enable the task source to begin operating.

232. A computer program product as defined in claim 231, the device also operating as the task source, the device including a task generator module configured to enable the computer to, under control of the control module, obtain information from which the tasks are to be generated and generate the tasks.

233. A computer program product as defined in claim 232 in which the interface module is further configured to enable the computer to transmit the tasks to at least one other device.

234. A computer program product as defined in claim 232 in which the device is not the task source, the control module being configured to enable the computer to enable the interface module to transmit a command to enable the task source to operate.

235. A computer program product as defined in claim 234 in which tasks are divided into a series of task sequences, and, in response to receipt of control information enabling execution of a sequence

of tasks currently being executed to be terminated , the control module is configured to enable the computer to enable the interface module to transmit a command to enable the task source to insert a resynchronize command in the series of tasks associated with the task sequence to enable termination of execution of the task sequence.

236. A computer program product as defined in claim 234 in which tasks are divided into a series of task sequences, and, in response to receipt of control information enabling execution of a sequence of tasks at least a portion of which have been received, but which are not being executed, the control module is configured to enable the computer to enable the interface module to transmit a command to enable the task source to insert a resynchronize command in the series of commands to disable execution of the task sequence.

237. A computer program product as defined in claim 228 in which the control module is configured to enable the computer to enable the user interface module interface module to transmit status information to the user interface module.

238. A computer program product as defined in claim 237 in which the status information includes the execution status of at least one task.

239. A computer program product as defined in claim 218 further including:

A. a migration information receiving module configured to enable the computer to receive migration information from the task source device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group, and

B. a migration control module configured to enable the computer to, after the migration information has been received.,

   i. distribute the series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the task source device, at which the devices comprising the synchrony group are to execute the respective task, and

   ii. notify the members of the synchrony group that it is to thereafter operate as the task source device.

240. A computer program product as defined in claim 239 in which the information respecting the tasks that are to be distributed is a source of streaming information, the migration control module

being configured to enable the computer to, after the migration information has been received, assemble the streaming information into tasks and associate each task with a said time stamp.

241. A computer program product as defined in claim 239 in which the information respecting the tasks that are to be distributed is in a file identified in the migration information, the migration control module being configured to enable the computer to obtain information from the identified file, assemble the information into tasks and associate each task with a said time stamp.

242. A computer program product as defined in claim 241 in which the information respecting the tasks that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control module being configured to enable the computer to assemble the information from a position in the file associated with the identified offset.

243. A computer program product as defined in claim 241 in which the information respecting the tasks that are to be distributed is in a series of files identified in the migration information, the migration control module being configured to enable the computer to obtain information from successive files in the series assemble the information into tasks and associate each task with as said time stamp.

244. A computer program product as defined in claim 218 further including a clock rate adjustment module configured to enable the computer to adjust the member device's clock rate in relation to a clock rate value maintained by the task source device's clock.

245. A computer program product as defined in claim 244 in which the clock rate adjustment module is configured to enable the computer to adjust the clock rate in relation to time stamps associated with respective ones of the tasks received from the task source device.

246. A computer program product for use in connection with a computer to provide a task source device for distributing a series of tasks for execution by at least one other device at respective execution times, the computer program product comprising a computer readable medium having encoded thereon:

A. a task acquisition module configured to enable the computer to obtain respective ones of the tasks;

B. a task execution time determination module configured to enable the computer to determine a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission module configured to enable the computer to transmit the series of tasks to said at least one other device; and

D. an execution control command generation module configured to enable the computer to , in response to a predetermined event, enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to modify the execution sequence of the tasks transmitted thereto.

247. A computer program product as defined in claim 246 which the task transmission module is configured to enable the computer to utilize a selected multi-cast transmission methodology.

248. A computer program product as defined in claim 246 in which the predetermined event comprises input indicia provided by an operator.

249. A computer program product as defined in claim 246 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to control execution of at least one task in relation to the time the predetermined event occurs.

250. A computer program product as defined in claim 249 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to terminate execution of a task that it is currently executing and/or to not execute at least one task that is to be executed subsequent to the predetermined event.

251. A computer program product as defined in claim 249 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation module being configured to enable the computer to , in response to said predetermined event,

(i) enable the task transmission module to insert into the series of tasks transmitted by the task transmission module, a command enabling the at least one other device to terminate execution of the currently executing task and/or to not execute at least one subsequent task, if any, in the task sequence including the currently executing task; and

(ii) enable

(a) the task execution time determination module to provide that time stamps associated with tasks of task sequences subsequent to the task sequence currently being executed will reflect the termination and/or non-execution of tasks from the task sequence of the task that is currently being executed, and

(b) the task transmission module to transmit the tasks of the subsequent task sequences to the at least one other device.

252. A computer program product as defined in claim 251 in which the task transmission module is configured to enable the computer to insert a command to enable the at least one other device to not execute all tasks in the task sequence subsequent to the terminated and/or non-executed task.

253. A computer program product as defined in claim 246 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to control execution of a task in relation to the task's position in the series of tasks.

254. A computer program product as defined in claim 253 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to cancel execution of a task in relation to the task's position in the series of tasks.

255. A computer program product as defined in claim 254 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted task in relation to the task's position in the series of tasks.

256. A computer program product as defined in claim 255 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the task transmission module to insert into the series of tasks transmitted by the task transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted task, in relation to the task's position in the series of tasks, for which execution has not begun.

257. A computer program product as defined in claim 253 in which the series of tasks includes a series of task sequences with each task sequence including a sub-set of the series of tasks, the execution control command generation module being configured to enable the computer to , in response to said predetermined event,

(i) enable the task transmission module to insert into the series of tasks transmitted by the task transmission module, a command enabling the at least one other device to cancel execution of a task sequence for which execution has not begun; and

(ii) enable

(a) the task execution time determination module to provide that time stamps associated with tasks of task sequences subsequent to the task sequence whose execution was cancelled will reflect the cancellation, and

(b) the task transmission module to transmit the tasks of the subsequent task sequences to the at least one other device.

258. A computer program product for use in connection with a computer to provide a task source device for distributing a series of tasks for execution by at least one other device at respective execution times, the computer program product comprising a computer readable medium having encoded thereon:

A. a task acquisition module configured to enable the computer to obtain respective ones of the tasks;

B. a task execution time determination module configured to enable the computer to determine a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time, the task execution time determination module being configured to enable the computer to provide, associated with at least one task, a time stamp indicating a task execution time later than the time currently indicated by a clock maintained by the task source device; and

C. a task transmission module configured to enable the computer to transmit the series of tasks to said at least one other device.

259. A computer program product as defined in claim 258 in which the task execution time determination module is configured to enable the computer to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one task distributed by the task source device can be executed by the member devices at the time indicated by the time stamp.

260. A computer program product as defined in claim 259 in which the task execution time determination module is configured to enable the computer to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least

one task distributed by the task source device can be distributed to the member devices for execution at the time indicated by the time stamp.

261. A computer program product as defined in claim 258 in which task acquisition module is configured to enable the computer to obtain information associated with the tasks from at least two types of information sources, the task execution time determination module being configured to enable the computer to select, for each task the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective task.

262. A computer program product for use in connection with a computer to provide a task source device for distributing a series of tasks for execution by at least one other device, the computer program product comprising a computer readable medium having encoded thereon:

A. a task acquisition module configured to enable the computer to obtain respective ones of the tasks;

B. a task execution time determination module configured to enable the computer to determine a time at which each respective task is to be executed, and to associate the respective task with a time stamp indicating its execution time;

C. a task transmission module configured to enable the computer to transmit the series of tasks to said at least one other device; and

D. a task source device migration control module configured to enable the computer to, in response to a predetermined event, provide migration information to another device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group.

263. A system comprising a plurality of devices,

A one device being enabled to operate as a task source device, the task source device being enabled to transmit a series of tasks to the other devices, each task being associated with execution time information;

B. at least one device being enabled to operate as a task execution device, the task execution device being enabled to receive the series of tasks from the network and executing each respective task at a time in relation to the timing information associated with the respective task; and

C. one device being enabled to operate as a control device, the control device being enabled to control operations of the task source device and the task execution device.

264. A system as defined in claim 263 in which the device that is enabled to operate as the control device is also enabled to operate as a task execution device.

265. A system as defined in claim 263 in which the device that is enabled to operate as the control device is also enabled to operate as the task source device.

266. A system as defined in claim 263 in which the device that is enabled to operate as the task source device is also enabled to operate as a task execution device.

267. A system for facilitating processing of data by device operating as a member of a synchrony group, synchronously with a clock maintained by a data source device that provides the data to the synchrony group,

A. the data source device being configured to divide the data into a series of frames for distribution to the synchrony group, each frame being associated with a time stamp indicating a time, relative to a clock maintained by the data source device, at which the devices comprising the synchrony group are to process the frame;

B each member device in the synchrony group being configured to:

(i) periodically obtain from the data source device an indication of a current time value indicated by the data source device's clock;

(ii) receive the series of frames from the data source device;

(iii) determine, from the time stamp associated with each respective frame and a time differential value representing a difference between the current time value indicated by the data source device's clock, a current time value indicated by the member device's respective clock, a time, relative to the member device's respective clock, at which the member device is to process the frame; and

(iv) process the respective frame at the determined time.

268. A system as defined in claim 267 in which at least one frame comprises audio data, each member device, while processing each frame, producing sound represented by the respective frame.

269. A system as defined in claim 267 comprising a plurality of member devices, the member devices processing each frame synchronously in relation to the clock maintained by the data source device so that they all produce sound represented by the respective frame synchronously.

270. A system as defined in claim 269 in which the data source device is configured to provide time stamps for at least two successive frames so that the member devices will play them with a selected time interval therebetween.

271. A system as defined in claim 270 in which the time interval is zero.

272. A system as defined in claim 269 in which one of the member devices operates as a master device for the synchrony group, any other member devices comprising respective slave devices, the master device being configured to perform at least one type of synchrony group management operation in connection with the member devices comprising the synchrony group.

273. A system as defined in claim 272 in which the master device is further configured to provide status information relating to the status of the synchrony group to the user interface module.

274. A system as defined in claim 273 in which the status information includes identifications of the devices comprising the synchrony group.

275. A system as defined in claim 273 in which each frame is associated with a frame sequence, each frame sequence comprising at least one frame, the frame or frames comprising each frame sequence being associated with an audio track, each audio track being associated with an identifier, the status information including the identification of the audio track currently being played.

276. A system as defined in claim 272, the system comprising at least one additional device, in which, in one type of synchrony group management operation, the master device is configured to enable the at least one additional device to join the synchrony group as a slave device.

277. A system as defined in claim 276 in which the data source device is configured to distribute frames to the member devices using a selected multi-cast transmission methodology, the member devices being configured to buffer the frames until they are to be played, and further in which, when the at least one additional device joins the synchrony group as a slave device, the data source device is enabled to transmit at least one previously distributed frame to the slave device using a selected unicast transmission methodology.

278. A system as defined in claim 272 in which, in at least one type of synchrony group management operation, the master device is configured to enable the data source device to join the synchrony group as a slave device, the data source device continuing to operate as the data source device for the synchrony group.

279. A system as defined in claim 272 in which, in at least one type of synchrony group management operation, the master device is configured to enable a slave device to disengage from the synchrony

group, the disengaged slave device thereafter not being a member device in the synchrony group and not playing frames distributed by the data source device.

280. A system as defined in claim 272 in which, in at least one type of synchrony group management operation, the master device is configured to control the series of frames to be distributed by the data source device.

281. A system as defined in claim 280 in which each frame is associated with a frame sequence, each frame sequence comprising at least one frame, the frame or frames comprising each frame sequence being associated with an audio track, the master device being configured to control the frames to be distributed by the data source device on a track-by-track basis.

282. A system as defined in claim 281 in which, in at least one type of synchrony group management operation, the master device is further configured to control playback by the member devices of the frames associated with the respective audio tracks that have been distributed by the data source device in response to control information from the user interface module.

283. A system as defined in claim 282 in which, in response to control information from the user interface module to facilitate control of playback by the member devices, the master device is configured to enable the data source device to distribute a playback control command to the member devices.

284. A system as defined in claim 283 in which a playback control command enables the member devices to terminate playback of a track that is currently being played back.

285. A system as defined in claim 284 in which the playback control command enables the member devices to terminate playback of a track with a frame subsequent to a frame currently being played.

286. A system as defined in claim 283 in which the playback control command enables the member devices to cancel playback of a track at least one of whose frames has been previously distributed, but for which playback has not begun.

287. A system as defined in claim 272 in which, in at least one type of synchrony group management operation, the member device operating as the master device is configured to enable the master device to migrate from itself to another member device in the synchrony group.

288. A system as defined in claim 272 in which, in at least one type of synchrony group management operation, the master device is configured to enable the data source device to migrate from one device to another device in the system.

289. A system as defined in claim 288 in which

A. the device operating as the data source device is configured to, after being enabled to migrate the data source device to said other device, provide migration information to the other device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices,

B. the other device being configured to, after receiving the migration information,

i. distribute the series of frames to the synchrony group, each frame being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the data source device, at which the devices comprising the synchrony group are to execute the respective frame, and

ii. notify the members of the synchrony group that it is to thereafter operate as the data source device.

290. A system as defined in claim 289 in which the information respecting the frames that are to be distributed is a source of streaming information, the other device being configured to, after receiving the migration information, assemble the streaming information into frames and associate each frame with a said time stamp.

291. A system as defined in claim 289 in which the information respecting the frames that are to be distributed is in a file identified in the migration information, the other device being configured to obtain information from the identified file, assemble the information into frames and associate each frame with a said time stamp.

292.A system as defined in claim 291 in which the information respecting the frames that are to be distributed begins at an offset, identified in the migration information, into the identified file, the other device being configured to assemble the information from a position in the file associated with the identified offset.

293. A system as defined in claim 291 in which the information respecting the frames that are to be distributed is in a series of files identified in the migration information, the other device being configured to obtain information from successive files in the series assemble the information into frames and associate each frame with as said time stamp.

294. A system as defined in claim 267 in which the data source device is further configured to distribute the frames using a selected multi-cast message transmission methodology.

295. A system as defined in claim 267 in which the data source device is configured to provide, associated with at least one frame, a time stamp indicating a frame execution time later than the time currently indicated by the clock maintained by the data source device.

296. A system as defined in claim 295 in which the data source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be executed by the member devices at the time indicated by the time stamp.

297. A system as defined in claim 296 in which the data source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be distributed to the member devices for execution at the time indicated by the time stamp

298. A system as defined in claim 295 in which data source device is configured to obtain information associated with the frames from at least two types of information sources, the data source device being further configured to select, for each frame the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective frame.

299. A system as defined in claim 267 in which at least one member device is further configured to adjust its clock rate in relation to a clock rate value maintained by the data source device's clock.

300. A system as defined in claim 299 in which the at least one member device is configured to adjust its clock rate in relation to time stamps associated with respective ones of the frames received from the data source device.

301. A system as defined in claim 267 in which at least one other device operates as a data source device configured to distribute frames to a second synchrony group, the device operating as the data source device for the first synchrony group also operating as a member device of a second synchrony group.

302. A system as defined in claim 267 in which at least one member device comprises:

A. an interface module configured to receive the series of frames, each frame being associated with a time stamp, each time stamp indicating a time value relative to a time indicated a clock maintained by the data source, at which the device is to process the respective frame;

B. a current time retrieval module configured to obtain, from the data source, a current time value as indicated by the data source's clock;

C. an processing time determination module configured to determine, from the time stamp associated with each respective frame and a time differential value representing a difference between the current time value obtained by the current time retrieval module and a current time value indicated by a clock maintained by the device, a time, relative to the device's respective clock, at which the frame is to be processed; and

D. a frame processing module configured to process each respective frame at the time determined by the processing time determination module.

303. A system as defined in claim 302 in which the at least one frame comprises audio data, each member device, while processing each frame, producing a signal representing sound represented by the respective frame.

304. A system as defined in claim 302 further including a control module for controlling execution of commands received by said interface module.

305. A device as defined in claim 304 further comprising a buffer configured to buffer frames that the interface module has received until they are to be processed, the frame processing module being configured to process the buffered frames at the times determined by the processing time determination module.

306. A system as defined in claim 305 in which the each frame is associated with a track, each track including at least one frame, the interface module being further configured to receive playback control commands for controlling processing of tracks, the control module being configured to execute at least one playback control command controlling track processing.

307. A system as defined in claim 306 in which at least one type of playback control command enables the control module to, in turn, enable the frame processing module to terminate playback of a track that is currently being played back.

308. A system as defined in claim 307 in which the playback control command enables the control module to terminate playback of a track with a frame subsequent to a frame currently being played.

309. A system as defined in claim 306 in which the playback control command enables the control module to cancel playback of a track at least one of whose frames has been previously distributed, but for which playback has not begun.

310. A system as defined in claim 306 in which the control module is configured to enable the frame processing module to resume processing of at least one frame associated with a subsequent track at a time determined by the processing time determination module.

311. A system as defined in claim 302 in which the member device includes:

A. a migration information receiving module configured to receive migration information from the task source device, the migration information including a source for information respecting the tasks to be distributed, timing information relative to the clock maintained by the task source device and identifications of the member devices of the synchrony group, and

B. a migration control module configure to, after the migration information has been received.,

i. distribute the series of tasks to the synchrony group, each task being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the task source device, at which the devices comprising the synchrony group are to execute the respective task, and

ii. notify the members of the synchrony group that it is to thereafter operate as the task source device.

312. A system as defined in claim 311 in which the information respecting the frames that are to be distributed is a source of streaming information, the migration control module being configured to, after the migration information has been received, assembling the streaming information into frames and associate each frame with a said time stamp.

313. A system as defined in claim 311 in which the information respecting the frames that are to be distributed is in a file identified in the migration information, the migration control module being configured to obtain information from the identified file, assemble the information into frames and associate each frame with a said time stamp.

314. A system as defined in claim 313 in which the information respecting the frames that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control module being configured to assemble the information from a position in the file associated with the identified offset.

315. A system as defined in claim 313 in which the information respecting the frames that are to be distributed is in a series of files identified in the migration information, the migration control module being configured to obtain information from successive files in the series assemble the information into frames and associate each frame with as said time stamp.

316. A system as defined in claim 302 in which the member device further includes a clock rate adjustment module configured to adjust the member device's clock rate in relation to a clock rate value maintained by the data source device's clock.

317. A system as defined in claim 316 in which the clock rate adjustment module is configured to adjust the clock rate in relation to time stamps associated with respective ones of the frames received from the data source device.

318. A system as defined in claim 267 in which the data source device comprises:

A. a frame acquisition module configured to obtain respective ones of the frames;

B. a frame execution time determination module configured to determine a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission module configured to transmit the series of frames to said at least one other device; and

D. an execution control command generation module configured to, in response to a predetermined event, enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to modify the execution sequence of the frames transmitted thereto.

319. A system as defined in claim 318 in which at least one frame comprises audio data, each member device, while processing each frame, producing generating a signal representing sound represented by the respective frame.

320. A system as defined in claim 318 in which the frame transmission module is configured to utilize a selected multi-cast transmission methodology.

321. A system as defined in claim 318 in which the predetermined event comprises input indicia provided by an operator.

322. A system as defined in claim 321 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to control execution of at least one frame in relation to the time the predetermined event occurs.

323. A system as defined in claim 322 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to terminate execution of a frame that it is currently executing and/or to not execute at least one frame that is to be executed subsequent to the predetermined event.

324. A data source device as defined in claim 323 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation module being configured to, in response to said predetermined event,

(i) enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module, a command enabling the at least one other device to terminate execution of the currently executing frame and/or to not execute at least one subsequent frame, if any, in the frame sequence including the currently executing frame; and

(ii) enable

(a) the frame execution time determination module to provide that time stamps associated with frames of frame sequences subsequent to the frame sequence currently being executed will reflect the termination and/or non-execution of frames from the frame sequence of the frame that is currently being executed, and

(b) the frame transmission module to transmit the frames of the subsequent frame sequences to the at least one other device.

325. A system as defined in claim 324 in which the frame transmission module is configured to insert a command to enable the at least one other device to not execute all frames in the frame sequence subsequent to the terminated and/or non-executed frame.

326. A system as defined in claim 267 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to control execution of a frame in relation to the frame's position in the series of frames.

327. A system as defined in claim 326 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to cancel execution of a frame in relation to the frame's position in the series of frames.

328. A system as defined in claim 326 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted frame in relation to the frame's position in the series of frames.

329. A system as defined in claim 328 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted frame, in relation to the frame's position in the series of frames, for which execution has not begun.

330. A system as defined in claim 327 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation module being configured to, in response to said predetermined event,

(i) enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module, a command enabling the at least one other device to cancel execution of a frame sequence for which execution has not begun; and

(ii) enable

(a) the frame execution time determination module to provide that time stamps associated with frames of frame sequences subsequent to the frame sequence whose execution was cancelled will reflect the cancellation, and

(b) the frame transmission module to transmit the frames of the subsequent frame sequences to the at least one other device.

331. A system as defined in claim 267 in which the data source device comprises:

A. a frame acquisition module configured to obtain respective ones of the frames;

B. a frame execution time determination module configured to determine a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission module configured to transmit the series of frames to said at least one other device; and

D. a data source device migration control module configured to, in response to a predetermined event, provide migration information to another device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices of the synchrony group.

332. A system as defined in claim 331 in which the data source device is configured to provide, associated with at least one frame, a time stamp indicating a frame execution time later than the time currently indicated by the clock maintained by the data source device.

333. A system as defined in claim 332 in which the data source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be executed by the member devices at the time indicated by the time stamp.

334. A system as defined in claim 333 in which the data source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be distributed to the member devices for execution at the time indicated by the time stamp

335. A system as defined in claim 333 in which data source device is configured to obtain information associated with the frames from at least two types of information sources, the data source device being further configured to select, for each frame the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective frame.

336. A device for processing a series of frames of data provided by a data source at times specified by the data source in relation to a clock maintained by the data source, the device comprising

A. an interface module configured to receive the series of frames, each frame being associated with a time stamp, each time stamp indicating a time value relative to a time indicated a clock maintained by the data source, at which the device is to process the respective frame;

B. a current time retrieval module configured to obtain, from the data source, a current time value as indicated by the data source's clock;

C. an processing time determination module configured to determine, from the time stamp associated with each respective frame and a time differential value representing a difference between the current time value obtained by the current time retrieval module and a current time value indicated by a clock maintained by the device, a time, relative to the device's respective clock, at which the frame is to be processed; and

D. a frame processing module configured to process each respective frame at the time determined by the processing time determination module.

337. A device as defined in claim 336 in which the frame comprises audio data, each member device, while processing each frame, producing a signal representing sound represented by the respective frame.

338. A device as defined in claim 336 further including a control module for controlling execution of commands received by said interface module.

333. A system as defined in claim 332 in which the data source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be executed by the member devices at the time indicated by the time stamp.

334. A system as defined in claim 333 in which the data source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be distributed to the member devices for execution at the time indicated by the time stamp

335. A system as defined in claim 333 in which data source device is configured to obtain information associated with the frames from at least two types of information sources, the data source device being further configured to select, for each frame the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective frame.

336. A device for processing a series of frames of data provided by a data source at times specified by the data source in relation to a clock maintained by the data source, the device comprising

A. an interface module configured to receive the series of frames, each frame being associated with a time stamp, each time stamp indicating a time value relative to a time indicated a clock maintained by the data source, at which the device is to process the respective frame;

B. a current time retrieval module configured to obtain, from the data source, a current time value as indicated by the data source's clock;

C. an processing time determination module configured to determine, from the time stamp associated with each respective frame and a time differential value representing a difference between the current time value obtained by the current time retrieval module and a current time value indicated by a clock maintained by the device, a time, relative to the device's respective clock, at which the frame is to be processed; and

D. a frame processing module configured to process each respective frame at the time determined by the processing time determination module.

337. A device as defined in claim 336 in which the frame comprises audio data, each member device, while processing each frame, producing a signal representing sound represented by the respective frame.

338. A device as defined in claim 336 further including a control module for controlling execution of commands received by said interface module.

339. A device as defined in claim 336 further comprising a buffer configured to buffer frames that the interface module has received until they are to be processed, the frame processing module being configured to process the buffered frames at the times determined by the processing time determination module.

340. A device as defined in claim 339 in which the each frame is associated with a track, each track including at least one frame, the interface module being further configured to receive playback control commands for controlling processing of tracks, the control module being configured to execute at least one playback control command controlling track processing.

341. A device as defined in claim 340 in which at least one type of playback control command enables the control module to, in turn, enable the frame processing module to terminate playback of a track that is currently being played back.

342. A device as defined in claim 341 in which the playback control command enables the control module to terminate playback of a track with a frame subsequent to a frame currently being played.

343. A device as defined in claim 340 in which the playback control command enables the control module to cancel playback of a track at least one of whose frames has been previously distributed, but for which playback has not begun.

344. A device as defined in claim 340 in which the control module is configured to enable the frame processing module to resume processing of at least one frame associated with a subsequent track at a time determined by the processing time determination module.

345. A device as defined in claim 338, the control module being configured to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive frames for processing by the frame processing module.

346. A device as defined in claim 345 in which the control module is configured to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive frames currently being transmitted using a selected multi-cast transmission methodology, and further to retrieve at least one previously distributed frame using a selected unicast transmission methodology.

347. A device as defined in claim 346 in which the control module is configured to, in receipt of a command to disengage from a synchrony group, disable the interface module from receiving frames and the frame processing module from processing previously received frames that have not yet been processed.

348. A device as defined in claim 336 further comprising a user interface module interface module configured to receive control information from a user interface module, the control module being configured to enable selected operations to be performed in response to the control information.

349. A device as defined in claim 348 in which, in response to control information to enable another device to become a member of the device's synchrony group, the control module enables the interface module to transmit a command to the other device to enable the other device to become a member of the synchrony group.

350. A device as defined in claim 349 in which, in response to control information to enable another device that is a member of the device's synchrony group to disengage from the synchrony group, the control module enables the interface module to transmit a command to the other device to enable the other device to disengage from the synchrony group.

351. A device as defined in claim 348 in which, in response to control information to enable the data source to begin operating, the control module performs selected operations to enable the data source to begin operating.

352. A device as defined in claim 351, the device also operating as the data source, the device including a frame generator module configured to, under control of the control module, obtain information from which the frames are to be generated and generate the frames.

353. A device as defined in claim 342 in which the interface module is further configured to transmit the frames to at least one other device.

354. A device as defined in claim 343 in which the device is not the data source, the control module being configured to enable the interface module to transmit a command to enable the data source to operate.

355. A device as defined in claim 354 in which each frame is associated with a track, each track including at least one frame, and, in response to receipt of control information enabling processing of a track currently being processed to be terminated , the control module is configured to enable the interface module to transmit a command to enable the data source to insert a resynchronize command in the series of frames associated with the track to enable termination of processing of the track.

356. A device as defined in claim 354 in which each frame is associated with a track, each track including at least one frame, and, in response to receipt of control information enabling processing of a track at least a portion of which has been received, but for which processing has not begun, the control module is configured to enable the interface module to transmit a command to enable the data

source to insert a resynchronize command in the series of commands to disable processing of the track.

357. A device as defined in claim 344 in which the control module is configured to enable the user interface module interface module to transmit status information to the user interface module.

358. A device as defined in claim 357 in which the status information includes the processing status of at least one frame.

359. A device as defined in claim 336 further including:

A. a migration information receiving module configured to receive migration information from the data source device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices of the synchrony group, and

B. a migration control module configured to, after the migration information has been received.,

   i. distribute the series of frames to the synchrony group, each frame being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the data source device, at which the devices comprising the synchrony group are to execute the respective frame, and

   ii. notify the members of the synchrony group that it is to thereafter operate as the data source device.

360. A device as defined in claim 359 in which the information respecting the frames that are to be distributed is a source of streaming information, the migration control module being configured to, after the migration information has been received, assembling the streaming information into frames and associate each frame with a said time stamp.

361. A device as defined in claim 359 in which the information respecting the frames that are to be distributed is in a file identified in the migration information, the migration control module being configured to obtain information from the identified file, assemble the information into frames and associate each frame with a said time stamp.

362. A device as defined in claim 361 in which the information respecting the frames that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control module being configured to assemble the information from a position in the file associated with the identified offset.

363. A device as defined in claim 361 in which the information respecting the frames that are to be distributed is in a series of files identified in the migration information, the migration control module being configured to obtain information from successive files in the series assemble the information into frames and associate each frame with as said time stamp.

364. A device as defined in claim 336 further including a clock rate adjustment module configured to adjust the member device's clock rate in relation to a clock rate value maintained by the data source device's clock.

365. A data source device for distributing a series of frames for execution by at least one other device at respective execution times, the data source device comprising:

A. a frame acquisition module configured to obtain respective ones of the frames;

B. a frame execution time determination module configured to determine a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission module configured to transmit the series of frames to said at least one other device; and

D. an execution control command generation module configured to, in response to a predetermined event, enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to modify the execution sequence of the frames transmitted thereto.

366. A data source device as defined in claim 365 in which the frame transmission module is configured to utilize a selected multi-cast transmission methodology.

367. A data source device as defined in claim 365 in which the predetermined event comprises input indicia provided by an operator.

368. A data source device as defined in claim 367 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to control execution of at least one frame in relation to the time the predetermined event occurs.

369. A data source device as defined in claim 368 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command

to enable the at least one other device to terminate execution of a frame that it is currently executing and/or to not execute at least one frame that is to be executed subsequent to the predetermined event.

370. A data source device as defined in claim 368 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation module being configured to, in response to said predetermined event,

(i) enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module, a command enabling the at least one other device to terminate execution of the currently executing frame and/or to not execute at least one subsequent frame, if any, in the frame sequence including the currently executing frame; and

(ii) enable

(a) the frame execution time determination module to provide that time stamps associated with frames of frame sequences subsequent to the frame sequence currently being executed will reflect the termination and/or non-execution of frames from the frame sequence of the frame that is currently being executed, and

(b) the frame transmission module to transmit the frames of the subsequent frame sequences to the at least one other device.

371. A data source device as defined in claim 370 in which the frame transmission module is configured to insert a command to enable the at least one other device to not execute all frames in the frame sequence subsequent to the terminated and/or non-executed frame.

372. A data source device as defined in claim 362 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to control execution of a frame in relation to the frame's position in the series of frames.

373. A data source device as defined in claim 372 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to cancel execution of a frame in relation to the frame's position in the series of frames.

374. A data source device as defined in claim 371 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command

to enable the at least one other device to cancel execution of a previously-transmitted frame in relation to the frame's position in the series of frames.

375. A data source device as defined in claim 374 in which, in response to said predetermined event, the execution control command generation module is configured to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted frame, in relation to the frame's position in the series of frames, for which execution has not begun.

376. A data source device as defined in claim 371 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation module being configured to, in response to said predetermined event,

(i) enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module, a command enabling the at least one other device to cancel execution of a frame sequence for which execution has not begun; and

(ii) enable

(a) the frame execution time determination module to provide that time stamps associated with frames of frame sequences subsequent to the frame sequence whose execution was cancelled will reflect the cancellation, and

(b) the frame transmission module to transmit the frames of the subsequent frame sequences to the at least one other device.

377. A data source device for distributing a series of frames for execution by at least one other device at respective execution times, the data source device comprising:

A. a frame acquisition module configured to obtain respective ones of the frames;

B. a frame execution time determination module configured to determine a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time, the frame execution time determination module being configured to provide, associated with at least one frame, a time stamp indicating a frame execution time later than the time currently indicated by a clock maintained by the data source device; and

C. a frame transmission module configured to transmit the series of frames to said at least one other device.

377. A system as defined in claim 377 in which the data source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient

to provide that the at least one frame distributed by the data source device can be executed by the member devices at the time indicated by the time stamp.

379. A system as defined in claim 378 in which the data source device is configured to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be distributed to the member devices for execution at the time indicated by the time stamp

380. A system as defined in claim 377 in which data source device is configured to obtain information associated with the frames from at least two types of information sources, the data source device being further configured to select, for each frame the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective frame.

381. A data source device for distributing a series of frames for execution by at least one other device, the data source device comprising:

A. a frame acquisition module configured to obtain respective ones of the frames;

B. a frame execution time determination module configured to determine a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission module configured to transmit the series of frames to said at least one other device; and

D. a data source device migration control module configured to, in response to a predetermined event, provide migration information to another device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices of the synchrony group.

382. A method for operating a system to facilitate processing of data by device operating as a member of a synchrony group, synchronously with a clock maintained by a data source device that provides the data to the synchrony group, the method including the steps of

A. enabling the data source device being to divide the data into a series of frames for distribution to the synchrony group, each frame being associated with a time stamp indicating a time, relative to a clock maintained by the data source device, at which the devices comprising the synchrony group are to process the frame;

B enabling each member device in the synchrony group to:

(i) periodically obtain from the data source device an indication of a current time value indicated by the data source device's clock;

(ii) receive the series of frames from the data source device;

(iii) determine, from the time stamp associated with each respective frame and a time differential value representing a difference between the current time value indicated by the data source device's clock, a current time value indicated by the member device's respective clock, a time, relative to the member device's respective clock, at which the member device is to process the frame; and

(iv) process the respective frame at the determined time.

383. A method as defined in claim 382 in which at least one frame comprises audio data, each member device, while processing each frame, being enabled to generate a signal representing sound represented by the respective frame.

384. A method as defined in claim 382, the system comprising a plurality of member devices, the method including the steps of enabling the member devices to process each frame synchronously in relation to the clock maintained by the data source device so that they all produce sound represented by the respective frame synchronously.

385. A method as defined in claim 384 including the step of enabling the data source device to provide time stamps for at least two successive frames so that the member devices will play them with a selected time interval therebetween.

386. A method as defined in claim 385 in which the time interval is zero.

387. A method as defined in claim 382 in which one of the member devices operates as a master device for the synchrony group, any other member devices comprising respective slave devices, the method including the step of enabling the master device to perform at least one type of synchrony group management operation in connection with the member devices comprising the synchrony group.

388. A method as defined in claim 387, the method including the step of enabling the master device to perform at least one type of synchrony group management operation in response to control information provided thereto by a user interface module.

389. A method as defined in claim 388 further including the step of enabling the master device to provide status information relating to the status of the synchrony group to the user interface module.

390. A method as defined in claim 389 in which the status information includes identifications of the devices comprising the synchrony group.

391. A method as defined in claim 389 in which each frame is associated with a frame sequence, each frame sequence comprising at least one frame, the frame or frames comprising each frame sequence being associated with an audio track, each audio track being associated with an identifier, the status information including the identification of the audio track currently being played.

392. A method as defined in claim 388 the method comprising at least one additional device, in which, in one type of synchrony group management operation, the master device is configured to enable the at least one additional device to join the synchrony group as a slave device.

393. A method as defined in claim 392 in which the data source device is configured to distribute frames to the member devices using a selected multi-cast transmission methodology, the member devices being configured to buffer the frames until they are to be played, the method including the step of, when the at least one additional device joins the synchrony group as a slave device, the enabling the data source device to transmit at least one previously distributed frame to the slave device using a selected unicast transmission methodology.

394. A method as defined in claim 387 in which, in at least one type of synchrony group management operation, the method includes the step of enabling the master device to enable the data source device to join the synchrony group as a slave device, the data source device continuing to operate as the data source device for the synchrony group.

395. A method as defined in claim 387 in which, in at least one type of synchrony group management operation, the method includes the step of enabling the master device to enable a slave device to disengage from the synchrony group, the disengaged slave device thereafter not being a member device in the synchrony group and not playing frames distributed by the data source device.

396. A method as defined in claim 387 in which, in at least one type of synchrony group management operation, the method includes the step of enabling the master device to control the series of frames to be distributed by the data source device.

397. A method as defined in claim 387 in which each frame is associated with a frame sequence, each frame sequence comprising at least one frame, the frame or frames comprising each frame sequence being associated with an audio track, the method including the step of enabling the master device being configured to control the frames to be distributed by the data source device on a track-by-track basis.

398. A method as defined in claim 387 in which, in at least one type of synchrony group management operation, the method includes the step of enabling the master device to control playback by the member devices of the frames associated with the respective audio tracks that have

been distributed by the data source device in response to control information from the user interface module.

399. A method as defined in claim 387 in which, in response to control information from the user interface module to facilitate control of playback by the member devices, the method includes the step of enabling the master device to enable the data source device to distribute a playback control command to the member devices.

400. A method as defined in claim 399 in which a playback control command enables the member devices to terminate playback of a track that is currently being played back.

401. A method as defined in claim 399 in which the playback control command enables the member devices to terminate playback of a track with a frame subsequent to a frame currently being played.

402. A method as defined in claim 399 in which the playback control command enables the member devices to cancel playback of a track at least one of whose frames has been previously distributed, but for which playback has not begun.

403. A method as defined in claim 387 in which, in at least one type of synchrony group management operation, the method includes the step of enabling the member device operating as the master device to enable the master device to migrate from itself to another member device in the synchrony group.

404. A method as defined in claim 403 in which, in at least one type of synchrony group management operation, the method includes the step of enabling the master device to enable the data source device to migrate from one device to another device in the system.

405. A method as defined in claim 404 in which

A. the device operating as the data source device is enable to perform the step of, after being enabled to migrate the data source device to said other device, providing migration information to the other device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices,

B. the other device is enabled to, after receiving the migration information, perform the steps of

   i. distributing the series of frames to the synchrony group, each frame being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the data source device, at which the devices comprising the synchrony group are to execute the respective frame, and

ii. notifying the members of the synchrony group that it is to thereafter operate as the data source device.

406. A method as defined in claim 405 in which the information respecting the frames that are to be distributed is a source of streaming information, the other device being enabled to, after receiving the migration information, perform the steps of assembling the streaming information into frames and associating each frame with a said time stamp.

407. A method as defined in claim 405 in which the information respecting the frames that are to be distributed is in a file identified in the migration information, the other device being enabled to perform the steps of obtaining information from the identified file, assembling the information into frames and associating each frame with a said time stamp.

408.A method as defined in claim 407 in which the information respecting the frames that are to be distributed begins at an offset, identified in the migration information, into the identified file, the other device being enabled to perform the step of assembling the information from a position in the file associated with the identified offset.

409. A method as defined in claim 407 in which the information respecting the frames that are to be distributed is in a series of files identified in the migration information, the other device being enabled to perform the step of obtaining information from successive files in the series assemble the information into frames and associate each frame with as said time stamp.

410. A method as defined in claim 382 including the step in which the data source is further configured to distribute the frames using a selected multi-cast message transmission methodology, the device being enabled to receive the frames using the selected multi-cast message transmission methodology.

411. A method as defined in claim 382 in which the data source device is enabled to perform the step of providing, provide, associated with at least one frame, a time stamp indicating a frame execution time later than the time currently indicated by the clock maintained by the data source device.

412. A method as defined in claim 411 in which the data source device is enabled to perform the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be executed by the member devices at the time indicated by the time stamp.

413. A method as defined in claim 411 in which the data source device is enabled to perform the step of selecting the delay interval between the time indicated by the time stamp and the current time so

to be sufficient to provide that the at least one frame distributed by the data source device can be distributed to the member devices for execution at the time indicated by the time stamp

414. A method as defined in claim 411 in which data source device is enabled to perform the steps of obtaining information associated with the frames from at least two types of information sources, and selecting, for each frame, the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective frame.

415. A method as defined in claim 382 in which at least one member device is further enabled to perform the step of adjusting its clock rate in relation to a clock rate value maintained by the data source device's clock.

416. A method as defined in claim 415 in which the at least one member device is enabled to perform the step of adjusting its clock rate in relation to time stamps associated with respective ones of the frames received from the data source device.

417. A method as defined in claim 382 in which at least one other device operates as a data source device configured to distribute frames to a second synchrony group, the method including the steps of enabling the device to operate both as the data source device for the first synchrony group and also operating as a member device of a second synchrony group.

418. A method as defined in claim 382 including, within a member device,

A. a frame receiving step of receiving the series of frames,

B. a current time retrieval step of obtaining, from the data source device, the current time value as indicated by the frame source's clock;

C. an execution time determination step of determining, from the time stamp associated with each respective frame and a time differential value representing a difference between the current time value obtained during the current time retrieval step and a current time value indicated by a clock maintained by the device, a time, relative to the member device's respective clock at which the frame is to be executed; and

D. a frame execution step of executing each respective frame at the time determined during the execution time determination step.

419. A method as defined in claim 418 further including a control step for controlling execution of commands received during the frame receiving step

420. A method as defined in claim 419 in which the at least one member device further comprises a buffer configured to buffer frames received from the data source device during the frame receiving

step until they are to be executed, the frame execution step including the step of executing the buffered frames at the times determined during the execution time determination step.

420. A method as defined in claim 420 in which the series of frames includes a series of frame sequences with each frame sequence including a subset of the series of frames, the frame receiving step being further enabled to receive frame execution control commands for controlling frame execution, the control step including the step of, in response to receipt of at least one frame execution control command for controlling frame execution, controlling the frame execution step to terminate execution of a frame sequence currently being executed, and further to not execute subsequent frames in the frame sequence that have been buffered.

422. A method as defined in claim 421 in which the control step includes the step of enabling the frame execution step to resume execution of at least one frame associated with a subsequent frame sequence at a time determined during the execution time determination step.

423. A method as defined in claim 420 in which the series of frames includes a series of frame sequences with each frame sequence including a subset of the series of frames, the frame receiving step including the step of receiving frame execution control commands for controlling frame execution, the control step including the step of, in response to receipt of at least one frame execution control command to cancel execution of a frame sequence for which at least one frame has been buffered but for which execution has not begun, providing that frames from that frame sequence that have been buffered not be executed.

424. A method as defined in claim 423 in which the control step includes the step of enabling execution to resume of at least one frame associated with a frame sequence subsequent to the frame sequence for which at least one frame has been buffered but for which execution has not begun.

425. A method as defined in claim 418 in which the member device performs:

A. migration information receiving step of receiving migration information from the data source device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices of the synchrony group, and

B. a migration control step of, after the migration information has been received.,

   i. distributing the series of frames to the synchrony group, each frame being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the data source device, at which the devices comprising the synchrony group are to execute the respective frame, and

ii. notifying the members of the synchrony group that it is to thereafter operate as the data source device.

426. A method as defined in claim 425 in which the information respecting the frames that are to be distributed is a source of streaming information, the migration control step includes the steps of, after the migration information has been received, assembling the streaming information into frames and associating each frame with a said time stamp.

427. A method as defined in claim 425 in which the information respecting the frames that are to be distributed is in a file identified in the migration information, the migration control step including the steps of obtaining information from the identified file, assembling the information into frames and associating each frame with a said time stamp.

428. A method as defined in claim 427 in which the information respecting the frames that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control step including the step of assembling the information from a position in the file associated with the identified offset.

429. A method as defined in claim 427 in which the information respecting the frames that are to be distributed is in a series of files identified in the migration information, the migration control step including the steps of obtaining information from successive files in the series assemble the information into frames and associating each frame with as said time stamp.

430. A method as defined in claim 418 in which the member device is further enabled to perform a clock rate adjustment step including the step of adjusting the member device's clock rate in relation to a clock rate value maintained by the data source device's clock.

431. A method as defined in claim 430 in which the clock rate adjustment step includes the step of determining adjust the clock rate in relation to time stamps associated with respective ones of the frames received from the data source device.

432. A method as defined in claim 382 of operating a device for distributing a series of frames for execution by at least one other device at respective execution times, the method comprising:

A. a frame acquisition step of obtaining respective ones of the frames;

B. a frame execution time determination step of determining a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission step of transmitting the series of frames to said at least one other device; and

D. an execution control command generation step of, in response to a predetermined event and during the frame transmission step, the insertion into the series of frames transmitted by the frame transmission module of a command to enable the at least one other device to modify the execution sequence of the frames transmitted thereto.

433. A method as defined in claim 432 in which the frame transmission step includes the step of utilizing a selected multi-cast transmission methodology.

434. A method as defined in claim 432 in which the predetermined event includes the step of receiving input indicia provided by an operator.

435. A method as defined in claim 434 in which the execution control command generation step includes the step of enabling, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to control execution of at least one frame in relation to the time the predetermined event occurs.

436. A method as defined in claim 435 in which, the execution control command generation step includes the step of enabling, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to terminate execution of a frame that it is currently executing and/or to not execute at least one frame that is to be executed subsequent to the predetermined event.

437. A method as defined in claim 434 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation step including the steps of, in response to said predetermined event,

(i) enabling during the frame transmission step the insertion, into the series of frames, of a command enabling the at least one other device to terminate execution of the currently executing frame and/or to not execute at least one subsequent frame, if any, in the frame sequence including the currently executing frame; and

(ii) enabling

(a) during the frame execution time determination step, that time stamps associated with frames of frame sequences subsequent to the frame sequence currently being executed will reflect the termination and/or non-execution of frames from the frame sequence of the frame that is currently being executed, and

(b) during the frame transmission step, the transmission of the frames of the subsequent frame sequences to the at least one other device.

438. A method as defined in claim 437 in which the frame transmission step includes the step of inserting a command into the frame sequence to enable the at least one other device to not execute all frames in the frame sequence subsequent to the terminated and/or non-executed frame.

439. A method as defined in claim 435 in which the execution control command generation step include the step of enabling, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to control execution of a frame in relation to the frame's position in the series of frames.

440. A method as defined in claim 439 in which the execution control command generation step includes the step of, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to cancel execution of a frame in relation to the frame's position in the series of frames.

441. A method as defined in claim 439 in which the execution control command generation step includes the step of enabling, in response to said predetermined event and during the frame transmission step, the inserting into the series of frames of a command to enable the at least one other device to cancel execution of a previously-transmitted frame in relation to the frame's position in the series of frames.

442. A method as defined in claim 441 in which, the execution control command generation step includes the step of, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to cancel execution of a previously-transmitted frame, in relation to the frame's position in the series of frames, for which execution has not begun.

443. A method as defined in claim 439 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation step including the steps of, in response to said predetermined event,

(i) enabling, during the frame transmission step, the insertion, into the series of frames, of a command enabling the at least one other device to cancel execution of a frame sequence for which execution has not begun; and

(ii) enabling

(a) during the frame execution time determination step, that time stamps associated with frames of frame sequences subsequent to the frame sequence whose execution was cancelled will reflect the cancellation, and

(b) during the frame transmission step, the transmission of the frames of the subsequent frame sequences to the at least one other device.

444. A method as defined in claim 382 in which the data source device is enabled to perform

A. a frame acquisition step of obtaining respective ones of the frames;

B. a frame execution time determination step of determining a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission step of transmitting the series of frames to said at least one other device; and

D. a data source device migration control step of, in response to a predetermined event, providing migration information to another device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices of the synchrony group.

445. A method as defined in claim 382 in which the data source device is enabled to perform the step of providing, associated with at least one frame, a time stamp indicating a frame execution time later than the time currently indicated by the clock maintained by the data source device.

446. A method as defined in claim 445 in which the data source device is enabled to perform the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be executed by the member devices at the time indicated by the time stamp.

447. A method as defined in claim 446 in which the data source device is enabled to perform the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be distributed to the member devices for execution at the time indicated by the time stamp.

448. A method as defined in claim 445 in which data source device is enabled to perform the step of obtaining information associated with the frames from at least two types of information sources, and selecting for each frame the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective frame.

449. A method of enabling a device to process a series of frames of data provided by a data source at times specified by the data source in relation to a clock maintained by the data source, the device comprising

A. an interface step of receiving the series of frames, each frame being associated with a time stamp, each time stamp indicating a time value relative to a time indicated a clock maintained by the data source, at which the device is to process the respective frame;

B. a current time retrieval step of obtaining, from the data source, a current time value as indicated by the data source's clock;

C. an processing time determination step of determining, from the time stamp associated with each respective frame and a time differential value representing a difference between the current time value obtained during the current time retrieval step and a current time value indicated by a clock maintained by the device, a time, relative to the device's respective clock, at which the frame is to be processed; and

D. a frame processing step of processing each respective frame at the time determined during the processing time determination step.

450. A method as defined in claim 449 in which at least one frame comprises audio data, each member device, while processing said at least one frame, generating a signal representing sound represented by the respective frame.

451. A method as defined in claim 449 further including a control step for controlling execution of commands received by said device.

452. A method as defined in claim 451 further comprising a buffer configured to buffer frames that the interface module has received until they are to be processed, the frame processing step including the step of processing the buffered frames at the times determined during the processing time determination step.

453 . A method as defined in claim 451 in which the each frame is associated with a track, each track including at least one frame, the interface module being further configured to receive playback control commands for controlling processing of tracks, the method including the step of enabling the control module to execute at least one playback control command controlling track processing.

454. A method as defined in claim 453 including the step of, in response to at least one type of playback control command, enabling the control module to, in turn, enable the frame processing module to terminate playback of a track that is currently being played back.

455. A method as defined in claim 454 including the step of, in response to the playback control command, enabling the control module to terminate playback of a track with a frame subsequent to a frame currently being played.

456. A method as defined in claim 452 including the step of, in response to the playback control command, enabling the control module to cancel playback of a track at least one of whose frames has been previously distributed, but for which playback has not begun.

457. A method as defined in claim 453 including the step of enabling the control module to enable the frame processing module to resume processing of at least one frame associated with a subsequent track at a time determined by the processing time determination module.

458. A method as defined in claim 449 including the step of enabling the control module to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive frames for processing by the frame processing module.

459. A method as defined in claim 458 including the step of enabling the control module to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive frames currently being transmitted using a selected multi-cast transmission methodology, and further to retrieve at least one previously distributed frame using a selected unicast transmission methodology.

460. A method as defined in claim 449 including the step of enabling the control module to, in receipt of a command to disengage from a synchrony group, disable the interface module from receiving frames and the frame processing module from processing previously received frames that have not yet been processed.

461. A method as defined in claim 449, further including the steps of receiving control information from a user interface module, and enabling the control module to enable selected operations to be performed in response to the control information.

462. A method as defined in claim 461 including the step of, in response to control information to enable another device to become a member of the device's synchrony group, enabling the control module to enable the interface module to transmit a command to the other device to enable the other device to become a member of the synchrony group.

463. A method as defined in claim 461 including the step of, in response to control information to enable another device that is a member of the device's synchrony group to disengage from the synchrony group, enabling the control module to enable transmission of a command to the other device to enable the other device to disengage from the synchrony group.

464. A method as defined in claim 461 including the step of, in response to control information to enable the data source to begin operating, enabling the control module to perform selected operations to enable the data source to begin operating.

465. A method as defined in claim 464, the device also operating as the data source, the method including a frame generator step of, under control of the control module, obtaining information from which the frames are to be generated and generating the frames.

466. A method as defined in claim 465 further comprising the step of transmitting the frames to at least one other device.

467. A method as defined in claim 465 in which the device is not the data source, the method including the step of enabling the control module to enable transmission of a command to enable the data source to operate.

468. A method as defined in claim 467 in which each frame is associated with a track, each track including at least one frame, and, in response to receipt of control information enabling processing of a track currently being processed to be terminated , the method including the step of enabling the control module to enable transmission of a command to enable the data source to insert a resynchronize command in the series of frames associated with the track to enable termination of processing of the track.

469. A method as defined in claim 467 in which each frame is associated with a track, each track including at least one frame, and, in response to receipt of control information enabling processing of a track at least a portion of which has been received, but for which processing has not begun, the method including the step of enabling the control module to enable transmission of a command to enable the data source to insert a resynchronize command in the series of commands to disable processing of the track.

470. A method as defined in claim 451 further including the step of enabling the control module to enable the user interface module interface module to transmit status information to the user interface module.

471. A method as defined in claim 470 in which the status information includes the processing status of at least one frame.

472. A method as defined in claim 469 in which the member device performs:

A. migration information receiving step of receiving migration information from the data source device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices of the synchrony group, and

B. a migration control step of, after the migration information has been received.,

i. distributing the series of frames to the synchrony group, each frame being associated with a time stamp indicating a time, relative to the timing information that it receives from the device operating as the data source device, at which the devices comprising the synchrony group are to execute the respective frame, and

ii. notifying the members of the synchrony group that it is to thereafter operate as the data source device.

473. A method as defined in claim 472 in which the information respecting the frames that are to be distributed is a source of streaming information, the migration control step includes the steps of, after the migration information has been received, assembling the streaming information into frames and associating each frame with a said time stamp.

474. A method as defined in claim 472 in which the information respecting the frames that are to be distributed is in a file identified in the migration information, the migration control step including the steps of obtaining information from the identified file, assembling the information into frames and associating each frame with a said time stamp.

475. A method as defined in claim 474 in which the information respecting the frames that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control step including the step of assembling the information from a position in the file associated with the identified offset.

476. A method as defined in claim 474 in which the information respecting the frames that are to be distributed is in a series of files identified in the migration information, the migration control step including the steps of obtaining information from successive files in the series assemble the information into frames and associating each frame with as said time stamp.

477. A method as defined in claim 449 in which the member device is further enabled to perform a clock rate adjustment step including the step of adjusting the member device's clock rate in relation to a clock rate value maintained by the data source device's clock.

478. A method as defined in claim 477 in which the clock rate adjustment step includes the step of determining adjust the clock rate in relation to time stamps associated with respective ones of the frames received from the data source device.

479. A method of operating a device for distributing a series of frames for execution by at least one other device at respective execution times, the method comprising:

A. a frame acquisition step of obtaining respective ones of the frames;

B. a frame execution time determination step of determining a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission step of transmitting the series of frames to said at least one other device; and

D. an execution control command generation step of, in response to a predetermined event and during the frame transmission step, the insertion into the series of frames transmitted by the frame transmission module of a command to enable the at least one other device to modify the execution sequence of the frames transmitted thereto.

480. A method as defined in claim 479 in which at least one frame comprises audio data.

481. A method as defined in claim 479 in which the frame transmission step includes the step of utilizing a selected multi-cast transmission methodology.

482. A method as defined in claim 479 in which the predetermined event includes the step of receiving input indicia provided by an operator.

483. A method as defined in claim 482 in which the execution control command generation step includes the step of enabling, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to control execution of at least one frame in relation to the time the predetermined event occurs.

484. A method as defined in claim 483 in which, the execution control command generation step includes the step of enabling, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to terminate execution of a frame that it is currently executing and/or to not execute at least one frame that is to be executed subsequent to the predetermined event.

485. A method as defined in claim 482 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation step including the steps of, in response to said predetermined event,

(i) enabling during the frame transmission step the insertion, into the series of frames, of a command enabling the at least one other device to terminate execution of the currently executing frame and/or to not execute at least one subsequent frame, if any, in the frame sequence including the currently executing frame; and

(ii) enabling

(a) during the frame execution time determination step, that time stamps associated with frames of frame sequences subsequent to the frame sequence currently being executed will reflect the termination and/or non-execution of frames from the frame sequence of the frame that is currently being executed, and

(b) during the frame transmission step, the transmission of the frames of the subsequent frame sequences to the at least one other device.

486. A method as defined in claim 485 in which the frame transmission step includes the step of inserting a command into the frame sequence to enable the at least one other device to not execute all frames in the frame sequence subsequent to the terminated and/or non-executed frame.

487. A method as defined in claim 479 in which the execution control command generation step include the step of enabling, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to control execution of a frame in relation to the frame's position in the series of frames.

488. A method as defined in claim 487 in which the execution control command generation step includes the step of, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to cancel execution of a frame in relation to the frame's position in the series of frames.

489. A method as defined in claim 487 in which the execution control command generation step includes the step of enabling, in response to said predetermined event and during the frame transmission step, the inserting into the series of frames of a command to enable the at least one other device to cancel execution of a previously-transmitted frame in relation to the frame's position in the series of frames.

490. A method as defined in claim 489 in which, the execution control command generation step includes the step of, in response to said predetermined event and during the frame transmission step, the insertion into the series of frames of a command to enable the at least one other device to cancel execution of a previously-transmitted frame, in relation to the frame's position in the series of frames, for which execution has not begun.

491. A method as defined in claim 488 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation step including the steps of, in response to said predetermined event,

(i) enabling, during the frame transmission step, the insertion, into the series of frames, of a command enabling the at least one other device to cancel execution of a frame sequence for which execution has not begun; and

(ii) enabling

(a) during the frame execution time determination step, that time stamps associated with frames of frame sequences subsequent to the frame sequence whose execution was cancelled will reflect the cancellation, and

(b) during the frame transmission step, the transmission of the frames of the subsequent frame sequences to the at least one other device.

492. A method of operating a device for distributing a series of frames for execution by at least one other device at respective execution times, the method comprising:

A. a frame acquisition step of obtaining respective ones of the frames;

B. a frame execution time determination step of determining a time at which each respective frame is to be executed, and to associating the respective frame with a time stamp indicating its execution time, the frame execution time determination step including the step of providing, associated with at least one frame, a time stamp indicating a frame execution time later than the time currently indicated by a clock maintained by the data source device; and

C. a frame transmission step of transmitting the series of frames to said at least one other device.

493. A method as defined in claim 492 in which at least one frame comprises audio data.

494. A method as defined in claim 492 in which the frame execution time determination step includes the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the device can be executed by at least one other device at the time indicated by the time stamp.

495. A method as defined in claim 494 in which the frame execution time determination step includes the step of selecting the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the device can be distributed to other devices over a network for execution by the other devices at the time indicated by the time stamp.

496. A method as defined in claim 492 in which frame acquisition step includes the step of obtaining information associated with the frames from at least two types of information sources, the frame execution time determination step including the step of selecting, for each frame the delay interval

between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective frame.

497. A method of operating a device for distributing a series of frames for execution by at least one other device, the method comprising:

A. a frame acquisition module configured to obtain respective ones of the frames;

B. a frame execution time determination module configured to determine a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission module configured to transmit the series of frames to said at least one other device; and

D. a data source device migration control module configured to, in response to a predetermined event, provide migration information to another device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices of the synchrony group.

498. A method as defined in claim 497 in which at least one frame comprises audio data.

499. A computer program product for use in connection with a computer to provide a device for processing a series of frames of data provided by a data source at times specified by the data source in relation to a clock maintained by the data source, the computer program product comprising a computer-readable medium having encoded thereon

A. an interface module configured to enable the computer to receive the series of frames, each frame being associated with a time stamp, each time stamp indicating a time value relative to a time indicated a clock maintained by the data source, at which the device is to process the respective frame;

B. a current time retrieval module configured to enable the computer to obtain, from the data source, a current time value as indicated by the data source's clock;

C. an processing time determination module configured to enable the computer to determine, from the time stamp associated with each respective frame and a time differential value representing a difference between the current time value obtained by the current time retrieval module and a current time value indicated by a clock maintained by the device, a time, relative to the device's respective clock, at which the frame is to be processed; and

D. a frame processing module configured to enable the computer to process each respective frame at the time determined by the processing time determination module.

500. A computer program product as defined in claim 499 in which at least one frame comprises audio data, the computer, while processing said frame, generating a signal representing sound represented by the respective frame.

501. A computer program product as defined in claim 499 further including a control module a control module for enabling said computer to control execution of commands received by the interface module.

502. A computer program product as defined in claim 501 further comprising a buffer configured to enable the computer to buffer frames that the interface module has received until they are to be processed, the frame processing module being configured to enable the computer to process the buffered frames at the times determined by the processing time determination module.

503. A computer program product as defined in claim 502 in which the each frame is associated with a track, each track including at least one frame, the interface module being further configured to enable the computer to receive playback control commands for controlling processing of tracks, the control module being configured to enable the computer to execute at least one playback control command controlling track processing.

504. A computer program product as defined in claim 503 in which at least one type of playback control command enables the control module to, in turn, enable the frame processing module to terminate playback of a track that is currently being played back.

505. A computer program product as defined in claim 504 in which the playback control command enables the control module to terminate playback of a track with a frame subsequent to a frame currently being played.

506. A computer program product as defined in claim 503 in which the playback control command enables the control module to cancel playback of a track at least one of whose frames has been previously distributed, but for which playback has not begun.

507. A computer program product as defined in claim 503 in which the control module is configured to enable the computer to enable the frame processing module to resume processing of at least one frame associated with a subsequent track at a time determined by the processing time determination module.

508. A computer program product as defined in claim 501, the control module being configured to enable the computer to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive frames for processing by the frame processing module.

509. A computer program product as defined in claim 508 in which the control module is configured to enable the computer to, in response to receipt of a command to become a member of a synchrony group, enable the interface module to receive frames currently being transmitted using a selected multi-cast transmission methodology, and further to retrieve at least one previously distributed frame using a selected unicast transmission methodology.

510. A computer program product as defined in claim 501 in which the control module is configured to enable the computer to, in receipt of a command to disengage from a synchrony group, disable the interface module from receiving frames and the frame processing module from processing previously received frames that have not yet been processed.

511. A computer program product as defined in claim 501 further comprising a user interface module interface module configured to receive control information from a user interface module, the control module being configured to enable the computer to enable selected operations to be performed in response to the control information.

512. A computer program product as defined in claim 511 in which, in response to control information to enable another device to become a member of the device's synchrony group, the control module enables the interface module to transmit a command to the other device to enable the other device to become a member of the synchrony group.

513. A computer program product as defined in claim 511 in which, in response to control information to enable another device that is a member of the device's synchrony group to disengage from the synchrony group, the control module enables the interface module to transmit a command to the other device to enable the other device to disengage from the synchrony group.

514. A computer program product as defined in claim 511 in which, in response to control information to enable the data source to begin operating, the control module performs selected operations to enable the data source to begin operating.

515. A computer program product as defined in claim 514, the device also operating as the data source, the device including a frame generator module configured to enable the computer to, under control of the control module, obtain information from which the frames are to be generated and generate the frames.

516. A computer program product as defined in claim 515 in which the interface module is further configured to enable the computer to transmit the frames to at least one other device.

517. A computer program product as defined in claim 515 in which the device is not the data source, the control module being configured to enable the computer to enable the interface module to transmit a command to enable the data source to operate.

518. A computer program product as defined in claim 517 in which each frame is associated with a track, each track including at least one frame, and, in response to receipt of control information enabling processing of a track currently being processed to be terminated , the control module is configured to enable the computer to enable the interface module to transmit a command to enable the data source to insert a resynchronize command in the series of frames associated with the track to enable termination of processing of the track.

519. A computer program product as defined in claim 517 in which each frame is associated with a track, each track including at least one frame, and, in response to receipt of control information enabling processing of a track at least a portion of which has been received, but for which processing has not begun, the control module is configured to enable the computer to enable the interface module to transmit a command to enable the data source to insert a resynchronize command in the series of commands to disable processing of the track.

520. A computer program product as defined in claim 517 in which the control module is configured to enable the computer to enable the user interface module interface module to transmit status information to the user interface module.

521. A computer program product as defined in claim 520 in which the status information includes the processing status of at least one frame.

522. A computer program product as defined in claim 499 further including:

A. a migration information receiving module configured to enable the computer to receive migration information from the data source device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices of the synchrony group, and

B. a migration control module configured to enable the computer to, after the migration information has been received.,

   i. distribute the series of frames to the synchrony group, each frame being associated with a time stamp indicating a time, relative to the timing information that it receives

from the device operating as the data source device, at which the devices comprising the synchrony group are to execute the respective frame, and

ii. notify the members of the synchrony group that it is to thereafter operate as the data source device.

523. A computer program product as defined in claim 522 in which the information respecting the frames that are to be distributed is a source of streaming information, the migration control module being configured to enable the computer to, after the migration information has been received, assemble the streaming information into frames and associate each frame with a said time stamp.

524. A computer program product as defined in claim 522 in which the information respecting the frames that are to be distributed is in a file identified in the migration information, the migration control module being configured to enable the computer to obtain information from the identified file, assemble the information into frames and associate each frame with a said time stamp.

525. A computer program product as defined in claim 524 in which the information respecting the frames that are to be distributed begins at an offset, identified in the migration information, into the identified file, the migration control module being configured to enable the computer to assemble the information from a position in the file associated with the identified offset.

526. A computer program product as defined in claim 524 in which the information respecting the frames that are to be distributed is in a series of files identified in the migration information, the migration control module being configured to enable the computer to obtain information from successive files in the series assemble the information into frames and associate each frame with as said time stamp.

527. A computer program product as defined in claim 499 further including a clock rate adjustment module configured to enable the computer to adjust the member device's clock rate in relation to a clock rate value maintained by the data source device's clock.

528. A computer program product as defined in claim 527 in which the clock rate adjustment module is configured to enable the computer to adjust the clock rate in relation to time stamps associated with respective ones of the frames received from the data source device.

529. A computer program product for use in connection with a computer to provide a data source device for distributing a series of frames for execution by at least one other device at respective execution times, the computer program product comprising a computer readable medium having encoded thereon:

A. a frame acquisition module configured to enable the computer to obtain respective ones of the frames;

B. a frame execution time determination module configured to enable the computer to determine a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission module configured to enable the computer to transmit the series of frames to said at least one other device; and

D. an execution control command generation module configured to enable the computer to , in response to a predetermined event, enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to modify the execution sequence of the frames transmitted thereto.

530. A computer program product as defined in claim 529 in which at least one frame comprises audio data.

531. A computer program product as defined in claim 529 in which the frame transmission module is configured to enable the computer to utilize a selected multi-cast transmission methodology.

532. A computer program product as defined in claim 529 in which the predetermined event comprises input indicia provided by an operator.

533. A computer program product as defined in claim 529 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to control execution of at least one frame in relation to the time the predetermined event occurs.

534. A computer program product as defined in claim 533 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to terminate execution of a frame that it is currently executing and/or to not execute at least one frame that is to be executed subsequent to the predetermined event.

535. A computer program product as defined in claim 533 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation module being configured to enable the computer to , in response to said predetermined event,

(i) enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module, a command enabling the at least one other device to terminate execution of the currently executing frame and/or to not execute at least one subsequent frame, if any, in the frame sequence including the currently executing frame; and

(ii) enable

(a) the frame execution time determination module to provide that time stamps associated with frames of frame sequences subsequent to the frame sequence currently being executed will reflect the termination and/or non-execution of frames from the frame sequence of the frame that is currently being executed, and

(b) the frame transmission module to transmit the frames of the subsequent frame sequences to the at least one other device.

536. A computer program product as defined in claim 535 in which the frame transmission module is configured to enable the computer to insert a command to enable the at least one other device to not execute all frames in the frame sequence subsequent to the terminated and/or non-executed frame.

537. A computer program product as defined in claim 529 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to control execution of a frame in relation to the frame's position in the series of frames.

538. A computer program product as defined in claim 537 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to cancel execution of a frame in relation to the frame's position in the series of frames.

539. A computer program product as defined in claim 537 in which, in response to said predetermined event, the execution control command generation module is configured to enable the computer to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted frame in relation to the frame's position in the series of frames.

540. A computer program product as defined in claim 539 in which, in response to said predetermined event, the execution control command generation module is configured to enable the

computer to enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module a command to enable the at least one other device to cancel execution of a previously-transmitted frame, in relation to the frame's position in the series of frames, for which execution has not begun.

541. A computer program product as defined in claim 537 in which the series of frames includes a series of frame sequences with each frame sequence including a sub-set of the series of frames, the execution control command generation module being configured to enable the computer to , in response to said predetermined event,

(i) enable the frame transmission module to insert into the series of frames transmitted by the frame transmission module, a command enabling the at least one other device to cancel execution of a frame sequence for which execution has not begun; and

(ii) enable

(a) the frame execution time determination module to provide that time stamps associated with frames of frame sequences subsequent to the frame sequence whose execution was cancelled will reflect the cancellation, and

(b) the frame transmission module to transmit the frames of the subsequent frame sequences to the at least one other device.

542. A computer program product for use in connection with a computer to provide a data source device for distributing a series of frames for execution by at least one other device at respective execution times, the computer program product comprising a computer readable medium having encoded thereon:

A. a frame acquisition module configured to enable the computer to obtain respective ones of the frames;

B. a frame execution time determination module configured to enable the computer to determine a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time, the frame execution time determination module being configured to enable the computer to provide, associated with at least one frame, a time stamp indicating a frame execution time later than the time currently indicated by a clock maintained by the data source device; and

C. a frame transmission module configured to enable the computer to transmit the series of frames to said at least one other device.

543. A computer program product as defined in claim 542 in which the data is audio data.

544. A computer program product as defined in claim 542 in which the frame execution time determination module is configured to enable the computer to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be executed by the member devices at the time indicated by the time stamp.

545. A computer program product as defined in claim 544 in which the frame execution time determination module is configured to enable the computer to select the delay interval between the time indicated by the time stamp and the current time so to be sufficient to provide that the at least one frame distributed by the data source device can be distributed to the member devices for execution at the time indicated by the time stamp.

546. A computer program product as defined in claim 542 in which frame acquisition module is configured to enable the computer to obtain information associated with the frames from at least two types of information sources, the frame execution time determination module being configured to enable the computer to select, for each frame the delay interval between the time indicated by the time stamp and the current time based on the information source type that is associated with the respective frame.

547. A computer program product for use in connection with a computer to provide a data source device for distributing a series of frames for execution by at least one other device, the computer program product comprising a computer readable medium having encoded thereon:

A. a frame acquisition module configured to enable the computer to obtain respective ones of the frames;

B. a frame execution time determination module configured to enable the computer to determine a time at which each respective frame is to be executed, and to associate the respective frame with a time stamp indicating its execution time;

C. a frame transmission module configured to enable the computer to transmit the series of frames to said at least one other device; and

D. a data source device migration control module configured to enable the computer to, in response to a predetermined event, provide migration information to another device, the migration information including a source for information respecting the frames to be distributed, timing information relative to the clock maintained by the data source device and identifications of the member devices of the synchrony group.

548. A computer program product as defined in claim 547 in which the data is audio data.

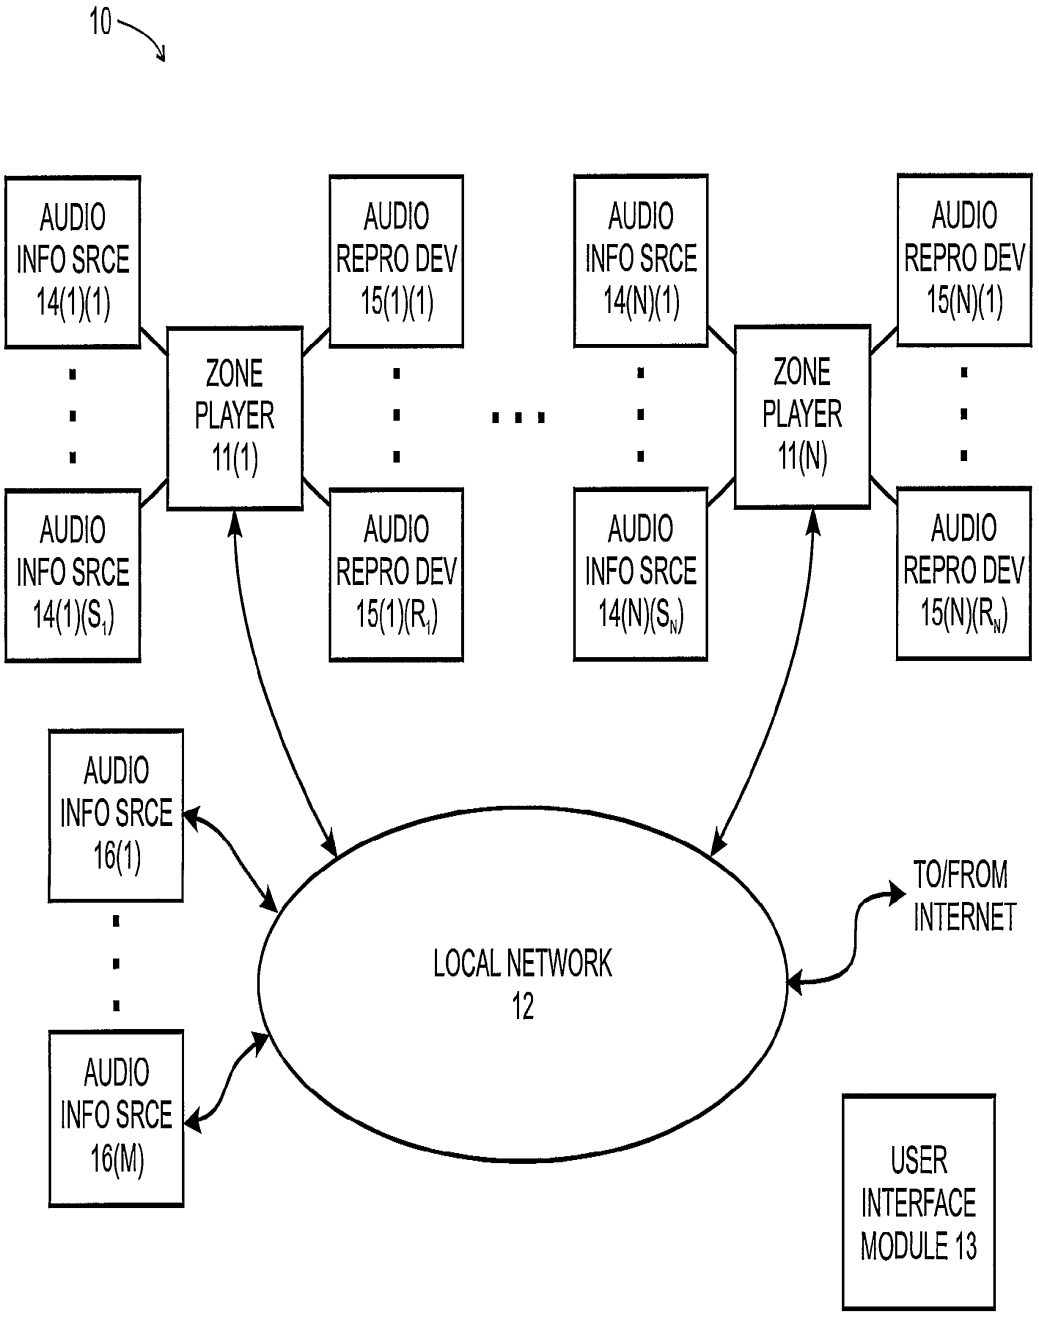

10
AUDIO INFO SRCE 14(1)(1)
AUDIO INFO SRCE 14(1)(S1)
ZONE PLAYER 11(1)
AUDIO REPRO DEV 15(1)(1)
AUDIO REPRO DEV 15(1)(R1)
AUDIO INFO SRCE 14(N)(1)
AUDIO INFO SRCE 14(N)(SN)
ZONE PLAYER 11(N)
AUDIO REPRO DEV 15(N)(1)
AUDIO REPRO DEV 15(N)(RN)
AUDIO INFO SRCE 16(1)
AUDIO INFO SRCE 16(M)
LOCAL NETWORK 12
TO/FROM INTERNET
USER INTERFACE MODULE 13

FIG. 1

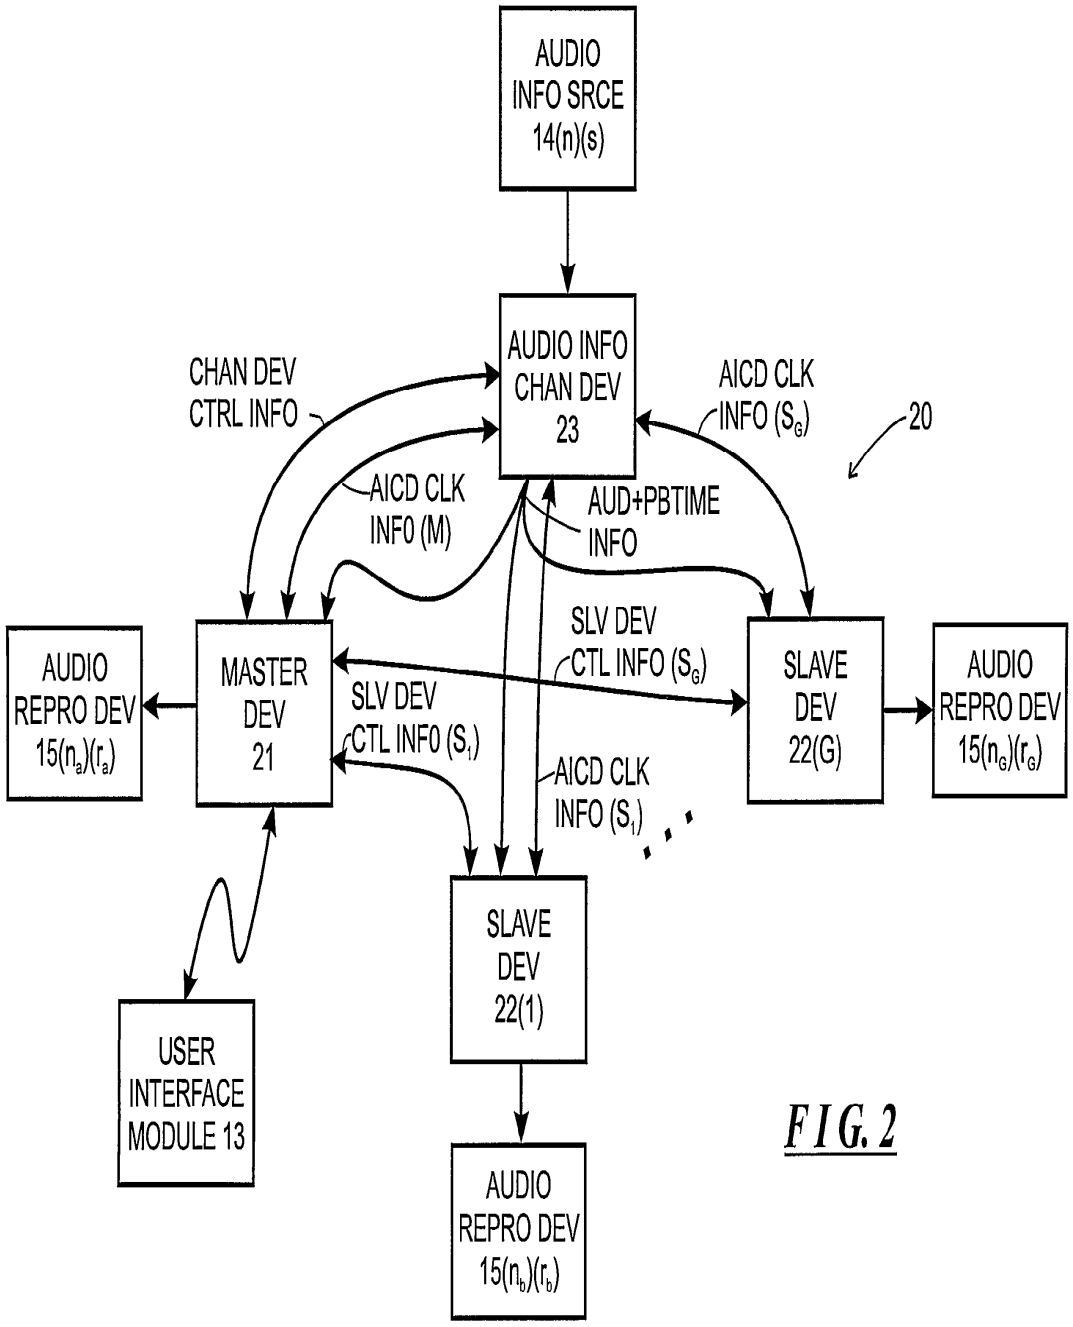

AUDIO INFO SRCE 14(n)(s)
AUDIO INFO CHAN DEV 23
CHAN DEV CTRL INFO
AICD CLK INFO (S_G)
20
AICD CLK INFO (M)
AUD+PBTIME INFO
AUDIO REPRO DEV 15(n_a)(r_a)
MASTER DEV 21
SLV DEV CTL INFO (S_G)
SLAVE DEV 22(G)
AUDIO REPRO DEV 15(n_G)(r_G)
SLV DEV CTL INFO (S_1)
AICD CLK INFO (S_1)
SLAVE DEV 22(1)
USER INTERFACE MODULE 13
AUDIO REPRO DEV 15(n_b)(r_b)

FIG. 2

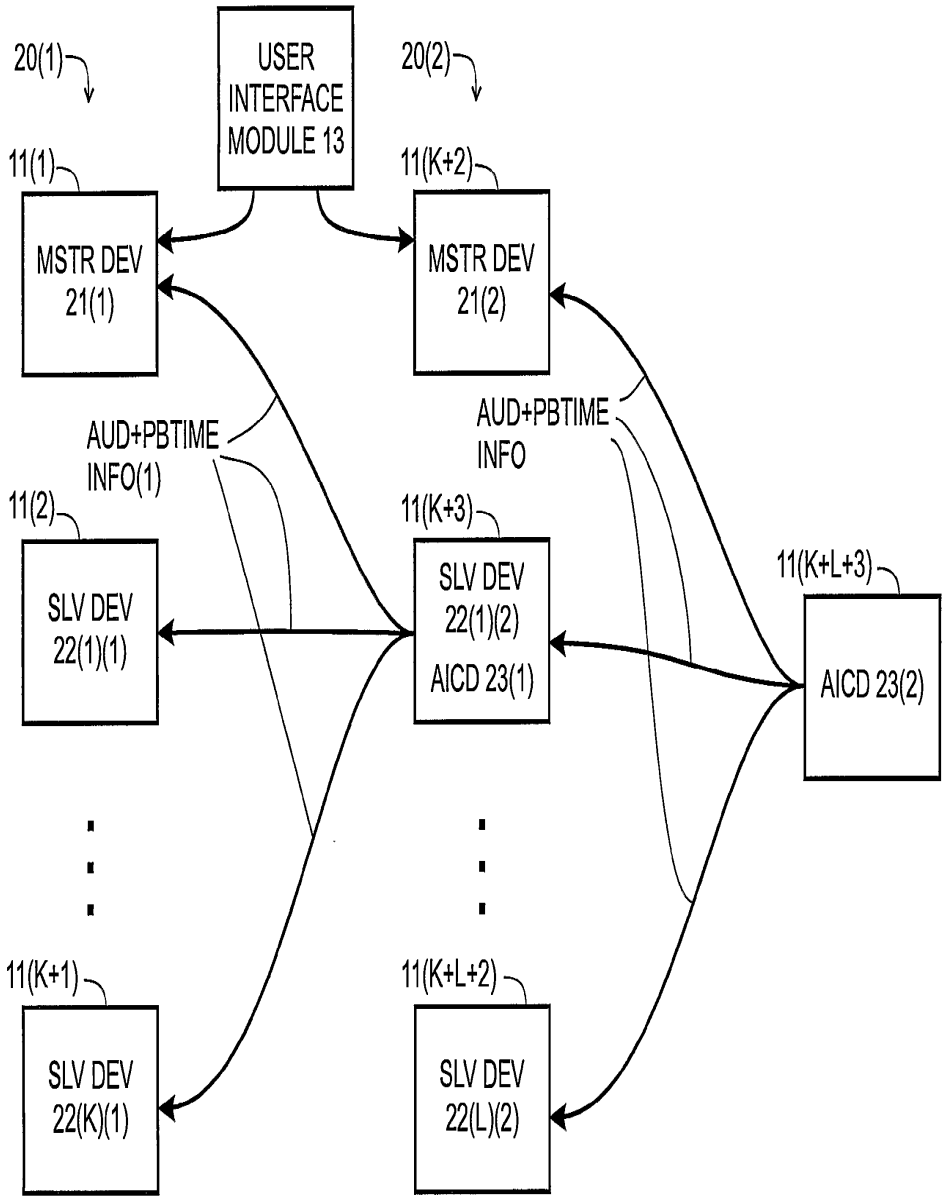

20(1)
USER INTERFACE MODULE 13
20(2)
11(1)
11(K+2)
MSTR DEV 21(1)
MSTR DEV 21(2)
AUD+PBTIME INFO(1)
AUD+PBTIME INFO
11(2)
11(K+3)
11(K+L+3)
SLV DEV 22(1)(1)
SLV DEV 22(1)(2) AICD 23(1)
AICD 23(2)
11(K+1)
11(K+L+2)
SLV DEV 22(K)(1)
SLV DEV 22(L)(2)

FIG. 2A

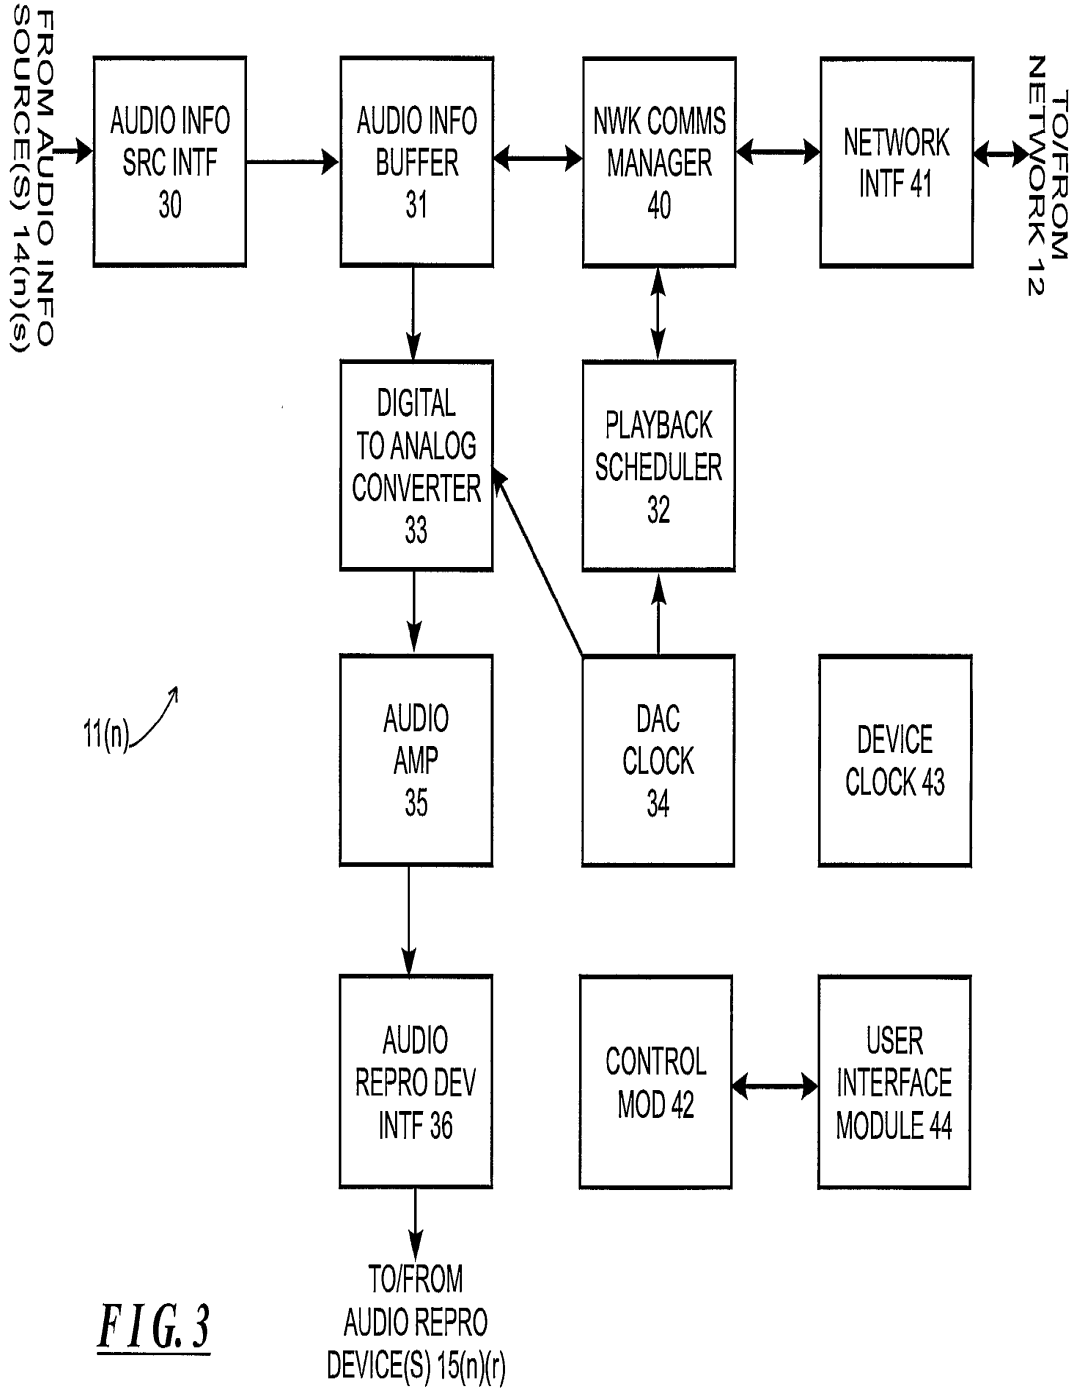

FROM AUDIO INFO SOURCE(S) 14(n)(s)
AUDIO INFO SRC INTF 30
AUDIO INFO BUFFER 31
NWK COMMS MANAGER 40
NETWORK INTF 41
TO/FROM NETWORK 12
DIGITAL TO ANALOG CONVERTER 33
PLAYBACK SCHEDULER 32
11(n)
AUDIO AMP 35
DAC CLOCK 34
DEVICE CLOCK 43
AUDIO REPRO DEV INTF 36
CONTROL MOD 42
USER INTERFACE MODULE 44
TO/FROM AUDIO REPRO DEVICE(S) 15(n)(r)

FIG. 3

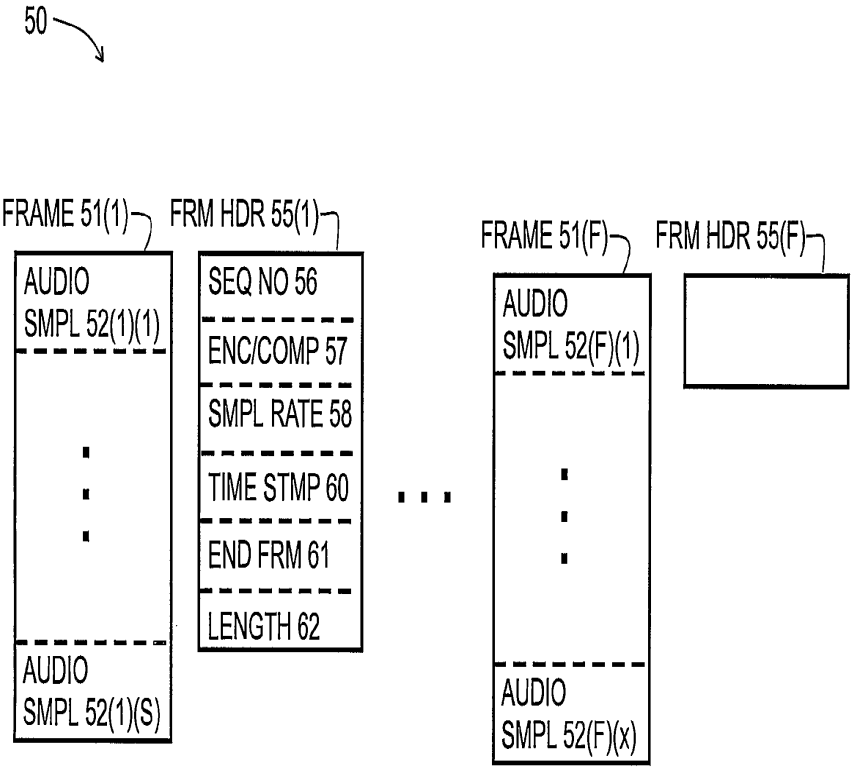

50
FRAME 51(1)
FRM HDR 55(1)
AUDIO SMPL 52(1)(1)
AUDIO SMPL 52(1)(S)
SEQ NO 56
ENC/COMP 57
SMPL RATE 58
TIME STMP 60
END FRM 61
LENGTH 62
FRAME 51(F)
FRM HDR 55(F)
AUDIO SMPL 52(F)(1)
AUDIO SMPL 52(F)(x)

FIG. 4

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference 13-0410-PCT	FOR FURTHER ACTION	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
International application No. **PCT/US2013/046372**	International filing date *(day/month/year)* **18 June 2013 (18.06.2013)**	(Earliest) Priority Date *(day/month/year)* 26 June 2012 (26.06.2012)
Applicant **SONOS, INC.**		

This International search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 3 sheets.

☐ It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**
   a. With regard to the **language**, the international search was carried out on the basis of:
      ☒ the international application in the language in which it was filed
      ☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))
   b. ☐ This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6*bis*(a)).
   c. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.
2. ☐ **Certain claims were found unsearchable** (See Box No. II)
3. ☐ **Unity of invention is lacking** (See Box No. III)
4. With regard to the **title**,
   ☒ the text is approved as submitted by the applicant.
   ☐ the text has been established by this Authority to read as follows:

5. With regard to the **abstract**,
   ☒ the text is approved as submitted by the applicant.
   ☐ the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority.

6. With regard to the **drawings**,
   a. the figure of the drawings to be published with the abstract is Figure No. 7
      ☐ as suggested by the applicant.
      ☒ as selected by this Authority, because the applicant failed to suggest a figure.
      ☐ as selected by this Authority, because this figure better characterizes the invention.
   b. ☐ none of the figure is to be published with the abstract.

Form PCT/ISA/210 (first sheet) (July 2009)

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/046372

A. CLASSIFICATION OF SUBJECT MATTER

G06F 15/16(2006.01)i

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F 15/16; G09B 5/00; G06F 17/00

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: media playback, link, queueplaylist, credentialauthorizeauthenticate

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2007-0048713 A1 (DANIEL PLASTINA et al.) 01 March 2007 See paragraph 8 - paragraph 48; paragraph 73 - paragraph 156;and figures 1-23.	1,4-11,14-20
A		2-3,12-13
A	US 2012-0030366 A1 (COLLART TODD et al.) 02 February 2012 See paragraph 5 - paragraph 35; paragraph 67 - paragraph 107; paragraph 115 - paragraph 163; and figures 1-7,11-19.	1-20
A	US 2006-0143236 A1 (MICHAEL WU) 29 June 2006 See paragraph 14 - paragraph 42; and figures 1-4.	1-20
A	US 2008-0091771 A1 (JEFFREY L. ALLEN et al.) 17 April 2008 See paragraph 28 - paragraph 48; claim 1 and figures 3-6.	1-20

☐ Further documents are listed in the continuation of Box C. ☒ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents,such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
23 August 2013 (23.08.2013)	**26 August 2013 (26.08.2013)**

Name and mailing address of the ISA/KR
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701, Republic of Korea
Facsimile No. +82-42-472-7140

Authorized officer
PARK Sang Hyun
Telephone No. +82-42-481-8263

Form PCT/ISA/210 (second sheet) (July 2009)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.
PCT/US2013/046372

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007-0048713 A1	01/03/2007	CN101243426 A	13/08/2008
		CN101243426 C0	13/08/2008
		KR 10-2008-0033399 A	16/04/2008
		US 2007-0048712 A1	01/03/2007
		US 2007-0048714 A1	01/03/2007
		US 2007-0083556 A1	12/04/2007
		US 7831605 B2	09/11/2010
		US 8140601 B2	20/03/2012
		WO 2007-021831 A1	22/02/2007
US 2012-0030366 A1	02/02/2012	AU 2008-335095 A1	18/06/2009
		AU 2008-335095 B2	28/03/2013
		CA 2708778 A1	18/06/2009
		EP 2232365 A2	29/09/2010
		EP 2232365 A4	31/07/2013
		JP 2011-520302 A	14/07/2011
		US 2009-150553 A1	11/06/2009
		US 8060631 B2	15/11/2011
		WO 2009-076474 A2	18/06/2009
		WO 2009-076474 A3	13/08/2009
US 2006-0143236 A1	29/06/2006	None	
US 2008-0091771 A1	17/04/2008	CN 101523495 A	02/09/2009
		CN 101997910 A	30/03/2011
		CN 101997910 B	13/03/2013
		CN101523495 B	06/04/2011
		EP 2089881 A1	19/08/2009
		EP 2089881 A4	16/01/2013
		JP 2010-507144A	04/03/2010
		KR 10-2009-0079885 A	22/07/2009
		MX 2009002806 A	31/03/2009
		WO 2008-048761 A1	24/04/2008

Form PCT/ISA/210 (patent family annex) (July 2009)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:
GEORGE CHRISTOPHER N.

HANLEY, FLIGHT & ZIMMERMAN, LLC 150 S.
WACKER DRIVE SUITE 2100 CHICAGO IL 60606 USA

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43bis.1)

Date of mailing *(day/month/year)*	**26 August 2013 (26.08.2013)**
Applicant's or agent's file reference 13-0410-PCT	**FOR FURTHER ACTION** See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2013/046372	**18 June 2013 (18.06.2013)**	26 June 2012 (26.06.2012)

International Patent Classification (IPC) or both national classification and IPC
G06F 15/16(2006.01)i

Applicant
SONOS, INC.

1. This opinion contains indications relating to the following items:

- [x] Box No. I Basis of the opinion
- [ ] Box No. II Priority
- [ ] Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- [ ] Box No. IV Lack of unity of invention
- [x] Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- [ ] Box No. VI Certain documents cited
- [ ] Box No. VII Certain defects in the international application
- [ ] Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1bis(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.
For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA/KR	Date of completion of this opinion	Authorized officer
Korean Intellectual Property Office 189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701, Republic of Korea Facsimile No. +82-42-472-7140	23 August 2013 (23.08.2013)	PARK, Sang Hyun Telephone No. +82-42-481-8263

Form PCT/ISA/237 (cover sheet) (July 2011)

WRITTEN OPINION OF THE
INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US2013/046372

Box No. I Basis of this opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed

   ☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43*bis*.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of:

   a. a sequence listing filed or furnished
      ☐ on paper
      ☐ in electronic form

   b. time of filing or furnishing
      ☐ contained in the international application as filed.
      ☐ filed together with the international application in electronic form.
      ☐ furnished subsequently to this Authority for the purposes of search.

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additioanl copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

WRITTEN OPINION OF THE
INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US2013/046372

Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Claims	2-3,12-13	YES
	Claims	1,4-11,14-20	NO
Inventive step (IS)	Claims	2-3,12-13	YES
	Claims	1,4-11,14-20	NO
Industrial applicability (IA)	Claims	1-20	YES
	Claims	NONE	NO

2. Citations and explanations :

(1) Reference is made to the following documents:

D1: US 2007-0048713 A1 (DANIEL PLASTINA et al.) 01 March 2007
D2: US 2012-0030366 A1 (COLLART TODD et al.) 02 February 2012
D3: US 2006-0143236 A1 (MICHAEL WU) 29 June 2006
D4: US 2008-0091771 A1 (JEFFREY L. ALLEN et al.) 17 April 2008

(2) Novelty & Inventive Steps

Concerning claims 1-10:
Claims 1,4-10
D1 discloses a community media playlist sharing system, wherein system users can upload media playlists which are automatically converted to a standardized format, media playlists can be shared with other users of the community media playlist sharing system, users can link to other playlists and obtain media from playlists automatically from third party points of purchase, users can provide recommendations, ratings and rankings relative to other playlists by system users and receive recommendations from other users and an editorial staff, participate in social networking functions including forums, discussions, weblogs, playlist topic groups and inter-user communications based on affinity for similar playlists, create playlists manually, automatically, or in real time, and wherein system users can obtain points for participation and popularity in the system and obtain profit by meta-affiliate sharing of profit for purchases of music made through playlists submitted by that user as well as through playlist toolbars distributed to third party websites.[paragraph 73 - paragraph 156, and figures 1-23] As all of the features of claims 1,4-10 are disclosed in D1, these claims are anticipated in D1. Therefore, claims 1,4-10 lack novelty under PCT Article 33(2) and are not considered as involving an inventive step under PCT Article 33(3).
Claims 2-3:
The subject matter of claims 2-3 is different from that of D1 in a guest mode. And it is

Continued on Supplemental Box

Supplemental Box

In case the space in any of the preceding boxes is not sufficient.
Continuation of: Box No. V

not obvious to a person skilled in the art by the documents, taken alone or in combination. Therefore, claims 2-3 are novel and involve an inventive step under PCT Article 33(2) and (3).

Concerning claims 11-20:
Claims 11,14-20:
D1 discloses a community media playlist sharing system, wherein system users can upload media playlists which are automatically converted to a standardized format, media playlists can be shared with other users of the community media playlist sharing system, users can link to other playlists and obtain media from playlists automatically from third party points of purchase, users can provide recommendations, ratings and rankings relative to other playlists by system users and receive recommendations from other users and an editorial staff, participate in social networking functions including forums, discussions, weblogs, playlist topic groups and inter-user communications based on affinity for similar playlists, create playlists manually, automatically, or in real time, and wherein system users can obtain points for participation and popularity in the system and obtain profit by meta-affiliate sharing of profit for purchases of music made through playlists submitted by that user as well as through playlist toolbars distributed to third party websites.[paragraph 73 - paragraph 156, and figures 1-23] As all of the features of claims 11,14-20 are disclosed in D1, these claims are anticipated in D1. Therefore, claims 11,14-20 lack novelty under PCT Article 33(2) and are not considered as involving an inventive step under PCT Article 33(3).
Claims 12-13:
The subject matter of claims 12-13 is different from that of D1 in a guest mode. And it is not obvious to a person skilled in the art by the documents, taken alone or in combination. Therefore, claims 12-13 are novel and involve an inventive step under PCT Article 33(2) and (3).

(3) Industrial Applicability

The subject matter of claims 1-20 is believed to be industrially applicable under PCT Article 33(4) because the claimed invention can be used in the computer industry.

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference 12-0402-PCT	**FOR FURTHER ACTION**	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
International application No. **PCT/US2013/046386**	International filing date *(day/month/year)* **18 June 2013 (18.06.2013)**	(Earliest) Priority Date *(day/month/year)* 26 June 2012 (26.06.2012)
Applicant **SONOS, INC.**		

This International search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 3 sheets.

☐ It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**
   a. With regard to the **language**, the international search was carried out on the basis of:
      ☒ the international application in the language in which it was filed
      ☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))
   b. ☐ This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6*bis*(a)).
   c. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.
2. ☐ **Certain claims were found unsearchable** (See Box No. II)
3. ☐ **Unity of invention is lacking** (See Box No. III)
4. With regard to the **title,**
   ☒ the text is approved as submitted by the applicant.
   ☐ the text has been established by this Authority to read as follows:
5. With regard to the **abstract,**
   ☒ the text is approved as submitted by the applicant.
   ☐ the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority.
6. With regard to the drawings,
   a. the figure of the **drawings** to be published with the abstract is Figure No. 1
      ☐ as suggested by the applicant.
      ☒ as selected by this Authority, because the applicant failed to suggest a figure.
      ☐ as selected by this Authority, because this figure better characterizes the invention.
   b. ☐ none of the figure is to be published with the abstract.

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/046386

A. CLASSIFICATION OF SUBJECT MATTER

G06Q 50/10(2012.01)i

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06Q 50/10; G06F 15/16; H04N 7/173; G06F 7/00; G06F 17/30; H04L 9/00; H04W 40/00

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords:playlist, authorization, participate, zone:

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 7483958 B1 (TAREK Z. ELABBADY et al.) 27 January 2009 See abstract, column 7, lines 61-66, column 11, lines 1-20, claims 1, 6 and figure 4.	1-21
Y	US 2010-0228740 A1 (ALAN CANNISTRARO et al.) 09 September 2010 See abstract, paragraphs [0009], [0036]-[0037], [0047]-[0057], [0089], claim 1 and figure 4.	1-21
A	US 2012-0129446 A1 (STEVE KO et al.) 24 May 2012 See abstract, paragraphs [0036]-[0037], claim 1 and figure 4.	1-21
A	US 2009-0031336 A1 (TIMOTHY R. CHAVEZ et al.) 29 January 2009 See abstract, paragraph [0021], claim 1 and figure 3.	1-21
A	US 2007-0299778 A1 (RYAN A. HAVESON et al.) 27 December 2007 See abstract, claim 1 and figure 3.	1-21

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents,such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
27 September 2013 (27.09.2013)	**30 September 2013 (30.09.2013)**

Name and mailing address of the ISA/KR
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701, Republic of Korea
Facsimile No. +82-42-472-7140

Authorized officer
KIM Sung Gon
Telephone No. +82-42-481-8746

Form PCT/ISA/210 (second sheet) (July 2009)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.
PCT/US2013/046386

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 7483958 B1	27/01/2009	None	
US 2010-0228740 A1	09/09/2010	None	
US 2012-0129446 A1	24/05/2012	AT 520238 T	15/08/2011
		CN 101346965 A	14/01/2009
		CN 102984041 A	20/03/2013
		EP 1964365 A1	03/09/2008
		EP 1964365 B1	10/08/2011
		EP 2299654 A1	23/03/2011
		HK 1122664 A1	18/11/2011
		JP 2009-521185 A	28/05/2009
		JP 2012-235471 A	29/11/2012
		US 2007-169115 A1	19/07/2007
		US 2010-049835 A1	25/02/2010
		US 2013-080599 A1	28/03/2013
		US 7702279 B2	20/04/2010
		US 8112032 B2	07/02/2012
		US 8315555 B2	20/11/2012
		WO 2007-075376 A1	05/07/2007
US 2009-0031336 A1	29/01/2009	None	
US 2007-0299778 A1	27/12/2007	None	

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:
JASPER JOSEPH T.

HANLEY, FLIGHT & ZIMMERMAN, LLC 150 S.
WACKER DRIVE, SUITE 2100 CHICAGO IL 60606 USA

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43bis.1)

Date of mailing *(day/month/year)*	**30 September 2013 (30.09.2013)**
Applicant's or agent's file reference 12-0402-PCT	**FOR FURTHER ACTION** See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2013/046386	**18 June 2013 (18.06.2013)**	26 June 2012 (26.06.2012)

International Patent Classification (IPC) or both national classification and IPC

G06Q 50/10(2012.01)i

Applicant
SONOS, INC.

1. This opinion contains indications relating to the following items:

- [x] Box No. I Basis of the opinion
- [ ] Box No. II Priority
- [ ] Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- [ ] Box No. IV Lack of unity of invention
- [x] Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- [ ] Box No. VI Certain documents cited
- [ ] Box No. VII Certain defects in the international application
- [ ] Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1bis(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.
For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA/KR	Date of completion of this opinion	Authorized officer
Korean Intellectual Property Office 189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701, Republic of Korea Facsimile No. +82-42-472-7140	27 September 2013 (27.09.2013)	KIM Sung Gon Telephone No. +82-42-481-8746

Form PCT/ISA/237 (cover sheet) (July 2011)

WRITTEN OPINION OF THE
INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US2013/046386

Box No. I Basis of this opinion

1. With regard to the **language**, this opinion has been established on the basis of :

   ☒ the international application in the language in which it was filed

   ☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43*bis*.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of:

   a. a sequence listing filed or furnished
      ☐ on paper
      ☐ in electronic form

   b. time of filing or furnishing
      ☐ contained in the international application as filed.
      ☐ filed together with the international application in electronic form.
      ☐ furnished subsequently to this Authority for the purposes of search.

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additioanl copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Form PCT/ISA/237 (Box No. I)(July 2011)

WRITTEN OPINION OF THE
INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US2013/046386

Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Claims	1-21	YES
	Claims	NONE	NO
Inventive step (IS)	Claims	NONE	YES
	Claims	1-21	NO
Industrial applicability (IA)	Claims	1-21	YES
	Claims	NONE	NO

2. Citations and explanations :

Reference is made to the following documents:

D1: US 7483958 B1 (TAREK Z. ELABBADY et al.) 27 January 2009
D2: US 2010-0228740 A1 (ALAN CANNISTRARO et al.) 09 September 2010

1. Novelty and Inventive Step

1.1 Claims 1-10

1.1.1 Independent claim 1

D1, which is considered to be the closest prior art to the subject matter of claim 1, discloses a method comprising: presenting media content information (e.g., music file, a play list) to the user, determining whether an original media playing license associated with the particular media contents file is available, generating a temporary media playing license from the original media playing license, the temporary media playing license specifying restrictions relating to playback of the particular media content file via the home network device; and sending the temporary media playing license to the home network device (See abstract, column 11, lines 1-20, claim 1, figure 4 in D1). Claim 1 differs from D1 in facilitating addition of content to the playlist by at least one participating user; and playback of content from the playlist in the restricted playback mode. However, D2 discloses a method comprising: receiving from at least one of the plurality of user devices a vote on a media item; and organizing the playlist based on the received votes (See abstract, paragraphs [0047]-[0057], claim 1, figure 4 in D2). Accordingly, claim 1 would have been obvious over D1 and D2. Therefore, claim 1 lacks an inventive step under PCT Article 33(3).

Continued on Supplemental Box

Supplemental Box

In case **the space in any of the preceding boxes is not sufficient.**
Continuation of : Box No. V

1.1.2 Dependent claims 2-10

The additional feature of claim 2 is identical to the feature of D1 in that media items in the playlist are automatically chosen and added to the playlist by a program on the playlist host device (See claim 6 in D1). Accordingly, claim 2 would have been obvious over D1 and D2. Therefore, claim 2 lacks an inventive step under PCT Article 33(3).

Claim 3 further specifies authenticating the at least one participating user to enable the at least one participating user to add content to the playlist in the restricted playback mode. However, the additional feature of claim 3 is virtually suggested by the feature of D2 considering that the host receives an opt-in command from an authorized user and the opt-in command instructs the host to allow remote access to the playlist (See abstract, paragraphs [0047]-[0057], figure 4 in D2). Accordingly, claim 3 would have been obvious over D1 and D2. Therefore, claim 3 lacks an inventive step under PCT Article 33(3).

The additional feature of claim 4 is identical to the feature of D2 in receiving from at least one of the plurality of user devices a vote on a media item; and organizing the playlist based on the received votes (See abstract, paragraphs [0047]-[0057], claim 1, figure 4 in D2). Accordingly, claim 4 would have been obvious over D1 and D2. Therefore, claim 4 lacks an inventive step under PCT Article 33(3).

Claims 5, 6 further specify that facilitating addition of content comprises accepting user selection of the content from a list of available content via a graphical user interface, and the list of available content includes a library of shared music available via a network. However, the additional features of claims 5, 6 are virtually suggested by the feature of D1 considering presenting media content information (e.g., metadata from the media library) to the user through a display screen, and selecting one or more music files, a play list, or other grouping of music files by the user (See abstract, column 11, lines 1-20, claim 1, figure 4 in D1). And D2 discloses a step of sending a playlist on a host device to a plurality of portable devices for display on a graphical user interface (See paragraph [0009] in D2). Accordingly, claims 5, 6 would have been obvious over D1 and D2. Therefore, claims 5, 6 lack an inventive step under PCT Article 33(3).

Claim 7 further comprises that the restricted playback mode is initiated and ended by a host user according to user authorization. However, the additional feature of claim 7 is merely a

Continued on The Next Page

Supplemental Box

In case the space in any of the preceding boxes is not sufficient.
Continuation of: Previous Page

matter of design option when the general knowledge in relevant field of the art is used. Accordingly, claim 7 would have been obvious over D1 and D2. Therefore, claim 7 lacks an inventive step under PCT Article 33(3).

Claim 8 further specifies configuring the restricted playback mode for a fixed time interval. However, the additional feature of claim 8 is virtually suggested by the feature of D2 considering ordering changes can be postponed until the interval since the last touch event has exceeded 2 seconds (See paragraph [0089] in D2). Accordingly, claim 8 would have been obvious over D1 and D2. Therefore, claim 8 lacks an inventive step under PCT Article 33(3).

Claim 9 further specifies that a restricted playback mode includes a specified configuration of zones. However, the additional feature of claim 9 is virtually suggested by the feature of D1 considering that the media content license may be associated with a specific entity or with multiple entities (e.g., groups) (See column 7, lines 61-66 in D1). Accordingly, claim 9 would have been obvious over D1 and D2. Therefore, claim 9 lacks an inventive step under PCT Article 33(3).

Claim 10 further comprises that the subset of playback functions includes limited user interaction with at least one of zone configuration and content playback parameters. However, the additional feature of claim 10 is merely a matter of design option when the general knowledge in relevant field of the art is used. Accordingly, claim 10 would have been obvious over D1 and D2. Therefore, claim 10 lacks an inventive step under PCT Article 33(3).

1.2 Claims 11-20

1.2.1 Independent claim 11

Claim 11 relates to a tangible computer readable storage medium including instructions which, when executed by a processor, are to implement at least a method. But it shares similar technical features with claim 1. Accordingly, the same reasoning applies to claim 11. Therefore, claim 11 lacks an inventive step under PCT Article 33(3).

1.2.2 Dependent claims 12-20

Claims 12-20 share the same technical features with claims 2-10, respectively. Accordingly, the

Continued on The Next Page

Supplemental Box

In case **the space in any of the preceding boxes is not sufficient.**
Continuation of : Previous Page

same reasoning applies to claims 12-20. Therefore, claims 12-20 lack an inventive step under PCT Article 33(3).

1.3 Independent claim 21

Claim 21 relates to a media playback system, but it shares similar technical features with claim 1. Accordingly, the same reasoning applies to claim 21. Therefore, claim 21 lacks an inventive step under PCT Article 33(3).

2. Industrial Applicability

Claims 1-21 are industrially applicable under PCT Article 33(4).

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant:	Nicholas A.J. Millington
Title:	METHOD AND APPARATUS FOR DYNAMIC MASTER DEVICE SWITCHING IN A SYNCHRONY GROUP
Appl. No.:	13/864,248
Filing Date:	April 17, 2013
Examiner:	Curtis Kuntz
Art Unit:	2656
Confirmation Number:	5095

SECOND SUPPLEMENTAL PRELIMINARY AMENDMENT

Sir:

Please enter the following amendments prior to examination of the application.

Amendments to the Specification are presented on page 2;

Remarks/Arguments begin on page 3 of this document.

Amendments to the Specification:

On page 1 of the specification, please amend the first paragraph after the section heading "Incorporation By Reference" with the new paragraph as follows:

This application claims priority under 35 USC §120 of US Application No. 13/297,000, filed November 15, 2011, entitled "System And Method For Synchronizing Operations Among A Plurality Of Independently Clocked Digital Data Processing Devices", which is a continuation of Patent Application Serial No. 10/816,217, now US Patent 8,234,395, filed on April 1, 2004, entitled "System and Method for Synchronizing Operations ~~Channel Handoff as~~ Among a Plurality of Independently Clocked Digital Data Processing Devices", which is a non-provisional application claiming priority under 35 USC §119 to Provisional Patent Application Serial No. 60/490,768, filed on July 28, 2003, entitled "Method for Synchronizing Audio Playback Between Multiple Networked Devices," now expired, all of which are assigned to the assignee of the present application and are incorporated herein by reference.

REMARKS

Applicants request that the above-presented amendments are entered prior to examination. The Examiner is invited to contact the undersigned if it would be helpful to advance the prosecution of the application.

Respectfully submitted,

Date: November 1, 2013

By: /Robert J. Irvine III/

TechLaw LLP
Telephone: (312) 834-4334

Robert J. Irvine III
Attorney for Applicant
Registration No. 41,865

Electronic Acknowledgement Receipt

EFS ID:	17290520
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 (13-0401)
Receipt Date:	01-NOV-2013
Filing Date:	17-APR-2013
Time Stamp:	19:58:52
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		SNS-71301US10-2ndSuppPreliminAmend.pdf	91210 d7d5c811f2ca38157f15d59dd79b906d97d95b79	yes	3

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Preliminary Amendment	1	1
	Specification	2	2
	Applicant Arguments/Remarks Made in an Amendment	3	3
Warnings:			
Information:			
	Total Files Size (in bytes):	91210	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 13/864,248	Filing Date 04/17/2013	☐ To be Mailed

ENTITY: ☒ LARGE ☐ SMALL ☐ MICRO

APPLICATION AS FILED – PART I

	(Column 1)	(Column 2)		
FOR	NUMBER FILED	NUMBER EXTRA	RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	

APPLICATION AS AMENDED – PART II

		(Column 1)		(Column 2)	(Column 3)		
	11/01/2013	CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
AMENDMENT	Total (37 CFR 1.16(i))	* 20	Minus	** 20	= 0	X $80 =	0
	Independent (37 CFR 1.16(h))	* 3	Minus	***3	= 0	X $420 =	0
	☐ Application Size Fee (37 CFR 1.16(s))						
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
						TOTAL ADD'L FEE	0

		(Column 1)		(Column 2)	(Column 3)		
		CLAIMS REMAINING AFTER AMENDMENT		HIGHEST NUMBER PREVIOUSLY PAID FOR	PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
AMENDMENT	Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))						
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
						TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

LIE
/PATRICIA F. LEWIS/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**
If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
13/864,248	04/17/2013	Nicholas A. J. Millington	SNS-71301.US.10 (13-0401)

CONFIRMATION NO. 5095

109473
TechLaw LLP
200 S. Wacker Drive
Suite 3100
Chicago, IL 60606

PUBLICATION NOTICE

OC000000063597882

Title:SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES

Publication No.US-2013-0231766-A1
Publication Date:09/05/2013

NOTICE OF PUBLICATION OF APPLICATION

The above-identified application will be electronically published as a patent application publication pursuant to 37 CFR 1.211, et seq. The patent application publication number and publication date are set forth above.

The publication may be accessed through the USPTO's publically available Searchable Databases via the Internet at www.uspto.gov. The direct link to access the publication is currently http://www.uspto.gov/patft/.

The publication process established by the Office does not provide for mailing a copy of the publication to applicant. A copy of the publication may be obtained from the Office upon payment of the appropriate fee set forth in 37 CFR 1.19(a)(1). Orders for copies of patent application publications are handled by the USPTO's Office of Public Records. The Office of Public Records can be reached by telephone at (703) 308-9726 or (800) 972-6382, by facsimile at (703) 305-8759, by mail addressed to the United States Patent and Trademark Office, Office of Public Records, Alexandria, VA 22313-1450 or via the Internet.

In addition, information on the status of the application, including the mailing date of Office actions and the dates of receipt of correspondence filed in the Office, may also be accessed via the Internet through the Patent Electronic Business Center at www.uspto.gov using the public side of the Patent Application Information and Retrieval (PAIR) system. The direct link to access this status information is currently http://pair.uspto.gov/. Prior to publication, such status information is confidential and may only be obtained by applicant using the private side of PAIR.

Further assistance in electronically accessing the publication, or about PAIR, is available by calling the Patent Electronic Business Center at 1-866-217-9197.

Office of Data Managment, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant:	Nicholas A.J. Millington
Title:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
Appl. No.:	13/864,248
Filing Date:	April 17, 2013
Examiner:	Curtis Kuntz
Art Unit:	2656
Confirmation Number:	5095

SUPPLEMENTAL PRELIMINARY AMENDMENT

Sir:

Please enter the following amendments prior to examination of the application.

Amendments to the Specification are presented on page 2;

Remarks/Arguments begin on page 3 of this document.

Please amend the application as follows:

Amendments to the Specification:

Please amend the title of the application as follows:

~~System and Method for Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices~~

Method and Apparatus for Dynamic Master Device Switching in a Synchrony Group

REMARKS

Applicants request that the above-presented amendments are entered prior to examination. Applicants have also submitted herewith a supplemental ADS indicating the changes to the title. The Examiner is invited to contact the undersigned if it would be helpful to advance the prosecution of the application.

Respectfully submitted,

Date: August 21, 2013

By: /Robert J. Irvine III/

TechLaw LLP
Telephone: (312) 834-4334

Robert J. Irvine III
Attorney for Applicant
Registration No. 41,865

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	SNS-71301-US.10
	Application Number	13/864,248
Title of Invention	System and Method for Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City		State/Province		Country of Residence	

Mailing Address of Inventor:

Address 1			
Address 2			
City		State/Province	
Postal Code		Country	

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

☐ **An Address is being provided for the correspondence Information of this application.**

Customer Number			
Email Address		Add Email	Remove Email

Application Information:

Method and Apparatus for Dynamic Master Device Switching in a Synchrony Group

Title of the Invention	~~System and Method for Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices~~		
Attorney Docket Number		Small Entity Status Claimed	☐
Application Type			
Subject Matter			
Total Number of Drawing Sheets (if any)		Suggested Figure for Publication (if any)	

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	
	Application Number	
Title of Invention	System and Method for Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices	

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)
☐ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number			

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

Prior Application Status		Remove	
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button.

Foreign Priority Information:

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(d). When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX)[i] the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(h)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

			Remove
Application Number	Country[i]	Filing Date (YYYY-MM-DD)	Access Code[i] (if applicable)

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button.

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	
	Application Number	
Title of Invention	System and Method for Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices	

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

☐ This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.
NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.

Authorization to Permit Access:

☐ Authorization to Permit Access to the Instant Application by the Participating Offices

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date o f filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	
	Application Number	
Title of Invention	System and Method for Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices	

Applicant 1

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Clear

◯ Assignee	◯ Legal Representative under 35 U.S.C. 117	◯ Joint Inventor
◯ Person to whom the inventor is obligated to assign.	◯ Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :

If the Applicant is an Organization check here. ☐

Prefix	**Given Name**	Middle Name	**Family Name**	Suffix

Mailing Address Information For Applicant:

Address 1			
Address 2			
City		**State/Province**	
Country		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Applicant Data may be generated within this form by selecting the Add button.

Non-Applicant Assignee Information:

Providing assignment information in this section does not subsitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Assignee 1

Complete this section only if non-applicant assignee information is desired to be included on the patent application publication in accordance with 37 CFR 1.215(b). Do not include in this section an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest), as the patent application publication will include the name of the applicant(s).

PTO/AIA/14 (03-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	
	Application Number	
Title of Invention	System and Method for Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices	

If the Assignee is an Organization check here. ☐

Prefix	**Given Name**	Middle Name	**Family Name**	Suffix

Mailing Address Information For Non-Applicant Assignee:

Address 1	
Address 2	

City		**State/Province**	
Country		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee Data may be generated within this form by selecting the Add button.

Signature:

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Signature	/Robert J. Irvine III/			Date (YYYY-MM-DD)	2013-08-19
First Name	Robert J.	Last Name	Irvine III	Registration Number	41865

Additional Signature may be generated within this form by selecting the Add button.

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent C o o p eration Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt	
EFS ID:	16643896
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 (13-0401)
Receipt Date:	22-AUG-2013
Filing Date:	17-APR-2013
Time Stamp:	10:35:18
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		SNS-71301US10-SuppPrelimAmend-title.pdf	88644 ab1d4f56854b93d415be4f9d80b6604ef2413748	yes	3

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Preliminary Amendment	1	1
	Specification	2	2
	Applicant Arguments/Remarks Made in an Amendment	3	3

Warnings:					
Information:					
2	Application Data Sheet	SNS-71301US10-SupplementalADS.pdf	161471 c8068522aa981c0990768f63c7534216311c5661	no	6
Warnings:					
Information:					
This is not an USPTO supplied ADS fillable form					
		Total Files Size (in bytes):	250115		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 13/864,248	Filing Date 04/17/2013	☐ To be Mailed

ENTITY: ☒ LARGE ☐ SMALL ☐ MICRO

APPLICATION AS FILED – PART I

	(Column 1)	(Column 2)		
FOR	NUMBER FILED	NUMBER EXTRA	RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	

APPLICATION AS AMENDED – PART II

AMENDMENT	08/22/2013	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 20	Minus	** 2	=	X $ =	
	Independent (37 CFR 1.16(h))	* 3	Minus	***3	=	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))						
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
						TOTAL ADD'L FEE	

AMENDMENT		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))						
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
						TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

LIE
/PEGGY YARBOROUGH/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**
If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
13/864,248	04/17/2013	Nicholas A. J. Millington	SNS-71301.US.10 (13-0401)

CONFIRMATION NO. 5095

109473
TechLaw LLP
200 S. Wacker Drive
Suite 3100
Chicago, IL 60606

POA ACCEPTANCE LETTER

OC000000061846500

Date Mailed: 06/17/2013

NOTICE OF ACCEPTANCE OF POWER OF ATTORNEY

This is in response to the Power of Attorney filed 06/07/2013.

The Power of Attorney in this application is accepted. Correspondence in this application will be mailed to the above address as provided by 37 CFR 1.33.

/zabraha/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

PTO/SB/80 (11-08)
Approved for use through 11/30/2011. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

POWER OF ATTORNEY TO PROSECUTE APPLICATIONS BEFORE THE USPTO

I hereby revoke all previous powers of attorney given in the application identified in the attached statement under 37 CFR 3.73(b).

I hereby appoint:

☑ Practitioners associated with the Customer Number: 109473

OR

☐ Practitioner(s) named below (if more than ten patent practitioners are to be named, then a customer number must be used):

Name	Registration Number	Name	Registration Number

as attorney(s) or agent(s) to represent the undersigned before the United States Patent and Trademark Office (USPTO) in connection with any and all patent applications assigned only to the undersigned according to the USPTO assignment records or assignment documents attached to this form in accordance with 37 CFR 3.73(b).

Please change the correspondence address for the application identified in the attached statement under 37 CFR 3.73(b) to:

☑ The address associated with Customer Number: 109473

OR

☐ Firm or Individual Name			
Address			
City		State	Zip
Country			
Telephone		Email	

Assignee Name and Address:
Sonos, Inc.
223 E. De La Guerra Street
Santa Barbara, CA 93101

A copy of this form, together with a statement under 37 CFR 3.73(b) (Form PTO/SB/96 or equivalent) is required to be filed in each application in which this form is used. The statement under 37 CFR 3.73(b) may be completed by one of the practitioners appointed in this form if the appointed practitioner is authorized to act on behalf of the assignee, and must identify the application in which this Power of Attorney is to be filed.

SIGNATURE of Assignee of Record

The individual whose signature and title is supplied below is authorized to act on behalf of the assignee

Signature	/Mark W. Triplett/	Date	September 14, 2012
Name	Mark W. Triplett	Telephone	
Title	Director, Intellectual Property		

This collection of information is required by 37 CFR 1.31, 1.32 and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/96 (08-12)
Approved for use through 01/31/2013. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

STATEMENT UNDER 37 CFR 3.73(c)

Applicant/Patent Owner: SONOS, INC.

Application No./Patent No.: 13/864,248 Filed/Issue Date: April 17, 2013

Titled: System and Method for Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices

SONOS, INC., a Corporation

(Name of Assignee) (Type of Assignee, e.g., corporation, partnership, university, government agency, etc.)

states that, for the patent application/patent identified above, it is (choose **one** of options 1, 2, 3 or 4 below):

1. ☑ The assignee of the entire right, title, and interest.

2. ☐ An assignee of less than the entire right, title, and interest (check applicable box):

   ☐ The extent (by percentage) of its ownership interest is _____%. Additional Statement(s) by the owners holding the balance of the interest must be submitted to account for 100% of the ownership interest.

   ☐ There are unspecified percentages of ownership. The other parties, including inventors, who together own the entire right, title and interest are:

   Additional Statement(s) by the owner(s) holding the balance of the interest must be submitted to account for the entire right, title, and interest.

3. ☐ The assignee of an undivided interest in the entirety (a complete assignment from one of the joint inventors was made). The other parties, including inventors, who together own the entire right, title, and interest are:

   Additional Statement(s) by the owner(s) holding the balance of the interest must be submitted to account for the entire right, title, and interest.

4. ☐ The recipient, via a court proceeding or the like (*e.g.*, bankruptcy, probate), of an undivided interest in the entirety (a complete transfer of ownership interest was made). The certified document(s) showing the transfer is attached.

The interest identified in option 1, 2 or 3 above (not option 4) is evidenced by either (choose **one** of options A or B below):

A. ☐ An assignment from the inventor(s) of the patent application/patent identified above. The assignment was recorded in the United States Patent and Trademark Office at Reel _____, Frame _____, or for which a copy thereof is attached.

B. ☑ A chain of title from the inventor(s), of the patent application/patent identified above, to the current assignee as follows:

1. From: Nicholas A.J. Millington To: Rincon Networks, Inc.

   The document was recorded in the United States Patent and Trademark Office at
   Reel 030384, Frame 0675, or for which a copy thereof is attached.

2. From: Rincon Networks, Inc. To: Sonos, Inc.

   The document was recorded in the United States Patent and Trademark Office at
   Reel 030518, Frame 0778, or for which a copy thereof is attached.

This collection of information is required by 37 CFR 3.73(b). The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450**.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/96 (08-12)
Approved for use through 01/31/2013. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

STATEMENT UNDER 37 CFR 3.73(c)

3. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

4. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

5. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

6. From: _____ To: _____

The document was recorded in the United States Patent and Trademark Office at

Reel _____, Frame _____, or for which a copy thereof is attached.

☐ Additional documents in the chain of title are listed on a supplemental sheet(s).

☑ As required by 37 CFR 3.73(c)(1)(i), the documentary evidence of the chain of title from the original owner to the assignee was, or concurrently is being, submitted for recordation pursuant to 37 CFR 3.11.

[NOTE: A separate copy (i.e., a true copy of the original assignment document(s)) must be submitted to Assignment Division in accordance with 37 CFR Part 3, to record the assignment in the records of the USPTO. See MPEP 302.08]

The undersigned (whose title is supplied below) is authorized to act on behalf of the assignee.

/Robert J. Irvine III/	June 4, 2013
Signature	Date
Robert J. Irvine III	41865
Printed or Typed Name	Title or Registration Number

[Page 2 of 2]

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	15956676
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 (13-0401)
Receipt Date:	07-JUN-2013
Filing Date:	17-APR-2013
Time Stamp:	10:24:07
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Power of Attorney	SONOS-POA.pdf	113691 bc339b23e44ae2271b93e7bdb8050526b8555675	no	1

Warnings:	
Information:	

2	Assignee showing of ownership per 37 CFR 3.73.	SNS-71301US10-373cStmt.pdf	79265 d183352dbfb2f377c42f399e7ce5698d208d68cb	no	3

Warnings:	
Information:	
Total Files Size (in bytes):	192956

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number: 13/864,248

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY FEE($)
BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A			N/A	280
SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A			N/A	600
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A			N/A	720
TOTAL CLAIMS (37 CFR 1.16(i))	20 minus 20 =	*			OR	x 80 =	0.00
INDEPENDENT CLAIMS (37 CFR 1.16(h))	3 minus 3 =	*				x 420 =	0.00
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).						0.00
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))							0.00
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL			TOTAL	1600

APPLICATION AS AMENDED - PART II

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT A	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE($)	ADDITIONAL FEE($)	OR	RATE($)	ADDITIONAL FEE($)
AMENDMENT B	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest found in the appropriate box in column 1.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
13/864,248	04/17/2013	Nicholas A. J. Millington	SNS-71301.US.10

CONFIRMATION NO. 5095

IMPROPER CPOA LETTER

109473
TechLaw LLP
200 S. Wacker Drive
Suite 3100
Chicago, IL 60606

OC000000061405402

Date Mailed: 05/30/2013

NOTICE REGARDING POWER OF ATTORNEY

This is in response to the power of attorney filed 04/17/2013. The power of attorney in this application is not accepted for the reason(s) listed below:

- The power of attorney from the inventors has not been accepted because the power of attorney must be signed by the applicant for patent. See 37 CFR 1.32(b)(4).

/ebekele/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY.DOCKET.NO	TOT CLAIMS	IND CLAIMS
13/864,248	04/17/2013	2484	1600	SNS-71301.US.10	20	3

CONFIRMATION NO. 5095

109473
TechLaw LLP
200 S. Wacker Drive
Suite 3100
Chicago, IL 60606

FILING RECEIPT



Date Mailed: 05/30/2013

Receipt is acknowledged of this non-provisional patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection. Please verify the accuracy of the data presented on this receipt. **If an error is noted on this Filing Receipt, please submit a written request for a Filing Receipt Correction. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections**

Inventor(s)

Nicholas A. J. Millington, Santa Barbara, CA;

Applicant(s)

SONOS, INC., Santa Barbara, CA

Assignment For Published Patent Application

SONOS, INC., Santa Barbara, CA

Power of Attorney: None

Domestic Priority data as claimed by applicant

This application is a CON of 13/297,000 11/15/2011
which is a CON of 10/816,217 04/01/2004 PAT 8234395
which claims benefit of 60/490,768 07/28/2003

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see http://www.uspto.gov for more information.) - None.
Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

If Required, Foreign Filing License Granted: 05/21/2013

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 13/864,248**

Projected Publication Date: 09/05/2013

Non-Publication Request: No

Early Publication Request: No

Title

SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES

Preliminary Class

386

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at http://www.uspto.gov/web/offices/pac/doc/general/index.html.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, http://www.stopfakes.gov. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4158).

LICENSE FOR FOREIGN FILING UNDER

Title 35, United States Code, Section 184

Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign AssetsControl, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop technology, manufacture products, deliver services, and grow your business, visit http://www.SelectUSA.gov or call +1-202-482-6800.

PTO/AIA/01 (06-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	System and Method for Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices

As the below named inventor, I hereby declare that:

This declaration is directed to:

- [x] The attached application, or
- [ ] United States application or PCT international application number _____

  filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Nicholas A.J. Millington Date (Optional): 2/25/13

Signature: Nicholas Millington

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	15769707
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 (13-0401)
Receipt Date:	15-MAY-2013
Filing Date:	
Time Stamp:	10:53:16
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Miscellaneous Incoming Letter	SNS-71301US10-DecComm.pdf	65110 072f36a4ea9db44dabef88ed7d5a3410809653cc	no	1
Warnings:					
Information:					

2	Oath or Declaration filed	SNS-71301US10-SignedDeclaration.pdf	227817 892f966644df68fe75e16a77d9e40e451cc7816d	no	2

Warnings:	
Information:	
Total Files Size (in bytes):	292927

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
(Attorney Docket No. SNS-71301.US.10)

In re Application of:	)	
	)	
Nicholas A.J. Millington	)	Examiner: TBD
	)	
Serial No.: 13/864,248	)	
	)	Art Unit: 2484
Filing Date: April 17, 2013	)	
	)	
For: System and Method For Synchronizing Operations Among a Plurality of Independently Clocked Digital Data Processing Devices	)	Confirmation No.: 5095

COMMUNICATION

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Applicant submits the attached executed Declaration and Power of Attorney for the above-referenced U.S. Patent Application.

If the Examiner has any questions, please contact the undersigned at (312) 834-4334.

Respectfully submitted,
TechLaw LLP

Date: May 14, 2013

By: /Robert J. Irvine III/
Robert J. Irvine III
Reg. 41,865

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (**Not for submission under 37 CFR 1.99)**		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

U.S.PATENTS						Remove
Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	4296278		1981-10-20	CULLISON; Steven B. et al.	
	2	4816989		1989-03-28	FINN; Alan M. et al.	
	3	5182552		1993-01-26	PAYNTING; Richard J.	
	4	5406634		1995-04-11	ANDERSON; Charles W. et al.	
	5	5440644		1995-08-08	FARINELLI; Robert et al.	
	6	5467342		1995-11-14	LOGSTON; Gary L. et al.	
	7	5491839		1996-02-13	SCHOTZ; Larry	
	8	5553222		1996-09-03	MILNE; Steve H. et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	9	5668884		1997-09-16	CLAIR; Jr. Roy Barry et al.	
	10	5673323		1997-09-30	SCHOTZ; Larry et al.	
	11	5815689		1998-09-29	SHAW; George H. J. et al.	
	12	5867691		1999-02-02	SHIRAISHI; Hajime	
	13	5875354		1999-02-23	CHARLTON; Paul et al.	
	14	5887143		1999-03-23	SAITO; Masahiko et al.	
	15	5946343		1999-08-31	SCHOTZ; Larry et al.	
	16	6009457		1999-12-28	MOLLER; Matthew Donaldson	
	17	6026150		2000-02-15	FRANK; Edward H. et al.	
	18	6031818		2000-02-29	LO; W. Steven et al.	
	19	6128318		2000-10-03	SATO; Takashi	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	20	6157957		2000-12-05	BERTHAUD; Jean-Marc	
	21	6175872		2001-01-16	NEUMANN; Eric K. et al.	
	22	6185737		2001-02-06	NORTHCUTT; J. Duane et al.	
	23	6195436		2001-02-27	SCIBORA; Marco et al.	
	24	6199169		2001-03-06	VOTH; Duane J.	
	25	6256554		2001-07-03	DILORENZO; Mark	
	26	6308207		2001-10-23	TSENG; Weinan William et al.	
	27	6324586		2001-11-27	JOHNSON; Aric R.	
	28	6332147		2001-12-18	MORAN; Thomas P. et al.	
	29	6351821		2002-02-26	VOTH; Duane J.	
	30	6430353		2002-08-06	HONDA; Susumu et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	31	6487296		2002-11-26	ALLEN; Steven W. et al.	
	32	6526325		2003-02-25	SUSSMAN; Robert et al.	
	33	6587127		2003-07-01	LEEKE; Steven D. et al.	
	34	6598172		2003-07-22	VANDEUSEN; Mark P. et al.	
	35	6611537		2003-08-26	EDENS; Glenn T. et al.	
	36	6631410		2003-10-07	KOWALSKI; John Michael et al.	
	37	6674803		2004-01-06	KESSELRING; William David	
	38	6757517		2004-06-29	CHANG; Chin-Chi	
	39	6778869		2004-08-17	CHAMPION; Mark	
	40	6826283		2004-11-30	WHEELER; Thomas M. et al.	
	41	6836788		2004-12-28	KIM; Sang-Gil et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	42	6898642		2005-05-24	CHAFLE; Girish Bhimrao et al.	
	43	6912610		2005-06-28	SPENCER; Thomas Vincent	
	44	6920373		2005-07-19	XI; Ning et al.	
	45	6934766		2005-08-23	RUSSELL; Paul J.	
	46	6985694		2006-01-10	DE; Bonet Jeremy S. et al.	
	47	7007106		2006-02-28	FLOOD; Mark et al.	
	48	7020791		2006-03-28	AWEYA; James et al.	
	49	7043651		2006-05-09	AWEYA; James et al.	
	50	7047308		2006-05-16	DESHPANDE; Sachin G.	
	51	7115017		2006-10-03	LAURSEN; Thomas et al.	
	52	7130608		2006-10-31	HOLLSTROM; Magnus et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	53	7162315		2007-01-09	GILBERT; Erik J.	
	54	7185090		2007-02-27	KOWALSKI; John Michael et al.	
	55	7187947		2007-03-06	WHITE; Russell W. et al.	
	56	7206367		2007-04-17	MOORE; Daniel L.	
	57	7209795		2007-04-24	SULLIVAN; Gary E. et al.	
	58	7312785		2007-12-25	TSUK; Robert W. et al.	
	59	7324857		2008-01-29	GODDARD; Mark D.	
	60	7333519		2008-02-19	SULLIVAN; Gary E. et al.	
	61	7372846		2008-05-13	ZWACK; Eduard	
	62	7392102		2008-06-24	SULLIVAN; Gary E. et al.	
	63	7643894		2010-01-05	BRAITHWAITE; Michael et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	64	7675943		2010-03-09	MOSIG; Rudiger et al.	
	65	7720096		2010-05-18	KLEMETS; Anders E.	
	66	7742740		2010-06-22	GOLDBERG; David et al.	
	67	7934239		2011-04-26	DAGMAN; Vadim	
	68	8014423		2011-09-06	THALER; Thomas et al.	
	69	8045952		2011-10-25	QURESHEY; Safi et al.	
	70	8050652		2011-11-01	QURESHEY; Safi et al.	
	71	8074253		2011-12-06	NATHAN; Guy	
	72	8086752		2011-12-27	MILLINGTON; Nicholas A.J. et al.	
	73	8103009		2012-01-24	MCCARTY; William A. et al.	
	74	8131390		2012-03-06	BRAITHWAITE; Michael et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	75	8234395		2012-07-31	MILLINGTON; Nicholas A.J.	
	76	8370678		2013-02-05	MILLINGTON; Nicholas A. J. et al.	
	77	8423659		2013-04-16	MILLINGTON; Nicholas A. J.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20010009604		2001-07-26	ANDO; Hideo et al.	
	2	20010022823		2001-09-20	RENAUD; Pierre	
	3	20010032188		2001-10-18	MIYABE; Yoshiyuki et al.	
	4	20020002039		2002-01-03	QURESHEY; Safi et al.	
	5	20020002562		2002-01-03	MORAN; Thomas P. et al.	
	6	20020034374		2002-03-21	BARTON; James M.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	7	20020042844		2002-04-11	CHIAZZESE; Giovanni	
	8	20020065926		2002-05-30	HACKNEY; Michael L.J. et al.	
	9	20020073228		2002-06-13	COGNET; Yves et al.	
	10	20020090914		2002-07-11	KANG; Hyun-Sook et al.	
	11	20020093478		2002-07-18	YEH; Chun Lin	
	12	20020112244		2002-08-15	LIOU; Shih-Ping et al.	
	13	20020129156		2002-09-12	YOSHIKAWA; Masato	
	14	20020143998		2002-10-03	RAJAGOPAL; Priya et al.	
	15	20020163361		2002-11-07	PARKIN; Michael W.	
	16	20020165921		2002-11-07	SAPIEYEVSKI; Jerzy	
	17	20030023741		2003-01-30	TOMASSETTI; Stephen Robert et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	18	20030035444		2003-02-20	ZWACK; Eduard	
	19	20030041173		2003-02-27	HOYLE; Stephen L.	
	20	20030041174		2003-02-27	WEN; Sayling et al.	
	21	20030066094		2003-04-03	VAN; Der Schaar Auke Sjoerd et al.	
	22	20030099212		2003-05-29	ANJUM; Farooq et al.	
	23	20030099221		2003-05-29	RHEE; Sokwoo	
	24	20030126211		2003-07-03	ANTTILA; Akseli et al.	
	25	20030195964		2003-10-16	MANE; Pravin D.	
	26	20030198257		2003-10-23	SULLIVAN; Gary E. et al.	
	27	20030210796		2003-11-13	MCCARTY; William A. et al.	
	28	20030231871		2003-12-18	USHIMARU; Chikara	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	29	20040001484		2004-01-01	OZGUNER; Tolga	
	30	20040001591		2004-01-01	MANI; Murali et al.	
	31	20040024925		2004-02-05	CYPHER; Robert E. et al.	
	32	20040027166		2004-02-12	MANGUM; Wayne G. et al.	
	33	20040170383		2004-09-02	MAZUR; Leszek	
	34	20040203378		2004-10-14	POWERS; Simon	
	35	20040249965		2004-12-09	HUGGINS; Guy Dwayne et al.	
	36	20040249982		2004-12-09	ARNOLD; Johann et al.	
	37	20040252400		2004-12-16	BLANK; William Thomas et al.	
	38	20050010691		2005-01-13	OYADOMARI; Randy et al.	
	39	20050013394		2005-01-20	RAUSCH; Mathias et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	40	20050058149		2005-03-17	HOWE; Wayne Richard	
	41	20050081213		2005-04-14	SUZUOKI; Masakazu et al.	
	42	20050114538		2005-05-26	ROSE; Steven W.	
	43	20050177643		2005-08-11	XU; Shugong	
	44	20050281255		2005-12-22	DAVIES; Trevor et al.	
	45	20050283820		2005-12-22	RICHARDS; Martin J. et al.	
	46	20050288805		2005-12-29	MOORE; Jeffrey C. et al.	
	47	20070038999		2007-02-15	MILLINGTON; Nicholas A.J. et al.	
	48	20070054680		2007-03-08	MO; Shaomin S. et al.	
	49	20070142944		2007-06-21	GOLDBERG; David et al.	
	50	20070271388		2007-11-22	BOWRA; Todd et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	51	20080120429		2008-05-22	MILLINGTON; Nicholas A.J. et al.	
	52	20090157905		2009-06-18	DAVIS; Bradford C.	
	53	20090228919		2009-09-10	ZOTT; Joseph A. et al.	
	54	20120029671		2012-02-02	MILLINGTON; Nicholas A. J. et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	0251584	EP	A2	1988-01-07	NORTHERN TELECOM LTD [CA]		☐
	2	0672985	EP	A1	1995-09-20	IBM [US]		☐
	3	1111527	EP	A2	2001-06-27	TOSHIBA KK [JP]		☐
	4	9525313	WO	A1	1995-09-21	IBM [US], et al.		☐
	5	9961985	WO	A1	1999-12-02	DOW CHEMICAL CO [US]		☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

NON-PATENT LITERATURE DOCUMENTS			Remove
Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T5
	1	AKYILDIZ I.F., et al., "Multimedia Group Synchronization Protocols for Integrated Services Networks," IEEE Journal on Selected Areas in Communications, 1996, Vol. 14 (1), pp. 162-173.	☐
	2	BENSLIMANE A., "A Multimedia Synchronization Protocol for Multicast Groups," Proceedings of the 26th Euromicro Conference, 2000, Vol. 1, pp. 456-463.	☐
	3	BIERSACK E., et al., "Intra- and Inter-Stream Synchronisation for Stored Multimedia Streams," IEEE International Conference on Multimedia Computing and Systems, 1996, pp. 372-381.	☐
	4	BRETL W.E., et al., MPEG2 Tutorial [online], 2000 [retrieved on 2009-01-13] Retrieved from the Internet:< http://www.bretl.com/mpeghtml/MPEGindex.htm>, pp. 1-23.	☐
	5	Final Office Action mailed 13 July 2009 for U.S. Application No. 10/816217 filed 01 April 2004.	☐
	6	Final Office Action mailed 21 October 2011 for U.S. Application No. 10/816217 filed 01 April 2004.	☐
	7	Final Office Action mailed 28 January 2011 for U.S. Application No. 10/816217 filed 01 April 2004.	☐
	8	Final Office Action mailed 30 June 2008 for U.S. Application No. 10/816217 filed 01 April 2004.	☐
	9	HUANG C.M., et al., "A Synchronization Infrastructure for Multicast Multimedia at the Presentation Layer," IEEE Transactions on Consumer Electronics, 1997, Vol. 43 (3), pp. 370-380.	☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	10	International Search Report for Application No. PCT/US04/23102, mailed on 01 August 2008, 5 pages.	☐
	11	ISHIBASHI Y., "A Group Synchronization Mechanism for Live Media in Multicast Communications," IEEE Global Telecommunications Conference, 1997, Vol. 2, pp. 746-752.	☐
	12	ISHIBASHI Y., "A Group Synchronization Mechanism for Stored Media in Multicast Communications," IEEE Information Revolution and Communications, 1997, Vol. 2, pp. 692-700.	☐
	13	JO J., et al., "Synchronized One-to-many Media Streaming with Adaptive Playout Control," Proceedings of SPIE, 2002, Vol. 4861, pp. 71-82.	☐
	14	MILLS D.L., "Network Time Protocol (Version 3) Specification, Implementation and Analysis," Network Working Group, March 1992, <http://www.ietf.org/rfc/rfc1305.txt>.	☐
	15	MILLS D.L., "Precision Synchronization of Computer Network Clocks," ACM SIGCOMM Computer Communication Review, 1994, Vol. 24 (2), pp. 28-43.	☐
	16	MOTOROLA., "Simplefi, Wireless Digital Audio Receiver, Installation and User Guide", 2001-12-31.	☐
	17	Non-Final Office Action mailed 18 January 2008 for U.S. Application No. 10/816217 filed 01 April 2004.	☐
	18	Non-Final Office Action mailed 21 June 2011 for U.S. Application No. 10/816217 filed 01 April 2004.	☐
	19	Non-Final Office Action mailed 22 January 2009 for U.S. Application No. 10/816217 filed 01 April 2004.	☐
	20	Non-Final Office Action mailed 25 June 2010 for U.S. Application No. 10/816217 filed 01 April 2004.	☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	21	Non-Final Office Action mailed 29 February 2012 for U.S. Application No. 13/297000 filed 15 November 2011.	☐
	22	Notice of Allowance mailed 27 December 2011 for U.S. Application No. 10/816217 filed 01 April 2004.	☐
	23	PARK S., et al., "Group Synchronization in MultiCast Media Communications," Proceedings of the 5th Research on Multicast Technology Workshop, 2003.	☐
	24	PRISMIQ; Inc., "PRISMIQ Media Player User Guide", 2003, 44 pages.	☐
	25	ROTHERMEL K., et al., "An Adaptive Stream Synchronization Protocol," 5th International Workshop on Network and Operating System Support for Digital Audio and Video, 1995.	☐
	26	SCHULZRINNE H., et al., "RTP: A Transport Protocol for Real-Time Applications, RFC 3550," Network Working Group, 2003, pp. 1-89.	☐
	27	The MPEG-2 Transport Stream. Retrieved from the Internet:< URL: http://www.coolstf.com/mpeg/#ts>.	☐
	28	Final Office Action mailed 13 September 2012 for U.S. Application No. 13/297000", 2012-09-13, 17 pages.	☐
	29	"UPnP; "Universal Plug and Play Device Architecture"; 08 June 2000; version 1.0; Microsoft Corporation; pp 1-54".	☐
	30	VOYETRA; Turtle Beach Inc., "AudioTron Quick Start Guide, Version 1.0", March, 2001, 24 pages.	☐
	31	VOYETRA; Turtle Beach Inc., "AudioTron Reference Manual, Version 3.0", May 2002, 70 pages.	☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

	32	VOYETRA; Turtle Beach Inc., "AudioTron Setup Guide, Version 3.0", May 2002, 38 pages.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	13864248
	Filing Date	2013-04-17
	First Named Inventor	Nicholas A. J. Millington
	Art Unit	2484
	Examiner Name	TBD
	Attorney Docket Number	SNS-71301.US.10 (13-0401)

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Robert J. Irvine III/	Date (YYYY-MM-DD)	2013-05-08
Name/Print	Robert J. Irvine III	Registration Number	41865

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Publication number: **0 251 584**
A2

(12) **EUROPEAN PATENT APPLICATION**

(21) Application number: 87305416.7

(51) Int. Cl.⁴: G06F 9/46

(22) Date of filing: 18.06.87

(30) Priority: 26.06.86 CA 512479

(43) Date of publication of application:
07.01.88 Bulletin 88/01

(84) Designated Contracting States:
AT DE FR GB IT NL SE

(71) Applicant: NORTHERN TELECOM LIMITED
600 de la Gauchetiere Street West
Montreal Quebec H3B 4N7(CA)

(72) Inventor: Goyer, Pierre
32 Terasse David
Gatineau Quebec J8V 1G4(CA)
Inventor: Selic, Branislav Vladeta
7 Whelan Drive
Nepean Ontario K2J 2A3(CA)

(74) Representative: Crawford, Andrew Birkby et al
A.A. THORNTON & CO. Northumberland
House 303-306 High Holborn
London WC1V 7LE(GB)

(54) **A synchronization service for a distributed operating system or the like.**

(57) A synchronization service which can be incorporated into a distributed operating system as a shared service. It allows the realization of different custom-built synchronization strategies for different applications. This approach is based on defining a general set of application-independent synchronization primitives. These are provided by the distributed operating system in the form of a synchronization service. By themselves the individual primities are insuffient to provide synchronization. However, they can be combined in different ways to realize customized synchronization strategies. This leaves the ultimate responsibility for synchronization with the application, but in a much simplified form. Application programs can combine these primitives to construct the most suitable form of synchronization.

EP 0 251 584 A2

A SYNCHRONIZATION SERVICE FOR A DISTRIBUTED OPERATING SYSTEM OR THE LIKE

This invention relates generally to the field of computers, and more specifically to a synchronization service for use with computers.

Background of the Invention

As computing tasks increase in size and complexity, one approach to speed up the execution of these tasks is to use distributed programs. A distributed program can be defined as a computer program which is partitioned into multiple concurrent components which execute on separate processing sites which do not share a common memory.

In this context, the term "program" is used to imply a global objective (i.e. common goal). Each component (or portion) of the program performs some portion of the overall activity required to attain this common goal. Thus, a distributed program represents a set of (functionally) tightly-coupled components operating in a (physically) loosely-coupled environment.

Asynchronous operation of concurrent cooperating activities results in the time-dependencies and race conditions which can lead to errors. For example, two processes attempting to simultaneously update a shared variable may interfere with each other so that an incorect value is assigned to the variable. The solution to such problems is through synchronization. Synchronization can be defined as the organization of actions and interactions of a system of concurrent asynchronous entities for the purpose of achieving some common objective.

One example of a distributed computer system is the case of a replicated database where each copy of the database is on a separate processing element. When a change is made to one copy then this change must be propagated to all the others if consistency is to be maintained. This involves synchronization. The situation may be complicated further if two or more conflicting changes are initiated simultaneously on different copies. In that case synchronization is required not only to ensure that all copies end up in the same state but also that the resulting state is valid. Other situations where synchronization is necessary include the restoration of the current state to new or recovering copies and the handling of failures.

Distributed synchronization can also be useful in standby schemes where redundant components are configured for greater system availability. In this case the compnoents have to agree as to which will be the active and which the standby components and must also arrange for proper switchover in case of failures.

The cited examples illustrate the diverse ways in which synchronization is used in distributed systems. As can be expected, different applications can have different demands on synchronization: some may require fast response while others may place more emphasis on reliability and fault tolerance. This indicates that the choice of the most suitable synchronization technique and tis implementation can only be made if the particular needs of the application are considered.

Unfortunately, in a large system supporting many different types of distributed application programs, leaving synchronization entirely to the application program could result in excessive duplication of effort, unreliable design, and suboptimal utilization of resources. Even worse, perhaps, is the possibility that the relatively complex issue of synchronization could dominate the design to such an extent that functional concerns are neglected. From that point of view a trusted, system-based, synchronization facility is preferred.

There are several important characteristics of distributed programs which make them significantly more difficult to design and implement compared to conventional non-distributed programs:

(1) Concurrent execution. This means that there is no single sequential control thread such as represented by the execution trace of a non-distributed program. Concurrency introduces timing dependencies among the system components which can lead to deadlocks or instability.

(2) Significant communication delays. The exchange of information between components of a distributed program involves non-negligible and randomly-distributed transmission delays. If these delays are comparable to the rate at which the components change state, the system may become unstable.

(3) Partial failure modes. Failures of distributed components require complex detection and recovery algorithms which are difficult to design and verify. Two types of partial failures exist:

- Communication path failures can result in the duplication, temporal reordering, or total loss of information being exchanged; and
-Processing component failures (hardware and software) lead to temporary loss of functionality.

The recovering action for each type of failure is quite different. Unfortunately, It is often difficult to distinguish them on the basis of the observed symptoms.

From the definition of synchronization it can be seen that the need for synchronization is determined by the shared objective of the cooperating distributed entities. This common objective places interdependencies on the individual entities so that a change in the state of one necessitates appropriate changes (reactions) in others. This can be expressed as a requirement to preserve certain application-dependent state consistency constraints. The problem of maintaining consistency is further complicated by the fact that each entity, in addition to internal interactions, is also exposed to independent interactions with the environment. (The environment consists of other distributed components which do not share the same objective as the synchronized system, but which use it for their own purposes). This means that the stimulus to change state can occur simultaneously in two or more synchronized entities. The synchronization problem can then be viewed as one of ordering concurrent interdepenent activities.

The simplest form of ordering which guarantees consistency is serialization: the execution of activities one at a time. Although synchronization strategies exist which are not based on serialization, they will not be considered here due to their relative complexity.

Two basic, and not necessarily exclusive, classes of strategies exist for achieving serialization in distributed systems;

(1) Centralized strategies.

In this case, the ordering of activities is performed by a unique distinguished entity. Synchronized entities, with externally induced work requests, first approach the distinguished entity for permission. This entity resolves concurrent requests by granting a right to only one of the competing entities. When that entity completes its work, the right is granted to another entity, and so on.

A major feature of this type of scheme is that there is a single point of control. This allows the implementation of relatively complex yet reliable and efficient scheduling algorithms. Examples of centralized strategies can be found in A Decentralized Control Method in a Distributed System by J.P. Cabanel et al, Proceedings 1st Ceonference, Distributed Proc. Systems, Huntsville, Al, 1979 and in A Failure Tolerant Centralized Mutual Exclusion Algorithm by G. N. Buckley et al, Proceedings 4th Conference, Distributed Computer Systems, San Francisco, Ca. 1984.

(2) Distributed strategies.

In this case, there is no central scheduler. Instead, ordering is accomplished through distributed agreement. Key to this scheme is a shared "clock" (logical or physical). This is generally a monotonically increasing numeric variable which is maintained consistently by all the synchronized entities. Work requests are timestamped with the clock value at the time of arrival and then processed in order. However, because two or more requests can be concurrent (i.e., they have the same timestamps), ties are resolved through group negotiation: a new work request is first broadcast to all other entities which respond either with a simple acknowledgement or a work request of their own. Once an entity is aware of all concurrent work requests within the group, it orders them according to some tie-breaking rule and then processes them. Since each entity uses the same ordering algorithm each will perceive the same sequence of events as all the others.

The distinguishing feature of distributed strategies is that operation does not depend on a single critical entity at any time. This makes them very fault-tolerant. However, they are generally less efficient than centralized strategies when the number of entities to be synchronized is large. Examples of distributed strategies can be found in Time, Clocks, and the Ordering of Events in a Distributed System by L. Lamport, Comm. ACM, (21,7), July 1978, in An Algorithm for Maintaining the Consistency of Multiple Copies by D. Herman et al, Proceedings, 1st Conference Distributed Proc. Systems, Huntsville, Al., 1979 and in Synchronization in Distributed Programs by F.B. Schneider, ACM Transactions on Prog. Lang. & Syst, (4,2), April, 1982.

Combinations of these two forms, such as the circulating sequencer proposed in Algorithms for Distributed Data Sharing Systems Which Use Tickets by G. Le Lann, Proc. 3rd Berkeley Workshop on Dist. Data, Aug. 1978, are possible. In that scheme, a centralized controller is used to control the clock used for timestamping. (Although the controller function is circulated among the distributed entities, at any given time it is performed by only one entity.) The ordering of activities is then done in a distributed fashion, based on timestamp values and a tie-breaking rule.

The following patents depict examples of distributed processing in general, and attention is directed to them: U.S. patent 3,411,139 dated November 12, 1968 by J.T. Lynch et al; U.S. patent 3,631,405 dated December 28, 1971 by G.S. Hoff et al; U.S. patent 3,771,137 dated November 6, 1973 by R.P. Barner et al; and U.S. patent 4,115,866 dated September 19, 1978 by J.L.G. Janssens et al.

Summary of the Invention

One objective of the Synchronization Service of the present invention is to provide a set of application-independent capabilities which would allow the construction of specific synchronization strategies belonging to the categories listed above. To do this it must incorporate the essential abstract features of those strategies. These are defined in the form of a general synchronization paradigm described in a following section.

Because of concurrent execution and the possibility of partial failures, it is necessary to closely synchronize the operation of the distributed components of a program. Synchronization can be defined as the ordering of actions and interactions of components in a distributed program so that the state of each component remains consistent with the common goal.

Experience with concurrent systems has shown that the synchronization problem is difficult to solve even for non-distributed situations; the number of possible component interactions is usually very large, increasing the probability of a design error.

A further difficulty is caused by the fact that no single synchronization strategy is adequate for all distributed programs. If multiple distributed programs are to be supported on a system, this means that the synchronization problem may have to be solved in many different ways.

Given the diversity of synchronization strategies and the difficulty of implementing them, is it possible to provide some assistance to designers of distributed programs to increase the reliability of their designs?

The approach to this problem, presented by the present invention, consists of providing a set of primitive synchronization operators at the level of a distributed operating system. Such operators can be used to construct more complex forms of synchronization customized to different applications. This approach has the following advantages:

-It provides a one-time trusted implementation of common mechanisms;

-It does not favour any particular synchronization strategy which would favour some applications but penalize others;

-It provides a systematic framework (programming model) for designing and implementing distributed programs.

The operating system component which implements the synchronization primitives (operators) is called the Synchronization Service.

The essential idea behind the Synchronization Service is that the synchronization problem can be tackled hierarchically. Each level in the hierarchy may have different synchronization mechanisms based on the synchronization facilities of the levels below. The lower levels of this hierarchy can be designed to be application-independent and can therefore be provided as a reliable system service. This, in turn, increases the reliability of programs and reduces development time.

This approach to distributed synchronization attempts to decompose the synchronization problem. At the lowest level of decomposition a general set of application-independent synchronization primitives is defined. These are provided by the distributed operating system in the form of a synchronization service 10. By themselves the primitives are insufficient to provide synchronization. However, they can be combined in different ways to realize customized synchronization strategies. This leaves the ultimate responsibility for synchronization with the application program, but in a much simplified form. The role of the synchronization service 10 is to hide many of the more basic housekeeping functions inherent in distributed synchronization. For instance, all fault-tolerant synchronization schemes require a monitoring function to keep track of the operational status of all relevant distributed components. The present invention consolidates such a function as a system service where it can be shared by many application programs.

Stated in other terms, the present invention is a general service, provided within a distributed operating system, which can be used by application and system programs to implement synchronization between program components that are physically distributed.

Stated in other terms, the present invention is a synchronization service for use with a computer having a distributed operation system, to allow the construction of a customized synchronization scheme, for synchronizing the constituent portions of a distributed program, the service comprising: a general set of application-independent synchronization primitives, whereby the construction of the customized synchronization scheme is achieved by the selective implementation of the application-independent synchronization primitives.

Stated in yet other terms, the present invention is a synchronization service for use with a computer having an operating system distributed over a plurality of processing elements, to allow the construction of a customized synchronization scheme, for synchronizing the constituent components of a distributed program, the service comprising the steps of:

a) joining a program component on a fist processing element to a group of existing program components on at least a second processing element so that each of the existing components is aware of the present and location of the joining components;

b) informing each member of the group of physically distributed program components when one or more components which are members of the group, depart from it;

c) selecting, as a distinguished member, one program component from a group of distributed program components such that, within the group, there is never more than one distinguished member; and

d) providing mutually exclusive rights to the group of distributed program components such that no more than one component can appropriate a given right at any time.

Stated in still other terms the present invention is a synchronization service, for use with a computer having an operating system distributed over a plurality of processing elements, to allow the construction of customized synchronization schemes for synchronizing the constituent components of a distributed program, the service including a synchronization master control comprising: master control means for activating the synchronization service; polling means for polling the processing elements associated with the components of the distributed program so as to monitor the status of the processing elements; control means for joining new members to the group, and for handling departures of members from the group; and a database means containing information representative of the current state of the synchronization service at a given point in time.

Brief Description of the Drawings

The present invention will now be described in more detail with reference to the accompany drawings, wherein like parts in each of the several figures are identified by the same reference character, and wherein:

Figure 1 depicts a simplified block diagram of the synchronization service of the present invention;

Figure 2a is similar to Figure 1 but is for one specific embodiment thereof;

Figure 2b is a variation on the embodiment of Figure 2a;

Figure 2c is similar to Figure 2b;

Figure 3a is a chart depicting the primitives and corresponding replies employed by the invention;

Figure 3b is a symbolic representation of the constituent tasks of synchronization master control of Figure 1;

Figure 3c is a symbolic representation of the constituent tasks of member agent 11 of Figure 1;

Figure 4 is a simplified functional flow diagram for a database;

Figure 5 is a simplified functional flow diagram for a database;

Figures 6 to 8, 9a, 9b, and 10 to 13 inclusive represent action sequences helpful for understanding the operation of the present invention; and

Figure 14 is a simplified representation of the useage dependencies helpful in understanding the operation of the present invention.

Detailed Description

Synchronization service 10 is based on a general distributed program paradigm. This paradigm is represented by the concept of synchronization groups. A synchronization group is a set of distributed program components called "members", and referred to by the reference character 18, which cooperate to achieve a common objective. Note that members 18 are not a part of synchronization service 10, but they use synchronization service 10.

In other words, the distributed operating system 15, to which synchronization service 10 is applied, will support both distributed application programs and distributed system programs. Both the distributed application and system programs consist of several program components (called members 18) which in turn consist of subcomponents called tasks. In synchronization service 10 there is one synchronization group for each distributed application or system program.

A primitive synchronization operator has effect only within the domain of a particular synchronization group. Synchronization groups, therefore, encapsulate units of tightly coupled distributed functionality. Of course, synchronization service 10 allows many synchronization group to coexist on a single distributed operating system 15.

The basic construct of synchronization service 10 is the synchronization group representing a set (i.e., a system) of distributed entities which are tightly coupled to each other in some way. The

state and action dependencies which bind these entities are not specified at this level so that synchronization groups are decoupled from application semantics.

Formally, a synchronization group is a set of components, called members 18, in which each group ideally has the following properties:

(1) Uniqueness: There can be any number of synchronization groups in a larger system but each synchronization group is distinguished from all others by a unique synchronization gruop identifier.

(2) Physical distribution: Each member of a synchronization group exists on a different processing element 12. (This is simply a matter of convenience: extending the concept of synchronization groups to logically distributed entities is possible). Note that there are no restrictions concerning the number of synchronization groups which may have members 18 on a particular processing element 12. This means that two or more synchronization groups can overlap in physical space.

(3) Reliable communication: Communication between any pair of members 18 is non-lossy, non-duplicating, and order-preserving. Furthermore, full connectivity is assumed; i.e., each member 18 can communicate directly with all other members 18. If the physical system does not have these properties then it is assumed that an underlying communication service exists which provides them. The intent here is to isolate communications issues from synchronization issues.

(4) Dynamic behavior: Members 18 can depart or join the synchronization group at any time and independently of each other. (The group exists as long as at least one member 18 exists.) Departures may be either application-driven or due to processing element 12 failure. This property captures the dynamic nature of real-world components.

(5) Mutual exclusion: Each synchronization group maintains a set of shared objects called rights, each of which can be either free or associated with at most one member 18. They are functionally equivalent to semaphores (reference: E.W. Dijkstra, Cooperating Sequential Processes, Technical Report EWD-123, Technological University, Eindhoven, 1965) but for a distributed environment. (However, a member 18 can hold more than one right at a time.) A departing member 18 cannot abscond with a right since any rights it holds are automatically freed. In essence, rights are a general mechanism for distinguishing between group members 18. The assignment of functional significance to rights is up to the application.

(6) Distinguished member: One and only one member 18 of every synchronization group is designated as its distinguished member. The appointment is made at random and is transferred to another member 18 if the current distinguished member 18 departs. This property is intended to serve those synchronization strategies which require a central coordinator although synchronization service 10 makes no assumptions regarding the functional significance of the distinguished member 18. (Note that the distinguished member feature is simply a special case of the mutual exclusion property but has been singled out purely for convenience.) Since the selection and preservation of a distinguished member 18 is by synchronization service 10, application programs need not implement their own election algorithms.

A synchronization group represents a unit of synchronization. The facilities of the synchronization service 10 (described later) are all limited in scope to the respective synchronization group.

The synchronization problem is often formulated as a problem of maintaining data consistency in a dynamic environment. From that point of view, the synchronization service 10 ensures consistency of the folowing information sent to members 18:

(1) current membership list;

(2) the identity of the distinguished member; and

(3) the status of all group rights.

This information is maintained consistently and correctly in the face of continual departures and arrivals of members 18.

The concept of synchronization groups does not encompass application program-level consistency; that is the responsibility of the application program. Instead, a synchronization group maintains a consistent view (on all its members 18) concerning the status of its objects: the list of active members 18, the status of rights, the distinguished member designation. These responsibilities are therefore removed from the view of the application program.

Figure 1 depicts a simplified block diagram of synchronization service 10 of the present invention. A distributed application program, structured as a synchronization group, typically has members 18 (i.e. distributed program components) which are physically distributed across two or more processing elements 12a....12n (referred to collectively as processing elements 12). In the implementation of Figure 1, the structure of synchronization service 10 matches the structure of the synchronization group by providing a local synchronization controller, i.e. member agent 11, for each group member 18. Thus, there is a separate implementation of synchronization service 10 for each application program; note, however, that there is only one synchronization master control 13 regardless of how

many implementations, and only one synchronization agent 14 per processing element 12. Each implementation is functionally independent of the others.

Note that the group of processing elements 12 together with processing element 19 form part of the distributed computing environment (i.e. distributed operating system 15) which synchronization service 10 is designed to synchronize. Not also that each processing element 12 may have a plurality of member agents 11, and that processing element 19 may be combined with one of the processing elements 12.

Member agents 11 provide the main interface to the synchronization service 10. Application program components (i.e. members 18) initiate synchronization activities by invoking the desired synchronization primitives (to be described later). This is communicated to the local member agent 11 which then interacts with other member agents 11 in order to effect the specified synchronization function. The member agent 11 also informs the members 18 of synchronization requests initiated by other members 18 as well as group events such as the failure of active members 18 and the joining of new ones.

Member agents 11 are dynamic entities which follow the dynamics of the application programs they serve. A member agent 11 is created (by the local synchronization agent 14) when an application program component (i.e. member 18) requests to be synchronized with other members 18 in a synchronization group. It is destroyed when the member 18 is unsynchronized.

To ensure coherent behaviour of synchronization service 10, control of the individual implementations of the service 10 is centralized. This is done through a three-level hierarchy with a unique master controller at the top (i.e. synchronization master control 13), an intermediate layer of controllers in the middle (i.e. synchronization agents 14a to 14n, referred to collectively as synchronization agents 14), and a layer of member agents 11 at the bottom. This hierarchy allows a decomposition of the control problem into smaller more comprehensive subproblems. Note from Figure 1 that there is one synchronization agent 14 for each processing element 12, and it controls all the member agents 11 in that processing element 12.

Figure 2a is similar to Figure 1, but depicts a specific embodiment of the synchronization service, referred to by reference character 100 as applied to distributed operating system 115. In Figure 2a there is a synchronization master control 13 on processing element 119, three processing elements 112a, 112b, and 112c, three synchronization agents 14a, 14b, and 14c, along with six members 18a to 18f along with their corresponding member agents 11a to 11f respectively. In the distributed computing example of Figure 2a, processing elements 112a, 112b, 112c and 119 are each an IBM PC-AT. Note that the members 18a to 18f inclusive are not part of synchronization service 100 while everything else shown in Figure 2a is. Members 18a to 18f inclusive use the synchronization service 100. Note also that there is another synchronization master control (not shown) on standby.

Figure 2b is similar to Figure 2a, but is further simplified and depicts only those items that constitute one implementation of synchronization service 100 (i.e. implementation 100a). That is, members 18a and 18e (Figures 2a and 2b) form one synchronized group. Members 18b, 18c, 18d, and 18f (Figure 2a) form at least one other synchronized group.

Figure 2c is a simplified application to exemplify synchronization service 100a of Figure 2b. In Figure 2c the hardware implementating synchronization service 100a is a group of IBM personal computers of the AT series, linked by an IBM LAN (local area network) 226. That is, processing element 112a is an IBM PC-AT computer 212a, processing element 112b is an IBM PC-AT computer 212b, and processing element 119 is an IBM PC-AT computer 219.

In Figure 2c, computer 212a is a telephone operator's workstation as is computer 212b. The application in Figure 2c is to maintain a telephone directory and to allow the user at both computers 212a and 212b to have access to the telephone directory, to access it to determine an individual's telephone number, and to be able to update the telephone directory as changes occur. Computer 219, in this example, handles the tasks of synchronization master control 13 and database 16 (Fig. 2b)

Returning now to the general case of Figure 1, the role of synchronization master control 13 is to provide internal synchronization between the components of the local synchronization service 10. In essence, it performs those functions where a consistent (but not necessarily correct) view of the system 15 is required. More precisely, synchronization master control 13 is responsible for:

(1) Activation of synchronization service 10.

This is done by activating the synchronization agents 14 as the processing elements 12 are restarted.

(2) Monitoring of processing elements 12.

This function involves observing (polling) the status of all processing elements 12 by communicating with local synchronization agents 14. Any changes in these states are detected by synchronization master control 13 and appropriate notifications are dispatched to the synchronization service components affected by the change.

(3) Management of synchronization groups.

synchronization master control 13 is the central arbiter for all synchronization groups in the local synchronization service 10. It is involved in handling transient conditions which occur in group operation:
-group establishment,
-joining of new members 18, and
-departures of joined members 18.
Note that synchronization master control 13 does not participate in the steady-state operation of synchronization groups and, consequently, is not normally a performance bottleneck.

Synchronization master control 13 must be highly fault-tolerant since synchronization service 10 may be used to implement standby schemes by applications. For that reason it is backed up by at least one other instance operating in standby mode. If the currently active synchronization master control 13 fails, the standby will take its place. Because this is the Synchronization Service, the selection of an active synchronization master control 13 from the set of instances must be done through an internal agreement (election). This is the only place in the entire system where the synchronization service 10 cannot be used for such a purpose. However, in this case, the problem occurs in a very specific context and can be solved in a specific way (for example, by using a bully algorithm for a distributed election as described in Elections in a Distributed Computing System by H. Garcia-Molina, IEEE Trans. on Computers, (C-31,1), Jan. 1982).

Once the active synchronization master control 13 has been selected, the standby resorts to a monitoring mode in which it periodically polls the active instance until a failure is detected.

Since a standby is used, following a failure of the synchronization master control 13, its previous state must be reconstructed on the standby, preferably without involving the application program. This can be achieved through the information kept by the synchronization agents 14. As a consequence, except for slightly extended service times due to the recovering process, application programs are unaware of synchronization master control 13 failure.

SYNC MASTER TASK

The Sync Master Task 20 is the root task (i.e. program) of the synchronization service 10 control hierarchy. It provides the central control point for all synchronization groups. It consists of four main subcomponents as depicted in Figure 3b and is located within synchronization master control 13. The four main subcomponents of the Sync Master Task 20 are as follows:

SYNC MASTER CONTROL 21 establishes and maintains the operational state of the Sync Master Task 20. This includes the Sync Master recovery algorithm. Sync Master Control 21 consists of the main procedure of the Sync Master Task 20.

POLLING CONTROL 22 is responsible for detecting failure of processing elements 12. This subcomponent sends periodic messages to all synchronization agents 14. If a reply is not received within a certain time interval (after several retries have been attempted) the corresponding processing element 12 is declared as failed and a notification is sent to all remaining synchronization agents 14. This subcomponent is implemented within the Sync Master Task 20.

SYNC AGENT CONTROL 23 deals with events which occur at the processing element 12 level. This subcomponent is responsible for activating newly-recovered synchronization agents 14 as well as for accepting notifications, from the synchronization agents 14, about the arrivals and departures of group members 18. These are then relayed to the appropriate Group Control 24. This subcomponent is also implemented within the Sync Master Task 20.

GROUP CONTROL 24 handles events which are relevant to one group. This includes the joining and departure of group members 18. The Group Control function is implemented by the Group Master Task 25. There is one such task 25 for each synchronization group. Tasks 25 are created dynamically by the Sync Master Task 20.

The tasks comprising the Sync Master Task 20 maintain a shared detabase 16 (Figure 1) which represents a snapshot of the current state of the synchronization service 10. This database is described later.

SYNCHRONIZATION AGENT

A synchronization agent 14 resides in the control program of each processing element 12 which requires synchronization service 10 and it is the sole representative of the Sync Master Task 20 in that processing element 12. The synchronization agent 14 has the following responsibilities:
-It accepts SYNCHRONIZE directives and creates corresponding member agents 11.
-It monitors the status of all active member agents 11 on its processing element 12 and detects their disappearance (spontaneous or planned).
-It notifies the synchronization master control 13 of all changes (arrivals and departures) of Member Agents 11 on its processing element 12.
The synchronization agent 14 is implemented by the Sync Agent Task which is part of the operating system 15 on the corresponding processing element 12.

The synchronization agents 14 are permanent representatives of synchronization master control 13 within their host processing element 12. They have three main purposes:

(1) Synchronization agents 14 are a focal point for controlling all member agents 11 within a single processing element 12. This reduces the load on synchronization master control 13 which simply sends common control information to synchronization agents 14 for distribution to local member agents 11.

(2) Synchronization agents 14 isolate member agents 11 from the effects of synchronization master control 13 failures. All communication between the synchronization master control 13 and Member Agents 11 is channeled through the synchronization agents 14. If the synchronization master control 13 is temporarily unavailable (due to failure), the synchronization agents 14 will hold member agent 11 messages destined for the synchronization master control 13 until the latter is reinstated. In this way failures of the synchronization master control 13 are masked from member agents 11 and hence the applications.

(3) Synchronization agents 14 participate in the recovery of the synchronization master control 13. When a synchronization master control 13 is being reinstated it can reconstruct its operational state simply by querying all the synchronization agents 14. This is much faster and more reliable than querying the member agents 11 since these are more dynamic and more numerous.

The synchronization master control 13 maintains a database 16 (Figure 1) which represents the current state of the synchronization service 10 within the system 15. The database can be accessed through two keys:
-by group identifier --for access to the data for a particular synchronization group, and
-by processing element identifier --for access to synchronization service components located on a particular processing element 12.

The basic structure used is the linked list of dynamically allocated control blocks, each block corresopnding to some synchronization service component. This represents a trade-off between the requirement to minimize storage costs and the need for fast access to the data.

The next section describes the operation of the internal mechanisms used to achieve the synchronization functions. In the following discussion the communication between member agents 11 is assumed to be reliable; i.e., it is non-lossy, non duplicating, and order preserving. If the communication medium is unreliable an underlying reliable communication service provided within the distributed operating system can be used.

Rights are a set of shared objects within each synchronization group; each right can be free or assicated with at most one member 18. One example of a right is a database lock whereby only one user at a time can write to a database and no one else can read or write at that time. See also the "Update" right referred to later.

Rights are distributed in a centralized fashion since that minimizes overhead and complexity. In principle, this can be done by any member agent 11. For convenience, the control and distribution of rights are performed by the distinguished member (one of the members 18). The distinguished member 18 already has the uniqueness and fault-tolerant properties which are also required by the controller for rights. Thus, the Member 18 selected as the distinguished member has to perform this special function in addition to its standard synchronization functions. The selection of a distinguished member is done, by the synchronization master control 13, at the time the group is established (see below).

When a member 18 requires a right, its member agent 11 directs the request to the distinguished member 18. If the right is available, the distinguished member 18 will grant the right and inform the requesting member agent 11. If the right is already appropriated, then depending on the type of request made, the request is either queued by the distinguished member 18 or it is refused. In the first case, requests are handled on a first-come first-served basis.

Should the current distinguished member 18 fail, a new one is appointed by the synchronization master control 13 (which is also resonsible for detecting the failure). Of course, until a new distinguished member 18 is appointed, rights cannot be distributed or retrieved, but all the other synchronization services are still available. In order to mini-

mize the effect of a distinguished member 18 failure, the state of rights is reconstructed to the point just prior to failure. Each member 18 keeps a list of all rights which it has appropriated as well as a list of all its outstanding rights requests. This information is then exchanged with the new distinguished member 18 which can then assume the same state as the previous distinguished member 18. The entire switchover process is transparent to the application program.

If a member 18 fails, the distinguished member will automatically release any rights held by that member 18 and also purge any queued requests generated by that member 18.

Member agent 11 is the main functional component of synchronization service 10 and is responsible for handling all directives initiated by the user. It performs four classes of functions as depicted in Figure 3c and as represented by the following:

The COMMUNICATIONS HANDLER 33 provides a reliable (order-preserving, non-lossy, non-duplucating) communications service between group members 18; in order to minimize deadlocks the communication mode used is asynchronous message passing. This function is required only if there is no reliable communications service present within the distributed operating system 15.

The GROUP STATE HANDLER 32 maintains a local version of the current state of all the other group members 18.

The DIRECTIVE HANDLER 31 provides the interface between user tasks (components of members 18) and the member agent 11.

The DM HANDLER 30 implements the distinguished member functionality and is active on only one member 18 of the group at a time. This member 18 is selected by the Group Master Task 25 (Figure 3b). The distinguished member 18 is responsible for allocation of rights as well as for broadcasting group status change notifications to all other members 18 of the group. (This information is received from the Group Master Task 25.)

Member agents 11 are created dynamically by the synchronization agent 14 in response to a SYNCHRONIZE directive (Primitive). They are also destroyed by the synchronization agent 14 after they have left the group or following a failure.

Broadcasts and Acknowledgements

When an application program initiates a broadcast (via the GROUP-BROADCAST primitive), its local member agent 11 distributes the information to all other active member agents 11. It then accumulates acknowledgements until all active member agents 11 have replied after which the application program is notified (via the GRP-ACK reply signal).

If an element 12 fails before its acknowledgement is dispatched, the broadcasting member agent 11 will assume an implicit acknowledgement from that member so that failures will not disrupt the application.

Group Establishment and Joining of New Members

A newly joining member 18 first informs the synchronization master control 13 (via its synchronization agent 14) of its intent to join the synchronization group. The synchronization master control 13 then determines if this is the first reported member of the group. If it is, then this Member 18 is designated as the distinguished member 18 and a notification is sent back. This establishes the group.

If the group is already established, synchronization master control 13 registers the new member 18 as being in the joining state and informs the group's distinguished member agent 11. Upon receiving this notification the distinguished member agent 11 broadcasts a join request to all member agents 11 on the list and waits for the corresponding group acknowledgement. The period between the broadcast of the join request and the full acknowledgement of that request by all joined member agents 11 is called the joining interval. During that time some member agents 11 will become aware of the new member agent 11 before others. This opens up the possibility that some messages broadcast within the group may bypass the partially synchronized member agent 11. If messages received by this member agent 11 are passed to the application program, then the application program function of this member 18 would not necessarily perceive the same sequence of group events as other members 18; it could miss some. Therefore, the new member agent 11 must acknowledge any messages received from other member agents 11 (in order to satisfy the acknowledgement requirement) but, once acknowledged, the messages are discarded; i.e. they are not passed on to the application (an exception is messages containing other joining or departure requests which are processed by the member agent 11 but still not relayed to the application). This mode of operation remains in effect until the join

request is finally acknowledged by the entire group. At that point, the new member 18 informs its application that it is fully joined and switches to normal operation. The overall effect, as perceived by the application, is that the joining operation is atomic.

The handling of messages that were discarded during the joining interval is no different to the application program than the handling of messages missed by the member 18 while it was down; that is, once synchronized with the gruop, the application program must proceed to upgrade its functional state to be consistent with the functional states of other members 18. The best method for achieving this depends on the application program.

Departure of Members

The departure of a member 18 from a synchronization group occurs when the member 18 decides to unsynchronize or when the host processing element 12 fails. In the former case, the departing procedure is as follows: the synchronization group (i.e. agent 11) notifies the Sync Master Task 20 of its intention. This event is relayed, via the appropriate Group Master Task 25 (Figure 3b), to the distinguished member 18 of the group. The distinguished member 18 then broadcasts this information to all other group members 18. Note that there is one Group Master Task 25 for every synchronization group defined in service 10.

In the case of a processing element 12 failure, the failure is detected by the Polling Control 22 within Sync Master Task 20 (Figure 3b) and the same sequence as described above is executed.

If the departed member 18 was a distinguished member, Group Master Task 25 will first select a new distinguished member 18 and then proceed in the same manner as above.

The synchronization agents 14 are intermediaries between synchronization master control 13 and the member agents 11. Synchronization agents 14 are created and dispatched when their host processing element 12 is initialized. Upon creation they wait to be contacted by the synchronization master 13, if one exists. Any application level requests for synchronization are queued until an acknowledgement is received from synchronization master control 13.

During normal operation, the synchronization agents 14 serve as a relay point for communication between the synchronization master control 13 and the member agents 11. All communication is buffered until acknowledged by the receiver so that the Member Agents 11 are protected from temporary failures of synchronization master control 13. The synchronization agents 14 also extract and store any information relevant to the reestablishment of the synchronization master control 13.

Most of the operation of the synchronization master control 13 has already been described above. The only aspect remaining is the monitoring function.

The monitoring of the existence of processing elements 12 is done by the Polling Control 22 which polls each individual synchronization agent 14. The failure of a processing element 12 implies that the corresponding synchronization agent 14 is down as well as all member agents 11 that were present on that processing element 12. When that happens the synchronization master control 13 notifies all affected Group Master Tasks 25. These, in turn, inform their distinguished member agents 11 which then broadcast this information to other member agents 11.

Before we go any further, it may be advantageous to introduce the primitives used with synchronization service 10. The primitives can be split into two categories:

(1) Synchronous Primitives are in the form of request-reply pairs; member agents 11 submit requests for some action to be performed on their behalf and synchronization service 10 eventually matches these with appropriate replies.

(2) Asynchronous Notifications are spontaneous signals informing a member agent 11 about changes in the status of its group or conveying a message sent by some other member agent 11.

There are only two types of asynchronous notifications that can be sent to a member agent 11:
-GROUP-CHANGE (group status) is sent when a new member 18 has joined or an active member 18 has departed from the group. The status information includes the complete new membership list and the id of the new distinguished member.
-GROUP-MSG (message) signals the arrival of a message from some other member 18 (broadcast or point-to-point).

The application program must allow forms of communication (i.e. synchronous and asynchronous) although it may choose to handle asynchronous communications in a synchronous manner by ignoring them until the current activity sequence is complete.

The synchronous primitives and corresponding replies are depicted in chart form in Figure 3a, to which attention is directed.

The primitives are:

-SYNCHRONIZE (group-id)

This is directive which is issued by a member 18 (via its member agent 11) when it wishes to become synchronized with the group specified by <group-id>. If no group exists at the time, one is established. The only signal expected in reply to this directive is the SYNCH-DONE signal.

-SYNC-DONE (group-status)

This is a signal from the synchronization service 10 (i.e. member agent 11) in response to a successful synchronization of a member 18 following the invocation of the SYNCHRONIZE directive. The return parameter, <group status>, contains the same information about the status of the group as the GROUP-CHANGE primitive described below. It includes a <dm-flag> parameter which informs the member 18 if it is the bearer of the distinguished member status.

- UNSYNC

This directive is used when a member 18 decides to depart from its group. It ensures orderly deactivation.

-UNSYNC-DONE

This signal is a confirmation that the member 18 has been removed from its synchronization group.

-GROUP-CHANGE (group-status)

This is an asynchronous signal which is generated by the member agent 11 whenever a new member 18 joins the group or when a member 18 departs from the group. If this mmeber 18 is the new distinguished member as a result of the change, a <dm-flag> parameter in the <group-status> data record will be set appropriately. The treatment of this situation is left to the application program. The new status of the group is also returned.

-REQ-RIGHT (right-id, mode)

This directive is issued when a member 18 needs exclusive access to a group right. If the right is available, it is guaranteed to be granted to only one requesting member 18 (there may be multiple simultaneous requests for the same right). If the right is not available, then if the <mode> parameter specifie a "queued" request, it is queued until it can be serviced. Alternatively, if the <mode> parameter specifies "immediate" the request is refused since the right has already been appropriated by another member 18 of the group.

- R-GRANTED (right-id)

This signal informs a member 18 that it has been granted the required right.

-R-REFUSED (right-id)

This signal informs a member 18 which has requested a right, with the "immediate reply" mode specified in the request, that the right is not available. (If a queued request was made then this signal will never be generated.)

-REL-RIGHT (right-id)

This directive is used to release an appropriated right.

-R-RELEASED (right-id)

This signal is the reply to the REL-RIGHT directive.

-QRY-RIGHTS

This is a directive which is used to obtain a snapshot of the distribution of group rights among group members.

-R-STATUS (rights-status)

This is a reply signal to the QRY-RIGHTS directive. The <rights-status> parameter lists, for each group right, the member-id of the member which owns it, it any.

Note that service 10 cannot guarantee the currency of the returned information since changes in the distribution of rights can occur at any time.

-GRP-BRDCST (message)

This directive is used to broadcast a synchronization event (message) to all synchronized members 18. It is the responsibility of the synchronization service 10 (via member agent 11) to ensure that all members 18 receive the message. The <message> parameter can be used to timestamp the synchronization event. The higher level software is responsible for supplying this parameter as well as interpreting its functional significance.

-GROUP-ACK

This is an acknowledgement signal for the GRP-BRDCST directive. It signifies that all members 11 have received the latest broadcast message.

-SND-TO-MEM (message)

This directive is used to send a point-to-point message to another group member 18.

-MSG-ACK

This is an acknowledgement that the latest point-to-point message has been received by the destination member 18.

Before the invention is described further, it may be of value to give some brief examples of the application of the primitives.

The first example is the control of a standby configuration. In this configuration there are two or more distributed program components (i.e. members 18) each on a different processing element 12, each of which is equally capable of providing the necessary function. Only one should be active at any given time while the others are standing by, ready to be activated should the active one fail. Assuming that they are all part of the same synchronization group that the algorithm which each member 18 executes is the same (the synchronization service primitives are highlighted in capitals):

SYNCHRONIZE;
Wait for SYNC-DONE signal;
If not selected as the distinguished member then
Repeat
Listen for SYNC-CHANGE signals;
until selected as the distinguished member;
Execute funciton;

If a member 18 is not selected as the distinguished member following synchronization with the group, then it simply waits until it is designated as the distinguished member.

The next exampled concerns the updating of a replicated database, i.e. the same example mentioned in the Background of the Invention. In this case there are multiple instances of a database, each of which can initiate an update request as a result of external activity. Such requests will be called external to distinguish them from "shadow requests". Shadow request are copies of an external request which a member 18 sends to all other members 18 so that they can make the appropriate changes to their copies of the database. For brevity, the handling of any other requests except update requests is ignored.

The solution shown below uses the mutual exclusion feature of the synchronization group. A right, called the Update right, is defined. The holder of this right is the member 18 whose request will be honored; all other members 18 must withhold their requests and perform the shadow request sent by the holder of the right:

Solution A

Repeat
Wait for next request;
If external request then
begin
REQ-RIGHT (update);
While waiting for R-GRANTED
Handle any incoming shadow requests;
GRP-BRDCST (external request);
Handle external request;
REL-RIGHT (Update);
end
else
Handle shadow request;
until termination;

Note that the application program need not be concerned with spontaneous failures of other members 18 since that is handled by the synchronization service 10.

An important problem which must be handled by this application (i.e. Solution A, above) is the addition of new or recovering instances. These will not necessarily have the same state as the others and therefore must be brought to the same functional level. The situation is complicated by the possibility that updates may be initiated at other instances whil the new instance is being upgraded. One method of dealing with this is for the new instance to appropriate the Update right to ensure that the state remains unchanged while it is being upgraded. The algorithm performed by a restarting instance is then:

SYNCHRONIZE;
Wait for SYNC-DONE signal;
REQ-RIGHT (Update);

While waiting for R-GRANTED
Discard any shadow requests received;
Obtain current copy of database;
REL-RIGHT (Update);

Following this, the normal request-processing algorithm described above (i.e. Solution A) is executed.

The current copy of the database is obtained from any other member 18 through an internal protocol using point-to-point messages (i.e. SND-TO-MEM directives). Instead of a copy of the entire database it may be more convenient to request an update log and then perform the updates missed while the member 18 instance was down.

Figure 4 is a functional flow diagram representing the synchronization service 10 database 16 when accessed through the processing element 12 identifier.

The head and tail pointers (AGT-LST-HD and AGT-LST-TL) respectively, point to a linked list of synchronization agent control blocks (tAGT-CB) for those synchronization agents 14 involved.

There is one synchronization agent control block tAGT-CB for each processing element 12 which requires synchronization service 10. It contains a link (AGT-LST-LNK) to other synchronization agent control blocks tAGT-CB. This chain enables quick scanning of affected processing elements 12 when an entire block of processing elements 12 fails. Each synchronization agent control block tAGT-CB also contains a pointer (MMCB-LST-HD) to a chain of member agent control blocks (tMEM-CB) which reside on that processing element 12. Through this chain it is possible to detect quickly all synchronization groups which are affected by the failure of a processing element 12. Whereas all other chains in the synchronization service database 16 remain unchanged once they are established, this chain follows the dynamics of member 18 joinings and departures.

In order to simplify searching and list maintenance, the last Sync Agent control block tAGT-CB in the list is a dummy block.

Each member agent control block tMEM-CB corresponds to one member 18 of one synchronization group. Among other data, this control block contains a pointer (not shown in the diagram) to the corresponding synchronization agent control block tAGT-CB. This link allows quick reconfiguring of the processing element-Member Agent chain when necessary.

Figure 5 is a functional flow diagram representing the synchronization service database 16 when accessed through the unique group identifier.

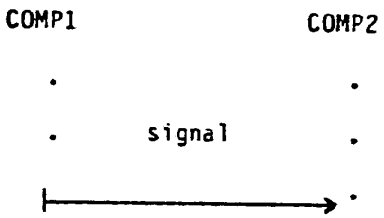
GRP-HDR [gropu id] is a static array of pointers. Each item of the array points to a circular list of member agent control blocks (tMEM-CB of which belong to the same synchronization group.

The member agent control blocks tMEM-CB are linked into a circular list to facilitate selection of a distinguished member. This list grows as members 18 are added to the synchronization group, each successive block identified by the next available positive integer (MEM-ID). This integer corresponds to the member 18 identifier.

ACTION SEQUENCES

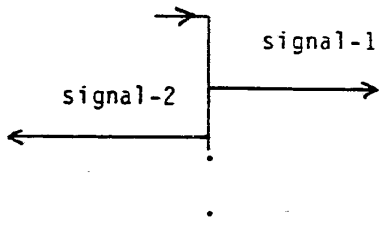
This section describes various action sequences within synchronization service 10. A diagrammatic representation is used to show these sequences. The following conventions are used.

- A full horizontal line indicates a message or rendezvous between two components (i.e. programs or tasks):



where signal is the name of the entry procedure in component COMP2 which accepts the signal. COMP1 is the component which sent the signal.

- a vertical line (1) following the reception of a signal indicates processing within the appropriate component which received the signal. This processing results in one or more signals being dispatched to other components:



- If an asterisk (*) appears next to a signal it implies that the signal may be repeated (to different destinations).

- A signal which is enclosed in braces (e.g.,<signal> indicates that the signal is not mandatory and may be omitted depending on circumstances.

• Bracketed numbers in the Figures (e.g. (1)) designate explanatory notes which contain textual descriptions pertaining to various signals. The notes are given in the text relating to the relevant Figure. As it is believed that the Figures are self-explanatory, only brief comments will be made regarding the figures.

Joining of New Members

An "empty" synchronization group is one in which no members 18 are active. When the first member 18 joins, it is designated as the distinguished member by default. Figure 6 depicts a member 18 joining an empty group. The sequence of events is as depicted in Figure 6, to which attention is directed. The abbreviations used in the Figures are as follows: APPL means an application task; SNYC AGT means the synchronization agent task; MEM AGT means the member agent task; SYNC MST menas the synchronization master task 20; GRP MST means the Group Master Tast 25; MEM AGT (DM) means the distinguished member agent task; APPL (DM) means the application task which corresponds to the distinguished member agent.

The following notes refer to the bracketed numbers in Figure 6.

Notes:

(1) The Sync Agent will create a member agent task only if it had not existed previously, otherwise it will REINIT a previously created task instance. The START-AGT signal which follows initialization is used to pass initial data to the member agent 11.

(2) The AGT-MST-MSG signal to the sync master 13 contains the complete information about all member agents 11 on this processing element 12, including the newly-created member 18. (This ensures state convergence even in the presence of design faults.) The MST-REPLY message is used for positive acknowledgement so that the Sync Agent can send the next message to the Sync Master if it has one. (Only one outstanding message is allowed between the Sync Master and a Sync Agent.)

(3) If this group has not previously existed, a new Group Master is created. In that case a STARTUP signal follows to pass initial data to the Group Master Task, and an ACTIVATE signal is used to force it into an operational mode. If the group had existed previously (but had lost all its members) the existing Group Master Task is used.

(4) Once the Group Master has been activated, a GRP-EVENT signal is sent by the Sync Master informing it of the joining of the first member.

(5) Upon receipt of the GRP-EVENT signal, the Group Master selects the newly-created member agent 11 as the distinguished member and sends it a GRP-STATE signal. This signal establishes a connection between the Group Master and the Distinguished Member. All subsequent GRP-STATE signals are sequenced to ensure proper event ordering as well as to guard against communication failures. The GRP-REPLY signal is used to acknowledge one or more GRP-STATE signals and provides reliable communication.

The GRP-STATE signal contains the complete new state of the group rather than just information about the changes. This ensures that the system will converge to the true state even in the presence of design faults.

(6) The CHECK-FAIL signal is used to poll the application task to detect unexpected failures of the application. The application task never receives this signal; however, should the task fail, the member agent task will be notified by the underlying operating system kernel.

(7) The SYNC-REPLY signal contains a reply code of SYNC-DONE.

Figure 7 depicts the sequence for joining an existing synchronization group; the Group Master Task already exists, and the distinguished member is used to notify (broadcast) the other members 18 of the presence of a new member 18.

Notes:

(8) If the application has so requested, a GROUP-CHANGE signal is sent by each member agent 11 to the application whenever it detects a change in status of the group. In the case of the distinguished member the status change is gleaned from the GRP-STATE signal.

(9) When the distinguished member receives a GRP-STATE signal which indicates a group change (not all do), it broadcasts the new state to all other group members 18 using a MEM-MSG signal. Each member 18 acknowledges such messages with an ACK signal to provide reliable communication.

(10) When the newly-joining member 18 receives its MEM-MSG from the distinguished member it will send a SYNC-REPLY signal to the application (instead of a GROUP-CHANGE signal).

The control flow for the departure of a member 18 is shown in Figure 8. Note that the case of a processing element 12 failure is not shown here but is instead treated separately.

Notes:

(11) The sequence shown here corresponds to a voluntary departure initiated by the application program issuing an UNSYNC directive. This results in the member agent task terminating which, in turn, sends a COMPLETE signal to the parent task, the Sync Agent. The sequence is similar in situ-

ations where the departure is not voluntary:
-When the application task fails, the member agent 11 is notified (through the failure of the CHECK-FAIL message) which results in the termination of the member agent task (and consequently, raising of the COMPLETE signal).
-If the member agent task itself fails, the Sync Agent is notified by the operating system kernel with a COMPLETE signal.

(12) Indicates a <GROUP-CHANGE> signal to an application task not shown in the Figure.

Figures 9a and 9b together depict the recovery of the synchronization master control 13 (i.e. Sync Master). The most general case is considered, i.e. the case of a running synchronization service 10 with active synchronization groups. This includes, as a subset, the case of a "cold" start.

Notes:

(13) Upon activation, Sync Master sends an ACTIVATE signal to each configured Sync Agent. The Sync Agent, whether they are already active or not, will respond with an AGT-REPLY message which includes a list of all member agents supported on that processing element. (The MST-REPLY signal is used for acknowledgement: refer to Figure 6.)

(14) Following activation of the Sync Agent, the Polling Control subcomponent 22 of Sync Master Control sends a POLL-AGT message to which the Sync Agent responds with an AGT-REPLY message. This exchange is repeated periodically to detect outages of the processing element.

(15) A Group Master is initiated only the first time a group is encountered. Refer to Figure 6 for further details on Group Master initiation.

(16) Before activating a Group Master, it is provided with data regarding the status of its members through GRP-DATA signals. Each signal contains the information for one member agent of one group. (This info is obtained from the AGT-REPLY messages.) In this way, the Group Master reconstructs the status of its group.

(17) After all Sync Agents have responded, the reconstruction is complete and an ACTIVATE signal is sent to all Group Masters. The Group Masters respond by sending a GRP-STATE signal to all distinguished members. Since this signal contains the complete group state, any group changes that might have occurred while the Sync Master was down are detected.

Figure 10 depicts the procedure for handling the failure of a processing element 12 which contains a synchronization agent 14 (Sync Agent).

Notes:

(18) The Sync Master detects a failure of a processing element when a TIME_OUT event is received. This means that a Sync Agent has not responded to a poll.

(19) For each group affected by the processing element failure, the Sync Master will send a GRP-EVENT signal to the respective Group Master.

Figure 11 depicts the procedure for the recovery of a processing element 12 which is part of synchronization service 10. Note that the recovery of processing element 12 does not extend to recovering member agent 11 tasks. It is assumed that these will be recovered when the application tasks (i.e. programs) which use them are restarted. Thus, the only action to recover a process element 12 is to integrate the sync agent 14 with the rest of synchronization service 10.

Notes:

(20) When a previously failed Sync Agent finally responds to a POLL_AGT signal, the Sync Master initiates the recovery procedure.

(21) The Sync Master registers the new processing element 12 and sends an ACTIVATE signal to the Sync Agent on that processing element 12 (Refer to Figure 9a for a more detailed description of the activation sequence).

Figure 12 depicts the procedure for member 18 to member 18 messages. This procedure (protocol) is used both for group broadcasts and point-to-point messages between members 18.

Notes:

(22) In case of a broadcast, a copy of the message is sent to each member 18. If the member 18 is not active, the message is not sent.

(23) Upon receiving a MEM-MSG signal which indicates an application-level message the message is relayed to the application task responsible for receiving asynchronous messages.

(24) A SYNC-REPLY (codes: GRP-ACK or MSG_ACK) signal is sent back to the originator.

Figure 1 depicts the procedure employed in the processing of all directives which require distinguished member intervention (rights handling directives).

Notes:

(25) If the request is made on the distinguished member site, then no message is sent.

(26) A reply signal (MEM-MSG followed by a SYNC-REPLY to the application).

In one implementation made by the inventors, the code for synchronization service 10 was contained in ten files which were distributed into six units, the useage dependencies (and compilation order) of which are shown in Figure 14.

Notes:

SYNCCTRL contains the stub of the Sync Master unit and directives to include three files (SYNCMST, SYNCGMST, and SYNCPOLL) which implement the Sync Master function. It also contains the definitions required for the master database 16. SYNCMST is an "include" file which

contains the code for the Sync Master Task.
SYNCGMST is an "include" file which contains the code for the Group Master Task.
SYNCPOLL contains the Polling Control 22.
SYNCLOCL contains the stub of the Sync Agent unit as well as a definition of data and procedure objects shared by the Sync Agent Task and the Member Agent Tasks. It also contains directives to include two files (SYNCAGT, SYNCMAGT).
SYNCRESX contains the definition of the SYNCHRONIZE primitive. This unit must be loaded with the application code which uses the synchronization service.
SYNCDEFI contains a set of internal compile-object definitions for the synchronization service. This includes the definition of all entries used for communication between synchronization service components which are hidden from user programs. Since this file contains only compile objects it is not loaded.
SYNCDEFX contains a definition of all compile-objects which are exported by the synchronization servie to its users. Since this file contains only compile objects it is not loaded.
Application tasks which use the synchronization service need to include SYNCRESX and SYNCDEFX in their useage lists.

Note that SYNCCTRI implements the function of synchronization master control 13 (Fig. 1), and that SYNCLOCL and SYNCRESX together implement the functions of member agent 11 and synchronization agent 14 (Fig. 1). The files SYNCDEFI and SYNCDEFX are not resident in the service; they can be thought of as tools used in the construction of the synchronization service but they are not themselves a part of it.

Simplified pseudocode listings for the main constituents of the invention follow as appendix I. They are believed to be self-explanatory. Any elaboration of the material is accomplished through the use of appended notes, to which attention is directed.

As a further aid to the understanding and to the use of the present invention the following (a copy of a "User's Reference" to the synchronization service of the present invention, as prepared by one of the inventors) is included as Appendix II. It will expand on the use of the present invention.

Claims

1. A synchronization service [10] for use with a computer having a distributed operating system, to allow the construction of a customized synchronization scheme, for synchronizing the constituent portions of a distributed program, said service comprising:

a general set of application-independent synchronization primitives, whereby the construction of said customized synchronization scheme is achieved by the selective implementation of said application-independent synchronization primitives.

2. The synchronization service of claim 1 wherein said application-independent primitives comprise the following functions: synchronize; synchronize done; and unsynchronize.

3. The synchronization service of claim 2 wherein said primitives further comprise the following functions: request right; right granted; right refused; release right; group broadcast; and group acknowledge.

4. The synchronization service of claim 3 wherein said primitives further comprise the following functions: unsynchronize done; send to member; and message acknowledge.

5. A synchronization service [10] for use with a computer having an operating system [15] distributed over a plurality of processing elements [12], to allow the construction of a customized synchronization scheme, for synchronizing the constituent portions [18] of a distributed program, said service comprising:
a common synchronization master control means [13];
a synchronization agent means [14] for each processing element;
a plurality of application program components [18], each component located on a different processing element, each said component having associated therewith a member agent [11], said member agent being a program for interfacing with said synchronization agent means, and said synchronization agent means interfacing between said master control means and said member agent, whereby a customized synchronization scheme can be constructed based upon a general set of application-independent synchronization primitives contained in both said synchronization agent means [14] and said member agent [11] and accessed via said synchronization agent means.

6. The synchronization service of claim 5 wherein said application-independent primitives comprise the following functions: synchronize; synchronize done; and unsynchronize.

7. The synchronization service of claim 6 wherein said application-independent primitives further comprise the following functions: request right; right granted; right refused; release right; group broadcast; and group acknowledge.

8. A synchronization service [10] for use with a computer having an operating system [15] distributed over a plurality of processing elements, to allow the construction of a customized synchroniza-

tion scheme, for synchronizing the constituent component of a distributed program, said service [10] comprising the steps of:

a) joining a program component [18] on a first processing element [12] to a group of existing program components [18] on at least a second processing element [12] so that each of the existing components is aware of the presence and location of the joining components;

b) informing each member of a group of physically distributed program components when one or more components which are members of said group, depart from it;

c) selecting, as a distinguished member, one program component from a group of distributed program components such that, within said group, there is never more than one said distinguished member; and

d) providing mutually exclusive rights to said group of distributed program components such that no more than one said component can appropriate a given right at any time.

9. The synchronization service of claim 8 further including the step of providing reliable point-to-point communication between said distributed program components on the basis of their internal group identifiers.

10. The synchronization service of claim 9 further including the step of providing a broadcast mechanism from any one program component to all other program components which are currently declared as being in the same group as the broadcasting component.

11. The synchronization service of claim 10 wherein said program components are components of an application program.

12. The synchronization service of claim 10 wherein said program components are components of an operating system program.

13. The synchronization service of claim 8 wherein said physical processing elements are logically distributed entities at one physical location.

14. A synchronization service [10], for use with a computer having an operating system [15] distributed over a plurality of processing elements [12], to allow the construction of customized synchronization schemes for synchronizing the constituent components [18] of a distributed program, said service comprising, as required, the steps of:

a) establishing a synchronization group for said distributed program, said group comprising at least one distributed program component [18];

b) joining a program component [18] to said group of existing program components so that each of the components is aware of the presence and the location of all the other components in said group;

c) informing each member of said group of distributed program components when one or more components which are members of said group, depart from it;

d) selecting, as a distinguished member for said group, one program component from said group of distributed program components such that, within said group, there is never more than one said distinguished member; and

e) providing mutually exclusive rights to said group of distributed program components such that no more than one said component can appropriate a given right at any time.

15. The synchronization service of claim 14 further including the step of providing full connectivity between all said distributed program components of said group.

16. The synchronization service of claim 15 wherein said distributed program is an application program.

17. The synchronization service of claim 15 wherein said distributed program is an operating system program.

18. The synchronization service of claim 15 wherein each said program component [18] is on a different processing element [12].

19. A synchronization service [10], for use with a computer having an operating system [15] distributed over a plurality of processing elements [12], to allow the construction of customized synchronization schemes for synchronizing the constituent components [18] of a distributed program, said service including a synchronization master control [13] comprising:

master control means [21] for activating said synchronization service;

polling means [22] for polling the processing elements [12] associated with said components of said distributed program so as to monitor the status of said processing element;

control means [24] for joining new members [18] to said group, and for handling departures of members [18] from said group; and

a database means [16] containing information representative of the current state of said synchronization service at a given point in time.

20. The synchronization service of claim 19 further including, at each said processing element, a synchronization agent [14] comprising:

means for accepting synchronization directives and for creating corresponding member agents; and

means for monitoring the status of all active member agents on said processing element and reporting same to said synchronization master control [13].

21. The synchronization service of claim 20 further including at each said processing element [12], a member agent [11] each synchronization

group, comprising:
communications means [33] for providing a reliable communications service between program components [18];
storage means [32] for maintaining a local version of the current state of all other program components [18];
handler means [31] for providing the interface between user tasks and said member agent [11]; and
distinguished member means [30] for implementing the distinguished member function on only one program component [18] at any given time.

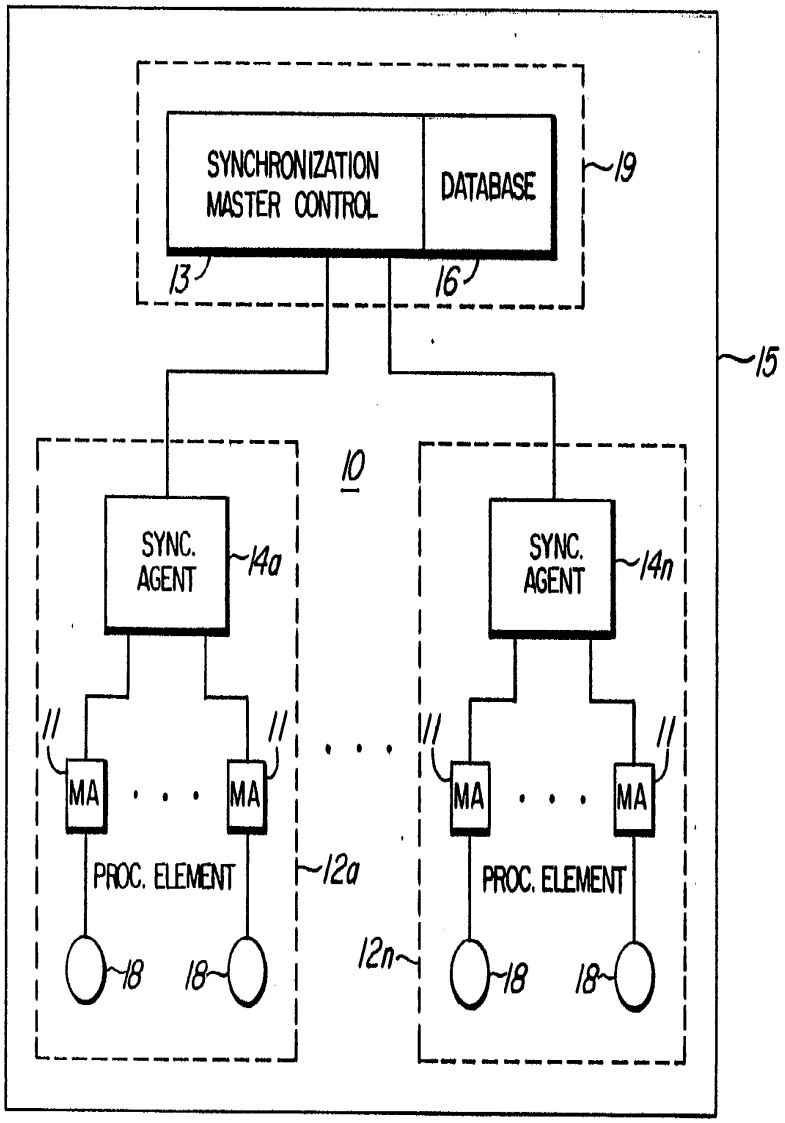


FIG. 1

0 251 584

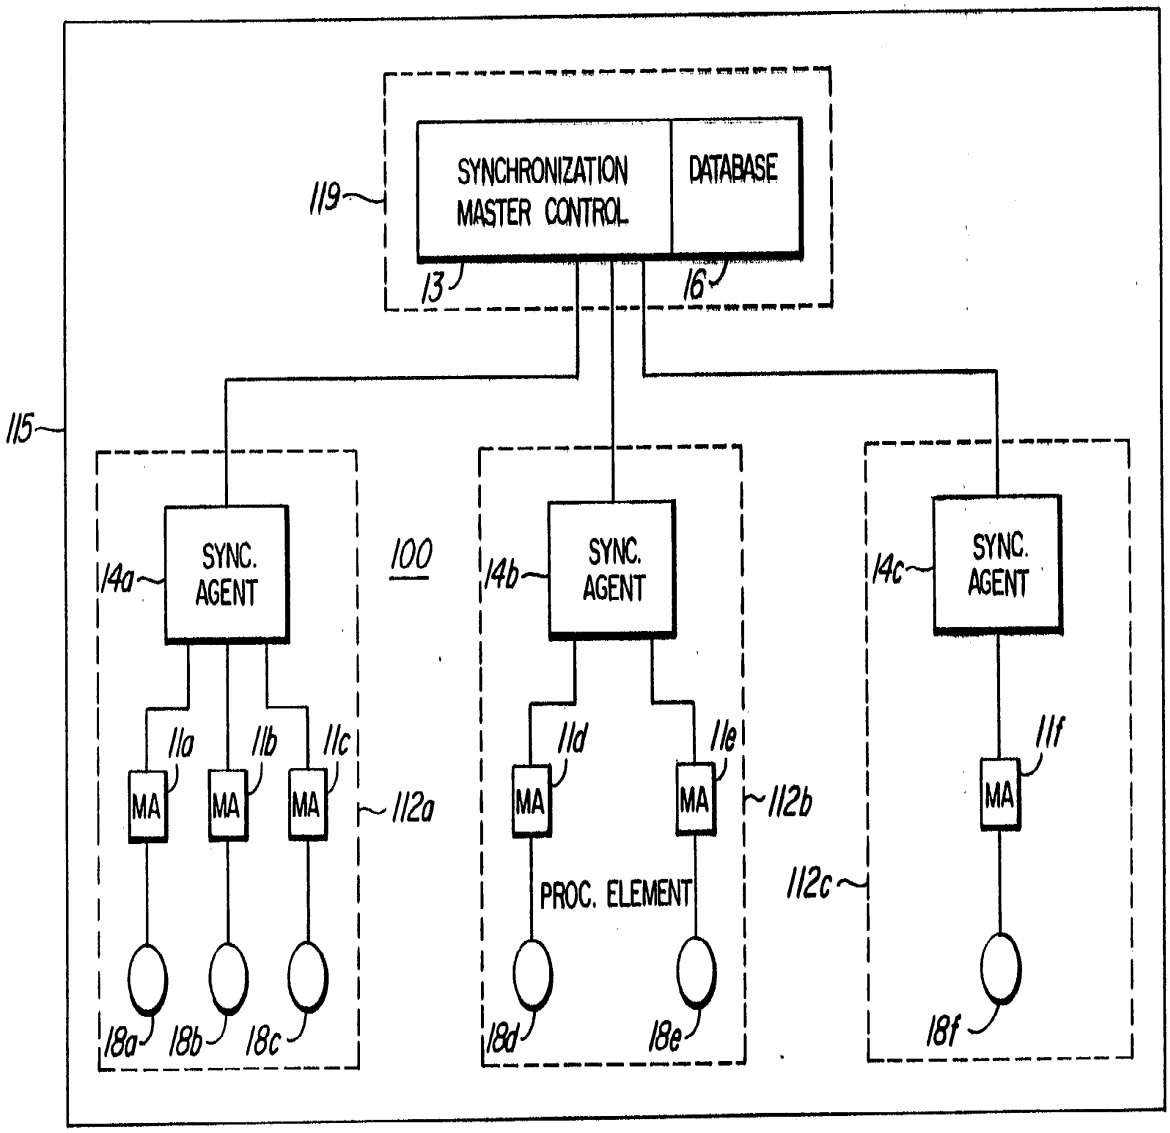


FIG. 2a

0 251 584

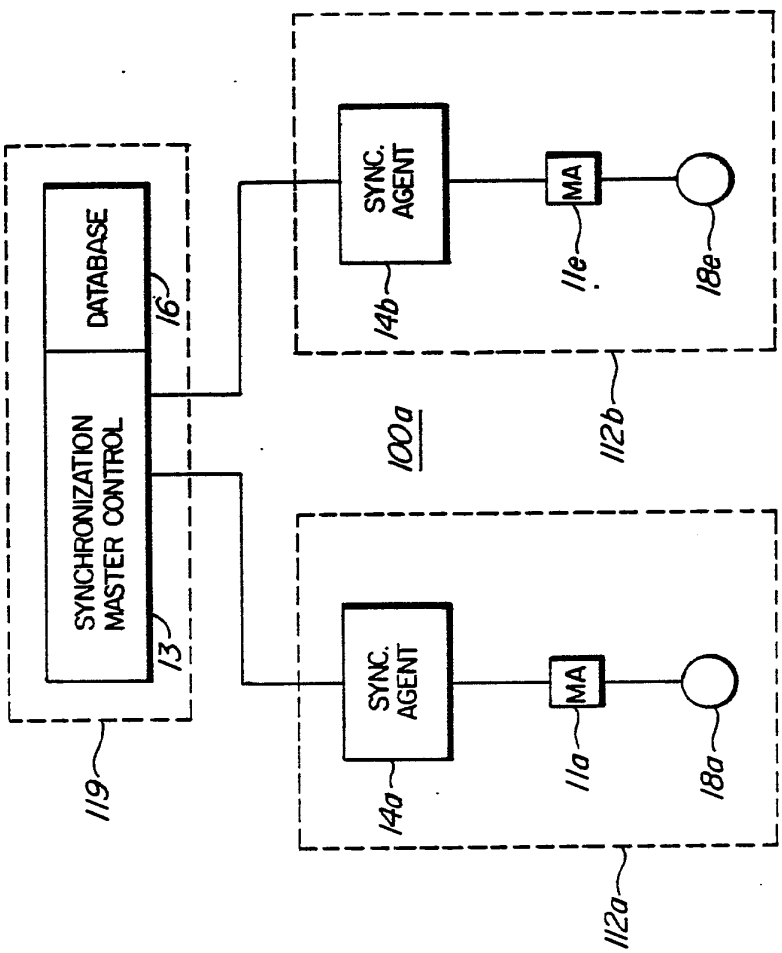

SYNCHRONIZATION MASTER CONTROL
DATABASE
13
16
119
SYNC. AGENT
MA
14a
11a
18a
112a
100a
SYNC. AGENT
MA
14b
11e
18e
112b

FIG. 2b

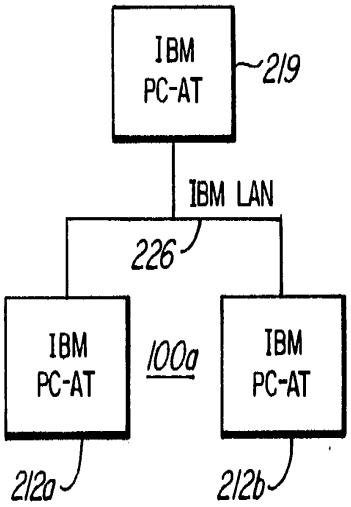


FIG. 2c

SYNCHRONOUS PRIMITIVES

REQUEST	REPLY
SYNCHRONIZE (group-id)	SYNC_DONE (group-status)
UNSYNC	UNSYNC_DONE
REQ_RIGHT (right-id, Immed)	R_GRANTED
	R_REFUSED
REQ_RIGHT (right-id, Queued)	R:GRANTED
REL_RIGHT (right-id)	R_RELEASED
QRY_RIGHTS	R_STATUS (rights-status)
GRP_BRDCST (message)	GRP_ACK
SND_TO_MEM (message)	MSG_ACK

ASYNCHRONOUS NOTIFICATIONS

GROUP-CHANGE (group-status)
GROUP-MSG (message)

FIG. 3a

0 251 584

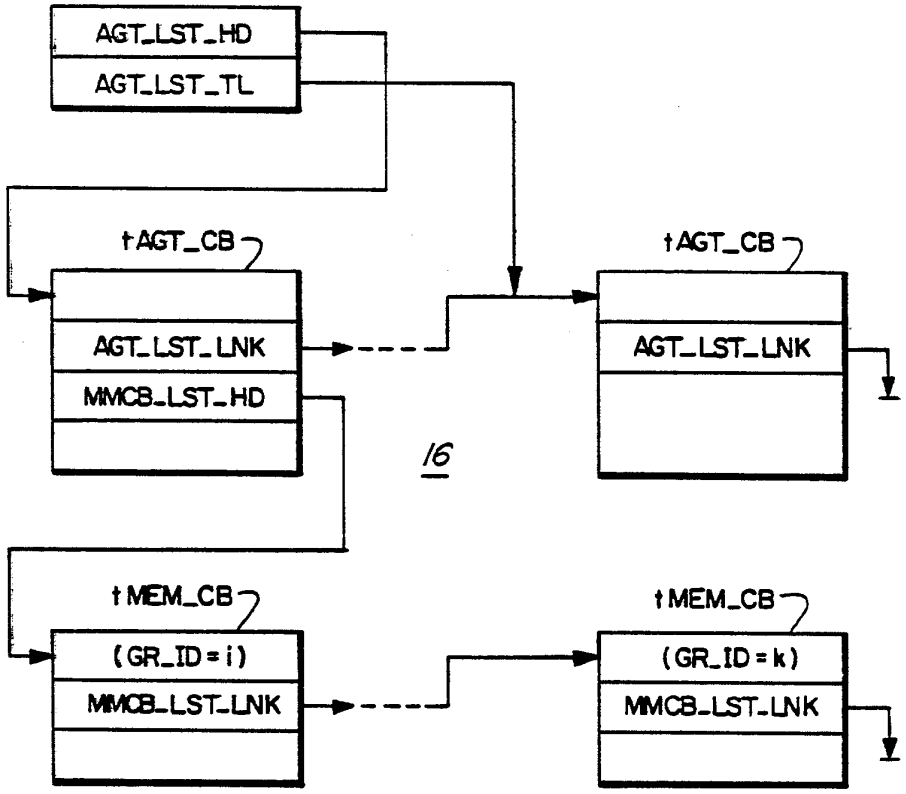

AGT_LST_HD
AGT_LST_TL
tAGT_CB
AGT_LST_LNK
MMCB_LST_HD
tAGT_CB
AGT_LST_LNK
16
tMEM_CB
(GR_ID = i)
MMCB_LST_LNK
tMEM_CB
(GR_ID = k)
MMCB_LST_LNK

FIG. 4

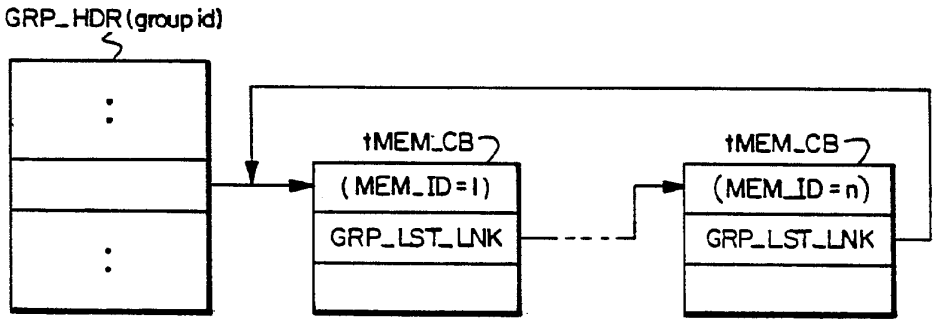

GRP_HDR (group id)
tMEM_CB
(MEM_ID = 1)
GRP_LST_LNK
tMEM_CB
(MEM_ID = n)
GRP_LST_LNK

FIG. 5

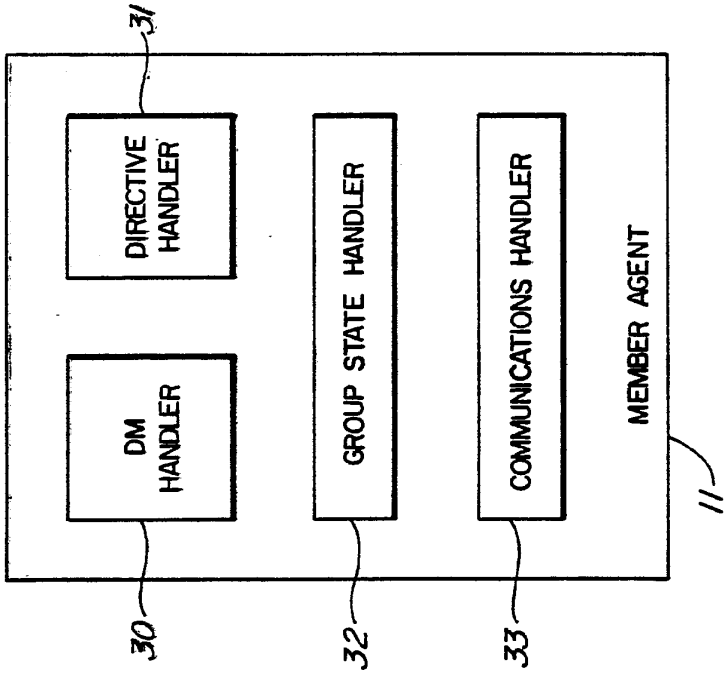

31
DIRECTIVE HANDLER
DM HANDLER
GROUP STATE HANDLER
COMMUNICATIONS HANDLER
MEMBER AGENT
11
30
32
33

FIG. 3c

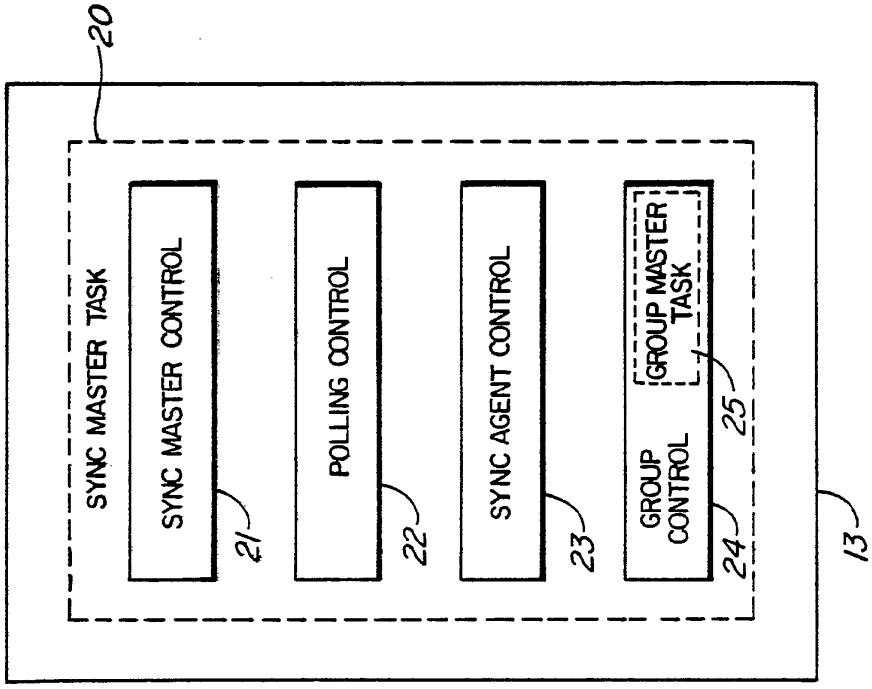
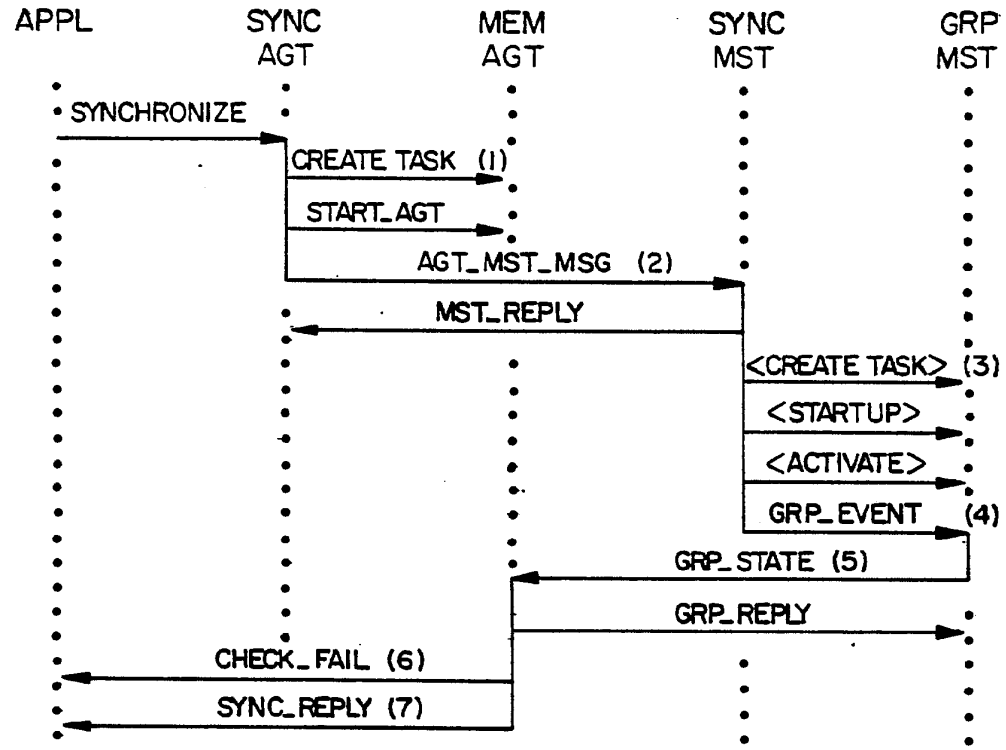

20
SYNC MASTER TASK
SYNC MASTER CONTROL
POLLING CONTROL
SYNC AGENT CONTROL
GROUP CONTROL
GROUP MASTER TASK
21
22
23
24
25
13

FIG. 3b



LEGEND
APPL = MEMBER 18
SYNC AGT = SYNCHRONIZATION AGENT 14
MEM AGT = MEMBER AGENT 11
SYNC MST = SYNCHRONIZATION MASTER 13
GRP MST = GROUP MASTER 25

FIG. 6

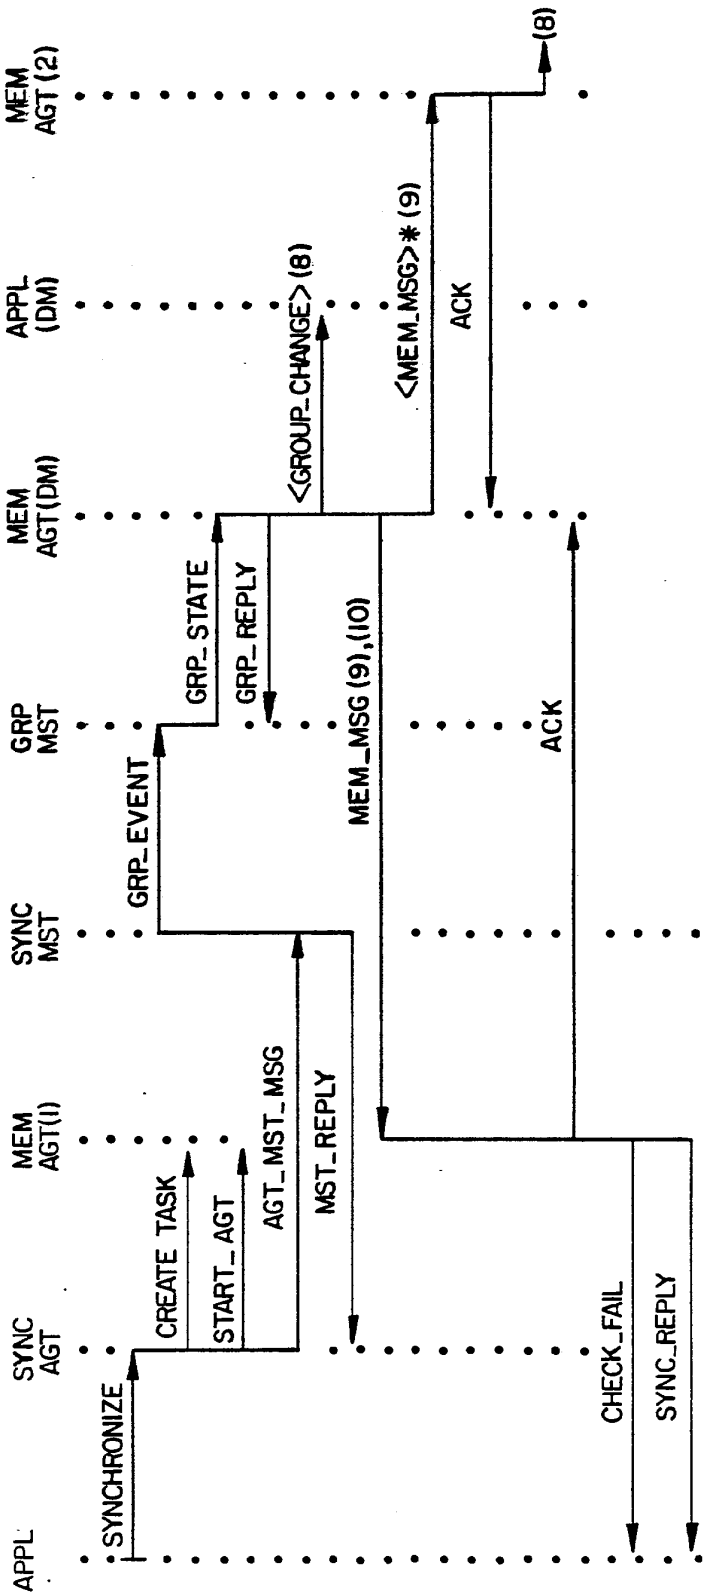

APPL
SYNC AGT
MEM AGT(I)
SYNC MST
GRP MST
MEM AGT(DM)
APPL (DM)
MEM AGT (2)
SYNCHRONIZE
CREATE TASK
START_AGT
AGT_MST_MSG
MST_REPLY
GRP_EVENT
GRP_STATE
GRP_REPLY
<GROUP_CHANGE> (8)
MEM_MSG (9),(10)
<MEM_MSG>* (9)
ACK
(8)
ACK
CHECK_FAIL
SYNC_REPLY

FIG. 7

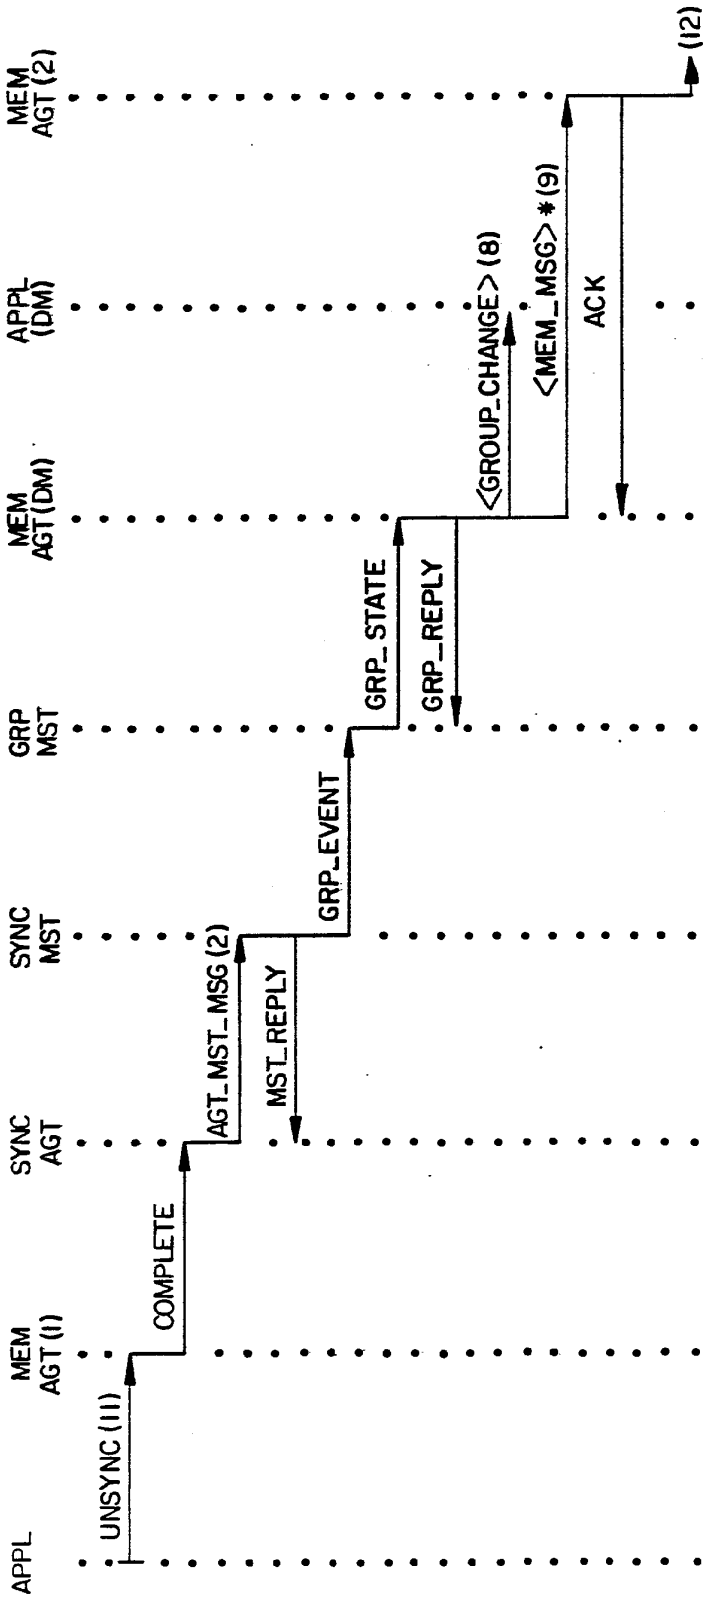

APPL
MEM AGT (1)
SYNC AGT
SYNC MST
GRP MST
MEM AGT (DM)
APPL (DM)
MEM AGT (2)
UNSYNC (11)
COMPLETE
AGT_MST_MSG (2)
MST_REPLY
GRP_EVENT
GRP_STATE
GRP_REPLY
<GROUP_CHANGE> (8)
<MEM_MSG> * (9)
ACK
(12)

FIG. 8

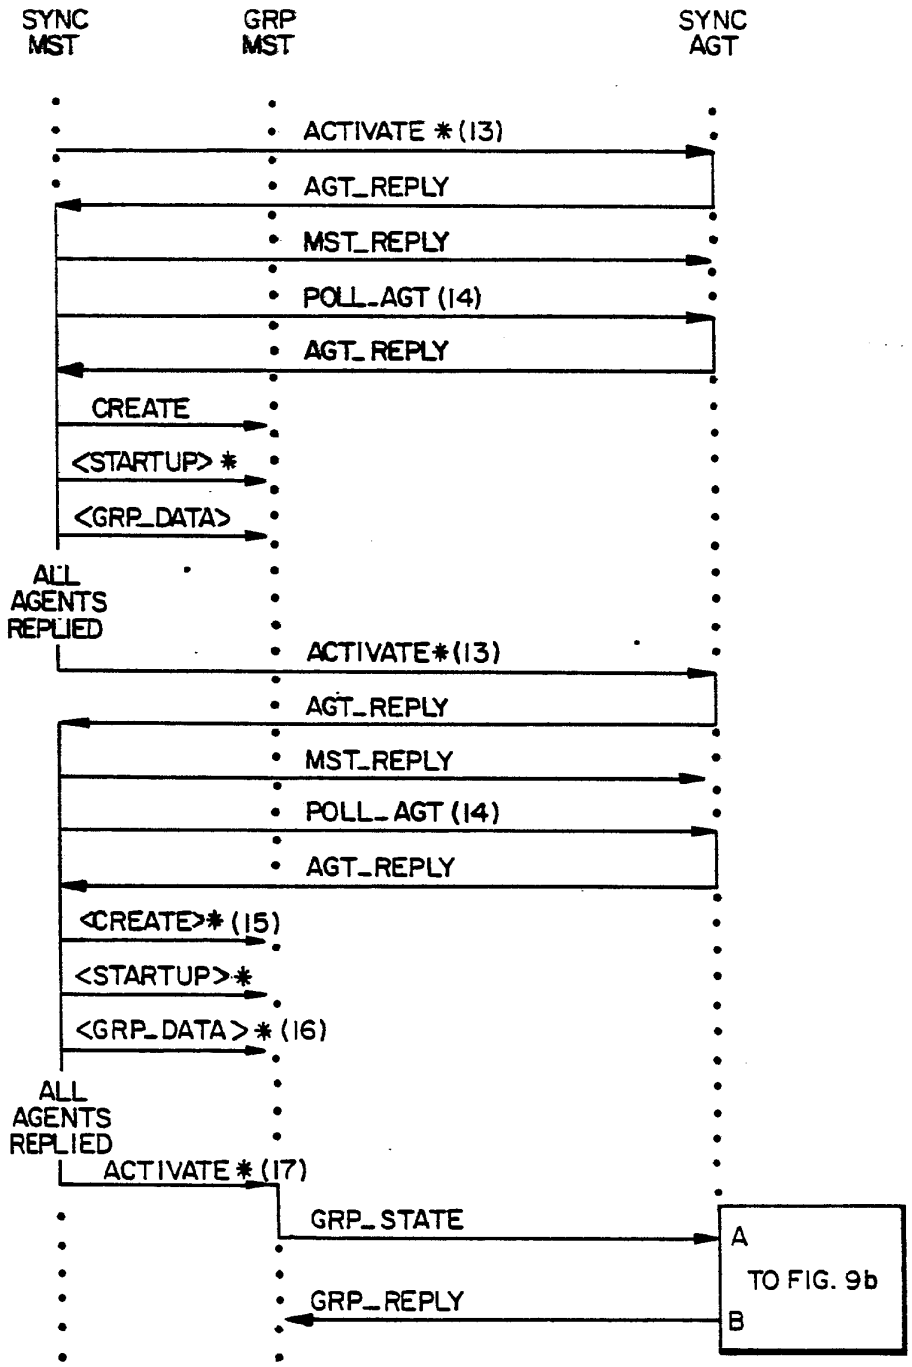

SYNC MST
GRP MST
SYNC AGT
ACTIVATE * (13)
AGT_REPLY
MST_REPLY
POLL_AGT (14)
AGT_REPLY
CREATE
<STARTUP> *
<GRP_DATA>
ALL AGENTS REPLIED
ACTIVATE*(13)
AGT_REPLY
MST_REPLY
POLL_AGT (14)
AGT_REPLY
<CREATE>* (15)
<STARTUP>*
<GRP_DATA>* (16)
ALL AGENTS REPLIED
ACTIVATE * (17)
GRP_STATE
GRP_REPLY
A
TO FIG. 9b
B

FIG. 9a

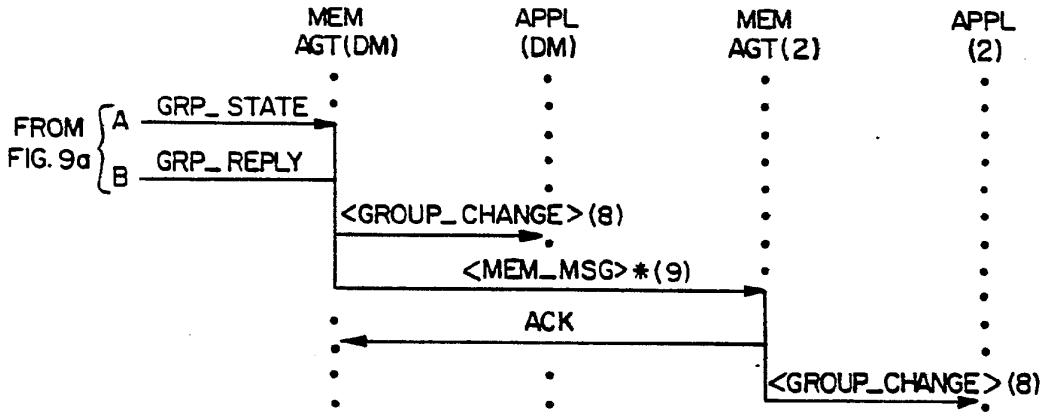

MEM AGT(DM)
APPL (DM)
MEM AGT(2)
APPL (2)
FROM FIG. 9a
A
B
GRP_ STATE
GRP_ REPLY
<GROUP_ CHANGE> (8)
<MEM_MSG>* (9)
ACK
<GROUP_CHANGE> (8)

FIG. 9b

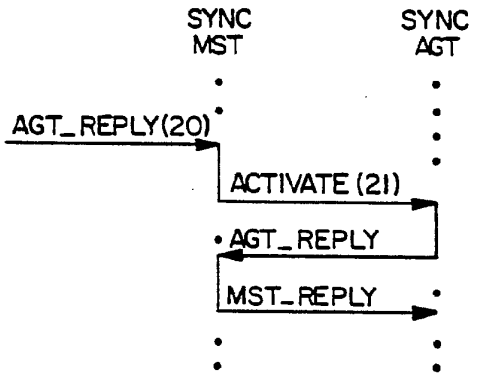

SYNC MST
SYNC AGT
AGT_REPLY(20)
ACTIVATE (21)
AGT_REPLY
MST_REPLY

FIG. 11

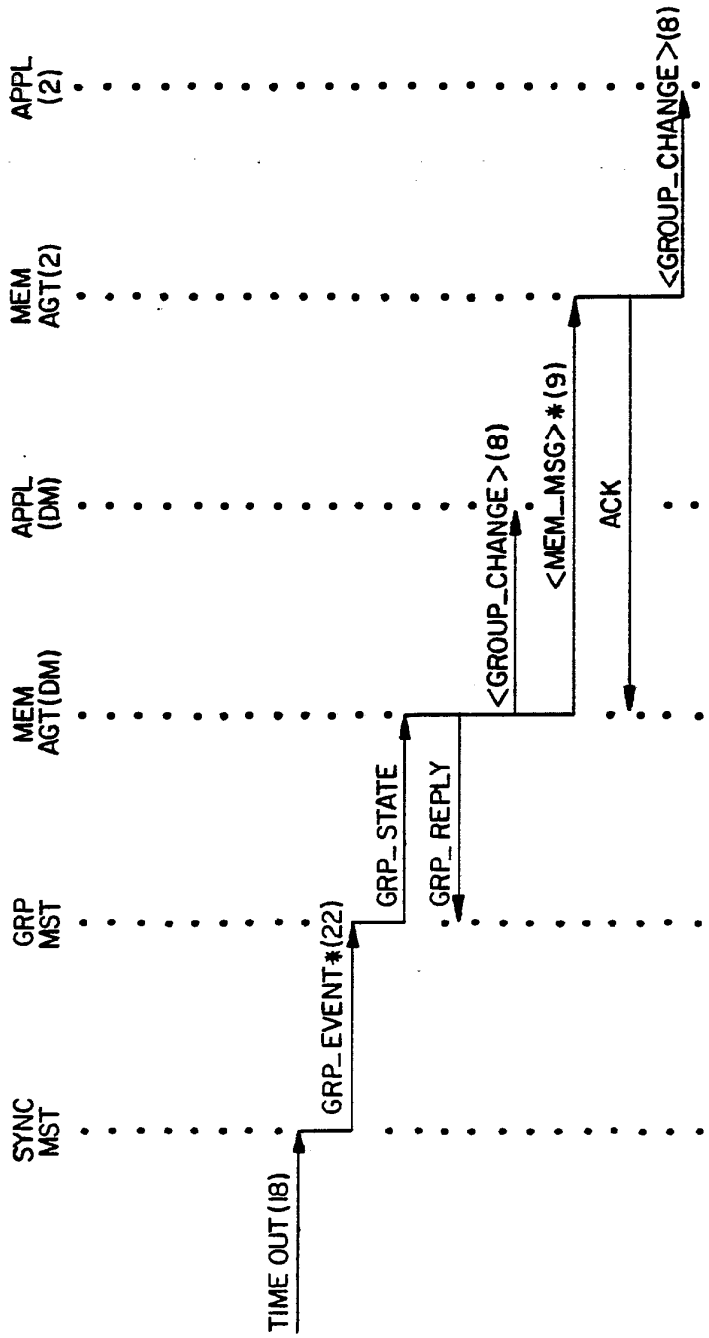

SYNC MST
GRP MST
MEM AGT(DM)
APPL (DM)
MEM AGT(2)
APPL (2)
TIME OUT(18)
GRP_EVENT*(22)
GRP_STATE
GRP_REPLY
<GROUP_CHANGE>(8)
<MEM_MSG>*(9)
ACK
<GROUP_CHANGE>(8)

FIG. IO

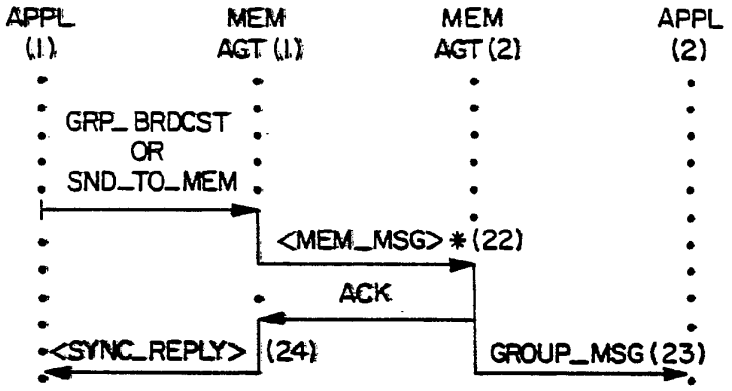

APPL (1)
MEM AGT (1)
MEM AGT (2)
APPL (2)
GRP_BRDCST OR SND_TO_MEM
<MEM_MSG> * (22)
ACK
<SYNC_REPLY> (24)
GROUP_MSG (23)

FIG. 12

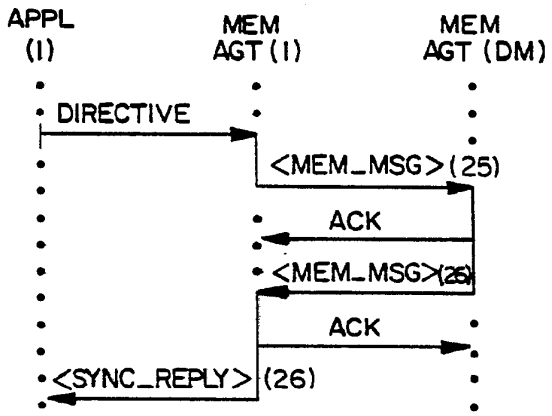

APPL (1)
MEM AGT (1)
MEM AGT (DM)
DIRECTIVE
<MEM_MSG> (25)
ACK
<MEM_MSG> (26)
ACK
<SYNC_REPLY> (26)

FIG. 13

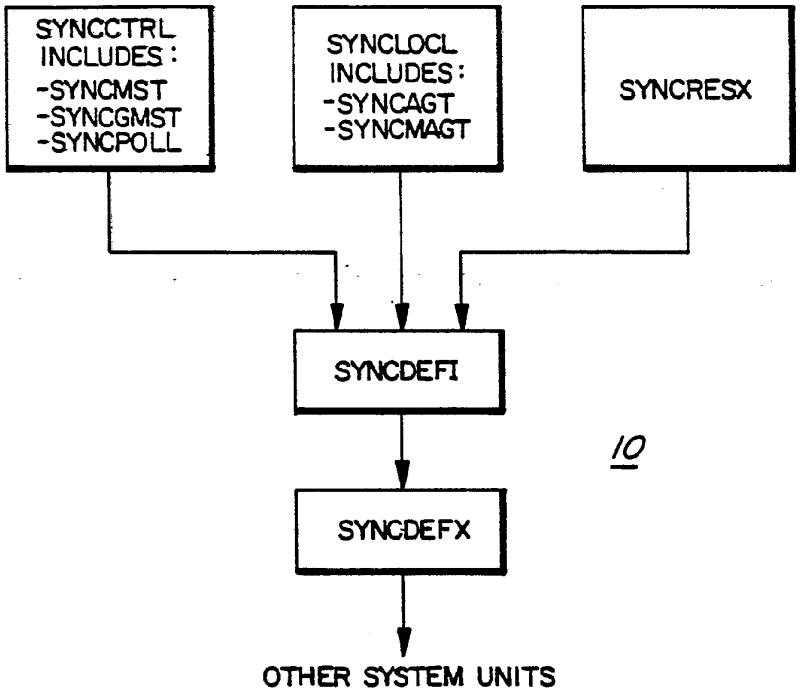

SYNCCTRL
INCLUDES:
-SYNCMST
-SYNCGMST
-SYNCPOLL
SYNCLOCL
INCLUDES:
-SYNCAGT
-SYNCMAGT
SYNCRESX
SYNCDEFI
10
SYNCDEFX
OTHER SYSTEM UNITS

FIG. 14

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Publication number : **0 672 985 A1**

(12) **EUROPEAN PATENT APPLICATION**

(21) Application number : **95300753.1**

(22) Date of filing : **07.02.95**

(51) Int. Cl.[6] : **G06F 11/14**, G06F 11/20

(30) Priority : **22.02.94 US 199444**

(43) Date of publication of application :
20.09.95 Bulletin 95/38

(84) Designated Contracting States :
DE FR GB

(71) Applicant : **International Business Machines Corporation**
Old Orchard Road
Armonk, N.Y. 10504 (US)

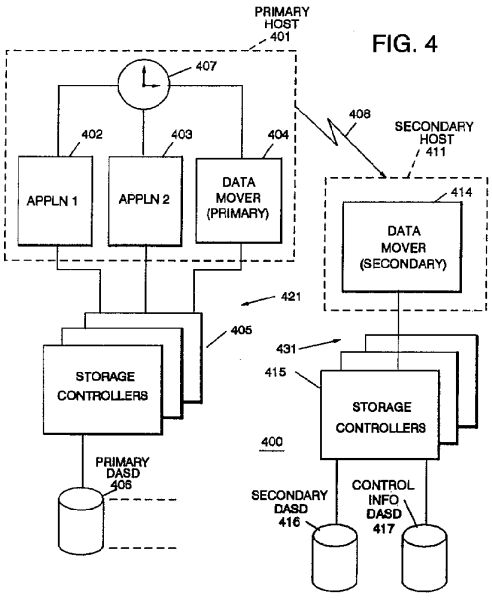
(72) Inventor : **Kern, Robert Frederic**
8338 East Calexico Street
Tucson, AZ 85730 (US)
Inventor : **Kern, Ronald Maynard**
761 North Colette Place
Tucson, AZ 85748 (US)
Inventor : **McBride, Gregory Edward**
8622 East Fairmount Place
Tucson, AZ 85715 (US)
Inventor : **Shackelford, David Michael**
1348 West Thurber Place Drive
Tucson, AZ 85705 (US)

(74) Representative : **Litherland, David Peter**
IBM United Kingdom Limited
Intellectual Property Department
Hursley Park
Winchester, Hampshire SO21 2JN (GB)

(54) **Asynchronous remote data duplexing.**

(57) A remote data shadowing system provides storage based, real time disaster recovery capability. Record updates at a primary site cause write I/O operations in a storage subsystem therein. The write I/O operations are time stamped and the time, sequence, and physical locations of the record updates are collected in a primary data mover. The primary data mover groups sets of the record updates and associated control information based upon a predetermined time interval, the primary data mover appending a prefix header to the record updates thereby forming self describing record sets. The self describing record sets are transmitted to a remote secondary site wherein consistency groups are formed such that the record updates are ordered so that the record updates can be shadowed in an order consistent with the order the record updates cause write I/O operations at the primary site.

FIG. 4



EP 0 672 985 A1

The present invention relates generally to disaster recovery techniques in data processing systems, and more particularly, to a system for real-time remote copying of data.

Data processing systems, in conjunction with processing data, typically are required to store large amounts of data (or records), which data can be efficiently accessed, modified, and re-stored. Data storage is typically separated into several different levels, or hierarchically, in order to provide efficient and cost effective data storage. A first, or highest level of data storage involves electronic memory, usually dynamic or static random access memory (DRAM or SRAM). Electronic memories take the form of semiconductor integrated circuits wherein millions of bytes of data can be stored on each circuit, with access to such bytes of data measured in nano-seconds. The electronic memory provides the fastest access to data since access is entirely electronic.

A second level of data storage usually involves direct access storage devices (DASD). DASD storage, for example, can comprise magnetic and/or optical disks, which store bits of data as micrometer sized magnetic or optical altered spots on a disk surface for representing the "ones" and "zeros" that make up those bits of the data. Magnetic DASD, includes one or more disks that are coated with remnant magnetic material. The disks are rotatably mounted within a protected environment. Each disk is divided into many concentric tracks, or closely spaced circles. The data is stored serially, bit by bit, along each track. An access mechanism, known as a head disk assembly (HDA), typically includes one or more read/write heads, and is provided in each DASD for moving across the tracks to transfer the data to and from the surface of the disks as the disks are rotated past the read/write heads. DASDs can store giga-bytes of data with the access to such data typically measured in milli-seconds (orders of magnitudes slower than electronic memory). Access to data stored on DASD is slower due to the need to physically position the disk and HDA to the desired data storage location.

A third or lower level of data storage includes tape and/or tape and DASD libraries. At this storage level, access to data is much slower in a library since a robot is necessary to select and load the needed data storage medium. The advantage is reduced cost for very large data storage capabilities, for example, terabytes of data storage. Tape storage is often used for back-up purposes, that is, data stored at the second level of the hierarchy is reproduced for safe keeping on magnetic tape. Access to data stored on tape and/or in a library is presently of the order of seconds.

Having a back-up data copy is mandatory for many businesses as data loss could be catastrophic to the business. The time required to recover data lost at the primary storage level is also an important recovery consideration. An improvement in speed over tape or library back-up, includes dual copy. An example of dual copy involves providing additional DASD's so that data is written to the additional DASDs (sometimes referred to as mirroring). Then if the primary DASDs fail, the secondary DASDs can be depended upon for data. A drawback to this approach is that the number of required DASDs is doubled.

Another data back-up alternative that overcomes the need to provide double the storage devices involves writing data to a redundant array of inexpensive devices (RAID) configuration. In this instance, the data is written such that the data is apportioned amongst many DASDs. If a single DASD fails, then the lost data can be recovered by using the remaining data and error correction procedures. Currently there are several different RAID configurations available.

The aforementioned back-up solutions are generally sufficient to recover data in the event that a storage device or medium fails. These back-up methods are useful only for device failures since the secondary data is a mirror of the primary data, that is, the secondary data has the same volume serial numbers (VOLSERs) and DASD addresses as the primary data. System failure recovery, on the other hand, is not available using mirrored secondary data. Hence still further protection is required for recovering data if a disaster occurs destroying the entire system or even the site, for example, earthquakes, fires, explosions, hurricanes, etc. Disaster recovery requires that the secondary copy of data be stored at a location remote from the primary data. A known method of providing disaster protection is to back-up data to tape, on a daily or weekly basis, etc. The tape is then picked up by a vehicle and taken to a secure storage area usually some kilometers away from the primary data location. A problem is presented in this back-up plan in that it could take days to retrieve the back-up data, and meanwhile several hours or even days of data could be lost, or worst, the storage location could be destroyed by the same disaster. A somewhat improved back-up method would be to transmit data to a back-up location each night. This allows the data to be stored at a more remote location. Again, some data may be lost between back-ups since back-up does not occur continuously, as in the dual copy solution. Hence, a substantial data amount could be lost which may be unacceptable to some users.

More recently introduced data disaster recovery solutions include remote dual copy wherein data is backed-up not only remotely, but also continuously. In order to communicate duplexed data from one host processor to another host processor, or from one storage controller to another storage controller, or some combination thereof, a substantial amount of control data is required for realizing the process. A high overhead, however, can interfere with a secondary site's ability to keep up with a primary site's processing, thus threatening the ability of the secondary site to be

able to recover the primary in the event a disaster occurs.

Accordingly it is desired to provide a method and apparatus for providing a real time update of data consistent with the data at a primary processing location using minimal control data, wherein the method and apparatus operates independently of a particular application data being recovered, that is, generic storage media based rather than specific application data based.

An aim of the present invention is to provide an improved system and method for shadowing DASD data to a secondary site for disaster recovery.

According to a first aspect of the present invention, a method for forming consistency groups provides for disaster recovery capability from a remote site. Data updates generated by one or more applications running in a primary processor are received by a primary storage subsystem, wherein the primary storage subsystem causes I/O write operations to write each data update therein. The primary storage subsystem is synchronized by a common timer, and a secondary system, remote from the primary processor, shadows the data updates in sequence consistent order such that the secondary site is available for disaster recovery purposes. Thee method comprising steps of: (a) time stamping each write I/O operation occurring in the primary storage subsystem; (b) capturing write I/O operation record set information from the primary storage subsystem for each data update; (c) generating self describing record sets from the data updates and the respective record set information, such that the self describing record sets are sufficient to re-create a sequence of the write I/O operations; (d) grouping the self describing record sets into interval groups based upon a predetermined interval threshold; and (e) selecting a first consistency group as that interval group of self describing record sets having an earliest operational time stamp, the individual data updates being ordered within the first consistency group based upon time sequences of the I/O write operations in the primary storage subsystem.

In another aspect of the present invention, a primary system has a primary processor running one or more applications, wherein the applications generating record updates, and the primary processor generating self describing record sets therefrom. Each self describing record set is sent to a secondary system remote from the primary system, wherein the secondary system shadows the record updates in sequence consistent order based upon the self describing record sets for real time disaster recovery purposes. The primary processor is coupled to a primary storage subsystem wherein the primary storage subsystem receives the record updates and causes I/O write operations for storing each record update therein. The primary processor comprises a sysplex timer for providing a common time source to the applications and to the primary storage subsystem for synchronization purposes, and a primary data mover, synchronized by the sysplex timer, prompts the primary storage subsystem for providing record set information to the primary data mover for each record update. The primary data mover groups a plurality of record updates and each corresponding record set information into time interval groups, and inserts a prefix header thereto. Each time interval group forms the self describing record sets.

A preferred embodiment of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

FIG. 1 is a block diagram of a disaster recovery system having synchronous remote data shadowing capabilities.

FIG. 2 is a flow diagram of a method for providing synchronous remote copy according to the disaster recovery system of FIG. 1.

FIG. 3 is a flow diagram of a method of an I/O error recovery program (I/O ERP) operation.

FIG. 4 is a block diagram of a disaster recovery system having asynchronous remote data shadowing capabilities.

FIG. 5 is a data format diagram showing a prefix header that prefixes read record sets from the primary site of FIG. 4.

FIG. 6 is a data format diagram describing fields making up a read record set.

FIG. 7 is a state table identifying volume configuration information.

FIG. 8 is a master journal as used by the secondary site of FIG. 4.

FIG. 9 is an example sequence for forming a consistency group.

FIG. 10 is a flow diagram showing a method of collecting information and read record sets for forming consistency groups.

FIG. 11 is a flow diagram showing a method of forming consistency groups.

FIG. 12 is a table indicating full consistency group recovery rules application for an ECKD architecture device for a given sequence of I/O operations to a DASD track.

FIG. 13 is a description of the rules to be used in the table of FIG. 12.

FIG. 14 is a flow diagram of a method of writing read record set copies to a secondary site with full consistency group recovery capability.

A typical data processing system may take the form of a host processor, such as an IBM System/360 or IBM System/370 processor for computing and manipulating data, and running, for example, data facility storage management subsystem/multiple virtual systems (DFSMS/MVS) software, having at least one IBM 3990 storage controller attached thereto, the storage controller comprising a memory controller

and one or more cache memory types incorporated therein. The storage controller is further connected to a group of direct access storage devices (DASDs) such as IBM 3380 or 3390 DASDs. While the host processor provides substantial computing power, the storage controller provides the necessary functions to efficiently transfer, stage/destage, convert and generally access large databases.

Disaster recovery protection for the typical data processing system requires that primary data stored on primary DASDs be backed-up at a secondary or remote location. The distance separating the primary and secondary locations depends upon the level of risk acceptable to the user, and can vary from several kilometers to thousands of kilometers. The secondary or remote location, in addition to providing a back-up data copy, must also have enough system information to take over processing for the primary system should the primary system become disabled. This is due in part because a single storage controller does not write data to both primary and secondary DASD strings at the primary and secondary sites. Instead, the primary data is stored on a primary DASD string attached to a primary storage controller while the secondary data is stored on a secondary DASD string attached to a secondary storage controller.

The secondary site must not only be sufficiently remote from the primary site, but must also be able to back-up primary data in real time. The secondary site needs to back-up primary data as the primary data is updated with some minimal delay. Additionally, the secondary site has to back-up the primary data regardless of the application program (e.g., IMS, DB2) running at the primary site and generating the data and/or updates. A difficult task required of the secondary site is that the secondary data must be order consistent, that is, secondary data is copied in the same sequential order as the primary data (sequential consistency) which requires substantial systems considerations. Sequential consistency is complicated by the existence of multiple storage controllers each controlling multiple DASDs in a data processing system. Without sequential consistency, secondary data inconsistent with primary data would result, thus corrupting disaster recovery.

Remote data duplexing falls into two general categories, synchronous and asynchronous. Synchronous remote copy involves sending primary data to the secondary location and confirming the reception of such data before ending a primary DASD input/output (I/O) operation (providing a channel end (CE) and device end (DE) to the primary host). Synchronous copy, therefore, slows the primary DASD I/O response time while waiting for secondary confirmation. Primary I/O response delay is increased proportionately with the distance between the primary and secondary systems - a factor that limits the remote distance to tens of kilometers. Synchronous copy, however, provides sequentially consistent data at the secondary site with relatively little system overhead.

Asynchronous remote copy provides better primary application system performance because the primary DASD I/O operation is completed (providing a channel end (CE) and device end (DE) to the primary host) before data is confirmed at the secondary site. Therefore, the primary DASD I/O response time is not dependent upon the distance to the secondary site and the secondary site could be thousands of kilometers remote from the primary site. A greater amount of system overhead is required, however, for ensuring data sequence consistency since data received at the secondary site will often not be in order of the primary updates. A failure at the primary site could result in some data being lost that was in transit between the primary and secondary locations.

Synchronous real time remote copy for disaster recovery requires that copied DASD volumes form a set. Forming such a set further requires that a sufficient amount of system information be provided to the secondary site for identifying those volumes (VOLSERs) comprising each set and the primary site equivalents. Importantly, the secondary site forms a "duplex pair" with the primary site and the secondary site must recognize when one or more volumes are out of sync with the set, that is, "failed duplex" has occurred. Connect failures are more visible in synchronous remote copy than in asynchronous remote copy because the primary DASD I/O is delayed while alternate paths are retried. The primary site can abort or suspend copy to allow the primary site to continue while updates for the secondary site are queued, the primary site marking such updates to show the secondary site is out of sync. Recognizing exception conditions that may cause the secondary site to fall out of sync with the primary site is needed in order that the secondary site be available at any time for disaster recovery. Error conditions and recovery actions must not make the secondary site inconsistent with the primary site.

Maintaining a connection between the secondary site and the primary site with secondary DASD present and accessible, however, does not ensure content synchronism. The secondary site may lose synchronism with the primary site for a number of reasons. The secondary site is initially out of sync when the duplex pair is being formed and reaches sync when an initial data copy is completed. The primary site may break the duplex pair if the primary site is unable to write updated data to the secondary site in which case the primary site writes updates to the primary DASD under suspended duplex pair conditions so that the updating application can continue. The primary site is thus running exposed, that is, without current disaster protection copy until the duplex pair is restored. Upon restoring the duplex pair, the secondary site is not immediately in sync. After applying now pending up-

dates, the secondary site returns to sync. The primary site can also cause the secondary site to lose sync by issuing a suspend command for that volume to the primary DASD. The secondary site re-syncs with the primary site after the suspend command is ended, duplex pair is re-established, and pending updates are copied. On-line maintenance can also cause synchronization to be lost.

When a secondary volume is out of sync with a primary volume, the secondary volume is not useable for secondary system recovery and resumption of primary applications. An out-of-sync volume at the secondary site must be identified as such and secondary site recovery-takeover procedures need to identify the out-of-sync volumes for denying application access (forcing the volumes off-line or changing their VOLSERs). The secondary site may be called upon to recover the primary site at any instant wherein the primary site host is inaccessible - thus the secondary site requires all pertinent information about a sync state of all volumes. The secondary storage subsystem, that is the secondary storage controllers and DASD, is unable to determine all conditions causing the primary site to break synchronism due to primary site-encountered exceptions. For example, the primary site may break a duplex pair if the primary site is unable to access the secondary peer due to a primary I/O path or link failure that the secondary site is unaware of. In this case the secondary site shows in-sync state while the primary site indicates the duplex pair is broken.

External communication may notify the secondary site that an out-of-sync duplex pair volume exists. This is realizable by employing a user systems management function. Primary I/O operations end with channel end/device end/unit check (CE/DE/UC) status and sense data indicates the nature of the error. With this form of I/O configuration an error recovery program (ERP) processes the error and send an appropriate message to the secondary processor before posting the primary application that I/O is complete. The user is then responsible to recognize the ERP suspend duplex pair message and secure that information at the secondary location. When the secondary site is depended upon to become operational in place of the primary site, a start-up procedure brings the secondary DASD on-line to the secondary host wherein sync status stored in the secondary DASD subsystem is retrieved for ensuring that out-of-sync volumes are not brought on-line for application allocation. This sync status merged with all ERP suspend duplex pair messages gives a complete picture of the secondary out-of-sync volumes.

Referring now to FIG. 1, a disaster recovery system 10 is shown having a primary site 14 and a secondary site 15, wherein the secondary site 15 is located, for example, 20 kilo-meters remote from the primary site 14. The primary site 14 includes a host processor or primary processor 1 having an application and system I/O and Error Recovery Program 2 running therein (hereinafter referred to as I/O ERP 2). The primary processor 1 could be, for example, an IBM Enterprise Systems/9000 (ES/9000) processor running DFSMS/MVS operating software and further may have several application programs running thereon. A primary storage controller 3, for example, an IBM 3990 Model 6 storage controller, is connected to the primary processor 1 via a channel 12. As is known in the art, several such primary storage controllers 3 can be connected to the primary processor 1, or alternately, several primary processors 1 can be attached to the primary storage controllers 3. A primary DASD 4, for example, an IBM 3390 DASD, is connected to the primary storage controller 3. Several primary DASDs 4 can be connected to the primary storage controller 3. The primary storage controller 3 and attached primary DASD 4 form a primary substorage system. Further, the primary storage controller 3 and the primary DASD 4 could be single integral units.

The secondary site 15 includes a secondary processor 5, for example, an IBM ES/9000, connected to a secondary storage controller 6, for example an IBM 3990 Model 3, via a channel 13. A DASD 7 is further connected to the secondary storage controller 6. The primary processor 1 is connected to the secondary processor 5 by at least one host-to-host communication link 11, for example, channel links or telephone T1/T3 line links, etc. The primary processor 1 may also have direct connectivity with the secondary storage controller 6 by, for example, multiple Enterprise Systems Connection (ESCON) links 9. As a result, the I/O ERP 2 can communicate, if required, with the secondary storage controller 6. The primary storage controller 3 communicates with the secondary storage controller 6 via multiple peer-to-peer links 8, for example, multiple ESCON links.

When a write I/O operation is executed by an application program running in the primary processor 1, a hardware status channel end/device end (CE/DE) is provided indicating the I/O operation completed successfully. Primary processor 1 operating system software marks the application write I/O successful upon successful completion of the I/O operation, thus permitting the application program to continue to a next write I/O operation which may be dependent upon the first or previous write I/O operation having successfully completed. On the other hand, if the write I/O operation was unsuccessful, the I/O status of channel end/device end/unit check (hereinafter referred to as CE/DE/UC) is presented to the primary processor 1 operating system software. Having presented unit check, the I/O ERP 2 takes control obtaining specific sense information from the primary storage controller 3 regarding the nature of the failed write I/O operation. If a unique error to a volume oc-

curs then a unique status related to that error is provided to the I/O ERP 2. The I/O ERP 2 can thereafter perform new peer-to-peer synchronization error recovery for maintaining data integrity between the primary storage controller 3 and the secondary storage controller 6, or in the worst case, between the primary processor 1 and the secondary processor 5.

Referring to FIGs 2 and 3, the error recovery procedure is set forth. In FIG. 2, a step 201 includes an application program running in the primary processor 1 sending a data update to the primary storage controller 3. At step 203 the data update is written to the primary DASD 4, and the data update is shadowed to the secondary storage controller 6. At step 205 the duplex pair status is checked to determine whether the primary and secondary sites are synchronized. If the duplex pair status is in a synchronized state, then the data update is written to the secondary DASD 7 at step 207 while processing then continues at the primary processor 1 via application programs running thereat.

In the case that the duplex pair is in a "failed" state, then at step 209 the primary storage controller 3 notifies the primary processor 1 that duplex pair has suspended or failed. The duplex pair can become "failed" due to communication failure between the primary storage controller 3 and the secondary storage controller 6 via communication links 8. Alternatively, duplex pair can become "failed" due to errors in either the primary or secondary subsystem. If the failure is in the communication links 8, then the primary storage controller 3 is unable to communicate the failure directly to the secondary storage controller 6. At step 211 the primary storage controller 3 returns I/O status CE/DE/UC to the primary processor 1. The I/O ERP 2 quiesces the application programs hence taking control of the primary processor 1 at step 213 for error recovery and data integrity before returning control to the application requesting the write I/O operation.

FIG. 3 represents steps performed by the I/O ERP 2. The I/O ERP 2 issues a sense I/O to the primary storage controller 3 at step 221. The sense I/O operation returns information describing the cause of the I/O error, that is, the data description information is unique to the storage controllers or duplex pair operation regarding specific errors. In the event that the data description information indicates that the peer-to-peer communication links 8 have failed between the primary storage controller 1 and the secondary storage controller 6, then at step 223 the I/O ERP 2 issues a storage controller level I/O operation against the primary storage controller 3 and the secondary storage controller 6 indicating that the affected volume is to be placed in failed synchronous remote copy state. This secondary storage controller 6 is able to receive the state of the affected volume from the I/O ERP 2 via the - multiple ESCON links 9 or the host-to-host communication link 11. Consequently, the current status of the duplex pair operation is maintained at both the primary processor 1 and the secondary processor 5 in conjunction with applications running in the primary processor 1. Consoles 18 and 19 are provided for communicating information from the primary processor 1 and secondary processor 4, respectively, wherein the I/O ERP posts status information to both consoles 18 and 19.

Data integrity has been maintained at step 225 upon successful completion of the failed synchronous remote copy I/O operation to the primary storage controller 3 and the secondary storage controller 6. Therefore, if a recovery is attempted at the secondary site 15 the secondary storage controller 6 identifies the volume marked "failed synchronous remote copy" as not being useable until data on that volume are synchronized with other data in that synchronization group by data recovery means (conventional data base logs and/or journals for determining the state of that data on the volume).

Step 227 tests to determine whether the I/O ERP 2 received successful completion of the I/O operations at the primary storage controller 3 and the secondary storage controller 6 on the failed synchronous remote copy status update. Upon successful completion, the I/O ERP 2 returns control to the Primary processor 1 at step 229. Otherwise step 231 performs a next level recovery notification which involves notifying an operator, via the console 18, of the failed volume and that a status of that volume at either the primary storage controller 3 or the secondary storage controller 6 may not be correct. The notification is shadowed to the secondary site 15, via the console 19 or a shared DASD data set, for indicating the specific volume status there.

An error log recording data set is updated at step 233. This update is written to either the primary DASD 4 or some other storage location and is shadowed to the secondary site 15. Having completed the error recovery actions, the I/O ERP 2, at step 235, posts to the primary application write I/O operation a "permanent error" for causing the primary application an error normal "permanent error" recovery for the failed write I/O operation. Once the error is corrected, the volume states can be recovered, first to pending (recopy changed data) and then back to full duplex. The data may later be re-applied to the secondary DASD 7 once duplex pair is re-established.

When establishing a duplex pair a volume can be identified as CRITICAL according to a customer's needs. For a CRITICAL volume, when an operation results in failing a duplex pair, a permanent error failure of the primary volume is reported irrespective of the actual error's location. With CRIT=Y, all subsequent attempts to write to the primary DASD 405 of the failed pair will receive a permanent error, ensuring that no data is written to that primary volume that can-

not also be shadowed to the paired secondary volume. This permits complete synchronization with the primary application actions and the I/O data operations when required.

Consequently, the disaster recovery system 10 described herein, introduces outboard synchronous remote copy such that a primary host process error recovery procedure having an I/O order (channel command word (CCW)) may change a status of a primary and secondary synchronous remote copy volume from duplex pair to failed duplex thereby maintaining data integrity for several types of primary and secondary subsystem errors. Storage based back-up, rather than application based back-up, wherein data updates are duplicated in real time has been provided. The disaster recovery system 10 also attempts several levels of primary/secondary status updates, including: (1) primary and secondary storage controller volume status updates; (2) primary and secondary host processor notification on specific volume update status via operator messages or error log recording common data sets; and (3) CRITICAL volume indication, future updates to the primary volume can be prevented if the volume pair goes failed duplex. Hence, real time, full error disaster recovery is accomplished.

Asynchronous remote data shadowing is used when it is necessary to further increase a distance between primary and secondary sites for reducing the probability that a single disaster will corrupt both primary and secondary sites, or when primary application performance impact needs to be minimized. While the distance between primary and secondary sites can now stretch across the earth or beyond, the synchronization of write updates across multiple DASD volumes behind multiple primary subsystems to multiple secondary subsystems is substantially more complicated. Record write updates can be shipped from a primary storage controller via a primary data mover to a secondary data mover for shadowing on a secondary storage subsystem, but the amount of control data passed therebetween must be minimized while still being able to re-construct an exact order of the record write updates on the secondary system across several storage controllers as occurred on the primary system across multiple DASD volumes behind several storage controllers.

FIG. 4 depicts an asynchronous disaster recovery system 400 including a primary site 421 and a remote or secondary site 431. The primary site 421 includes a primary processor 401, for example, an IBM ES/9000 running DFSMS/MVS host software. The primary processor 401 further includes application programs 402 and 403, for example, IMS and DB2 applications, and a primary data mover (PDM) 404. A common sysplex clock 407 is included in the primary processor 401 for providing a common reference to all applications (402, 403) running therein, wherein all system clocks or time sources (not shown) synchronize to the sysplex clock 407 ensuring all time dependent processes are properly timed relative to one another. The primary storage controllers 406, for example, synchronize to a resolution appropriate to ensure differentiation between record write update times, such that no two consecutive write I/O operations to a single primary storage controller 404 can exhibit the same time stamp value. The resolution, and not the accuracy, of the sysplex timer 407 is critical. The PDM 404, though shown connected to the sysplex timer 407, is not required to synchronize to the sysplex timer 407 since write I/O operations are not generated therein. A sysplex timer 407 is not required if the primary processor 401 has a single time reference (for example, a single multi-processor ES/9000 system).

A plurality of primary storage controllers 405, for example, IBM 3990 Model 6 storage controllers, are connected to the primary processor 401 via a plurality of channels, for example, fiber optic channels. Connected to each primary storage controller 405 is at least one string of primary DASDs 406, for example, IBM 3390 DASDs. The primary storage controllers 405 and the primary DASDs 406 form a primary storage subsystem. Each storage controller 405 and primary DASD 406 need not be separate units, but may be combined into a single drawer.

The secondary site 431, located for example, some thousands of kilometers remote from the primary site 421, similar to the primary site 421, includes a secondary processor 411 having a secondary data mover (SDM) 414 operating therein. Alternatively, the primary and secondary sites can be the same location, and further, the primary and secondary data movers can reside on a single host processor (secondary DASDs may be just over a fire-wall). A plurality of secondary storage controllers 415 are connected to the secondary processor 411 via channels, for example, fiber optic channels, as is known in the art. Connected to the storage controllers 415 are a plurality of secondary DASDs 416 and a control information DASD(s) 417. The storage controllers 415 and DASDs 416 and 417 comprise a secondary storage subsystem.

The primary site 421 communicates with the secondary site 431 via a communication link 408. More specifically, the primary processor 401 transfers data and control information to the secondary processor 411 by a communications protocol, for example, a virtual telecommunications access method (VTAM) communication link 408. The communication link 408 can be realized by several suitable communication methods, including telephone (T1, T3 lines), radio, radio/telephone, microwave, satellite, etc.

The asynchronous data shadowing system 400 encompasses collecting control data from the primary storage controllers 405 so that an order of all

data writes to the primary DASDs 406 is preserved and applied to the secondary DASDs 416 (preserving the data write order across all primary storage subsystems). The data and control information transmitted to the secondary site 431, must be sufficient such that the presence of the primary site 421 is no longer required to preserve data integrity.

The applications 402, 403 generate data or record updates, which record updates are collected by the primary storage controllers 405 and read by the PDM 404. The primary storage controllers 405 each grouped its respective record updates for an asynchronous remote data shadowing session and provides those record updates to the PDM 404 via non-specific primary DASD 406 READ requests. Transferring record updates from the primary storage controllers 405 to the PDM 404 is controlled and optimized by the PDM 404 for minimizing a number of START I/O operations and time delay between each read, yet maximizing an amount of data transferred between each primary storage controller 405 and the primary processor 401. The PDM 404 can vary a time interval between non-specific READs to control this primary storage controller-host optimization as well as a currency of the record updates for the secondary DASDs 416.

Collecting record updates by the PDM 404, and transmitting those record updates to the SDM 414, while maintaining data integrity, requires the record updates to be transmitted for specific time intervals and in appropriate multiple time intervals with enough control data to reconstruct the primary DASDs 406 record WRITE sequence across all primary storage subsystems to the secondary DASDs 416. Re-constructing the primary DASDs 406 record WRITE sequences is accomplished by passing self describing records from the PDM 404 to the SDM 414. The SDM 414 inspects the self describing records for determining whether any records for a given time interval have been lost or are incomplete.

FIGs 5 and 6 show a journal record format created by the PDM 404 for each self describing record, including a prefix header 500 (FIG. 5), and a record set information 600 (FIG. 6) as generated by the primary storage controller 405. Each self describing record is further journaled by the SDM 414 for each time interval so that each self describing record can be applied in time sequence for each time interval to the secondary DASDs 416.

Referring now to FIG. 5, the prefix header 500, which is inserted at the front of each record set, includes a total data length 501 for describing the total length of the prefix header 500 and actual primary record set information 600 that is transmitted to the SDM 414 for each record set. An operational time stamp 502 is a time stamp indicating a start time for the operational set that the PDM 404 is currently processing. The operational time stamp 502 is generated by the PDM 404 (according to the sysplex timer 407) when performing a READ RECORD SET function to a set of the primary storage controllers 405. An I/O time 610 (FIG. 6) of the primary DASDs 406 write is unique for each primary storage controller 405 READ RECORD SET. The operational time stamp 502 is common across all storage controllers.

A READ RECORD SET command is issued by the PDM 404 and can be predicated upon one of the following conditions:

(1) Primary storage controller 405 attention interrupt based upon that primary storage controller predetermined threshold;

(2) Primary processor 401 timer interrupt based upon a predetermined time interval; or

(3) Record set information indicates additional information on outstanding record sets available but not yet read.

Condition (2) uses a timer interval to control how far behind the secondary system 431 executes during periods of low activity. Condition (3) occurs when the PDM 404 fails to drain all record sets during a processing interval which drives further activity for ensuring that the PDM 404 keeps up with primary storage controller 405 activity.

A time interval group number 503 is supplied by the PDM 404 to identify a time interval (bounded by operational time stamp 502 and a records read time 507) for which the current record sets belong (sets of records across all primary storage controllers 405 for a given time interval group form consistency groups). A sequence number within group 504 is derived based upon a hardware provided identification (to the PDM 404) of a write sequence order of application WRITE I/Os for primary storage controller 405 for each record set within a given time interval group 503. A primary SSID (substorage identification) 505 uniquely identifies the specific primary storage controller of the primary storage controllers 405 for each record set. A secondary target volume 506 is assigned by either the PDM 404 or the SDM 414 depending upon performance considerations. A records read time 507 supplies an operational time stamp that is common to all primary storage controllers 405 indicating an end time for the PDM 404 read record set process current interval.

The operational time stamp 502 and the records read time 507 are used by the PDM 404 to group sets of read record sets from each of the primary storage controllers 405. Time synchronization for grouping sets of read record sets is key only to the PDM 404 and as such, the PDM 404 could be synchronized to a central processing unit (CPU) clock running only the PDM 404 not attached to the syspelx timer 407. The PDM 404 does not write record updates, but the record updates, as stated previously, must be synchronized to a common time source.

Referring now to FIG. 6, the record set informa-

tion 600 is generated by the primary storage controllers 405 and collected by the PDM 404. Update Specific Information 601-610, includes a primary device unit address 601 of each record indicating the actual primary DASD 406 that the record update occurred on. A cylinder number/head number (CCHH) 602 indicates a location on primary DASD 406 for each record update. Primary SSID 603, the primary storage controller session identifier is the same as primary SSID 505. Status flags 604 provide status information regarding whether specific data records 620 follow. Sequence numbers 605 and 630 assign a number to each record for indicating whether the entire record set has been read (all data transferred to the PDM 404). Primary DASD write I/O type 606 is an operation indicator identifying the type of write operation performed on each record, the operation indicators including: update write; format write; partial track records follow; full track data follows; erase command performed; or write any performed. Search argument 607 indicates initial positioning information for the first read record set data record 620. A sector number 608 identifies that sector that the record was updated at. Count field 609 describes a number of specific record data fields 620 that follow. A host application time when the primary DASD 406 write update occurred is recorded in time of updates 610. Specific record data 620 provides a count/key/data (CKD) field of each record update. Lastly, the sequence number 630 is compared to the sequence number 605 for indicating whether the entire read record set was transferred to the PDM 404.

The update records are handled in software groups called consistency groups so that the SDM 414 can copy the record updates in the same order they were written at the primary DASDs 406. The information used for creating the consistency groups (across all record sets collected from all storage controllers 405) includes the: operational time stamp 502; time interval group number 503; sequence number within group 504; primary controller SSID 505; records read time 507; primary device address 601; the primary SSID 603; and the status flags 604. The information used for determining whether all records for a time interval group have been received for each storage controller 405 at the SDM 414 includes the: time interval group number 503; sequence number within group 504; physical controller ID 505; the primary SSID 603; and a total number of read record sets returned from each primary storage controller 405 for each operational time interval. The information necessary to place record updates on the secondary DASDs 416 equivalently to the primary DASDs 406 record updates with full recover possible includes the: secondary target volume 506; CCHH 602; primary DASD write I/O type 606; search argument 607; sector number 608; count 609; time of updates 610; and the specific record data 620.

FIGs 7 and 8 show a state table 700 and a master journal 800, respectively, for describing a current journal contents, which simplifies recovery time and journal transfer time. The state table 700 provides configuration information, collected by and common to the PDM 404 and SDM 414, and includes primary storage controller session identifiers (SSID numbers) and the volumes therein, and the corresponding secondary storage controller session identifiers and the corresponding volumes. Thus the configuration information tracks which primary volumes 710 or primary DASD extents map to secondary volumes 711 or secondary DASD extents. With a simple extension to the state table 700 indicating partial volume extents 712 (CCHH to CCHH), partial volume remote copy can be accomplished using the same asynchronous remote copy methods described herein, but for a finer granularity (track or extent) than full volume.

The master journal 800 includes: consistency group number; location on journal volumes; and operational time stamp. The master journal 800 further maintains specific record updates as grouped in consistency groups. The state table 700 and master journal 800 support disaster recovery, and hence must be able to operate in a stand-alone environment wherein the primary system 401 no longer exists.

A time stamp control is placed at the front and back of each master journal 800 to ensure that the entire control entry was successfully written. The time stamp control is further written to the secondary DASDs 417. The control elements include dual entries (1) and (2), wherein one entry is always a current entry, for example:

(1) Timestamp control | Control Info | Timestamp Control

(2) Timestamp Control | Control Info | Timestamp Control.

At any point in time either entry (1) or (2) is the current or valid entry, wherein a valid entry is that entry with equal timestamp controls at the front and back. Disaster recovery uses the valid entry with the latest timestamp to obtain control information. This control information, along with state information (environmental information regarding storage controllers, devices, and applied consistency groups), is used for determining what record updates have been applied to the secondary storage controllers 415.

After all read record sets across all primary storage controllers 405 for a predetermined time interval are received at the secondary site 431, the SDM 414 interprets the received control information and applies the received read record sets to the secondary DASDs 416 in groups of record updates such that the record updates are applied in the same sequence that those record updates were originally written on the primary DASDs 406. Thus, all primary application order (data integrity) consistency is maintained at the secondary site 431. This process is hereinafter refer-

red to as forming consistency groups. Forming consistency groups is based on the following assumptions: (A) application writes that are independent can be performed in any order if they do not violate controller sequence order; (B) application writes that are dependent must be performed in timestamp order, hence an application cannot perform a dependent write number two before receiving control unit end, device end from write number one; and (C) a second write will always be either (1) in a same record set consistency group as a first write with a later timestamp or (2) in a subsequent record set consistency group.

Referring to FIG. 9, an example of forming a consistency group (the consistency group could be formed at either the primary site 421 or secondary site 431), for example, for storage controllers SSID 1, SSID 2, and SSID 3 is shown (any number of storage controllers can be included but three are used in this example for clarity). Time intervals T1, T2 and T3 are assumed to occur in ascending order. An operational time stamp 502 of time interval T1 is established for storage controllers SSID 1, SSID 2 and SSID 3. The PDM 404 obtains record set data from storage controller SSIDs 1, 2, and 3 for time interval T1-T3. The record sets for SSIDs 1, 2, and 3 for time interval T1 are assigned to time interval group 1, G1 (time interval group number 503). The sequence number within group 504 is shown for each SSID 1, 2, and 3, wherein SSID has three updates as 11:59. 12:00, and 12:01, SSID 2 has two updates at 12:00 and 12:02, and SSID 3 has three updates at 11:58, 11:59, and 12:02. Record sets of time intervals T2 and T3 are listed but example times of updates are not given for simplicity.

Consistency group N can now be formed based upon the control information and record updates received at the secondary site 431. In order to ensure that no record update in time interval group number one is later than any record update of time interval group number two, a min-time is established which is equal to a the earliest read record set time of the last record updates for each storage controller SSID 1, 2, and 3. In this example then, min-time is equal to 12:01. Any record updates having a read record set time greater than or equal to min-time is included in the consistency group N + 1. If two record update times to a same volume were equal, though unlikely given sufficient resolution of the sysplex timer 407, the record update having the earlier sequence number within the time interval group N is kept with that group for consistency group N. The record updates are now ordered based upon read record set times. Record updates having equal times will cause the record update having the lower sequence number to be place before the later sequence numbered record update. Alternatively, record updates having equal time stamps, but to differing volumes, may be ordered arbitrarily as long as they are kept in the same consistency group.

If a primary storage controller 405 fails to complete a response to a read record set during a specified time interval, then a consistency group cannot be formed until that primary storage controller 405 completes. In the event that the primary storage controller 405 fails to complete its operation, then a missing interrupt results causing a system missing interrupt handler to receive control and the operation will be terminated. On the other hand, if the primary storage controller 405 timely completes the operation then the I/O will be driven to completion and normal operation will continue. Consistency group formation expects that write operations against the primary storage controllers 405 will have time stamps. Some programs, however, will cause writes to be generated without time stamps, in which case the primary storage controller 405 will return zeros for the time stamp. Consistency group formation can bound those records without time stamps based upon the timestamp that the data was read. If too many record updates without time stamps occur over a time interval such that the record updates are not easily bounded by consistency group times, then an error that the duplex volumes are out of synchronization may result.

FIGs 10 and 11 are flow diagrams presenting the method of forming consistency groups. Referring to FIG. 10, the process starts at step 1000 with the primary site 421 establishing remote data shadowing to occur. At step 1010 all application I/O operations are time stamped using the sysplex timer 407 as a synchronization clock (FIG. 4). The PDM 404 starts a remote data shadowing session with each primary storage controller 405 at step 1020 which includes identifying those primary volumes that will have data or records shadowed. Record set information 600 is trapped by the primary storage controllers 405 for each application WRITE I/O operation (see FIG. 6) by step 1030.

Step 1040 involves the PDM 404 reading the captured record set information 600 from each primary storage controller 405 according to a prompt including an attention message, a predetermined timing interval, or a notification of more records to read as described earlier. When the PDM 404 begins reading record sets, at step 1050, the PDM 404 prefixes each record set with a prefix header 500 (see FIG. 5) for creating specific journal records (a journal record includes the prefix header 500 and the record set information 600). The journal records contain the control information (and records) necessary for forming consistency groups at the secondary site 431 (or at the primary site 421).

At step 1060 the PDM 404 transmits the generated journal records to the SDM 414 via the communications link 408 (or within the same data mover system if the consistency groups are formed therein). The SDM 414 uses the state table 700 at step 1070

to gather the received record updates by group and sequence numbers for each time interval group and primary storage controller 405 established for the data shadowing session. The SDM 414 inspects the journal records at step 1080 to determine whether all record information has been received for each time interval group. If the journal records are incomplete, then step 1085 causes the SDM 414 to notify the PDM 404 to resend the required record sets. If the PDM 404 is unable to correctly resend, then the duplex volume pair is failed. If the journal records are complete, then step 1090 is performed which encompasses the SDM 414 forming the consistency groups.

Referring to FIG. 11, steps 1100-1160 representing step 1090 (FIG. 10) for forming consistency groups is shown. Consistency group formation starts at step 1100 wherein each software consistency group is written to an SDM 414 journal log ("hardened") on the secondary DASD 417 (FIG. 4). Step 1110 performs a test for determining whether the time interval groups are complete, that is, each primary storage controller 405 must have either presented at least one read record set buffer or have confirmation from the PDM 404 that no such record updates were placed in the record set buffer, and all read record set buffers with data (or null) must have been received by the SDM 414. If a time interval group is incomplete, then step 1110 retries reading the record sets from the primary storage controller 405 until the required data is received. If errors occur, a specific duplex volume pair or pairs may be failed. Having received complete time interval groups, step 1120 determines a first consistency group journal record. The first (or current) consistency group journal record is that record which contains the earliest operational time stamp 502 and the earliest time of update 610 of all records having equal operational time stamps 502.

Step 1130 inspects the records contained in the current consistency group journal record to determine which record will be the last record to be included therein (some records will be dropped and included in the next consistency group journal record). The last record in the current consistency group journal record is determined as a minimum update time (min-time) of the maximum update times for each primary storage controller 405 (that is, the last update of each primary storage controller 405 is compared and only the earliest of these remains in the current consistency group journal record).

Those remaining record updates in the current consistency group journal record are ordered according to time of update 610 and sequence number within group 504 by step 1140. A primary storage controller 405 that had no record updates does not participate in the consistency group. At step 1150, the remaining record updates of the current consistency group (having update times later than min-time) are passed to the next consistency group. Each sequence number within a group 504 should end with a null buffer indicating that all read record sets have been read for that operational time interval. If the null buffer is absent, then the step 1120 of defining the last record in the current software consistency group, coupled with the records read time 507 and time of update 610 can be used to determine the proper order of the application WRITE I/O operations across the primary storage controllers 405.

Step 1160 represents a back-end of the remote data shadowing process wherein specific write updates are applied to secondary DASDs 416 under full disaster recovery constraint. If when writing the updates to the secondary DASDs 416 an I/O error occurs, or the entire secondary site 431 goes down and is re-initialized, then the entire consistency group that was in the process of being written can be re-applied from the start. This permits the remote shadowing to occur without having to track which secondary DASDs 416 I/Os have occurred, which I/Os have not occurred, and which I/Os were in process, etc.

A key component of step 1160 is that the PDM 414 causes the records to be written efficiently to the secondary DASDs 416 so that the secondary site 431 can keep pace with the primary site 421. The requisite efficiency is accomplished, in part, by concurrently executing multiple I/O operations to different secondary DASDs 416. Serially writing one secondary DASD 416 at a time would cause the secondary site 431 to fall too far behind the primary site 421. Yet more efficiency is gained at the secondary site 431 by writing the records for each consistency group destined for a single secondary device via a single channel command word (CCW) chain. Within each single CCW chain, the I/O operations therein can be further optimized as long as those I/O operations to each secondary DASD 416 data track are maintained in the order of occurrence on the primary volumes.

Optimizing secondary I/O operations for specific consistency groups and within single CCW chains is based in part upon the pattern of primary write I/O operations, and in part upon the physical characteristics of the secondary DASDs 416. Optimization may vary somewhat depending upon whether secondary DASD 415 is count/key/data (CKD), extended count/key/data (ECKD), fixed block architecture (FBA), etc. Consequently, a number of WRITE I/Os (m) to a primary DASD 406 volume during a given time interval can be reduced to a single START I/O operation to a secondary DASD 416 volume. This optimization of the number of START I/Os to the secondary storage controllers 415 of m:1 can allow the secondary DASDs 416 to catch up with and thereby closer shadow the record updates at the primary site 421.

A key to successful remote data shadowing, and hence secondary I/O optimization, is minimizing unrecoverable errors in any of the concurrent multiple I/O operations to secondary DASDs 416 so that con-

sistent copies are available for recovery. A failure in a given secondary write could permit a later dependent write to be recorded without the conditioning write (e.g., a log entry indicating that a data base record has been updated when in reality the actual update write for the data base had failed violates the sequence integrity of the secondary DASD 416 copy).

A failed secondary 416 copy is unusable for application recovery until that failure to update has been recovered. The failed update could be corrected by having the SDM 414 request a current copy from the PDM 404. In the mean time the secondary data copy is inconsistent and hence unusable until the PDM 404 responds with the current update and all other previous updates are processed by the PDM 414. The time required to recover the failed update typically presents an unacceptably long window of non-recovery for adequate disaster recovery protection.

Effective secondary site 431 I/O optimization is realized by inspecting the data record sets to be written for a given consistency group and building chains based upon rules of the particular secondary DASD 416 architecture, for example, an ECKD architecture. The optimization technique disclosed herein simplifies recovery from I/O errors such that when applying a consistency group, if an I/O error occurs, the CCW chain can be re-executed, or in the event of a secondary initial program load (IPL) recovery, the entire consistency group can be re-applied without data loss.

FIG. 12 summarizes full consistency group recovery (FCGR) rules for building CCW chains for all WRITE I/O combinations for an ECKD architecture, wherein CCHHR record format is used (cylinder number, head number, record number). FIG. 12 is created by inspecting each possible combination of WRITE I/O operations to a DASD track within a consistency group's scope. The FCGR rules of FIG. 12, described in FIG.s 13A and 13B, are then followed to govern data placement (secondary DASD 416 I/O write CCW chains) for yielding full recovery for an error in applying a consistency group. The FCGR rules depicted in FIG. 12 would be extended appropriately as new WRITE I/O operations are added. These rules can exist in hardware or software at the secondary site 431. The FCGR rules advantageously reduces READ record set to a same DASD track analysis to an inspection of the primary DASD 406 WRITE I/O type, search argument, and count and key fields.

If a DASD track is written without inspecting the consistency group write operations as shown in FIG. 12, then previously written data records potentially cannot be re-written. For example, assume that a chain includes:

WRITE UPDATE to record five; and

FORMAT WRITE to record one,

wherein record one and record five occur on the same DASD track with record one preceding record five. Record five is updated by an UPDATE WRITE CCW and a FORMAT WRITE I/O CCW updates record one erasing a remainder of the track thus deleting record five. If this chain had to be re-executed, a LOCATE RECORD CCW that will position to the beginning of record five will no longer have a positioning point (record five no longer exists), and the chain is not fully recoverable from the beginning. Since the write operations have already been successful at the primary site 421, always being able to apply an entire consistency group on the secondary DASD 416 is required to maintain data consistency and integrity.

FIG. 14, steps 1410 through 1470, provides more detail as to the process represented by step 1160 of FIG. 11, while using the FCGR rules defined in FIG. 12. At step 1410 the SDM 414 divides the records of the current consistency group into two categories. A first category includes I/O orders directed to a same secondary DASD volume, and a second category includes I/O orders of those records in the first category that are directed to a same CCHH (i.e., records being updated to a same DASD track).

Having categorized the records of the current consistency group, step 1420 conforms application WRITE I/Os and SDM 414 WRITE I/Os to the architecture of the secondary DASDs 416, for example, to ECKD architecture FCGR rules (see FIG. 12) for identifying data placement on a track and track/record addressing. The SDM 414 groups secondary DASD WRITE I/O operations to the same volumes into single I/O CCW chains at step 1430. Step 1440 involves moving the head disk assembly (HDA) of each secondary DASD 416 according to search arguments and specific record data (CKD fields) for the actual secondary DASD 416 writes.

Step 1450 compares READ SET BUFFERS one and two for those records making up the second categories (there typically will be a plurality of second categories, one for each track receiving records), using the FCGR rules of FIG. 12 for determining whether a subsequent write operation invalidates a previous write operation or DASD search argument (positioning at a record that is now erased, etc.). The READ SET BUFFERS one and two contain adjacent read record sets. Following the FCGR rules ensures that the SDM 414 can re-write an entire consistency group, in the event of an error, without re-receiving record updates from the primary site 421. After the SDM 414 applies the current consistency group to the secondary DASD 416, step 1460 updates the state table (FIG. 7) and the master journal (FIG. 8).

The remote copy process continues in real time as step 1470 gets a next consistency group (which becomes the current consistency group) and returns processing to step 1410. The remote copy process will stop if the primary site 421 to secondary site 431 communication terminates. The communication may terminate if volume pairs are deleted from the process by the PDM 404, the primary site is destroyed

(disaster occurs), an orderly shutdown is performed, or a specific takeover action occurs at the secondary site 431. Consistency groups journaled on the secondary site 431 can be applied to the secondary DASD 416 during a takeover operation. The only data lost is that data captured by the primary site 421 that has not been completely received by the SDM 414.

In summary, synchronous and asynchronous remote data duplexing systems have been described. The asynchronous remote data duplexing system provides storage based, real time data shadowing. A primary site runs applications generating record updates, and a secondary site, remote from the primary site, shadows the record updates and provides disaster recovery for the primary site. The asynchronous remote data duplexing system comprises a sysplex timer for synchronizing time dependent processes in the primary site, and a primary processor at the primary site for running the applications, the primary processor having a primary data mover therein. A plurality of primary storage controllers are coupled to the primary processor for issuing write I/O operations for each record update, each primary storage controller DASD write I/O operation being synchronized to the sysplex timer. A plurality of primary storage devices receive the write I/O operations and store the record updates therein accordingly. The primary data mover collects record set information from the plurality of primary storage controllers for each record update and appends a prefix header to a predetermined group of record set informations. The prefix header and predetermined group of record set informations form the self describing record sets. Each record set information includes a primary device address, a cylinder number and head number (CCHH), a record update sequence number, a write I/O type, a search argument, a sector number, and a record update time. The prefix header includes a total data length, an operational time stamp, a time interval group number, and a records read time. A secondary processor at the secondary site has a secondary data mover, the secondary data mover receiving the self describing record sets from the primary site. A plurality of secondary storage controllers are coupled to the secondary processor, and a plurality of secondary storage devices are coupled to the secondary storage controllers for storing the record updates copies. The secondary data mover determines whether the transmitted self describing record sets are complete and forms consistency groups from the self describing record sets and provides the record updates from each consistency group to the plurality of secondary storage controllers for writing to the plurality of secondary storage devices in an order consistent with a sequence that the record updates were written to the plurality of primary storage devices.

While the invention has been particularly shown and described with reference to preferred embodiments thereof, it will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the scope of the invention as defined in the appended claims. For example, the consistency groups have been described as being formed by the secondary data mover based upon received self describing record sets, however, the consistency groups could be formed at the primary site based upon write record sets or elsewhere in the secondary site. The formats the storage devices at the primary and secondary sites need not be identical. For example, CKD records could be converted to fixed block architecture (FBA) type records, etc. Nor are the storage devices meant to be limited to DASD devices.

Claims

1. A method of providing asynchronous data duplexing, wherein data updates generated by one or more applications running in a primary processor are received by a primary storage subsystem, the primary storage subsystem causing I/O write operations to write each data update therein, each write I/O operation being time stamped, the time stamps synchronized by a common timer, and wherein a secondary system, whether local to or remote from the primary processor, shadows the data updates in sequence consistent order such that the secondary site is available for disaster recovery purposes, the method comprising the steps of:

   (a) time stamping each write I/O operation occurring in the primary storage subsystem;

   (b) capturing write I/O operation record set information from the primary storage subsystem for each data update;

   (c) generating self describing record sets from the data updates and respective record set informations, the self describing record sets containing sufficient control information to enable recreation of a sequence of the write I/O operations solely by the secondary system;

   (d) grouping the self describing record sets into interval groups, each interval group being measured from an operational time stamp start time and continuing for a predetermined interval threshold; and

   (e) selecting a current consistency group as that interval group of self describing record sets having an earliest operational time stamp, the individual data updates being ordered within the current consistency group based upon time sequences of the I/O write operations in the primary storage subsystem.

2. The method as claimed in claim 1 wherein the step (b) further includes initiating sessions with the primary storage subsystem based upon the operational time stamps for identifying a starting time for each interval group, each interval group being bounded by consecutive operational time stamps.

3. The method as claimed in claim 1 or claim 2 wherein the step (d) includes adding a prefix header describing each interval group.

4. The method as claimed in any preceding claim further comprising a step (f) transmitting the interval groups of self describing record sets to the secondary site .

5. The method as claimed in claim 4 further comprising a step (g) determining at the secondary site whether each received self describing record set is complete.

6. The method as claimed in claim 5 wherein the step (g) further includes the secondary site requesting the primary site to re-transmit any missing data updates if the secondary site determined a self describing record set is incomplete.

7. The method as claimed in claim 6 further comprising a step (h) determining at the secondary site whether each time interval group is complete.

8. The method as claimed in claim 7 wherein the step (h) further includes the secondary site requesting the primary site to re-send a missing record set if the secondary site determined that an interval group was incomplete.

9. The method as claimed in claim 8 further comprising a step (i) writing the received data updates at the secondary site to the secondary storage subsystem according to the sequence of the corresponding write I/O operations at the primary site as ordered in the consistency groups.

10. A method for providing remote data shadowing for disaster recovery purposes in a data processing system including a primary site having a primary processor running a primary data mover and applications generating record updates, the primary processor coupled to a primary storage subsystem having storage devices for storing the record updates according to write I/O operations issued by the primary processor to the primary storage subsystem, the primary site further including a common system timer for synchronizing time dependent operations in the primary site, the system further including a secondary site having a secondary processor communicating with the primary processor, and a secondary storage subsystem for storing copies of the record updates in sequence consistent order, the method comprising the steps of:
    (a) time stamping each write I/O operation in the primary storage subsystem;
    (b) establishing a session with each storage device in the primary storage subsystem;
    (c) capturing record set information from each storage device in the primary storage subsystem;
    (d) reading record sets and respective record set information in the primary data mover;
    (e) prefixing each record set with a header and creating self describing record sets therefrom;
    (f) transmitting the self describing record sets to the secondary processor in time interval groups according to predetermined time intervals;
    (g) forming consistency groups from the self describing record sets; and
    (h) shadowing the record updates of each consistency group to the secondary storage subsystem in a sequence consistent order.

11. The method as claimed in claim 10 wherein the record sets are transmitted to the secondary processor asynchronously.

12. The method as claimed in claim 10 or claim 11 wherein the step (g) is performed at the secondary site.

13. The method as claimed in any of claims 10 to 12 wherein the step (f) further includes determining at the secondary site whether each received self describing record set is complete.

14. The method as claimed in claim 13 wherein the step (f) further includes requesting the primary site to re-transmit any missing record updates if the primary site determined a received self describing record set is incomplete.

15. The method as claimed in any of claims 10 to 14 further comprising a step (i) determining at the secondary site whether each time interval group is complete.

16. The method as claimed in claim 15 wherein the step (i) further includes requesting the primary site to re-send a missing record set if the secondary site determined that a interval group was incomplete.

17. The method as claimed in any of claims 10 to 16 wherein the step (c) identifies, in the record set information, a physical location on the primary storage devices where each record update is stored.

18. The method as claimed in claim 17 wherein the step (c) identifies, in the record set information, a sequence and time of update of each record update stored to the primary storage devices within the session.

19. The method as claimed in any of claims 10 to 18 wherein the step (e) identifies, in the prefix header, an interval group number for the session and sequence within group for each record update referred to therein.

20. A primary data processing system having a primary processor running one or more applications, the one or more applications generating record updates, the primary processor generating self describing record sets therefrom, the self describing record sets being sent to a secondary system, the secondary system shadowing the record updates in sequence consistent order based upon the self describing record sets for real time disaster recovery purposes, the primary processor being coupled to a primary storage subsystem wherein the primary storage subsystem receives the record updates and executes write I/O operations for storing each record update therein, the primary processor comprising:

a timer for providing a common time source to the applications and primary storage subsystem for synchronization purposes; and

primary data mover means prompting the primary storage subsystem for providing record set information to the primary data mover means for each record update, the primary data mover means grouping a plurality of record updates and each corresponding record set information into time interval groups, and inserting a prefix header thereto, each time interval group the self describing record sets. interval groups, and inserting a prefix header thereto, each time interval group the self describing record sets.

21. The primary system as claimed in claim 20 wherein the primary storage subsystem comprises:

a plurality of primary storage controllers, the plurality of primary storage controllers issuing the write I/O operations; and

a plurality of primary storage devices coupled to the plurality of primary storage controllers.

22. The primary system as claimed in claim 20 or claim 21 wherein the primary data mover means collects record set information for each write I/O operation for each primary storage controller of the plurality of primary storage controllers participating with each time interval group.

23. The primary system as claimed in any of claims 20 to 22 wherein each write I/O operation is time-stamped in the primary processor relative to the sysplex timer, each write I/O operation being issued to a primary storage controller of the plurality of primary storage controllers, each primary storage controller preserving the time-stamp and returning that time-stamp in the corresponding read record set to the primary data mover means.

24. The primary system as claimed in claim 23 wherein each record set information identifies a corresponding record update's physical location on a primary storage device of the plurality of primary storage devices.

25. The primary system as claimed in claim 23 or claim 24 wherein each record set information identifies a corresponding record update's primary subsystem identification, primary device address, cylinder number, and head number.

26. The primary system as claimed in claim 23 wherein the primary data mover means identifies a relative sequence of each write I/O update across all primary storage controllers participating in a time interval group.

27. The primary system as claimed in claim 26 wherein the primary data mover means creates a state table for journaling record updates and cross referencing a storage location of each record update on the primary system and the secondary system.

28. The primary system as claimed in claim 26 wherein the primary data mover means communicates the state table to the secondary system.

29. A remote data shadowing system including a primary site and a secondary site, the secondary site asynchronously shadowing record updates of the primary site in real time for disaster recovery purposes, the record updates generated by applications running at the primary site, the primary site comprising:

a sysplex timer;

a primary processor running the applications generating the record updates and issuing a corresponding write I/O operation for each record update, the primary processor having a pri-

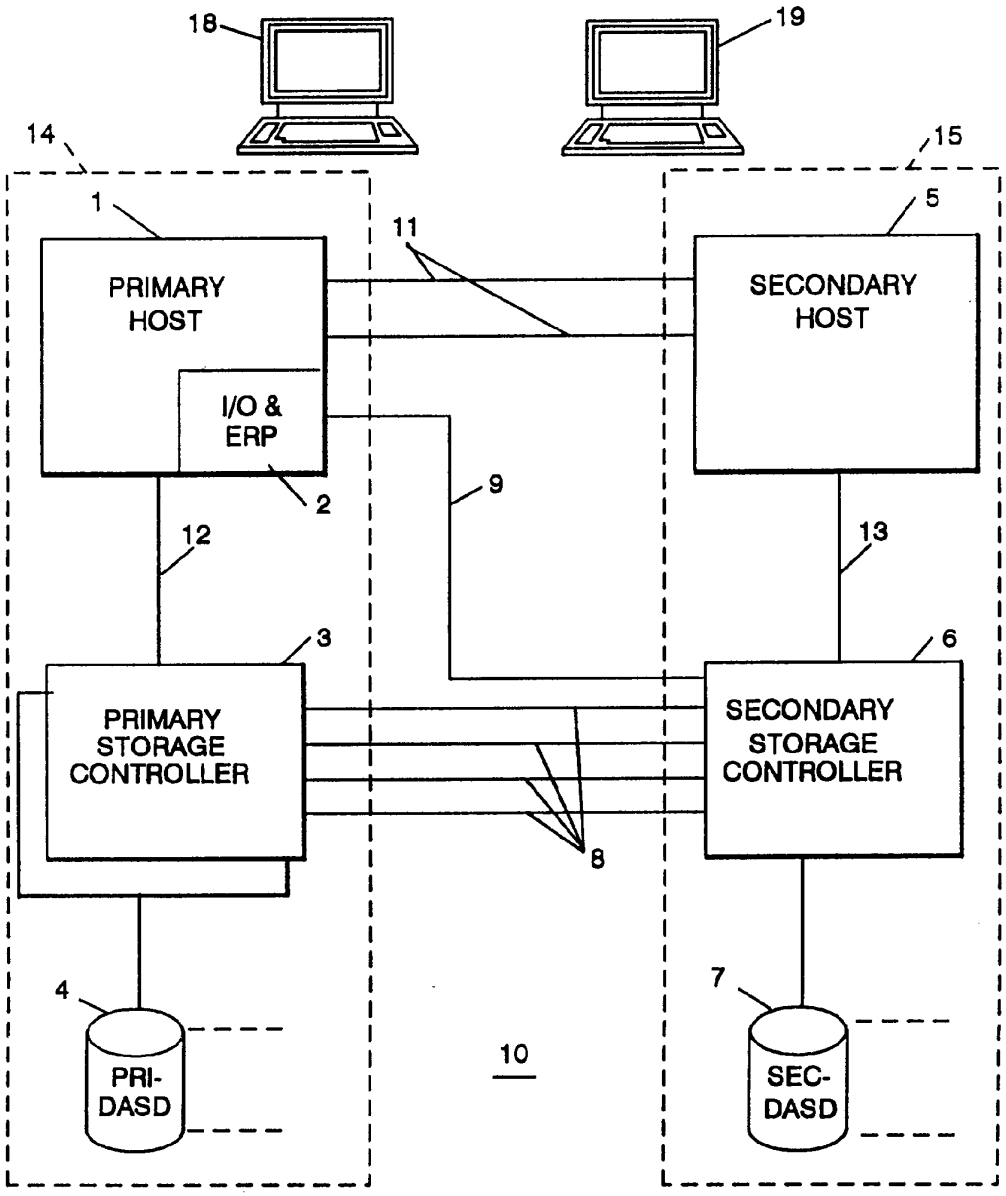
mary data mover therein;

a plurality of primary storage controllers directed to store the record updates, the plurality of primary storage controllers executing the issued write I/O operation for each record update; and

a plurality of primary storage devices receiving and storing the record updates therein according to the corresponding write I/O operations,

wherein the primary processor and each write I/O are time-stamped by the primary processor, as synchronized by the sysplex timer, such that write I/O operations are accurately sequence ordered relative to each other, the primary data mover collecting sets of record updates and combining each record set information as provided by each one of the plurality of primary storage controllers with the corresponding record update, each record set information including a relative sequence and time of each corresponding write I/O operation, the primary data mover collecting record updates into time interval groups and inserting a prefix header to each time interval group, wherein the prefix header includes information identifying the record updates included in each time interval group, each record set information and prefix header combined for creating self describing record sets, the self describing record sets being transmitted to the secondary site, wherein the self describing record sets provide information adequate for the secondary site to shadow the record updates therein in sequence consistent order without further communications from the primary site.

FIG. 1


18
19
14
15
1
5
11
PRIMARY HOST
SECONDARY HOST
I/O & ERP
2
9
12
13
3
6
PRIMARY STORAGE CONTROLLER
SECONDARY STORAGE CONTROLLER
8
4
7
PRI-DASD
10
SEC-DASD

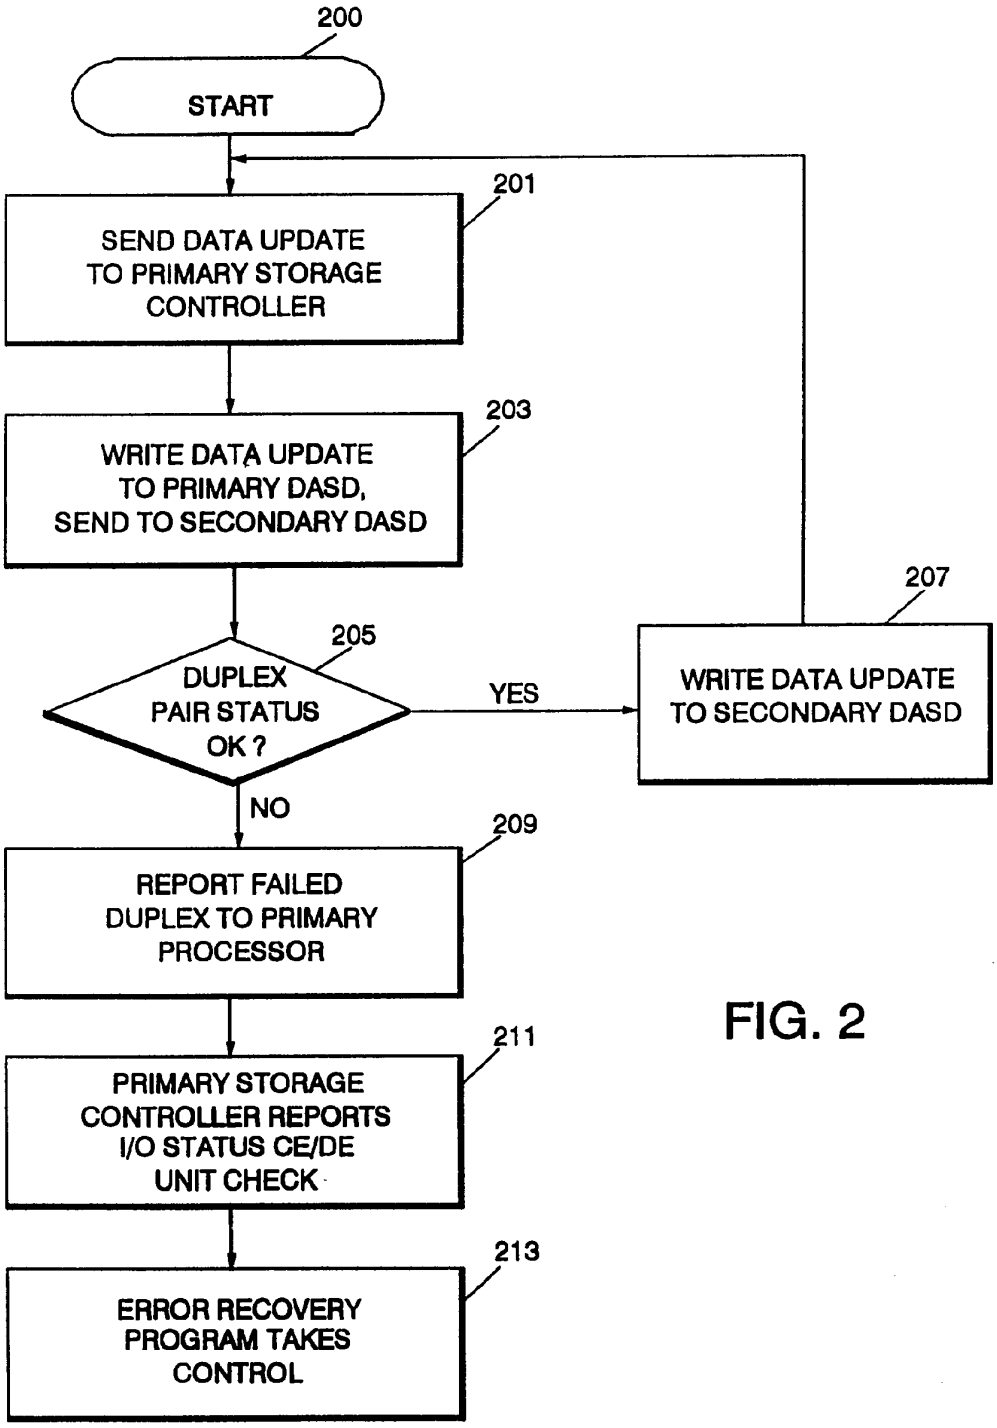

200
START
201
SEND DATA UPDATE TO PRIMARY STORAGE CONTROLLER
203
WRITE DATA UPDATE TO PRIMARY DASD, SEND TO SECONDARY DASD
205
DUPLEX PAIR STATUS OK ?
YES
207
WRITE DATA UPDATE TO SECONDARY DASD
NO
209
REPORT FAILED DUPLEX TO PRIMARY PROCESSOR
211
PRIMARY STORAGE CONTROLLER REPORTS I/O STATUS CE/DE UNIT CHECK
213
ERROR RECOVERY PROGRAM TAKES CONTROL

FIG. 2

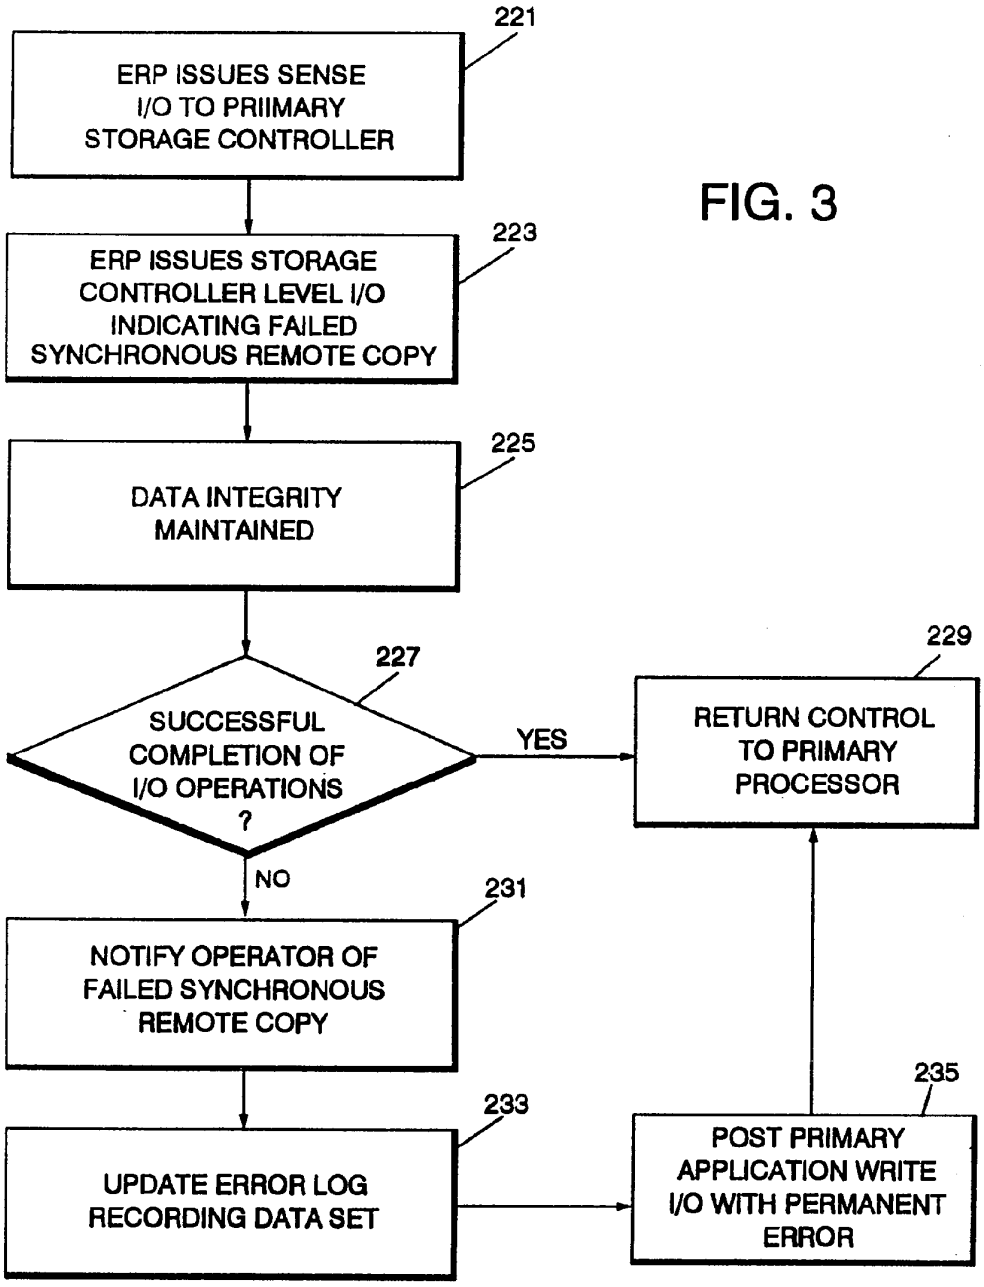

221
ERP ISSUES SENSE I/O TO PRIIMARY STORAGE CONTROLLER
223
ERP ISSUES STORAGE CONTROLLER LEVEL I/O INDICATING FAILED SYNCHRONOUS REMOTE COPY
225
DATA INTEGRITY MAINTAINED
227
SUCCESSFUL COMPLETION OF I/O OPERATIONS ?
YES
229
RETURN CONTROL TO PRIMARY PROCESSOR
NO
231
NOTIFY OPERATOR OF FAILED SYNCHRONOUS REMOTE COPY
233
UPDATE ERROR LOG RECORDING DATA SET
235
POST PRIMARY APPLICATION WRITE I/O WITH PERMANENT ERROR

FIG. 3

FIG. 4

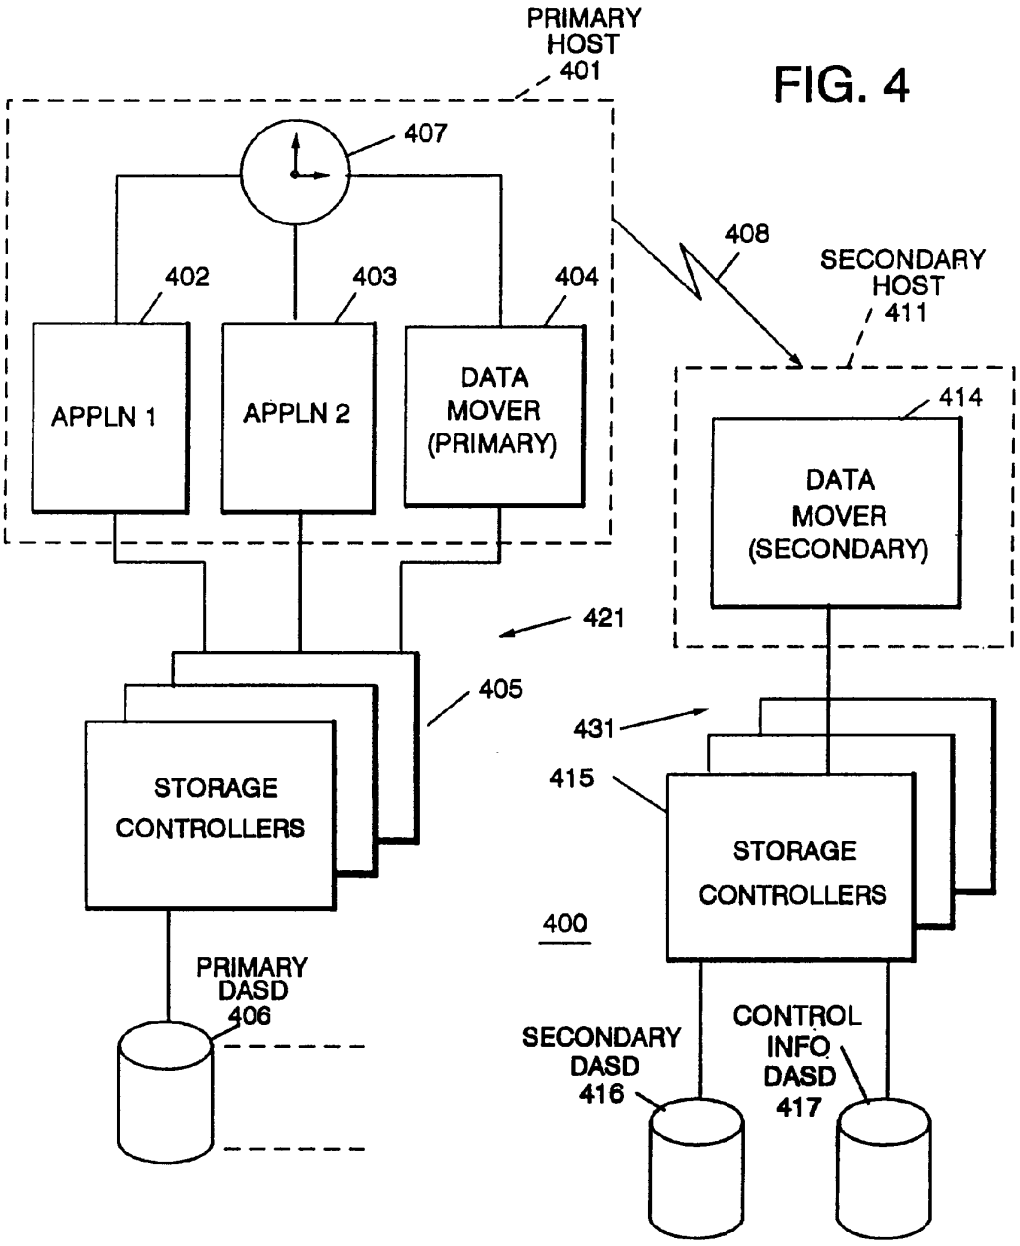
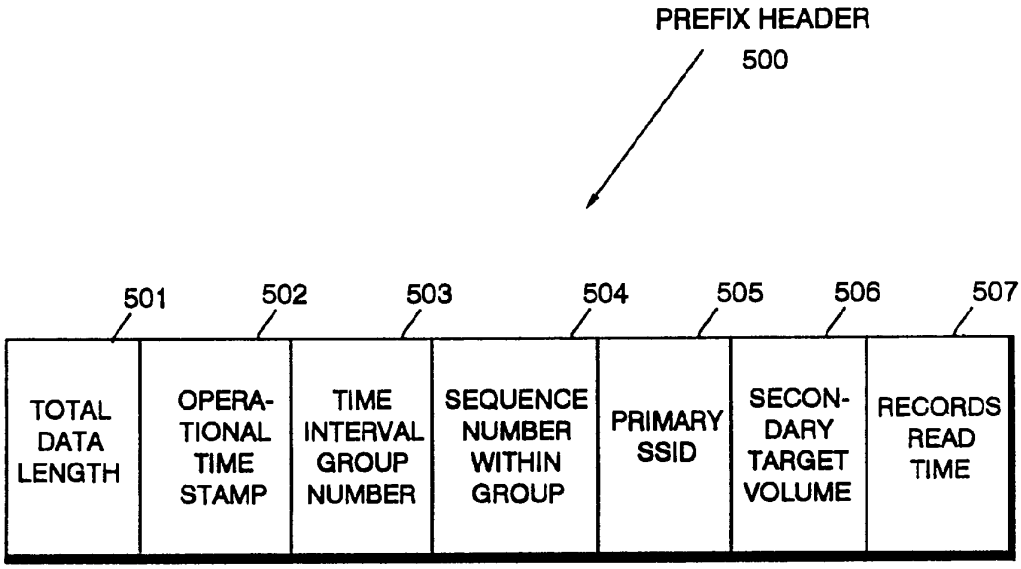

PRIMARY HOST
401
407
408
SECONDARY HOST
411
402
403
404
APPLN 1
APPLN 2
DATA MOVER (PRIMARY)
414
DATA MOVER (SECONDARY)
421
405
431
415
STORAGE CONTROLLERS
STORAGE CONTROLLERS
400
PRIMARY DASD
406
SECONDARY DASD
416
CONTROL INFO DASD
417

FIG. 5


PREFIX HEADER
500
501
502
503
504
505
506
507
TOTAL DATA LENGTH
OPERA-TIONAL TIME STAMP
TIME INTERVAL GROUP NUMBER
SEQUENCE NUMBER WITHIN GROUP
PRIMARY SSID
SECON-DARY TARGET VOLUME
RECORDS READ TIME

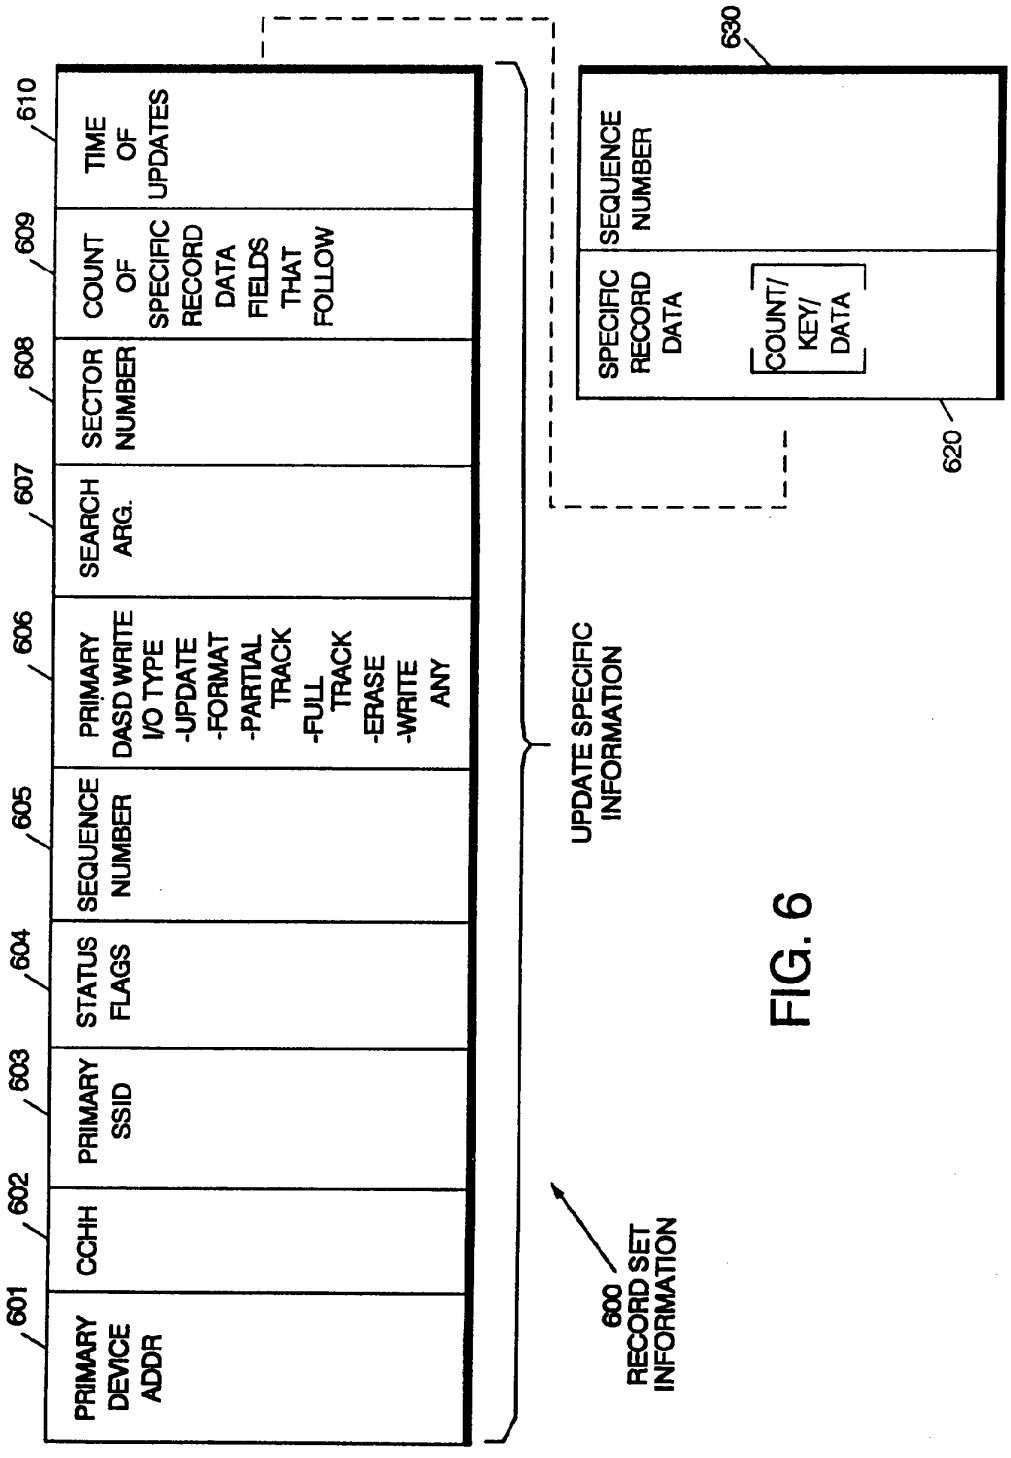
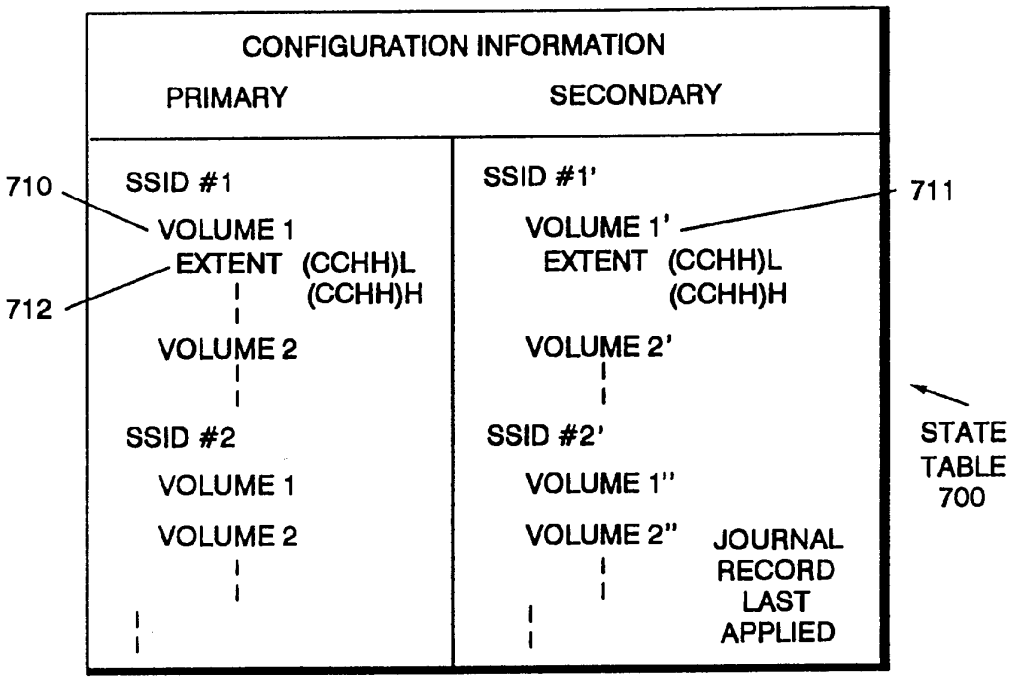

PRIMARY DEVICE ADDR
CCHH
PRIMARY SSID
STATUS FLAGS
SEQUENCE NUMBER
PRIMARY DASD WRITE I/O TYPE -UPDATE -FORMAT -PARTIAL TRACK -FULL TRACK -ERASE -WRITE ANY
SEARCH ARG.
SECTOR NUMBER
COUNT OF SPECIFIC RECORD DATA FIELDS THAT FOLLOW
TIME OF UPDATES
601
602
603
604
605
606
607
608
609
610
600 RECORD SET INFORMATION
UPDATE SPECIFIC INFORMATION
SPECIFIC RECORD DATA
COUNT/ KEY/ DATA
SEQUENCE NUMBER
620
630

FIG. 6

FIG. 7


CONFIGURATION INFORMATION
PRIMARY
SECONDARY
710
SSID #1
VOLUME 1
EXTENT (CCHH)L
(CCHH)H
712
VOLUME 2
SSID #2
VOLUME 1
VOLUME 2
SSID #1'
711
VOLUME 1'
EXTENT (CCHH)L
(CCHH)H
VOLUME 2'
SSID #2'
VOLUME 1''
VOLUME 2''
JOURNAL RECORD LAST APPLIED
STATE TABLE 700

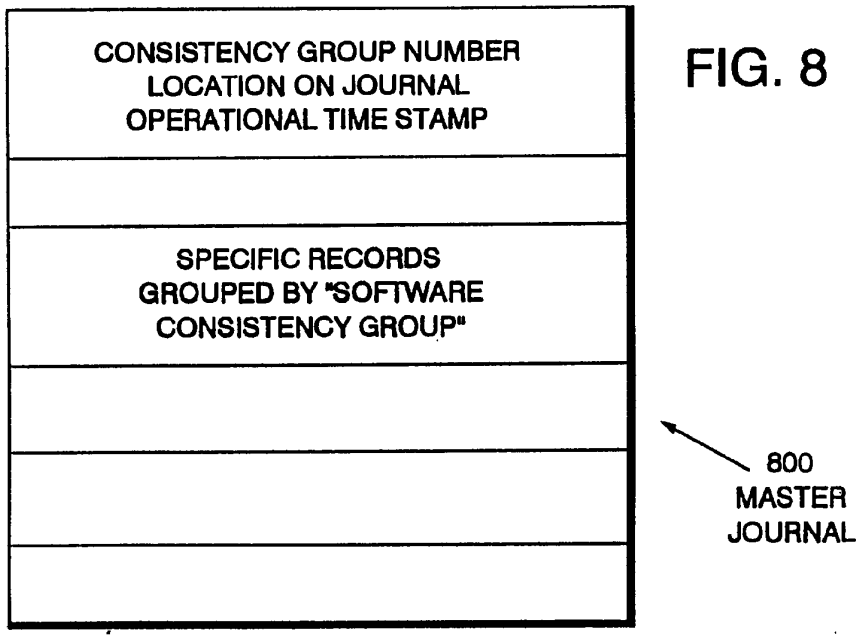

CONSISTENCY GROUP NUMBER
LOCATION ON JOURNAL
OPERATIONAL TIME STAMP
SPECIFIC RECORDS GROUPED BY "SOFTWARE CONSISTENCY GROUP"
800
MASTER JOURNAL

FIG. 8

FIG. 9

PHYSICAL CONTROL-LER ID	OPER-ATIONAL TIME STAMP	TIME INT GROUP #	READ RECORD SET TIME OF UPDATE / CONTROLLER		
			SEQ. 1 OF 3	SEQ. 2 OF 3	SEQ. 3 OF 3
SSID1	T_1	G_1	11:59 (2)	12:00 (5)	12:01 (6)
SSID2	T_1	G_1	12:00 (4)	12:02 * (7)	———
SSID3	T_1	G_1	11:58 (1)	11:59 (3)	12:02 * (8)
SSID1	T_2	G_2			
SSID2	T_2	G_2			
SSID3	T_2	G_2			
	T_3	G_3			

CONSISTENCY GROUP #1

(1) 11:58

(2) 11:59

(3) 11:59

(4) 12:00

(5) 12:00

(6) 12:01

EARLIEST OPERATIONAL TIME T_1, EARLIEST TIME OF UPDATE ACROSS SSID, ORDERED IN READ SEQUENCE

MIN TIME OF THE MAX TIMES OF UPDATE ACROSS ALL SSIDS

CONSISTENCY GROUP #2

(7) - - - - -

(8) - - - - -

FIG. 10

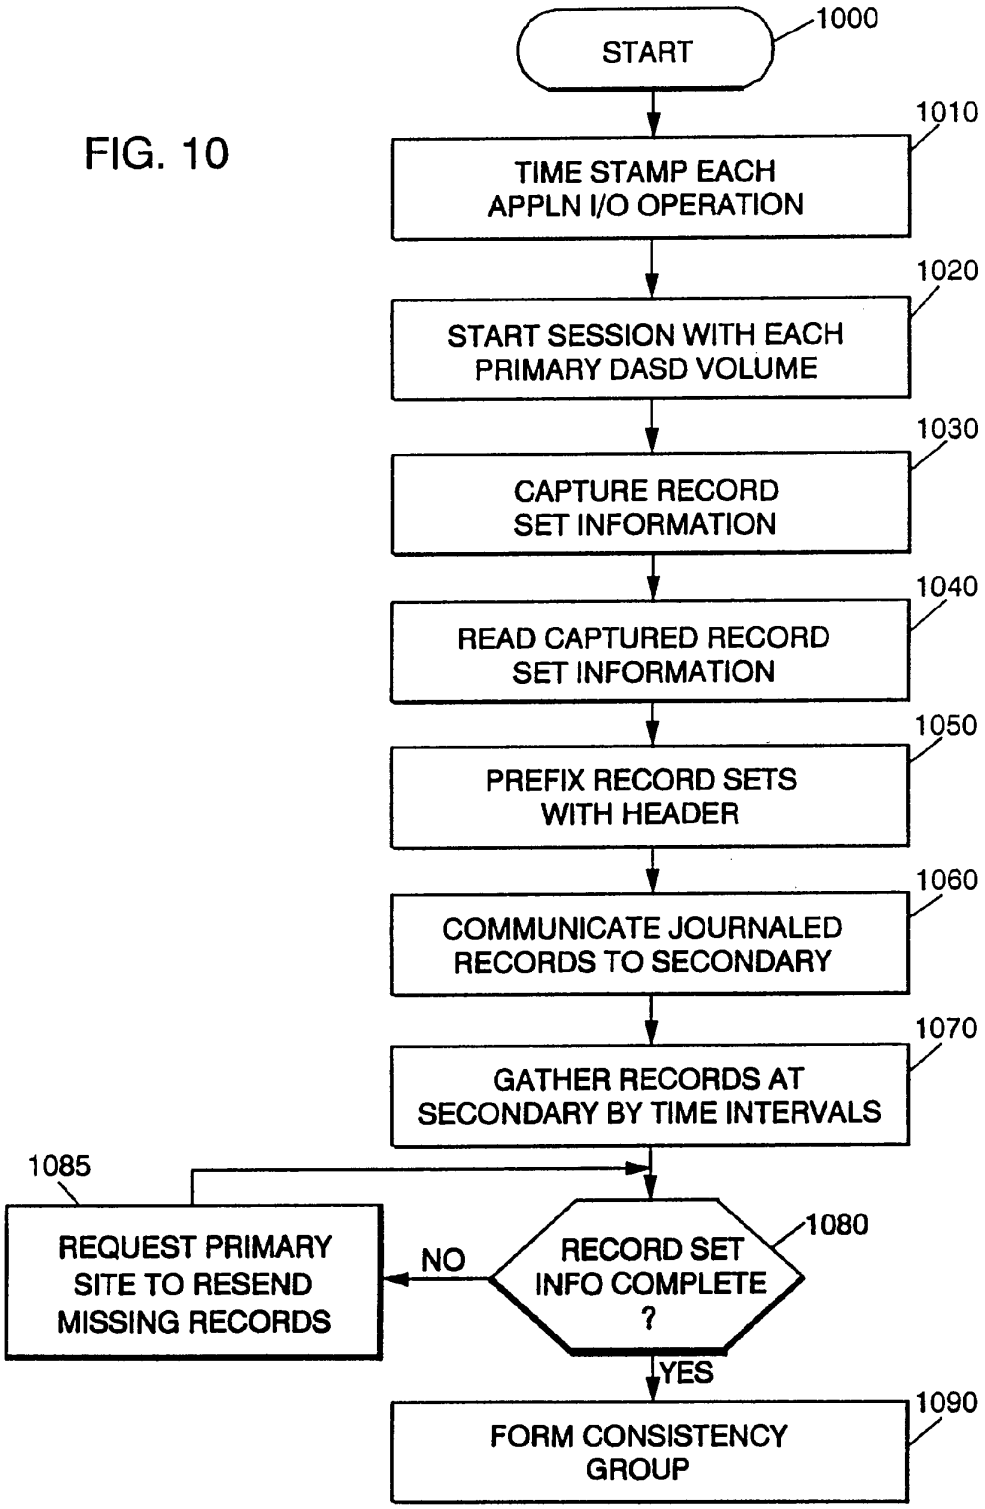

START
1000
TIME STAMP EACH APPLN I/O OPERATION
1010
START SESSION WITH EACH PRIMARY DASD VOLUME
1020
CAPTURE RECORD SET INFORMATION
1030
READ CAPTURED RECORD SET INFORMATION
1040
PREFIX RECORD SETS WITH HEADER
1050
COMMUNICATE JOURNALED RECORDS TO SECONDARY
1060
GATHER RECORDS AT SECONDARY BY TIME INTERVALS
1070
RECORD SET INFO COMPLETE ?
1080
NO
REQUEST PRIMARY SITE TO RESEND MISSING RECORDS
1085
YES
FORM CONSISTENCY GROUP
1090

FIG. 11

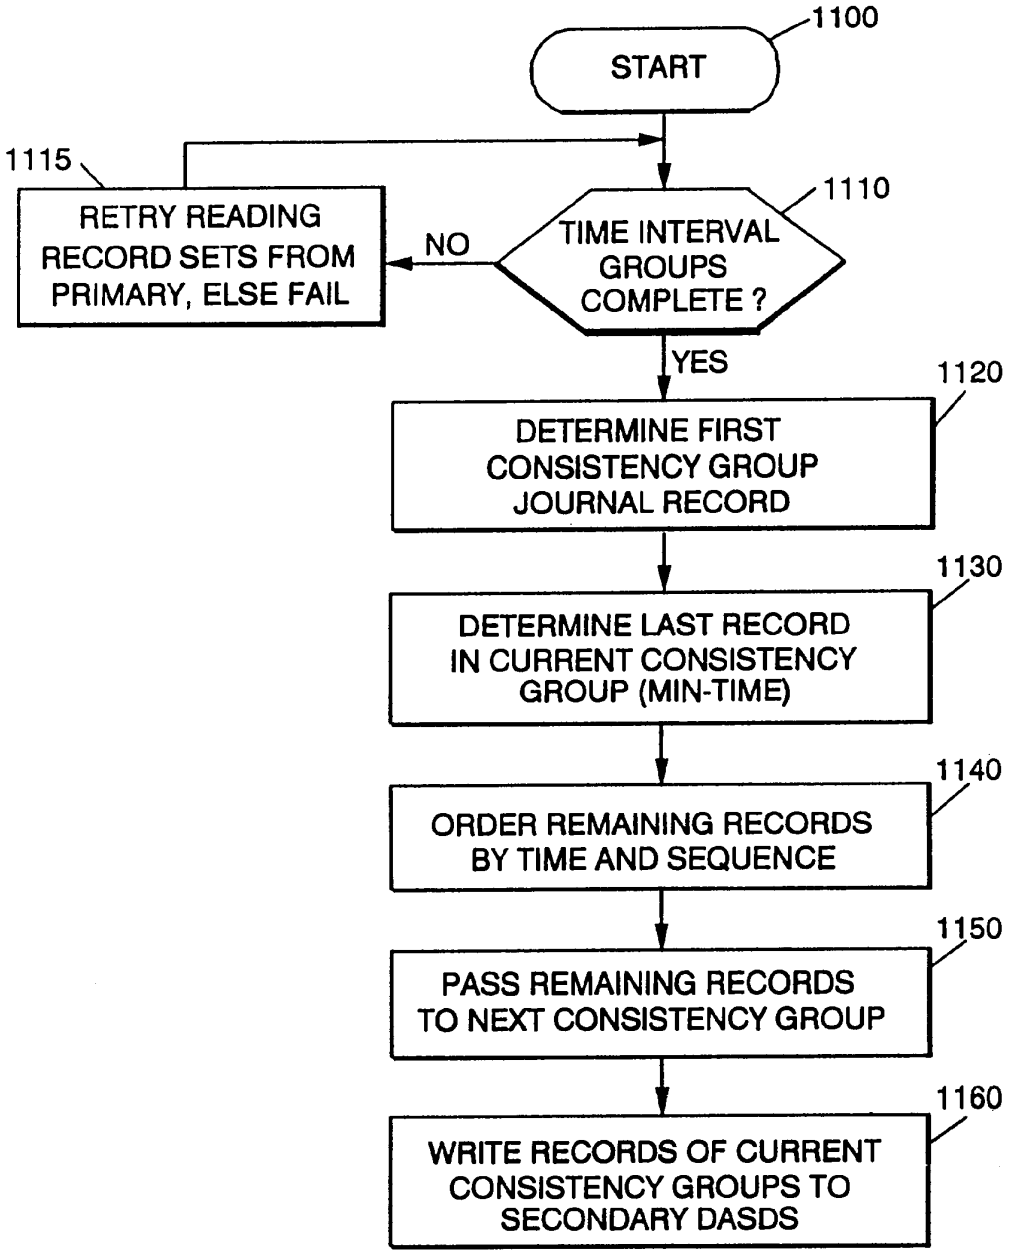

1100
START
1115
RETRY READING RECORD SETS FROM PRIMARY, ELSE FAIL
NO
1110
TIME INTERVAL GROUPS COMPLETE ?
YES
1120
DETERMINE FIRST CONSISTENCY GROUP JOURNAL RECORD
1130
DETERMINE LAST RECORD IN CURRENT CONSISTENCY GROUP (MIN-TIME)
1140
ORDER REMAINING RECORDS BY TIME AND SEQUENCE
1150
PASS REMAINING RECORDS TO NEXT CONSISTENCY GROUP
1160
WRITE RECORDS OF CURRENT CONSISTENCY GROUPS TO SECONDARY DASDS

FIG. 12

FULL CONSISTENCY GROUP RECOVERY RULES

READ RECORD SET BUFFER #2	READ RECORD SET BUFFER #1							
TYPE I/O WRITE OPERATION								
	UPDATE WRITE KL = 0	UPDATE WRITE KL ≠ 0	FORMAT WRITE FULL	FORMAT WRITE PARTIAL	ERASE FULL	ERASE PARTIAL	WRITE ANY KL = 0	WRITE ANY KL ≠ 0
UPDATE WRITE KL = 0	W*	E*	N	J	D	K	W	E*
UPDATE WRITE KL ≠ 0	E*	W*	N	J	D	K	E*	W
FORMAT WRITE FULL	T	T	R	T	R	T	R	R
FORMAT WRITE PARTIAL	C	C	N	H	F	L	W	W
ERASE FULL	T	T	R	R	R	T	T	T
ERASE PARTIAL	B	B	N	M	E	G	W	W
WRITE ANY KL = 0	W	E*	W	W	E	W	W	E*
WRITE ANY KL ≠ 0	E*	W	W	W	E	W	E*	W

FIG. 13

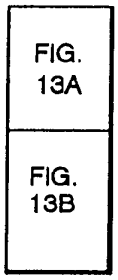


FIG. 13A

B - IF #1'S SEARCH ARG IS HIGHER THAN THE SEARCH ARG FOR #2, THEN THROW #1, ELSE DO BOTH.

C - IF #1'S RECORD IS EQUAL TO OR HIGHER THAN THE FIRST RECORD IN #2, THEN THROW #1, ELSE DO BOTH.

D - IF #2 IS UPDATING R0, THEN DO BOTH, ELSE ERROR.

E - ERROR (SHOULD NEVER HAPPEN).

E*- ERROR IF #1 AND #2 ARE THE SAME RECORD. (SHOULD NEVER HAPPEN WITHOUT A FORMAT WRITE IN BETWEEN).

F - IF FIRST RECORD IN #2 IS R1, THEN WRITE BOTH, ELSE ERROR.

G -IF #1'S SEARCH ARG IS EQUAL TO OR HIGHER THAN THE SEARCH ARG FOR #2, THEN THROW #1, ELSE ERROR.

H -IF #1'S SEARCH ARG IS HIGHER THAN THE LAST RECORD FOR #2. THEN THROW #1. ELSE IF #2'S SEARCH ARG IS HIGHER THAN THE LAST RECORD IN #1, THEN ERROR, ELSE WRITE BOTH.

TO OPTIMIZE FURTHER, CAN DO THE FOLLOWING INSTEAD:

```
IF (1'S SEARCH IS EQUAL TO OR HIGHER THAN THE LAST RECORD
      IN #2) OR (THE LAST RECORD IN #1 IS EQUAL TO OR
      HIGHER THAN THE LAST RECORD IN #2) AND (2'S SEARCH
      IS LESS THAN OR EQUAL TO THE LAST RECORD IN #1)
   THEN THROW 1
   ELSE IF (#2'S SEARCH ARG IS HIGHER THAN THE LAST RECORD
            IN #1)
            THEN ERROR
            ELSE WRITE BOTH
```



J - IF #2'S RECORD (OR SEARCH) IS HIGHER THAN THE LAST RECORD IN #1, THEN ERROR, ELSE WRITE BOTH.

K -IF #2'S RECORD (OR SEARCH) IS HIGHER THAN #1'S SEARCH, THEN ERROR, ELSE WRITE BOTH.

L -IF #1'S SEARCH ARG IS EQUAL TO OR HIGHER THAN #2'S SEARCH ARG THEN EITHER WRITE BOTH OR OK TO THROW #1, ELSE ERROR.

M -IF (#1'S SEARCH ARG IS EQUAL TO OR HIGHER THAN #2'S SEARCH ARG) THEN THROW 1
ELSE IF (#2'S SEARCH ARG IS HIGHER THAN THE LAST RECORD IN #1)
THEN ERROR
ELSE WRITE BOTH.

N -IF #2'S SEARCH ARG IS HIGHER THAN THE LAST RECORD IN #1, THEN ERROR, ELSE WRITE BOTH.

R -OK TO THROW #1.

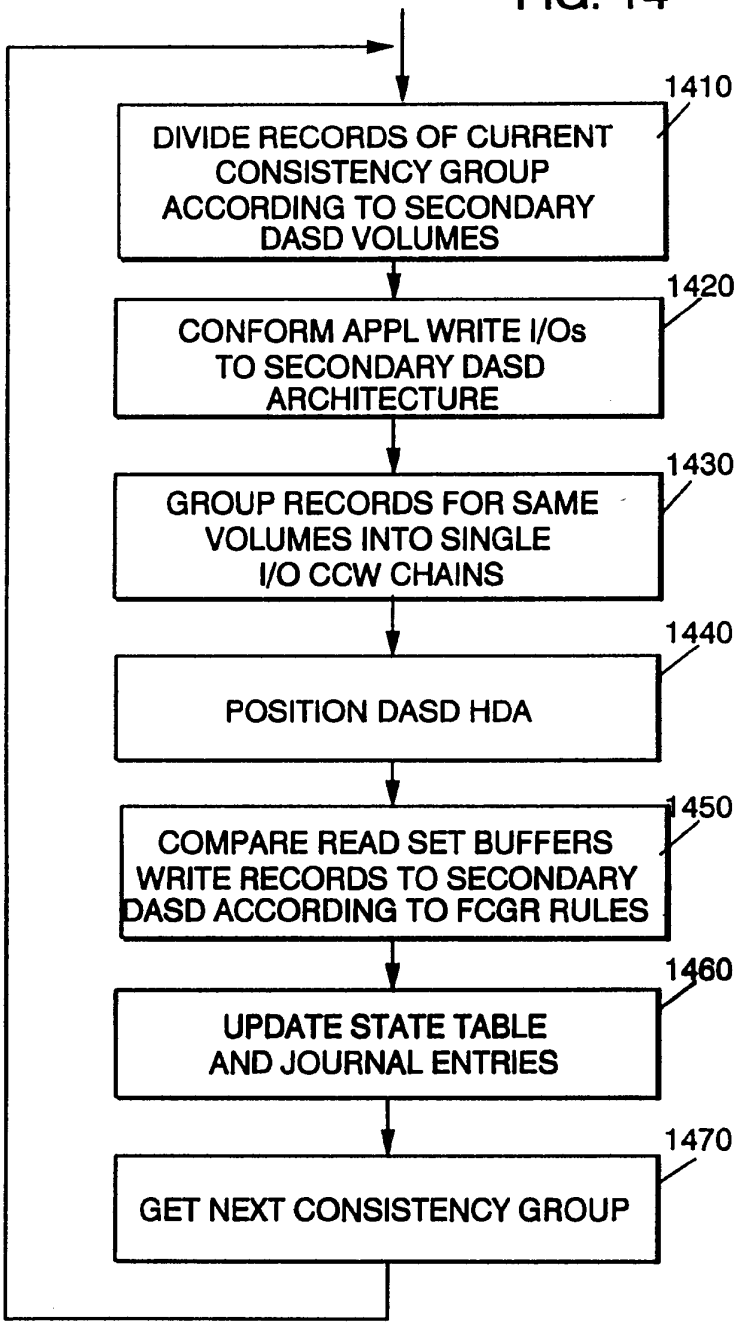
T -MUST THROW 1.

W -WRITE BOTH.

W*-IF #1 AND #2 HAVE THE SAME RECORDS, THEN THROW #1, ELSE DO BOTH OR MERGE RECORDS AND DO ONE WRITE.

FIG. 13B

FIG. 14


1410
DIVIDE RECORDS OF CURRENT CONSISTENCY GROUP ACCORDING TO SECONDARY DASD VOLUMES
1420
CONFORM APPL WRITE I/Os TO SECONDARY DASD ARCHITECTURE
1430
GROUP RECORDS FOR SAME VOLUMES INTO SINGLE I/O CCW CHAINS
1440
POSITION DASD HDA
1450
COMPARE READ SET BUFFERS WRITE RECORDS TO SECONDARY DASD ACCORDING TO FCGR RULES
1460
UPDATE STATE TABLE AND JOURNAL ENTRIES
1470
GET NEXT CONSISTENCY GROUP

European Patent Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 30 0753

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	WO-A-91 20034 (STORAGE TECHNOLOGY CORP) 26 December 1991 * abstract * * page 13, line 6 - page 14, line 20 * * page 18, line 34 - page 19, line 6 * * page 20, line 1 - line 35 * * figure 3 * ---	1,10,20, 21,29	G06F11/14 G06F11/20
A	PROCEEDINGS OF THE ACM SIGMOD INTERNATIONAL CONFERENCE ON MANAGEMEN OF DATA, SAN DIEGO, JUNE 2 - 5, 1992, vol. 21, STONEBRAKER M, pages 236-245, XP 000393592 ANUPAM BHIDE ET AL 'AN EFFICIENT SCHEME FOR PROVIDING HIGH AVAILABILITY' * abstract * * page 237, left column, line 36 - right column, line 52 * -----	1,11,20, 29	

TECHNICAL FIELDS SEARCHED (Int.Cl.6)

G06F

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
BERLIN	27 June 1995	Masche, C

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 111 527 B1

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
22.03.2006 Bulletin 2006/12

(51) Int Cl.:
G06Q 90/00 (2006.01) *H04L 29/06* (2006.01)
G06F 17/30 (2006.01)

(21) Application number: **00310244.9**

(22) Date of filing: **17.11.2000**

(54) **Communication method for data sychronization processing and electronic device therefor**

Kommunikationsverfahren für Datensynchronisierung und elektronische Vorrichtung dafür

Procédé de communication pour la synchronisation de données et dispositif électronique à cet effet

(84) Designated Contracting States:
DE FR GB

(30) Priority: **19.11.1999 JP 33023599**

(43) Date of publication of application:
27.06.2001 Bulletin 2001/26

(73) Proprietor: **KABUSHIKI KAISHA TOSHIBA**
Kawasaki-shi,
Kanagawa-ken 210-8572 (JP)

(72) Inventors:
- **Kizu, Sachio,**
  Kabushiki Kaisha Toshiba
  Tokyo 105-8001 (JP)
- **Natsubori, Shigeyasu,**
  Kabushiki Kaisha Toshiba
  Tokyo 105-8001 (JP)
- **Imai, Toru,**
  Kabushiki Kaisha Toshiba
  Tokyo 105-8001 (JP)
- **Amano, Shinya,**
  Kabushiki Kaisha Toshiba
  Tokyo 105-8001 (JP)
- **Muranaga, Tetsuro,**
  Kabushiki Kaisha Toshiba
  Tokyo 105-8001 (JP)
- **Takabatake, Yoshiaki,**
  Kabushiki Kaisha Toshiba
  Tokyo 105-8001 (JP)

(74) Representative: **Midgley, Jonathan Lee**
Marks & Clerk
90 Long Acre
London WC2E 9RA (GB)

(56) References cited:
EP-A- 0 926 608

- **LI GONG: "ENCLAVES: ENABLING SECURE COLLABORATION OVER THE INTERNET" IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, IEEE INC. NEW YORK, US, vol. 15, no. 3, 1 April 1997 (1997-04-01), pages 567-575, XP000683948 ISSN: 0733-8716**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Printed by Jouve, 75001 PARIS (FR)

EP 1 111 527 B1

Description

[0001] The present invention relates to a communication method for synchronously processing data between a plurality of devices and an electronic device therefor. In particular, the present invention relates to a communication method and an electronic device that have been modified so as to synchronize data in a wireless or wired network such that a network configuration is dynamically changed.

[0002] In recent years, apart from a personal computer, in a portable information processing terminal such as a portable cellular phone, a palm top computer or a wrist watch, personal data such as a telephone number (address) directory, schedule data, a task list, an E-mail can be managed and stored. Using this function, data is shared by using duplicate data in order to use the same data at home or outside as well as in office between electronic devices such as a personal computer or a portable information processing terminal. As a technique for keeping coincidence of a plurality of shared data, a data synchronization processing is known. For example, in the case where an attempt is made to share personal data such as an address directory or a schedule between a personal computer and a portable information processing terminal, duplication of data is created, and original data is held in the host and the duplicate data is held in each device other than the host. In the case where a change occurs with these items of data, the changed data is exchanged between these devices, whereby data consistency is maintained. This method is referred to as data synchronization.

[0003] As a typical computer program for data synchronization, Intelli Sync (Intelli Sync User's Guide Rev 3.1 Puma TECHNOLOGY 1999) is known. This enables data synchronization between personal computers or between a personal computer and a portable information processing terminal. This program enables data synchronization between two personal computers, one personal computer and one portable information processing terminal, or one personal computer and two portable information processing terminals.

[0004] In the case of this system, however, a device serving as a server must be participated in data synchronization processing. In the case where three or more devices share data, there applies a restriction that data synchronization cannot be performed between two devices other than server. Therefore, for example, in the case where a personal computer connected to a network is used as a server, thereby causing two portable information processing terminals to include the duplicate data, even if an attempt is made to use duplicate data after updated in one portable information processing terminal by the other portable information processing terminal outside, data synchronization processing cannot be performed by these portable information processing terminals other than server. Stated another way, the conventional data synchronization processing is performed between the server and the client in a one-to-one manner. It is not possible to perform the data synchronization processing in a one-to-plural manner.

[0005] Furthermore, in general, it is required to provide communication settings for a data synchronization party, settings regarding type of synchronization data, or setting of a method of solving a problem during data conflict. These settings are conventionally provided by a user operating an input device of a personal computer or portable information processing terminal through the setting screen displayed on a display device of the personal computer or portable information processing terminal and inputting data. This means that at least the display device and input device of the portable terminal device are required. Therefore, in the conventional method, because of a restriction specific to a device such as a wrist watch including device size, for example, it is very difficult to provide settings of data synchronisation in a device having only a limited input device or display device.

[0006] Accordingly, it is an object of the present invention to provide a communication method and electronic device that enables data synchronization between arbitrary devices without requiring a server.

[0007] EP-A-0926608 discloses a method of maintaining mutual consistency between copies of data stored at nodes in a network, the nodes comprising a server node, and a number of client nodes at remote locations, such as laptop computers linked by cellular telephone connections, the nodes being connected intermittently.

[0008] Following a reconnection between the remote client node and another of the nodes, timestamp information on objects cached at the other node, is sent to the remote client node. At the remote client node, the timestamp information is compared with corresponding information relating to the local copy, to determine if each of the cached objects is more up to date than the local copy. The cached objects are updated based on the local copy, according to the result of the comparison. This peer to peer resynchronisation means that the server need not monitor and control the process, so the system is more easily scaled up to handle more clients.

[0009] The document "Enclaves enabling secure collaboration over the internet", Li Gong, IEEE journal of selected areas in communications, IEEE Inc. New York, vol. 15, n. 3, 1 April 1997, discloses a group of collaborators which can dinamically form a secure virtual subnet wherein priorities are assigned to the group members for scure co-edition of files.

[0010] The invention provides a data synchronisation method as defined in Claims 1 and 22, and an electronic device and a slave device as defined in Claims 12 and 13 respectively.

[0011] Advantageous embodiments thereof are defined in the respective dependent claims.

[0012] According to the present invention, there is provided a communication method for synchronising data between a plurality of arbitrary devices using a wireless or wired network, the communication method comprising

a member registration step and a data synchronisation step. In the member registration step, a device specified as a master by a user transmits to a device specified as a slave by the user a certificate indicating that the slave belongs to a data synchronisation group that performs a data synchronisation processing and a priority to be used when the slave performs the data synchronisation processing, and then, the slave is registered as a member of the data synchronisation group. In the data synchronisation step, it is determined using the certificate whether or not one slave and one or more other slaves capable of making wireless or wired communication belong to the same data synchronisation group, and the data synchronisation processing is performed through wireless or wired communication in accordance with the priority between the slaves which are determined to belong to the same data synchronization group.

[0013] In this communication method, apart from the data synchronization step of actually processing data synchronization, the member registration step is provided for the purpose of specifying a target device for data synchronization. At this member registration step, the certificate and priority are transmitted from a device specified as the master to a device specified as the slave, whereby the slave can be registered as a member of the data synchronization group. The certificate and priority are provided to each member, whereby data synchronization can be easily performed between slaves even if the master is absent, as long as these devices are registered in the same data synchronization group.

[0014] In this case, as to whether or not the slaves capable of making wireless or wired communication belong to the same group, it is determined by exchanging the certificate between the slaves. Thus, even if a master does not exist, data synchronization can be performed reliably and safely between the slaves. In the case where a conflict occurs during data synchronization processing between slaves, such conflict can be solved by the priority assigned to each slave at the member registration step. Therefore, each slave may hold only the certificate and priority assigned to each slave itself, making it unnecessary to set data or the like concerning a communication party for data synchronization. Hence, a device having only its limited input device and display device can be used as a device targeted for synchronization. Further, as in a wired or wireless network such that a network configuration is dynamically changed, even in a case of use under environment that makes it difficult to provide data for statically specifying a communication party for data synchronization, a device that is such data synchronization party can be reliably specified.

[0015] The data synchronization step can be performed between one master and one or more arbitrary slaves having a certificate indicating that these slaves belong to the same group as well as between slaves.

[0016] In addition, there is no need to specify a master and a slave by devices, and an arbitrary device can be specified as a master or slave of the data synchronization group by data targeted for data synchronization. One device may function as a slave of the data synchronization group for one set of data or may function as a master of the data synchronization group for another set of data. Alternatively, one device may function as a slave of two data synchronization groups corresponding to two sets of different data.

[0017] In the case where a master and a slave are thus specified individually by data targeted for synchronization, the member registration step is also executed between a device specified as a master of a data synchronization group for the data targeted for data synchronization processing and a device specified as a slave of the master. Thus, for transmission of the certificate and priority as well, a master that is a transmission source and a slave that is a transmission destination are set by data, and the certificate and priority are set by data accordingly.

[0018] In addition to the aforementioned member registration step and data synchronization step, a master privilege transfer step may be provided, whereby only the master's function can be transferred to another device while the function and configuration of the current data synchronization group are maintained intact. Hence, a newly purchased device can be used as a master, and a slave can be temporarily used as a master outside.

[0019] In addition, it is preferable that a priority exchange step of exchanging a priority between slaves be further provided. In this manner, the priority registered in the member registration step can be changed later, and data synchronization processing can be achieved more flexibly between slaves.

[0020] The master privilege transfer step and priority exchange step as well are preferably performed after it has been determined that there does not exist any device other than two devices set in a mode for executing the above steps.

[0021] In order to easily change a member of the data synchronization group or the like, it is preferable to further provide a member registration releasing step of releasing a master or a slave from the data synchronization group. The slave releasing step may be performed by that slave device solely. However, the master releasing step is desirably executed together with the member registration releasing step of a slave communicable with the master. In this manner, a plurality of devices including a master can be released from the data synchronization group with one operation.

[0022] In addition, a conventional data synchronizing method assumes one-to-one cable connection, one-to-one connection with the server by modem, or one-to-one connection with the server by infrared-ray radiation. This is because a communication party for data synchronization can be uniquely specified.

[0023] However, if the data synchronization is performed in a bus type wired network or a wireless network free of directivity, unlike a case of one-to-one connection,

it is required to uniquely identify a party for data synchronization from the devices in a network by any method. In addition, in the case of a wireless or wired network such that a network configuration is dynamically changed, the network address of each device is not fixed, and a network address is often assigned during network connection or the like. Thus, the network address of a communication party cannot be used as static identification information for specifying a communication party for data synchronization. Therefore, in a wire or wireless network such that a configuration is dynamically changed as well, it is required to introduce a new technology for making it possible to uniquely identify a party for data synchronization.

[0024] Although it is not presumed in one-to-one connection, in a bus type wired network or a wireless network free of directivity, there is a possibility that a device other than device for data synchronization, for example, a malicious third party exists in the same network, especially in a wireless network. Thus, it is required to ensure data security, and at the same time, to prevent data leak due to computer hacking or the like. In particular, in a wireless network, there is high possibility that a device is merely disposed within the reach of wireless communication, whereby the device can be easily participated in the network. Thus, data security is important. In addition, without maliciousness, there is a danger that data contained in a device other than that targeted for synchronization, the device existing within the reach of wireless communication, is mistakenly or unintentionally changed by data synchronization processing between devices via wireless communication.

[0025] According to the present invention, such an undesired data synchronization is prevented from being occurred since the party for data synchronization can be easily identified by exchanging the certificate between the devices.

[0026] This summary of the invention does not necessarily describe all necessary features so that the invention may also be a sub-combination of these described features.

[0027] The invention can be more fully understood from the following detailed description when taken in conjunction with the accompanying drawings, in which:

FIG. 1 is a view illustrating a member registration step and a data synchronization step that are employed in a communication method according to one embodiment of the present invention;

FIG. 2 is a view showing an example of data set of each device employed according to the present embodiment;

FIG. 3 is a block diagram depicting the configuration of each device employed according to the present embodiment;

FIG. 4 is a view showing an example when functions provided to each device employed according to the present embodiment are classified by communication mode;

FIG. 5 is a view showing master and slave data exchange during data synchronization member registration according to the present embodiment;

FIG. 6 is a flow chart showing procedures for master processing during data synchronization member registration according to the present embodiment;

FIG. 7 is a flow chart showing procedures for slave processing during data synchronization member registration according to the present embodiment;

FIG. 8 is a view showing an example of communication mode during certificate data exchange and synchronizing process in the present embodiment;

FIG. 9 is a view showing another example of communication mode during certificate data exchange and synchronizing process in the present embodiment;

FIG. 10 is a view showing an example of processing at the time of data conflict in the present embodiment;

FIG. 11 is a flow chart showing procedures for master declaration processing in the present embodiment;

FIG. 12 is a flow chart showing procedures for slave processing in data synchronization member release processing according to the present embodiment;

FIG. 13 is a flow chart showing procedures for master processing in data synchronization member release processing according to the present embodiment;

FIG. 14 is a view showing master and slave data exchange in master privilege transfer processing according to the present embodiment;

FIG. 15 is a flow chart showing procedures for master processing during master privilege transfer processing according to the present embodiment;

FIG. 16 is a flow chart showing procedures for slave processing during master privilege transfer processing according to the present embodiment;

FIG. 17 is a view showing data exchange between devices during priority exchange in the present embodiment;

FIG. 18 is a flow chart showing procedures for device processing in priority exchange processing according to the present embodiment;

FIG. 19 is a timing chart showing set mode and execute mode applied to the present embodiment;

FIG. 20 is a view showing an example of an interface of a wrist watch to which the communication method according to the present embodiment is applied; and

FIG. 21 is a view showing an example of an interface of a portable information terminal to which the communication method according to the present embodiment is applied.

[0028] A preferred embodiment of a communication method for data synchronization processing according to the present invention will now be described with reference to the accompanying drawings.

First Embodiment

[0029] Hereinafter, an embodiment of a method and apparatus for synchronously processing data according to the present invention will be described with reference to the accompanying drawings.

Data Synchronization Group

[0030] Data synchronization denotes processing for, in the case where part of duplicated data is asynchronously updated independently relevant to plural sets of the duplicated data, entirely synchronizing the plural sets of these duplicated data. For example, assume that a set of data (data set) that a device A has is duplicated in a device B. Synchronization processing corresponds to processing for synchronizing data sets of the devices A and B each other. If a partial data field of the data set that the device A has is changed, a partial data field of the data set that the device B has is updated during next synchronization processing.

[0031] If the same data field of the data sets of the devices A and B is changed during and after the previous synchronization processing, in order to synchronize these data fields, it is required to select which of the data field values has a higher priority. When the data in the same data field is changed with a plurality of devices during or after synchronization processing, it is referred to as data conflict.

[0032] In the present invention, unlike a conventional method, data synchronization processing between information devices is not performed by one-to-one processing between a server and a client. This processing is performed based on a data synchronization group constructed exclusively for synchronization processing. Of the information devices capable of making data communication wirelessly or using a wire, a group configured between devices that performs data synchronization is referred to as a data synchronization group (or network). Namely, a data synchronization group composed of devices that should perform data synchronization processing corresponds to a data synchronization group. In this case, such data synchronization group may not be always configured by devices, and an arbitrary data synchronization group can be configured by each data set that enables data synchronization. Namely, each information device can configure the data synchronization group by the number of data sets that can be managed.

[0033] In addition, information devices may be a device employed for limited use such as wrist watch or may be a device having a plurality of functions such as personal computer. These information devices are not restricted by their sizes or the types or functions of input/output devices or display devices and the like. The data synchronization group is configured by a plurality of information devices. A device configuring one data synchronization group is referred to as a member of the data synchronization group.

[0034] A data synchronization group is a permanent concept. Data synchronization processing is performed by a member of the data synchronization group that is configured in advance, unlike a case in which a data synchronization group is configured every time data synchronization processing is performed.

Data Synchronization Group Communication

[0035] A communication method in a data synchronization group comprises a member registration step and a data synchronization step. FIG. 1 shows an example.

[0036] The member registration step is a step to be performed when a member is added to a data synchronization group. At this member registration step, a certificate and a priority are transmitted from a device specified as a master by a user to a device that the user want to newly register as a slave. The certificate indicates that a device receiving it belongs to the same data synchronization group (data synchronization network) as the master. That is, the certificate corresponds to a group ID or the like. The priority is used when a device receiving it performs data synchronization processing, and is employed to solve data conflict.

[0037] The data synchronization step is a step that performs data synchronization processing between a plurality of arbitrary members configuring a data synchronous group.

[0038] In FIG. 1, M denotes a master of a data synchronization group. The master M is a core of the data synchronization group. Only one master exist in one data synchronization group, and has a privilege of adding a member to the data synchronization group. S_1 to S_4 denote slaves of the data synchronization group. One or more slaves exist in one data synchronization group. Data synchronization member registration is performed between the slave and the master M in the data synchronization group, whereby the slave becomes a member of the data synchronization group.

[0039] A master and a slave may not be always specified by devices, and an arbitrary device can be specified as a master or slave of the data synchronization group corresponding to the data sets. Thus, one device may function as a slave of the data synchronization group relevant to one set of data or may function as a master of the data synchronization group relevant to another set of data. Alternatively, one device may function as each of the slaves of the two data synchronization groups corresponding to two different types of data.

Member Registration Step

[0040] The member registration step is executed by making one-to-one communication between a device that is a master of a data synchronization group and a slave that newly participates in the data synchronization group. Two members configure the data synchronization group at a stage at which a first slave S_1 is registered in

the master M. Subsequently, the member registration step is sequentially performed similarly between the master M and each of the slaves S_2 to S_4 to be newly registered, whereby the number of members configuring the data synchronization group increases one by one. The member registration step is not performed every data synchronization. Once a device is registered as the member, such registered device is continuously a member of the data synchronization group unless the device has been released explicitly. Data synchronization processing can be executed any time between the members in the same data synchronization group.

Data Synchronization Step

[0041] At the data synchronization step, data is synchronized between the members that belong to a data synchronization group. Here, all the members may perform data synchronization processing at one time or arbitrary members may perform data synchronization processing. In addition, the master M may be or may not be participated in data synchronization processing. According to the present invention, data synchronization processing can be performed between slaves. A plurality of synchronizing processes configured by other members can be performed independently at an arbitrary time, and a plurality of these synchronizing processes may be performed simultaneously in another place.

[0042] For example, as shown in FIG. 1, in the case where only the master M and the slave S_1 configure a data synchronization group, the data synchronization.step is executed between these two devices, namely between the master M and the slave S_1. Alternatively, in a state in which a slave S_2 is newly registered, the data synchronization step can be executed between the master M and slave S_1, between the master M and slave S_2, between the slave S_1 and the slave S_2, or between the master M and the slaves S_1 and S_2. Hereinafter, the synchronization step is executed similarly between arbitrary two devices irrespective of the master or slave.

Member Configuration

[0043] Now, the data and functional structures set to each of the master and slave devices will be described with reference to FIG. 2 and FIG. 3.

[0044] Set data for setting the data synchronization group that each device stores, includes master/slave specification data, identification data, certificate data, priority data, and a counter value.

[0045] The master/slave specification data is like a flag, and its initial value is blanked. When a device is registered as a master or slave, the corresponding flag is set.

[0046] The identification data denotes data for uniquely specifying a member in one data synchronization group, and is used as a key for generating certificate data. The synchronization group is defined by data type, and thus, the identification data differs depending on type of data set in one device. The identification data is generated in each device when one device is registered as a slave of a data synchronization group or when a device is operated as a master by a master declaration described later.

[0047] The certificate data corresponds to the aforementioned certificate, and is employed to certify that one device belongs to one data synchronization group.

[0048] For example, during data synchronization assume that it is authenticated by a public key encryption system as to whether or not devices belong to the same data synchronization group and encryption of synchronization data after authenticated. The identification data and the certificate data are employed as follows in an authentication and encryption system.

[0049] First, each of the master and slave generates a secret key and a public key. A pair of these secret key and public key corresponds to the identification data of each device. When the identification data is transmitted to another device, only the public key that is part of the identification data is transmitted. During data synchronization group registration, when a slave passes a public key as the identification data to a master, the master returns this public key of the slave as a certificate after the public key has been encrypted by own secret key. In this manner, the slave has own secret key and public key; and a certificate certified by the master and the public key of the master. Here, this certificate and the public key of the master are referred to as certificate data.

[0050] During actual authentication and encryption, when a device performing authentication receives a certificate from the counterpart, this certificate is decrypted by own public key of the master. If the certificate is correctly decrypted, it means that the device of the counterpart belongs to the same data synchronization group. Then, the public key of the counterpart is extracted from the decrypted data, and the synchronization data is encrypted by the extracted public key of the counterpart, and is transmitted, whereby synchronization data can be transmitted safely.

[0051] The certificate data may certify reliably that each device belongs to a data synchronization group. The data may be generated by other well known various methods without being limited to the aforementioned method.

[0052] The priority data is used to solve conflict of data synchronization as described previously. The data is used to determine a precedence (namely, priority) for solving such data conflict.

[0053] The counter value is data used to count the number of slaves registered by the master and determine the priority of that slave. The highest priority of the master is "1". The counter value is counted up every time a slave is registered, and the priorities "2", "3", ... are assigned to the slaves in order of registration. This counter is used by the master only.

[0054] As described above, the data synchronization

group is configured by each data set. Thus, the master/slave specification data, certificate data, identification data, priority data, and counter value are provided and managed by each data set. That is, in a device having a plurality of data sets, the certificate data, identification data, priority data, and counter value are not managed in one-dimensional manner, and are manages in two-dimensional manner, namely, by each data set, as shown in FIG. 2. FIG. 2 shows an example when a device belongs to three data synchronization groups corresponding to three types of data sets, i.e., "address directory", "time", and "E-mail", respectively. In this case, in a device having these three types of data sets, whether the device operates as a master or slave by each data set is managed. The counter value is valid only for a data set of the data synchronization group that operates a master. In addition, the aforementioned member registration step is performed by data synchronization group, the certificate data, identification data, priority data, and counter value are independent of each other by each data set, respectively. The secret key and public key configuring identification data are generated by the member registration step by employing random numbers, for example. The value of the identification data is continuously maintained until it has been removed from the data synchronization group in which a device is registered as a member at the member registration step.

[0055] In this manner, in the case where set data on a data synchronization group is managed by each data set, when one performs each of the functions including data synchronization member registration, data synchronization processing, master declaration, data synchronization member releasing, master privilege transfer, and priority exchange, a step is executed such that a user can select a target data set from among the data sets in the row direction shown in FIG. 2 in order to specify what is a data set targeted for processing.

[0056] FIG. 3 shows the functional configuration of each device. This network configuration is provided as an example when a device can be configured as a master and a slave.

[0057] As shown in FIG. 3, each device includes a communication device 11, a communication controller 12, a device input/output device 13, a device input/output controller 14, an authentication unit 15, a data synchronization group setting and management unit 16, a data synchronization processing unit 17, and a synchronous data storage device 18.

[0058] The communication device 11 makes communication with another device via a wireless or wired communication. This device 11 transmits data to another device, and receives data from such another device. As a wireless communication, there can be used an electric wave communication, free of directivity such as "Bluetooth", for example. In addition, a bus type LAN such as "Ethernet" can be used as a wired network. In either of the wireless and wired communication, the network configuration may not be statically determined. The present invention is applicable to a network such that the network configuration is dynamically changed.

[0059] The communication controller 12 controls data reception and/or transmission. The controller 12 includes a function for checking data and a request received by the communication device 11; a function for transmitting and receiving a processing request and response from the device input/output controller 14; a function for making a data transmission and reception request for the communication device 11; and a function for transmitting and receiving a request and a response among the authentication unit 15, data synchronization group setting and management unit 16, and data synchronization processing unit 17, respectively.

[0060] The device input/output device 13 is a user interface, and is an input device or a display device of own device.

[0061] The device input/output controller 14 controls data request and transmission and/or reception relevant to the device input/output device 13. The controller 14 includes a function for checking data and a request received by the device input/output device 13; a function for making a processing request for the communication controller 12; a function for receiving the processing request from the communication controller 12; a management function for writing, reading, deleting data relevant to the synchronization data storage device 18; and a function making a request for and a response to the data synchronization group setting and management unit 16.

[0062] The authentication unit 15 authenticates certificate data received from another device in response to an authentication request from the communication controller 12 by correlating the data with certificate data of the own device, the data being owned by a certificate data management controller 161 through the data synchronization group setting and management unit 16. The unit 15 determines whether or not another device belongs to its data synchronization group, i.e., whether or not both of the own and other devices belong to the same data synchronization group.

[0063] The data synchronization group setting and management unit 16 generally manages settings of the data synchronization group. The unit 16 includes a certificate data management controller 161 and a priority management controller 162. The data synchronization group setting and management unit 16 manages own settings by using the certificate data management controller 161 and the priority management controller 162. The unit 16 transfers data or a processing request to the communication controller 12, authentication unit 15, device input/output controller 14, data synchronization processing unit 17, priority management controller 162, and certificate data management controller 161.

[0064] The certificate data management controller 161 has a region for storing certificate data certified by a master of the own device and a region for storing identification data that is a source of the certificate data. The controller 161 manages certificate data and identification data for

generating such certificate data. The controller 161 includes functions for generating, deleting, and setting the certificate data and the identification that is a source of the certificate data according to a request from the data synchronization group setting and management unit 16.

[0065] The priority management controller 162 has a region for storing priority data on own device; a region for storing a counter value changed when a priority is issued in a master; and a counter device. The controller 162 includes a function for processing requests for generating, setting, deleting, and transferring priority data from the data synchronization group setting and management unit 16; and a function for resetting and incrementing the counter device.

[0066] The data synchronization controller 17 performs synchronization processing for a data set targeted for synchronization (synchronization data). The controller 17 includes a function for reading, writing, and deleting data relevant to the synchronization data storage device 18; a function for solving data conflict in view of the priority of own device and that of another synchronization device; and a function for processing a request from the communication controller 12 and delivering data.

[0067] The synchronization data storage device 18 stores synchronization data, and has a region for storing data.

[0068] The configuration of this device is provided as an example when a device can be configured as a master and a slave. This configuration may include a device (member) that can be configured as a slave only or a device (member) that can be configured as a master only. In this case, there can be provided a configuration in which functions or configurations required for a slave and a master are eliminated. For example, a certificate data producing function and a counter device and its control function and the like may be provided to only a device that can operate as a master, and is not required for a device that functions as a slave only.

[0069] Now, data synchronization group management and operation will be described here.

[0070] First, all functions owned by each device in a data synchronization group will be described with reference to FIG. 4.

[0071] A device that can configure a data synchronization group includes a total of six functions. The device includes a function to be processed without making communication with another device; a function requiring one-to-one communication between one device and another device; and a function requiring communication between one device to a plurality of devices (actually, made of a plurality of one-to-one communications). The "master declaration" function is a stand alone function that is processed without making communication with another device. The "data synchronization member registration", "master privilege transfer", and "priority exchange" functions are processed through one-to-one communication. The "data synchronization" function is performed between two or more devices. The "data synchronization member releasing" function is a stand along function in a slave, and a function that is processed through communication with one or more slaves in a master. Hereinafter, these functions will be described in detail. These functions can be performed by data synchronization group. When a function is read, it is required to specify an identification code of the data synchronization group.

(1) Data Synchronization Member Registration

[0072] In the case where a new member is added to a data synchronization group, "data synchronization member registration" processing is performed. This processing is adopted to achieve the aforementioned member registration step. The processing is performed by making data communication between a master of a data synchronization group and a member to be newly added as a slave. The master and newly registered member must be communicable with each other via a wireless or wired communication channel while power is supplied. The user calls the "data synchronization member registration" function each in the master and the member to be registered. Subsequently, each of the master and the member to be registered is placed in a member registration mode. In order to enter the member registration mode, the user interaction and exclusive control for another device are used such that the user holds down the registration button in each of the master and the new member, for example, and registration is permitted only when only two devices in the member registration mode exist, and security authentication during registration is provided. It is possible to determine whether or not only two devices in the member registration mode exist by checking a response to a broadcast from the master, for example. This prevents accidental registration of a device owned by a malicious third party within the reach of electric waves.

[0073] The effective communication range means that a device is within the reach of wireless communication in a wireless network. In a wired network, this range means a limited network area such as sub-network, for example, in which broadcasting is effective. Namely, in the range which has a possibility that a third party receives the communication, under the condition that there exists any device other than only two devices, i.e., a device targeted for member registration and a master, which are set in the member registration mode, data synchronization member registration is started.

[0074] Next, a function for registering a data synchronization member will be described in detail. In FIG. 5, there is shown master and slave communication processing to be performed together with synchronization member registration.

[0075] First, both of the master M and the slave S transmit to the counterpart a request for registering a data synchronization member, and then, the master notifies its identity to the slave. In this state, the master has all set data, but the slave has only identification data. When the slave receives the notification, it transmits the iden-

tification data to the master. The master transmits to the slave the certificate data for this slave generated using the identification data and the priority data, and communication processing is completed.

[0076] Now, the processing step in each device will be described in detail. In FIG. 6, there is shown the step of master processing when a master device registers a slave device as a member of the data synchronization group. A description is given with reference to FIG. 3.

[0077] First, when a user call of the "data synchronization member registration" function is transmitted to the device input/output controller 14 (step S11), the current setting of the data synchronization group (whether the own device is a master or slave) is determined at the data synchronization group setting and management unit 16 (step S12). If the own device is set as a master by a master declaration function described later, the user is prompted for a request for executing registration processing. When it is set as a slave, processing is terminated. In the case where nothing is set (i.e., blank data is specified), the device operates as a newly registered member described later (FIG. 7).

[0078] In the case where the own device is set as a master, when the device input/output device 13 receives a request for executing registration processing from the user (step S13), the communication device 11 transmits a request for registering a data synchronization member in a broadcasting manner (step S14). Then, the communication device 11 waits for a registration request from a new member. Although not shown, in the case where a registration request is not made for a predetermined time, processing terminates with time out. When a registration request from a new member is received (step S15), the registered device determines whether or not only two devices requires the registration, i.e., the own device and a new member (step S16). Namely, after a request for registering a data synchronization member has been transmitted between the own device and the counterpart, if only one counterpart transmits such request for registering a data synchronization member, it is determined that only two device, i.e., the own device and new member participate the registration.

[0079] In the case where any device other than only two devices, i.e., the own device and new member participate the registration (NO at the step S16), processing is terminated. In the case where only two devices participate the registration (YES at step S16), the communication device 11 notifies its identity to the new member (step S17). Then, when identification data that is a source of certificate data from the new member is received by the communication device 11 (step S18), the data synchronization group setting and management unit 16 generates the certificate data and the priority data to be passed to the new member, based on this data (steps S19 and S20). Further, the value of the counter that counts the number of registered devices is incremented by one (step S21).

[0080] Last, the communication device 11 transmits the generated certificate data and the priority data to the new member (steps S22 and S23). Although the counter value is used as the priority data, the priority may be determined in any system as long as a precedence can be defined by generating different values every issue, thereby making it possible to uniquely judge the priority among these generated values.

[0081] Now, in FIG. 7, there is shown the step of processing a new member when a desired device (member) newly participating in a data synchronization group is registered as a data synchronization group member by a master device. A description will be described with reference to FIG. 3.

[0082] When the device input/output device 13 receives the user request for "data synchronization member registration", as in the case with FIG. 6, the current data synchronization group setting is determined by the data synchronization group setting and management unit 16. In the case where the own device is set as a master, the aforementioned "data synchronization member registration" (FIG. 6) is performed. In the case where the own device is set as a slave, processing terminates. If a data synchronization group is not set (blank data is specified as a master or slave), the user is prompted for request for executing registration processing. When the user makes a request for executing registration processing to the device input/output device 13 (step S31), the communication device 11 receives a registration request from the master (step S32), and transmits such registration request to the master (step S33). Although not shown, in the case where there exists no response from the master within a predetermined period of time, processing terminates with time out. After the registration request has been transmitted to the master, the registered device determines whether or not only two devices, i.e., the own device and master participate the registration (step S34). In the case where only the two devices participate the registration, the own device waits for the counterpart master notification. In this manner, it is determined that the counterpart is a master (step S35), the data synchronization setting and management unit 16 generates identification data that is a source of the certificate data (step S36), and transmits the data to the master (step S37).

[0083] When the communication device 11 waits for the certificate data and priority data from the master, and receives them, the certificate data and priority data are set through the data synchronization group setting and management unit 16, and registration processing completes (steps S38 and S39).

[0084] In registration processing of a data synchronization member, it is required to share a data set of the same type between the master issuing the certificate data and priority and the new member to be registered. In the case where an attempt is made to register data sets of their different types, it is required to disable registration. This collation step can be performed by checking whether or not identifiers 01, 02, and 03 are assigned in advance

to an address directory, a time, and schedule, for example, during "transmission of a request for registering a data synchronization member" (step S14) and during "transmission of a request for registering a data synchronization member" (step S33), transmitting these identifiers between the master and the new member, and checking whether or not the identifiers are identical to each other.

[0085] Once it is guaranteed that a master and a slave that belong to a data synchronization group and a slave have a data set of the same type, it is authenticated as to whether or not the master and the slave belong to the same data synchronization group when another function is used. Thus, there is no possibility that data sets of the different types coexist. Even in a device having a plurality of data sets, the data sets in that device configure, respectively, other completely independent data synchronization groups, the data sets of different types and processing do not coexist.

(2) Data synchronization processing

[0086] Now, data synchronization processing will be described here.

[0087] When a master in a data synchronization group transmits certificate data and priority data to each member by data synchronization member registration processing, actual data synchronization processing can be performed between arbitrary members (even without intervening a master) by using at least these two sets of data.

[0088] During synchronization processing, the members belonging to the same data synchronization group first authenticate their identities each other, and then, perform data synchronization. Sets of data that do not conflict with each other are merged, respectively. In the case where a conflict occurs with some sets of data, the data with the highest priority is employed as data for solving a conflict.

[0089] As shown in FIG. 8, assume that a data synchronization group is configured by five members A, B, C, D, and E, wherein the member A is a master.

[0090] In the case where data synchronization processing is performed when all the members are present, one member transmits own certificate data to another member, receives certificate data from the counterpart, and authenticates their identity with each other. In transmission and authentication of certificate data, there is provided a method for all members transmitting/receiving certificate data to/from all the counterparts, thereby individual members using their received certificate data to authenticate their identities. Namely, among the members A to E, certificate data transmission/reception and authentication may be performed according to the number of combinations of different two members.

[0091] Alternatively, as shown in FIG. 9, certificate data transmission/reception and authentication may be performed only between a member with the highest priority (the member A here) and another member. Namely, there is provided a method in which the pairs of members such as A and B, A and C, and A and D, and A and E (the member A serves as a server), thereby performing authentication. After each of the members A to E has been authenticated to be a member of the same data synchronization group, in actual data synchronization processing, a member with the highest priority (the member A here) transmits/receives synchronization data in a data field to/from all other members, checks a conflict, solves such conflict based on the priority, and changes data so that the member can sequentially process the synchronization data in the next data field.

[0092] A conflict during data synchronization processing occurs, as shown in FIG. 10, in the case where the members A, C, and E have not been changed relevant to a data field 'x' but data owned by the members B and D is updated since previous synchronization processing. Means for solving such data conflict is to employ data owned by the member B with the highest priority. This is because, although the member A has the highest priority of the five, data is changed in the members B and D only. Even if data is not changed, a conflict during synchronization processing may be solved using the data owned by the member A with the highest priority.

[0093] Next, assume that only the members A, B, and C in this data synchronization group perform data synchronization. Such situation can occur when the members D and E are not turned ON or when only the members A, B, and C are carried to the outside, and data synchronization is performed outside. Authentication and data synchronization are performed in the same manner as described above. Then, assume that data owned by the members B and C conflict with each other (that is, assume that the data in the same data field owned by the members B and C has been updated during and after the previous synchronization). Such conflict is solved by using the data of a member having the highest priority of the members B and C. As a result, the data owned by the members A, B, and C are synchronized with each other. Assume that the remaining members D and E performs data synchronization in another place. As is the case with the members A, B, and C, if a conflict occurs, such conflict is solved by using the data with the highest priority of the members D and E, and the members D and E undergoes data synchronization.

[0094] As described above, data synchronization processing can be performed between slaves as well as between one master and one or more arbitrary slaves, each of which has a certificate indicating that they belong to the same data synchronization group that performs data synchronization processing. In this case as well, certificate data is transmitted between devices capable of making wireless or wired communication, whereby it is determined that these devices belong to the same data synchronization group. Thus, data synchronization can be performed safely by using only the certificate data without having specific data for uniquely specifying the

counterpart.

[0095] Data synchronization processing can be performed as follows, for example.

(a) A device set in a data synchronization processing mode transmits a data synchronization processing request in a broadcasting manner to other devices that exist within a communicable range;
(b) A certificate is transmitted between these devices, and it is authenticated as to whether or not they belong to the same data synchronization group each other; and
(c) When the above devices belong to the same data synchronization group, data synchronization of the corresponding data sets is executed in accordance with each priority.

[0096] In the case of data synchronization processing, unlike registration processing of a data synchronization member as described in FIG. 6 and FIG. 7, there does not apply a restriction that only two devices, i.e., a master and a member targeted for new registration. Data synchronization processing can be performed between two or more arbitrary devices that exist within a communicable range provided they belong to the same data synchronization group. All the devices to be synchronized may be set in a data synchronization processing mode; only one device may be set in a data synchronization processing mode by the user, and another device receiving a data synchronization processing request from that device may automatically enter the data synchronization processing mode according to that request.

[0097] In addition, in actual data synchronization processing, the user may specify a data set targeted for data synchronization processing in the data sets for which data synchronization groups are set. Alternatively, it is determined as to whether devices belong to the same data synchronization group relevant to all the data sets for which the data synchronization groups are set, whereby these synchronizing processes are executed simultaneously relevant to the data sets which belong to the same data synchronization group. Authentication data exists by each data set for which the data synchronization group is set. Thus, all the authentication data is exchanged each other, for example, thereby making it possible to check a data synchronization group to which both of the counterpart device and the own device belong. Of course, the identifier assigned to each data set is transmitted, whereby a data synchronization group to be authenticated may be specified.

(3) Master Declaration

[0098] Now, a "master declaration" function for setting the own device as a master of a data synchronization group will be described here.

[0099] In the case where a data synchronization group is newly constructed, it is required to determine a member as a master of the network. The "master declaration" function is provided to achieve this. The "master declaration" can provide setting such that only a member having a data set, which does not belong to any data synchronization group, can be a master of a data synchronization group that the member will construct. That is, a device that has already been set as a slave cannot be changed to a master. This setting is provided by a stand alone, and communication with another member is not performed.

[0100] In the existing data synchronization group, the "master declaration" cannot be performed by changing to a master a member that does not belong to the data synchronization group or a member that has been set as a slave or temporarily executing "data synchronization member releasing" for a slave to remove the slave from the data synchronization group, thereafter adding it as a master again. This is required to ensure security. In the latter case, there is defined another new data synchronization group in which a member performing "master declaration" is set as a master. The "master declaration" is used any time only when a new data synchronization group is constructed.

[0101] In FIG. 11, there is shown the processing step of "master declaration" in which a new member is set as a master of a new data synchronization group by the "master declaration".

[0102] When the user first makes an operation for calling the "master declaration" function for a data set through the device input/output device 13 (step S41), the data synchronization group setting and management unit 16 checks the current setting of the own device in the data synchronization group (step S42). In the case where the device has been set as a master or slave, processing is terminated. If no data synchronization group is defined (i.e., if blank data is specified as a master or slave), the data synchronization group setting and management unit 16 generates identification data on a data synchronization group (step S43), and generates authentication data based on the generated identification data (step S44).

[0103] Further, in the data synchronization group setting and management unit 16, a priority counter is reset, and the initial value of the counter is set as priority data indicative of own priority (steps S45 and S46). Then, the counter is incremented, and processing is terminated.

(4) Data Synchronization Member Releasing

[0104] "Data synchronization member releasing" is performed when a registered device is released from a data synchronization group. This operation differs depending on a master and a slave, and thus, a description will be given separately.

[0105] First, a description will be given with respect to a case in which a slave is released from a data synchronization group. This function is used when a slave participating in a data synchronization group is used while the slave is excluded from the member of the synchro-

nization group or when a slave device is discarded. The "data synchronization member releasing" is a stand-alone operation, and communication with a master or another device is not required. The existing data synchronization group is not affected by this function. The priority data owned by the released slave will be vacant because it is not reused in "data synchronization member registration". However, this does not mean that consistency of the data synchronization group is not destroyed.

[0106] In FIG. 12, there is shown the processing step when "data synchronization member releasing" is performed in a slave.

[0107] First, when a user make a request for "data synchronization member releasing" for the device input/output device 13 relevant to a data set (step S51), the data synchronization group setting and management unit 16 determines the current setting of the own device in the data synchronization group (step S52). When the own device is a slave, the data synchronization group setting and management unit 16 deletes the corresponding priority data, certificate data, and identification data (steps S53, S54, and S55). This releasing is performed by each synchronization group, and thus, network identification data is also specified when the releasing function is called.

[0108] In the case where a master is released from a data synchronization group, if a slave currently communicable with the master exists when "data synchronization member releasing" is performed, this slave is also released from the data synchronization group. This is because the master generally dismembers a slave without transferring its master privilege when the master desired to release the slave from the data synchronization group. In addition, only a communicable slave is released from the data synchronization group simultaneously because it is not assumed that all members in the data synchronization group are always present by the fact that the master cannot grasp all the members in the data synchronization group and the fact the slave is placed in another place due to purchase of a new device, discarding, faults, theft or the like. In this case, a slave that does not exist on the spot still participates in the previous data synchronization group, and data synchronization can be performed between the slaves that do not exist on the spot. However, in this case, since a master does not exist any longer, a data synchronization group member cannot be increased by newly performing data synchronization member registration.

[0109] In FIG. 13, there is shown the processing step when "data synchronization member releasing" is performed in a master. When a user makes a request for "data synchronization member releasing" relevant to a data set to the device input/output device 13, as described in FIG. 12, the data synchronization group setting and management unit 16 determines the current data synchronization group. In the case where the own device is a master, this device searches for a currently communicable slave (step S61). This may be determined according to a response from the slave relevant to data transmitted in a broadcasting manner. Then, the own device transmits a request for "data synchronization member releasing" and the master certificate data to a slave within the currently communicable range (steps S62 and S63).

[0110] The slave receiving this request performs authentication by using the master certificate data, and then, resets own setting as in "data synchronization member releasing" in the slave.

[0111] The master is then deletes the priority data, certificate data, identification data, and counter values of the own device, the counter is reset, and processing is terminated (steps S64 to S67).

(5) Master Privilege Transfer

[0112] The "master privilege transfer" function allows a data synchronization group to continuously exist intact in the existing data synchronization group, and allows the current master to transfer only the master's privilege to another member (slave). This function is used when one attempt to transfer a master privilege to another device while maintaining the functions and configuration of the data synchronization group, including cases in which where one attempts to specify a newly purchased device as a master of the existing data synchronization group; a master device is discharged; one attempts to change the current master device to a new device; and, when a desk top personal computer that is not mobile is specified as a master, one attempt to temporarily perform "data synchronization member registration" outside. This means that the slave and master privileges are exchanged in the existing data synchronization group. Alternatively, the master privilege may be transferred to a member that is not included in a data synchronization group by automatically exchange the privileges in combination with "data synchronization member registration".

[0113] In this process, the "master privilege transfer" function is called by a master and each slave device that transfers the master privilege. In a method similar to the "data synchronization member registration", the set data held by the master is transferred to a slave, and the set data held by the slave is transferred to the master. The data held by each of the master and slave is exchanged intact. Individual members do not have data concerning another device each other, and thus, these individual members do not affect those in another data synchronization group. Thus, processing can be performed safely and simply as in the "data synchronization member registration".

[0114] Now, the "master privilege transfer" function will be described in detail,

[0115] Here, master and slave communication processing to be performed together with the "master privilege transfer" processing will be described with reference to FIG. 14.

[0116] First, both of the master and the slave transmit a master privilege transfer request to each other. When both of them receive this request, they transmit certificate data for device authentication. After authentication processing, the master transmits the identification data, priority data, and counter value to the slave. Alternatively, the slave transmits the priority data and identification data, and communication processing is completed.

[0117] Now, the processing steps in each device will be described in detail. FIG. 15 shows the processing steps in a master device having the "master privilege transfer" function. A description will be given with reference to FIG. 3.

[0118] When the device input/output device 13 receives the "master privilege transfer" request from the user (step S71), the data synchronization group setting and management unit 16 judges the current setting of the own device in the data synchronization group (step S72). When the device is a master, the device input/output device 13 continuously waits for a request for executing processing from the user. When the device input/output device 13 receives a request for executing processing (step S73), the communication device 11 transmits a request for master privilege transfer processing in a broadcasting manner (step S74). Then, the communication device 11 waits a request for such transfer processing from the counterpart device. When such request from the counterpart device arrives (step S75), the communication device 11 and the communication controller 12 determine whether or not only the own device and the counterpart device are executable the master privilege transfer processing (step S76). In the case of affirmative determination (YES at the step S76), the communication device 11 transmits certificate data to the counterpart device (step S77). The sequence of the steps of transferring certificate data and determining the number of executable devices may be reversed. When the communication device 11 receives certificate data from the counterpart device (step S78), the authentication unit 15 and data synchronization group setting and management unit 16 authenticate the counterpart device, and determines that the counterpart device exists in the same data synchronization group as the master (step S79). If authentication fails, processing terminates (NO at the step S79). When authentication is successful (YES at the step S79), the data synchronization group setting and management unit 16, communication controller 12, and communication device 11 transmit the identification data, priority data, and counter value owned by the own device to the slave (steps S80, S81, and S82). Then, the communication device 11 receives the priority data and identification data from the slave (steps S83 and S84). The data synchronization group setting and management unit 16 sets the priority data, certificate data, and identification data received from the slave as own priority data, certificate data, and identification data (step S85), and resets the counter. Then, processing terminates (step S86). In this manner, a device that has been a master is specified as a slave device.

[0119] The timings of transmitting the priority data, counter value, certificate data, and identification data owned by the own device (steps S77, S80, S81, and S82) and those of receiving the priority data, certificate data, and identification data from a slave (steps S78, S83, and S84) may be in accordance with an arbitrary sequence.

[0120] Slave operations are generally similar to master operation. The processing steps are shown in FIG. 16.

[0121] When the device input/output device 13 receives a "master privilege transfer" request from the user, as described in FIG. 15, the data synchronization group setting and management unit 16 judges the current setting of the own device in the data synchronization group. In the case where this setting is not provided, processing terminates. When the device is a master, the aforementioned master operation is performed. In the case where it is determined as a slave, the device input/output device 13 waits for a request for executing master privilege transfer processing from the user. When the user executes master privilege transfer processing for the device input/output device 13 (step S91), the communication device 11 receives a master privilege transfer request from a master (step S92), and transmits the master privilege transfer request to the master (step S93). Then, at the step S94, the communication device 11 and the communication controller 12 check the number of devices that are executable for master privilege transfer processing. In the case where only one device other than own device is executable (YES at the step S94), it goes to the following steps.

[0122] First, the certificate data is exchanged with the master (steps S95 and S96). Then, the authentication unit 15 and the data synchronization group setting and management unit 16 authenticates the counterpart (step S97). After authentication, in the case where the communication counterpart is not a master of the data synchronization group to which the own device belong (NO at the step S97), processing terminates. In the case where the counterpart is the master (YES at the step S97), the priority data, counter value, and identification data owned by the master is received from the communication device 11 (steps S99, S100, and S101). Then, after the priority data and identification data owned by the own device have been transferred from the communication device 11 (steps S101 and S102), the data synchronization setting and management unit 16 sets the received master's identification data, certificate data, priority data, and counter value as own data (steps S103 and S104), and processing terminates.

[0123] The data synchronization group master can be set by each data set. In the case where a master device has a plurality of data sets, the master privilege transfer processing may be performed for only set data concerning a specific data set specified by the user. Of course, if a slave device specified as a transfer destination has a plurality of data sets, the master privileges concerning all the data sets owned by the master device can be trans-

ferred to one slave device in all.

(6) Priority Exchange

[0124] A "priority exchange" function exchanges the priorities owned by two slaves with each other. If a conflict occurs during data synchronization, such conflict is solved by using data on a member with its higher priority. The priority is assigned from the master in registration order. Thus, there can occur a case in which the user wants to cause the data on a slave registered later to precede the data on a slave registered previously. Although the priority may be selected during data synchronization member registration, the master or slave requires a display interface for selecting the priority. Here, devices are numbered in registration order for the purpose of simplification of processing so that priorities can be exchanged. At the time of "priority exchange" as well, as is the case with "data synchronization member registration" or "master privilege transfer", processing is performed only when only two devices are set.

[0125] Now, the communication processing step when the priorities are exchanged between slaves will be described with reference to FIG. 17.

[0126] First, slaves that want to exchange priorities transmit a priority exchange request to each other. When one slave receives such request, both of them transmit certificate data to each other. After it has been determined that the slaves belong to the same data synchronization group, both exchange priority data, and communication processing terminates.

[0127] In FIG. 18, there is shown the processing step of "priority exchange" in each slave device that exchanges the priority.

[0128] When the device input/output device 13 receives a "priority exchange" request from the user (step S101), the data synchronization group setting and management unit 16 judges the current setting of the own device in the data synchronization group. In the case where the setting is not provided or in the case of a master, processing terminates. In the case of a slave, the device input/output device 13 waits for a request for executing processing from the user. When the device input/output device 13 receives a request for executing processing (step S103), the communication device 11 transmits a request for processing "priority exchange" in a broadcasting manner, and waits for such request for processing "priority exchange" from the counterpart that exchange the priority. Then, when the communication device 11 receives a request for processing "priority exchange" from the counterpart (step S105), the communication controller 12 determines the number of devices that are currently executable for "priority exchange" (step S105) by using the data from the communication device 11 (step S106). In the case where only one device other than the own device is executable for "priority exchange" (YES at the step S106), the data synchronization group setting and management unit 16, communication controller 12, and communication device 11 receive and/or transmit certificate data from and/or to each other (steps S107 and S108). If three or more devices are executable within a communicable range (NO at the step S106), processing is terminated.

[0129] After certificate data has been transmitted and/or received, the authentication unit 15 and the data synchronization group setting and management unit 16 authenticate the communication counterpart based on the authentication data received at the communication device 11 (step S109). When it is authenticated that the counterpart device is a member of the same synchronization group (YES at the step S109), the own device transmits own priority data, and receives the counterpart's priority data (steps S110 and S111). Then, the data synchronization group setting and management unit 16 rewrites its own priority with the received priority, and terminates processing (step S112).

[0130] In the case where a device functions as a slave of the data synchronization group, the device having a plurality of data sets, the processing shown in FIG. 18 may be performed for a data set in which priority exchange is specified by the user.

(7) Method for Setting a Device

[0131] During each of the "data synchronization member registration", "priority exchange", and "master privilege transfer" processes, it is required to ensure safety of data communication to be performed together with setting. Therefore, prior to exchange actually set data, two devices targeted for setting are moved to be requested for executing processing. Then, after it is determined as to whether or not only two devices are requested for executing processing in a communicable range, it is required to determine whether or not such setting is continued. A series of these processes will be described in more detail.

[0132] In connection to these processes, each device has a set mode and an execute mode. The set mode serves to specify call functions including "data synchronization member registration", "priority change", and "master privilege transfer". The execute mode serves to start actual processing of a function called by the set mode, and make data communication with the set counterpart as required. Therefore, this execute mode means a request for executing processing.

[0133] In the case where a user attempts to perform "data synchronization member registration", "priority change", or "master privilege transfer", as shown in FIG. 19, the user moves two devices targeted for setting (devices 1 and 2) to the set mode. Then, each device determines which function is called, and determines whether or not the called function is valid in current setting of the own device. For example, in the case where the user attempts to perform "data synchronization member registration", the user determines that the called function is the "data synchronization member registration", and de-

termines whether or not the own device is a mater, slave or none. When the called function is valid, it can go to the execute mode. User's operation is required for the current mode to go to the execute mode.

[0134] The user simultaneously moves two devices to the execute mode relevant to the same function (strictly, an error Δt of about 1 second is permitted). At a stage before communication with the counterpart is established in the execute mode, each device waits for data indicating the execute mode from another device. After each device has received the data, each device transmits it to a transmission channel. During this data transmission, a data broadcasting transmission may be repeated with predetermined time intervals, for example. When each device receives data indicating the execute mode of the counterpart, after it is determined that there exists no other device that enters this mode and the devices 1 and 2 communicate with each other, communication of set data is started. In communication of the set data, data is encrypted to ensure security by using an encryption system such as the aforementioned public key encryption system. In the case where there exists any other device entering the execute mode, communication regarding the set data is not performed, and processing is interrupted and terminated.

[0135] A method of configuring a data synchronization group and the functions of each device has been described hereinbefore. Now, a specific example of providing devices configuring a data synchronization group will be described here.

[0136] FIG. 20 shows an example when a system for this communication method is applied to a wrist watch, and FIG. 21 shows an example when this system is applied to a portable information processing terminal.

[0137] In the wrist watch shown in FIG. 20, sections associated with data synchronization are two buttons that exist at the left of a character board and four display sections at the upper part of a clock display section. The two buttons include an upper left button employed to instruct a request for executing each function for setting a data synchronization group and for instructing the start of data synchronization processing; and a lower left button for selecting each function for setting a data synchronization group. The display sections include a mode display section for displaying each function for setting a data synchronization group; a display section for displaying a character M or S that represents a master or slave; a display section indicating a numeric value of priority data; and a display section indicating whether or not a data synchronization group is currently set. For a description of use, procedures for executing two typical functions, "data synchronization member registration" and "data synchronization processing" will be discussed here.

[0138] First, the "data synchronization member registration" will be described here. The lower left button is used to call a "data synchronization member registration" mode. The mode display section displays that the "data synchronization member registration" function has been called. When this call is valid, the user holds down the upper left button, and executes a request for "data synchronization member registration." In this example, if this wrist watch is registered as a member by a given master, the similar procedures are performed in the master as well. When the aforementioned "data synchronization member registration" processing is successful while both devices execute the request, registration processing completes. In the case of "data synchronization processing", the upper left button is first pressed, whereby a function for starting synchronization processing is executed. If one or more devices belong to the same synchronization group that exists in a communicable range, and that can execute synchronization processing, data synchronization processing is actually started. The start of data synchronization processing may be executed by any one device calling the corresponding function or may be executed by all the devices calling the data synchronization processing start functions simultaneously, all of these devices being actually targeted for data synchronization processing at this time.

[0139] A data set that the wrist watch has is generally limited to time data (clock data). Thus, a device targeted for data synchronization processing with the wrist watch is a device that belongs to a data synchronization group concerning time data (clock data) in the devices that exist in a range communicable with the wrist watch.

[0140] In FIG. 21, there is shown an example of the screen for setting data synchronization of a portable information processing terminal. Shown here is an example of the screen for setting data synchronization concerning time data. This screen is composed of a "set menu" for calling each function provided at the upper left part; a display section that is set to display the current settings; and an execute button for starting data synchronization processing, the button being provided at the lower left. In the case where an attempt is made to perform data synchronization with a device that exists in the range communicable with the portable information processing terminal, belongs to the same data synchronization group capable of executing synchronization processing, a synchronization button at the lower right is pressed, whereby the data synchronization processing is called. In addition, the "data synchronization member registration" is selected from the set menu, whereby the data synchronization member registration screen is displayed. The execute button provided on this screen is held down, whereby the function for data synchronization member registration can be executed in the same manner as in the foregoing example of the wrist watch.

[0141] As has been described above, according to the communication method of the present embodiment, there is no need to provide set data concerning a communication counterpart for data synchronization. Basically, data synchronization processing can be performed safely and reliably by a certificate (certificate data) and a priority. Thus, data synchronization processing can be performed between arbitrary slaves without requiring a

master. In particular, the present invention is preferable to application to a device that makes wireless communication outside. Moreover, synchronization processing of a data set that belongs to the same data synchronization group is automatically executed with another device that exists in a communicable range by instructing the start of synchronization processing. Therefore, in a device having only a limited user interface such as wrist watch or portable cellular phone, the device being difficult in setting concerning a communication counterpart as well, data synchronization can be easily performed by data synchronization group such that a group member is dynamically changed.

[0142] For example, although each embodiment has been described singly, a plurality of embodiments may be combined with each other as required. The communication method according to the present embodiment can be achieved by a computer program. Thus, the computer program is recorded in a computer readable recording medium such as CD-ROM, and then, is introduced into the computer or the computer program is recorded into the recording medium such as ROM, and then, is merely incorporated into an electronic device, making it possible to achieve an effect similar to the present embodiment using a general electronic device.

[0143] As has been described above, according to the present invention, it is possible to provide a communication method and electronic device that enables data synchronization between arbitrary devices without requiring a server. It is possible to perform data synchronization with a device that does not have a limited input device and a display device. It is possible to make communication for data synchronization safely between arbitrary devices in a wired or wireless network such that a network configuration dynamically changes. In exchanging set data between devices, only specific two devices perform setting operations simultaneously after the presence of two set devices has been guaranteed, whereby setting can be done safely.

Claims

1. A data synchronisation method comprising the steps of:

   transmitting from a master device to a slave device a certificate indicating that the slave device belongs to a data synchronisation group to which the master device belongs and a priority to be used for solving conflict of data during data synchronisation, thereby registering the slave device as a member of the data synchronisation group to which the master device belongs;
   determining whether or not a first slave device and a second slave device which is capable of communicating with the first slave device belong to the same data synchronisation group by using the certificate;
   and performing data synchronisation between the first slave device and the second slave device based on the priority if the first slave device and the second slave device belong to the same data synchronisation group.

2. The method according to Claim 1, wherein said registering is performed after it is confirmed that there is no device other than said master device and said slave device capable of communicating with said master device and set in a registration mode.

3. The method according to Claim 1, wherein said data synchronisation is performed between one master device and one or more slave devices which belong to the same synchronisation group or between at least two slave devices which belong to the same synchronisation group.

4. The method according to Claim 1, wherein said master device and slave device store plural types of data and belong to plural data synchronisation groups defined for each of said plural types of data.

5. The method according to Claim 4, wherein said registering is performed by transmitting said certificate and said priority from the master device belonging to a given synchronisation group defined for a given type of data to the slave device belonging to the given synchronisation group, said certificate and said priority being set depending on the given type of data.

6. The method according to Claim 1, further comprising transmitting data required to operate as the master device from said master device to a slave device which is targeted to transfer a master privilege, thereby transferring the master privilege to the slave device.

7. The method according to Claim 6, wherein said master privilege transferring is performed after it is confirmed that there is no slave device other than said master device and said slave device capable of communicating with said master device and set in a master privilege transfer mode.

8. The method according to Claim 1, further comprising exchanging the priority between slave devices that belong to the same data synchronisation group.

9. The method according to Claim 8, wherein said priority exchanging is performed after it is confirmed that there is no slave device other than said master device and said slave device capable of communicating with said master device and set in a priority exchanging mode.

10. The method according to Claim 1, further comprising releasing said master device or said slave device from the data synchronisation group, wherein, if said master device is released, said master device transmits a releasing instruction to a slave device which is capable of communicating with said master device to release said slave device as well as said master device.

11. A method according to claim 4, wherein said registering is performed by transmitting said certificate and said priority from the master device set by a user for a given type of data to the slave device set by the user for the given type of data, said certificate and said priority being set depending on the type of data.

12. An electronic device comprising:

a master unit, when an own device is set by a user as a master, configured to transmit to a slave device a certificate indicating that the slave device belongs to a data synchronization group to which the master device belongs and a priority to be used for solving conflict of data during data synchronisation, thereby registering the slave device as a member of the data synchronization group to which the master device belongs;
a slave unit, when the own device is set by the user as the slave, configured to receive from the master unit the certificate and the priority; and
a data synchronization unit (16, 17) configured to determine whether or not the own device and another device which is capable of communicating with the own device belong to the same data synchronization group by using the certificate, and perform data synchronization between the own device and the other device based on the priority if the own device and the other device belong to the same data synchronization group.

13. A slave device comprising:

a receiver (11, 12) configured to receive, from a master device for a data synchronization group, a certificate indicating that the slave device belongs to the data synchronization group to which the master device belongs and a priority to be used for solving conflict of data during data synchronisation;
a memory (18) configured to store the received certificate and the priority;
a certification unit (15) configured to determine, when a synchronization request from another device is received, whether or not the other device and the slave device belong to the same data synchronization group by using the certificate stored in the memory; and
a synchronization unit (17) configured to perform data synchronization between the other device and the slave device based on the priority stored in the memory when the certification unit determines that the other device and the slave device belong to the same data synchronization group.

14. The device according to claim 13, wherein said memory (18) is configured to store the received certificate and the priority after it is confirmed that there is no device other than said master device and said slave device, capable of communicating with said master device and set in a registration mode.

15. The device according to claim 13, wherein said other device comprises a master device or one or more slave devices.

16. The device according to claim 13, wherein said master device and slave device are adapted to store plural types of data and belong to plural data synchronization groups defined for each of said plural types of data.

17. The device according to claim 16, wherein said memory (18) is configured to store said certificate and said priority from the master device belonging to a given synchronization group defined for a given type of data, said certificate and said priority being set depending on the given type of data.

18. The device according to claim 13, wherein the receiver (11, 12) is configured to receive data required to operate as the master device from said master device, thereby transferring a master privilege to the slave device.

19. The device according to claim 18, wherein said master privilege transferring is performed after it is confirmed that there is no slave device other than said master device and said slave device capable of communicating with said master device and set in a master privilege transfer mode.

20. The device according to claim 13, further comprising an exchanging unit configured to exchange the priority between another slave device that belongs to the same data synchronization group.

21. The device according to claim 20, wherein said priority exchanging is performed after it is confirmed that there is no slave device other than said master device and said slave device capable of communicating with said master device and set in a priority exchanging mode.

22. A data synchronization method for a slave device comprising the steps of:

receiving from a master device for a data synchronization group, a certificate indicating that the slave device belongs to the data synchronization group to which the master device belongs and a priority to be used for solving conflict of data during data synchronization;
storing the received certificate and the priority;
determining, when a synchronization request from another device is received, whether or not the other device and the slave device belong to the same data synchronization group by using the stored certificate; and
performing data synchronization between the other device and the slave device based on the stored priority when it is determined that the other device and the slave device belong to the same data synchronization group.

23. The method according to claim 22, comprising the step of storing the received certificate and the priority after it is confirmed that there is no device other than said master device and said slave device capable of communicating with said master device and set in a registration mode.

24. The method according to claim 22, wherein said other device comprises: a master device or one or more slave devices.

25. The method according to claim 22, wherein said master device and said slave device belong to plural data synchronization groups and comprising storing, in said master device and said slave device, plural types of data defined for each of said plural data synchronization groups.

26. The method according to claim 25, comprising storing said certificate and said priority from the master device belonging to a given synchronization group defined for a given type of data, said certificate and said priority being set depending on the given type of data.

27. The method according to claim 22, comprising receiving data required for a slave device to operate as the master device from said master device, for transferring a master privilege to the slave device.

28. The method according to claim 27, wherein transferring of said master privilege is performed after it is confirmed that there is no slave device other than said master device and said slave device capable of communicating with said master device and set in a master privilege transfer mode.

29. The method according to claim 22, further comprising exchanging the priority between another slave device that belongs to the same data synchronization group.

30. The method according to claim 29, wherein exchanging of priority is performed after it is confirmed that there is no slave device other than said master device and said slave device capable of communicating with said master device and set in a priority exchanging mode.

31. A carrier medium carrying computer readable instructions for controlling a computer to carry out a method of any one of claims 1 to 11 and 22 to 30.

Revendications

1. Procédé de synchronisation de données, comprenant les étapes consistant à :

transmettre d'un dispositif maître à un dispositif esclave un certificat indiquant que le dispositif esclave appartient à un groupe de synchronisation de données auquel appartient le dispositif maître et une priorité à utiliser pour résoudre un conflit de données pendant la synchronisation de données, enregistrant ainsi le dispositif esclave comme un membre du groupe de synchronisation de données auquel appartient le dispositif maître ;
déterminer si oui ou non un premier dispositif esclave et un deuxième dispositif esclave, qui est capable de communiquer avec le premier dispositif esclave, appartiennent au même groupe de synchronisation de données en utilisant le certificat ; et
effectuer la synchronisation de données entre le premier dispositif esclave et le deuxième dispositif esclave sur la base de la priorité si le premier dispositif esclave et le deuxième dispositif esclave appartiennent au même groupe de synchronisation de données.

2. Procédé selon la revendication 1, dans lequel ledit enregistrement est effectué après qu'il a été confirmé qu'il n'existe aucun dispositif autre que ledit dispositif maître et ledit dispositif esclave capable de communiquer avec ledit dispositif maître et paramétré en mode d'enregistrement.

3. Procédé selon la revendication 1, dans lequel ladite synchronisation de données est effectuée entre un seul dispositif maître et un ou plusieurs dispositifs esclaves qui appartiennent au même groupe de synchronisation ou entre au moins deux dispositifs esclaves qui appartiennent au même groupe de synchronisation.

4. Procédé selon la revendication 1, dans lequel ledit

dispositif maître et ledit dispositif esclave stockent plusieurs types de données et appartiennent à plusieurs groupes de synchronisation de données définis pour chacun desdits plusieurs types de données.

5. Procédé selon la revendication 4, dans lequel ledit enregistrement est effectué en transmettant ledit certificat et ladite priorité du dispositif maître appartenant à un groupe de synchronisation donné défini pour un type donné de données au dispositif esclave appartenant au groupe de synchronisation donné, ledit certificat et ladite priorité étant paramétrés en fonction du type donné de données.

6. Procédé selon la revendication 1, comprenant en outre la transmission de données nécessaires pour fonctionner comme le dispositif maître, dudit dispositif maître à un dispositif esclave qui est ciblé pour transférer un privilège de maître, transférant ainsi le privilège de maître au dispositif esclave.

7. Procédé selon la revendication 6, dans lequel ledit transfert de privilège de maître est effectué après qu'il a été confirmé qu'il n'existe aucun dispositif esclave autre que ledit dispositif maître et ledit dispositif esclave capable de communiquer avec ledit dispositif maître et paramétré en mode de transfert de privilège de maître.

8. Procédé selon la revendication 1, comprenant en outre l'échange de la priorité entre des dispositifs esclaves qui appartiennent au même groupe de synchronisation de données.

9. Procédé selon la revendication 8, dans lequel ledit échange de priorité est effectué après qu'il a été confirmé qu'il n'existe aucun dispositif esclave autre que ledit dispositif maître et que ledit dispositif esclave, capable de communiquer avec ledit dispositif maître et paramétré en mode d'échange de priorité.

10. Procédé selon la revendication 1, comprenant en outre la libération dudit dispositif maître ou dudit dispositif esclave du groupe de synchronisation de données, dans lequel, si ledit dispositif maître est libéré, ledit dispositif maître transmet une instruction de libération à un dispositif esclave qui est capable de communiquer avec ledit dispositif maître pour libérer ledit dispositif esclave ainsi que ledit dispositif maître.

11. Procédé selon la revendication 4, dans lequel ledit enregistrement est effectué en transmettant ledit certificat et ladite priorité du dispositif maître paramétré par un utilisateur pour un type donné de données au dispositif esclave paramétré par l'utilisateur pour le type donné de données, ledit certificat et ladite priorité étant paramétrés en fonction du type de données.

12. Dispositif électronique, comprenant :

une unité de maître, quand un dispositif lui-même est paramétré par un utilisateur comme maître, configurée pour transmettre à un dispositif esclave un certificat indiquant que le dispositif esclave appartient à un groupe de synchronisation de données auquel appartient le dispositif maître et une priorité à utiliser pour résoudre un conflit de données pendant la synchronisation de données, enregistrant ainsi le dispositif esclave comme un membre du groupe de synchronisation de données auquel appartient le dispositif maître ;
une unité d'esclave, quand le dispositif lui-même est paramétré par l'utilisateur comme esclave, configurée pour recevoir de l'unité de maître le certificat et la priorité ; et
une unité de synchronisation de données (16, 17) configurée pour déterminer si oui ou non le dispositif lui-même et un autre dispositif, qui est capable de communiquer avec le dispositif lui-même, appartiennent au même groupe de synchronisation de données en utilisant le certificat, et pour effectuer une synchronisation de données entre le dispositif lui-même et l'autre dispositif sur la base de la priorité si le dispositif lui-même et l'autre dispositif appartiennent au même groupe de synchronisation de données.

13. Dispositif esclave, comprenant :

un récepteur (11, 12) configuré pour recevoir, d'un dispositif maître pour un groupe de synchronisation de données, un certificat indiquant que le dispositif esclave appartient au groupe de synchronisation de données auquel appartient le dispositif maître et une priorité à utiliser pour résoudre un conflit de données pendant la synchronisation de données ;
une mémoire (18) configurée pour stocker le certificat reçu et la priorité ;
une unité de certification (15) configurée pour déterminer, quand une demande de synchronisation d'un autre dispositif est reçue, si oui ou non l'autre dispositif et le dispositif esclave appartiennent au même groupe de synchronisation de données en utilisant le certificat stocké dans la mémoire ; et
une unité de synchronisation (17) configurée pour effectuer la synchronisation de données entre l'autre dispositif et le dispositif esclave sur la base de la priorité stockée dans la mémoire quand l'unité de certification détermine que l'autre dispositif et le dispositif esclave appar-

tiennent au même groupe de synchronisation de données.

14. Dispositif selon la revendication 13, dans lequel ladite mémoire (18) est configurée pour stocker le certificat reçu et la priorité après qu'il a été confirmé qu'il n'existe aucun dispositif autre que ledit dispositif maître et ledit dispositif esclave, capable de communiquer avec ledit dispositif maître et paramétré en mode d'enregistrement.

15. Dispositif selon la revendication 13, dans lequel ledit autre dispositif comprend un dispositif maître ou un ou plusieurs dispositifs esclaves.

16. Dispositif selon la revendication 13, dans lequel lesdits dispositifs maître et esclave sont adaptés pour stocker plusieurs types de données et appartiennent à plusieurs groupes de synchronisation de données définis pour chacun desdits plusieurs types de données.

17. Dispositif selon la revendication 16, dans lequel ladite mémoire (18) est configurée pour stocker ledit certificat et ladite priorité provenant du dispositif maître appartenant à un groupe de synchronisation donné défini pour un type donné de données, ledit certificat et ladite priorité étant paramétrés en fonction du type donné de données.

18. Dispositif selon la revendication 13, dans lequel le récepteur (11, 12) est configuré pour recevoir des données nécessaires pour fonctionner comme le dispositif maître dudit dispositif maître, transférant ainsi un privilège de maître au dispositif esclave.

19. Dispositif selon la revendication 18, dans lequel ledit transfert de privilège de maître est effectué après qu'il a été confirmé qu'il n'existe aucun dispositif esclave autre que ledit dispositif maître et ledit dispositif esclave capable de communiquer avec ledit dispositif maître et paramétré en mode de transfert de privilège de maître.

20. Dispositif selon la revendication 13, comprenant en outre une unité d'échange configurée pour échanger la priorité entre un autre dispositif esclave qui appartient au même groupe de synchronisation de données.

21. Dispositif selon la revendication 20, dans lequel ledit échange de priorité est effectué après qu'il a été confirmé qu'il n'existe aucun dispositif esclave autre que ledit dispositif maître et ledit dispositif esclave capable de communiquer avec ledit dispositif maître et paramétré en mode d'échange de priorité.

22. Procédé de synchronisation de données pour un dispositif esclave comprenant les étapes consistant à :

recevoir d'un dispositif maître pour un groupe de synchronisation de données, un certificat indiquant que le dispositif esclave appartient au groupe de synchronisation de données auquel appartient le dispositif maître et une priorité à utiliser pour résoudre un conflit de données pendant la synchronisation de données ;
stocker le certificat reçu et la priorité ;
déterminer, quand une demande de synchronisation d'un autre dispositif est reçue, si oui ou non l'autre dispositif et le dispositif esclave appartiennent au même groupe de synchronisation de données en utilisant le certificat stocké ; et
effectuer la synchronisation de données entre l'autre dispositif et le dispositif esclave sur la base de la priorité stockée quand il est déterminé que l'autre dispositif et le dispositif esclave appartiennent au même groupe de synchronisation de données.

23. Procédé selon la revendication 22, comprenant l'étape consistant à stocker le certificat reçu et la priorité après qu'il a été confirmé qu'il n'existe aucun dispositif autre que ledit dispositif maître et que ledit dispositif esclave, capable de communiquer avec ledit dispositif maître et paramétré en mode d'enregistrement.

24. Procédé selon la revendication 22, dans lequel ledit autre dispositif comprend : un dispositif maître ou un ou plusieurs dispositifs esclaves.

25. Procédé selon la revendication 22, dans lequel ledit dispositif maître et ledit dispositif esclave appartiennent à plusieurs groupes de synchronisation de données et comprenant le stockage, dans ledit dispositif maître et ledit dispositif esclave, de plusieurs types de données définis pour chacun desdits plusieurs groupes de synchronisation de données.

26. Procédé selon la revendication 25, comprenant le stockage dudit certificat et de ladite priorité dudit dispositif maître appartenant à un groupe de synchronisation donné défini pour un type donné de données, ledit certificat et ladite priorité étant paramétrés en fonction du type donné de données.

27. Procédé selon la revendication 22, comprenant la réception de données nécessaires pour qu'un dispositif esclave fonctionne comme le dispositif maître dudit dispositif maître, pour transférer un privilège de maître au dispositif esclave.

28. Procédé selon la revendication 27, dans lequel le transfert dudit privilège de maître est effectué après

qu'il a été confirmé qu'il n'existe aucun dispositif esclave autre que ledit dispositif maître et que ledit dispositif esclave capable de communiquer avec ledit dispositif maître et paramétré en mode de transfert de privilège de maître.

29. Procédé selon la revendication 22, comprenant en outre l'échange de la priorité entre un autre dispositif esclave qui appartient au même groupe de synchronisation de données.

30. Procédé selon la revendication 29, dans lequel l'échange de priorité est effectué après qu'il a été confirmé qu'il n'existe aucun dispositif esclave autre que ledit dispositif maître et que ledit dispositif esclave capable de communiquer avec ledit dispositif maître et paramétré en mode d'échange de priorité.

31. Moyen de support supportant des instructions lisibles par ordinateur pour commander un ordinateur pour qu'il exécute un procédé selon l'une quelconque des revendications 1 à 11 et 22 à 30.

Patentansprüche

1. Ein Datensynchronisationsverfahren umfassend die Schritte:

   Übertragen eines Zertifikats von einem Haupt- bzw. Master-Gerät zu einem Folge- bzw. Slave-Gerät, wobei das Zertifikat andeutet, dass das Slave-Gerät zu einer Datensynchronisationsgruppe gehört, zu der das Master-Gerät gehört, sowie einer Priorität, die während der. Datensynchronisation zum Lösen von Datenkonflikten einzusetzen ist, wodurch das Slave-Gerät als ein Mitglied der Datensynchronisationsgruppe, zu der das Master-Gerät gehört, registriert wird;
   Bestimmen ob oder nicht ein erstes Slave-Gerät und ein zweites Slave-Gerät, das in der Lage ist, mit dem zu derselben Datensynchronisationsgruppe gehörenden, ersten Slave-Gerät unter Verwendung des Zertifikats zu kommunizieren;
   und Ausführen einer Datensynchronisation zwischen dem ersten Slave-Gerät und dem zweiten Slave-Gerät gestützt auf die Priorität, falls das erste Slave-Gerät und das zweite Slave-Gerät zu derselben Datensynchronisationsgruppe gehören.

2. Das Verfahren nach Anspruch 1, wobei das Registrieren ausgeführt wird, nachdem bestätigt worden ist, dass es kein anderes Gerät als das erste Master-Gerät und das Slave-Gerät gibt, das in der Lage ist, mit dem Master-Gerät zu kommunizieren, und in einen Registrierungsmode eingestellt ist.

3. Das Verfahren nach Anspruch 1, wobei die Datensynchronisation zwischen einem Master-Gerät und einem oder mehreren Slave-Geräten, die zu derselben Synchronisationsgruppe gehören, oder zwischen wenigstens zwei Slave-Geräten, die zu derselben Synchronisationsgruppe gehören, ausgeführt wird.

4. Das Verfahren nach Anspruch 1, wobei das Master-Gerät und Slave-Geräte vielfache Datentypen speichern und zu mehreren Datensynchronisationsgruppen, die jeweils für die vielfachen Datentypen definiert sind, gehören.

5. Das Verfahren nach Anspruch 4, wobei das Registrieren ausgeführt wird, indem das Zertifikat und die Priorität von dem Master-Gerät, das zu einer gegebenen Synchronisationsgruppe gehört, die für einen gegebenen Datentyp definiert ist, zu dem Slave-Gerät, das zu der gegebenen Synchronisationsgruppe gehört, wobei das Zertifikat und die Priorität in Abhängigkeit vom gegebenen Datentyp eingestellt werden.

6. Das Verfahren nach Anspruch 1, ferner umfassend das Übertragen von Daten, die zum Operieren als das Master-Gerät des Master-Geräts erforderlich sind, zu einem Slave-Gerät, das darauf gerichtet ist, ein Master-Privileg zu übertragen, wodurch das Master-Privileg auf das Slave-Gerät übertragen wird.

7. Das Verfahren nach Anspruch 6, wobei das Übertragen des Master-Privilegs ausgeführt wird, nachdem bestätigt worden ist, dass es kein anderes Slave-Gerät verschieden von dem Master-Gerät und dem Slave-Gerät gibt, das in der Lage ist, mit dem Master-Gerät zu kommunizieren, und in einen Master-Privileg-Übertragungsmodus eingestellt ist.

8. Das Verfahren nach Anspruch 1, ferner umfassend Austauschen der Priorität zwischen zu der gleichen Datensynchronisationsgruppe gehörenden Slave-Geräten.

9. Das Verfahren nach Anspruch 8, wobei das Austauschen der Priorität ausgeführt wird, nachdem bestätigt worden ist, dass es kein Slave-Gerät verschieden von dem Master-Gerät und dem Slave-Gerät gibt, das in der Lage ist, mit dem Master-Gerät zu kommunizieren, und in einen Prioritätsaustauschmodus eingestellt ist.

10. Das Verfahren nach Anspruch 1, ferner umfassend Freigeben des Master-Geräts oder Slave-Geräts aus der Datensynchronisationsgruppe, wobei, wenn das Master-Gerät freigegeben wird, das Master-Ge-

rät eine Freigabeanweisung an ein Slave-Gerät überträgt, welches in der Lage ist, mit dem Master-Gerät zu kommunizieren, um sowohl das Slave-Gerät als auch das Master-Gerät freizugeben.

11. Ein Verfahren nach Anspruch 4, wobei das Registrieren ausgeführt wird, indem das Zertifikat und die Priorität von dem durch einen Benutzer für einen gegebenen Datentyp eingestellten Master-Gerät übertragen wird, an das von dem Benutzer für einen gegebenen Datentyp eingestellte Slave-Gerät, wobei das Zertifikat und die Priorität in Abhängigkeit von dem Datentyp eingestellt werden.

12. Ein elektronisches Gerät umfassend:

eine Haupt- bzw. Master-Einheit, wenn ein eigenes Gerät durch einen Benutzer als ein Master eingestellt ist, konfiguriert zum Übertragen eines Zertifikat zu einem Slave-Gerät, wobei das Zertifikat andeutet, dass das Slave-Gerät zu einer Datengruppe gehört, zu der das Master-Gerät gehört, sowie einer Priorität, die zum Lösen von Datenkonflikten während der Datensynchronisation einzusetzen ist, wodurch das Slave-Gerät als ein Mitglied der Datensynchronisationsgruppe, zu der das Master-Gerät gehört, registriert wird;
eine Slave-Einheit, wenn das eigene Gerät durch den Benutzer als der Slave eingestellt wird, konfiguriert, zum Empfangen des Zertifikats und der Priorität von der Master-Einheit; und
eine Datensynchronisationseinheit (16, 17), konfiguriert zum Bestimmen unter Verwendung des Zertifikats, ob oder nicht das eigene Gerät und ein anderes Gerät, das in der Lage ist, mit dem eigenen Gerät zu kommunizieren, zu der gleichen Datensynchronisationsgruppe gehört, und Ausführen der Datensynchronisation zwischen dem eigenen Gerät und dem anderen Gerät gestützt auf die Priorität, falls das eigene Gerät und das andere Gerät zu der gleichen Datensynchronisationsgruppe gehören.

13. Ein Slave-Gerät umfassend:

einen Empfänger (11, 12), von einem Master-Gerät für eine Datensynchronisationsgruppe konfiguriert zum Empfangen eines Zertifikats, das andeutet, dass das Slave-Gerät zu der Datensynchronisationsgruppe gehört, zu der das Master-Gerät gehört, sowie einer Priorität, die zum Lösen von Datenkonflikten während der Datensynchronisation einzusetzen ist;
einen Speicher (18), konfiguriert zum Speichern des empfangenen Zertifikats und der Priorität;
eine Zertifizierungseinheit (15), konfiguriert zum Bestimmen, unter Verwendung des in dem Speicher gespeicherten Zertifikats, wenn eine Synchronisationsanfrage von einem anderen Gerät empfangen wird, ob oder nicht das andere Gerät und das Slave-Gerät zu derselben Datensynchronisationsgruppe gehören; und
eine Synchronisationseinheit (17), konfiguriert zum Ausführen von Datensynchronisation, gestützt auf der in dem Speicher gespeicherten Priorität, zwischen dem anderen Gerät und dem Slave-Gerät, wenn die Zertifizierungseinheit feststellt, dass das andere Gerät und das Slave-Gerät zu der gleichen Datensynchronisationsgruppe gehören.

14. Das Gerät nach Anspruch 13, wobei der Speicher (18) konfiguriert ist zum Speichern des empfangenen Zertifikats und der Priorität, nachdem bestätigt worden ist, dass es kein Gerät verschieden von dem Master-Gerät und dem Slave-Gerät gibt, das in der Lage ist, mit dem Master-Gerät zu kommunizieren, und in einen Registrierungsmodus eingestellt ist.

15. Das Gerät nach Anspruch 13, wobei das andere Gerät ein Master-Gerät oder eine oder mehrere Slave-Geräte umfasst.

16. Das Gerät nach Anspruch 13, wobei das Master-Gerät und das Slave-Gerät dazu angepasst sind, vielfältige Datentypen zu speichern und zu vielfältigen, jeweils für die vielfältigen Datentypen definierten Datensynchronisationsgruppen zu gehören.

17. Das Verfahren nach Anspruch 16, wobei der Speicher (18) konfiguriert ist zum Speichern des Zertifikats und der Priorität von dem Master-Gerät, das zu einer gegebenen, für einen gegebenen Datentypen definierten Synchronisationsgruppe gehört, wobei das Zertifikat und die Priorität in Abhängigkeit von dem gegebenen Datentyp eingestellt werden.

18. Das Gerät nach Anspruch 13, wobei der Empfänger (11, 12) konfiguriert ist zum Empfangen von Daten, die zum Operieren als das Master-Gerät des Master-Geräts benötigt werden, wodurch ein Master-Privileg auf das Slave-Gerät übertragen wird.

19. Das Gerät nach Anspruch 18, wobei das Übertragen des Master-Privilegs ausgeführt wird, nachdem bestätigt worden ist, dass es kein Slave-Gerät verschieden von dem Master-Gerät gibt, und dass das Slave-Gerät in der Lage ist, mit dem Master-Gerät zu kommunizieren, und in einen Master-Privileg-Übertragungsmodus eingestellt ist.

20. Das Gerät nach Anspruch 13, ferner umfassend eine Austauscheinheit, konfiguriert zum Austauschen der Priorität zwischen einem anderen, zu der gleichen

Datensynchronisationsgruppe gehörenden Slave-Gerät.

21. Das Gerät nach Anspruch 20, wobei das Austauschen der Priorität ausgeführt wird, nachdem bestätigt worden ist, dass es kein Slave-Gerät verschieden von dem Master-Gerät und dem Slave-Gerät gibt, das in der Lage ist, mit dem Master-Gerät zu kommunizieren, und in einen Prioritätsaustauschmodus eingestellt ist.

22. Ein Datensynchronisationsverfahren für ein Slave-Gerät, umfassend die Schritte:

Empfangen eines Zertifikats von einem Haupt- bzw. Master-Gerät für eine Datensynchronisationsgruppe, wobei das Zertifikat andeutet, dass das Slave-Gerät zu der Datensynchronisationsgruppe gehört, zu der das Master-Gerät gehört, sowie einer Priorität, die zum Lösen von Datenkonflikten während der Datensynchronisation einzusetzen ist;
Abspeichern des empfangenen Zertifikats und der Priorität;
Bestimmen, unter Verwendung des gespeicherten Zertifikats, wann eine Synchronisationsanfrage von einem anderen Gerät empfangen wird, ob oder nicht das andere Gerät und das Slave-Gerät zu der gleichen Datensynchronisationsgruppe gehören; und Ausführen einer Datensynchronisation zwischen dem anderen Gerät und dem Slave-Gerät gestützt auf die gespeicherte Priorität, wenn bestimmt worden ist, dass das andere Gerät und das Slave-Gerät zu dergleichen Datensynchronisationsgruppe gehören.

23. Das Verfahren nach Anspruch 22, umfassend den Schritt des Abspeicherns des empfangenen Zertifikats und der Priorität, nachdem bestätigt worden ist, dass es kein Gerät anders als das Master-Gerät und das Slave-Gerät gibt, das dazu in der Lage ist, mit dem Master-Gerät zu kommunizieren, und in einen Registrierungsmodus eingestellt ist.

24. Das Verfahren nach Anspruch 22, wobei das andere Gerät ein Master-Gerät und eine oder mehrere Slave-Geräte umfasst.

25. Das Verfahren nach Anspruch 22, wobei das Master-Gerät und die Slave-Geräte zu mehreren Datensynchronisationsgruppen gehören, und das umfasst Speichern in dem Master-Gerät und dem Slave-Gerät von vielfältigen, für jedes der vielfältigen Datensynchronisationsgruppen definierten Typen von Daten.

26. Das Verfahren nach Anspruch 25, umfassend Speichern des Zertifikats und der Priorität, die von dem zu einer gegebenen, für einen gegebenen Datentyp definierten Synchronisationsgruppe gehörenden Master-Gerät, wobei das Zertifikat und die Priorität in Abhängigkeit von dem gegebenen Datentyp eingestellt werden.

27. Das Verfahren nach Anspruch 22, umfassend Empfangen von Daten, die für ein Slave-Gerät benötigt werden zum Operieren als das Master-Gerät von dem Master-Gerät, zum Übertragen eines Master-Privilegs auf das Slave-Gerät.

28. Das Verfahren nach Anspruch 27, wobei das Übertragen des Master-Privilegs ausgeführt wird, nachdem bestätigt worden ist, dass es kein anderes Slave-Gerät verschieden von dem Master-Gerät und dem Slave-Gerät gibt, das in der Lage ist, mit dem Master-Gerät zu kommunizieren, und in einen Master-Privileg-Übertragungsmodus eingestellt ist.

29. Das Verfahren nach Anspruch 22, ferner umfassend das Austauschen der Priorität zwischen einem anderen, zu der gleichen Datensynchronisationsgruppe gehörenden Slave-Gerät.

30. Das Verfahren nach Anspruch 29, wobei das Austauschen der Priorität ausgeführt wird, nachdem bestätigt worden ist, dass es kein Slave-Gerät gibt verschieden von dem Master-Gerät und dem Slave-Gerät, das in der Lage ist, mit dem Master-Gerät zu kommunizieren, und in einen Prioritätsaustauschmodus eingestellt ist.

31. Ein Speichermedium, das computerlesbare Anweisungen trägt zum Steuern eines Computers, um ein Verfahren nach einem der Ansprüche 1 bis 11 und 22 bis 30 auszuführen.

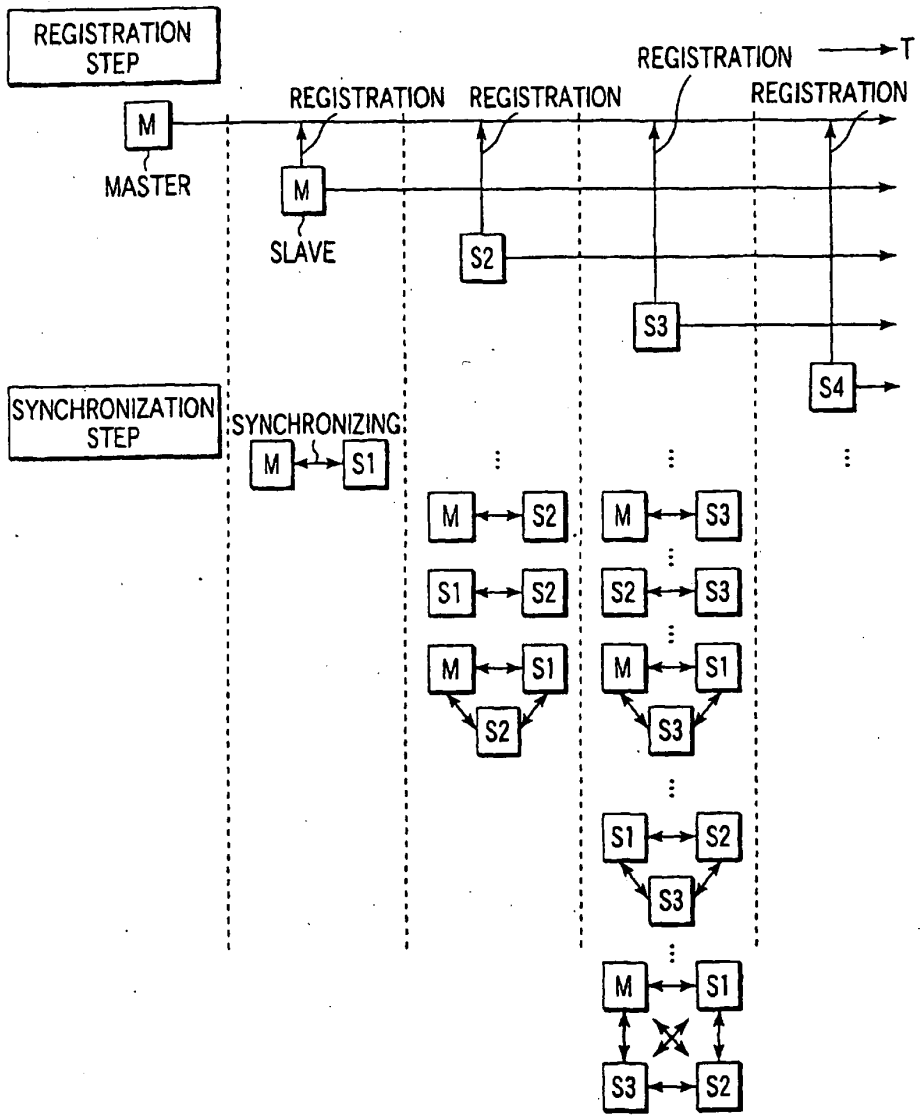

REGISTRATION STEP
REGISTRATION
REGISTRATION
REGISTRATION
REGISTRATION
T
M
MASTER
M
SLAVE
S2
S3
S4
SYNCHRONIZATION STEP
SYNCHRONIZING
M
S1
M
S2
M
S3
S1
S2
S2
S3
M
S1
M
S1
S2
S3
S1
S2
S3
M
S1
S3
S2

FIG. 1

TYPE OF DATA SET	MASTER/ SLAVE	PRIORITY DATA	CERTIFICATE DATA	IDENTIFICATION DATA	COUNTER VALUE
ADDRESS DIRECTORY	MASTER	1	01101110	10101010	12
TIME	SLAVE	5	00111010	01010101	———
SCHEDULE	———	———	———	———	———
TASK LIST	———	———	———	———	———
E-MAIL	SLAVE	8	11010011	11110100	———
⋮	⋮	⋮	⋮	⋮	⋮

FIG. 2

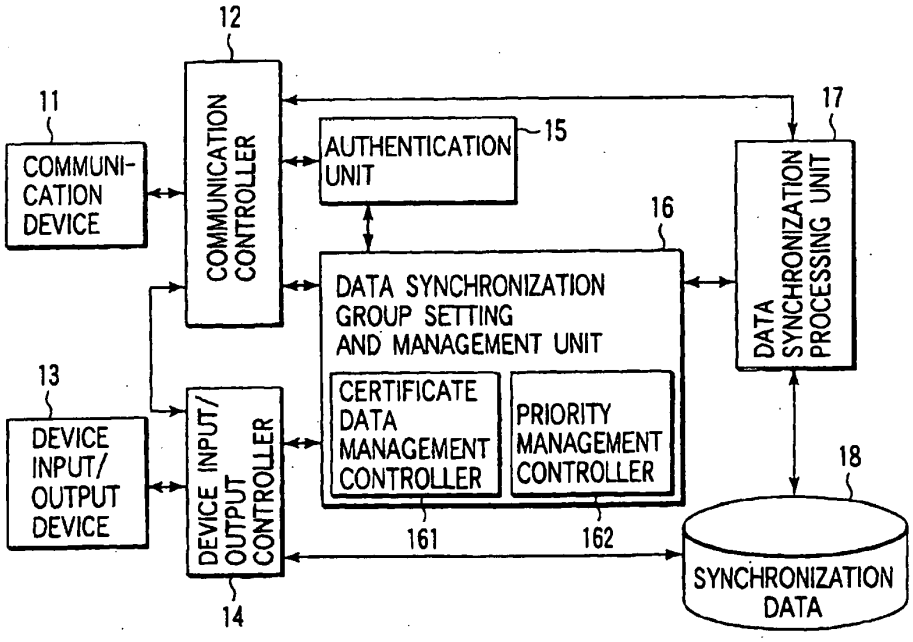


FIG. 3

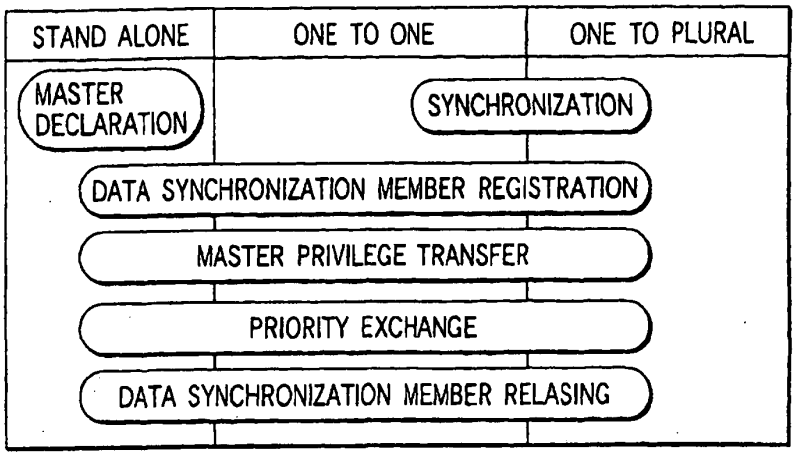

STAND ALONE
ONE TO ONE
ONE TO PLURAL
MASTER DECLARATION
SYNCHRONIZATION
DATA SYNCHRONIZATION MEMBER REGISTRATION
MASTER PRIVILEGE TRANSFER
PRIORITY EXCHANGE
DATA SYNCHRONIZATION MEMBER RELASING

FIG. 4

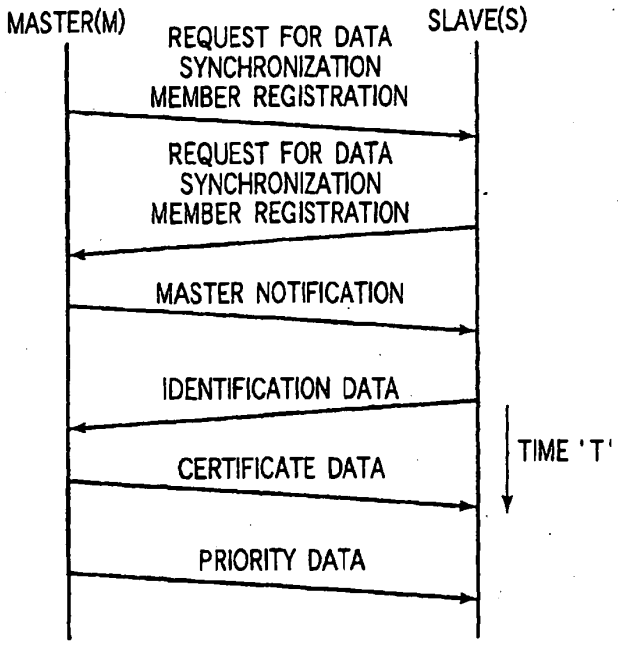

MASTER(M)
SLAVE(S)
REQUEST FOR DATA SYNCHRONIZATION MEMBER REGISTRATION
REQUEST FOR DATA SYNCHRONIZATION MEMBER REGISTRATION
MASTER NOTIFICATION
IDENTIFICATION DATA
CERTIFICATE DATA
PRIORITY DATA
TIME 'T'

FIG. 5

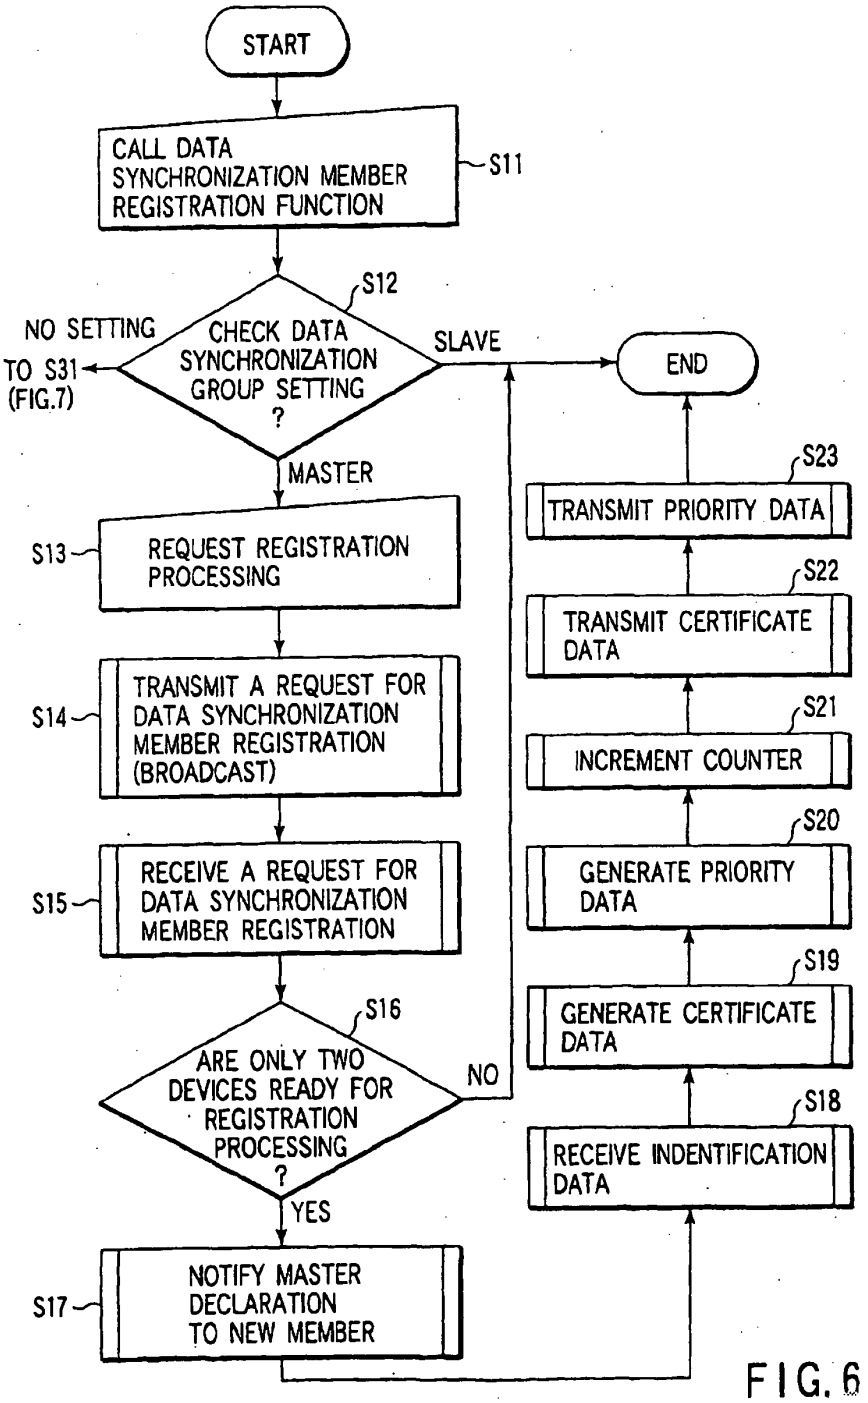

START
CALL DATA SYNCHRONIZATION MEMBER REGISTRATION FUNCTION
S11
S12
CHECK DATA SYNCHRONIZATION GROUP SETTING ?
NO SETTING
TO S31 (FIG.7)
SLAVE
END
MASTER
S13
REQUEST REGISTRATION PROCESSING
S14
TRANSMIT A REQUEST FOR DATA SYNCHRONIZATION MEMBER REGISTRATION (BROADCAST)
S15
RECEIVE A REQUEST FOR DATA SYNCHRONIZATION MEMBER REGISTRATION
S16
ARE ONLY TWO DEVICES READY FOR REGISTRATION PROCESSING ?
NO
YES
S17
NOTIFY MASTER DECLARATION TO NEW MEMBER
S18
RECEIVE INDENTIFICATION DATA
S19
GENERATE CERTIFICATE DATA
S20
GENERATE PRIORITY DATA
S21
INCREMENT COUNTER
S22
TRANSMIT CERTIFICATE DATA
S23
TRANSMIT PRIORITY DATA

FIG. 6

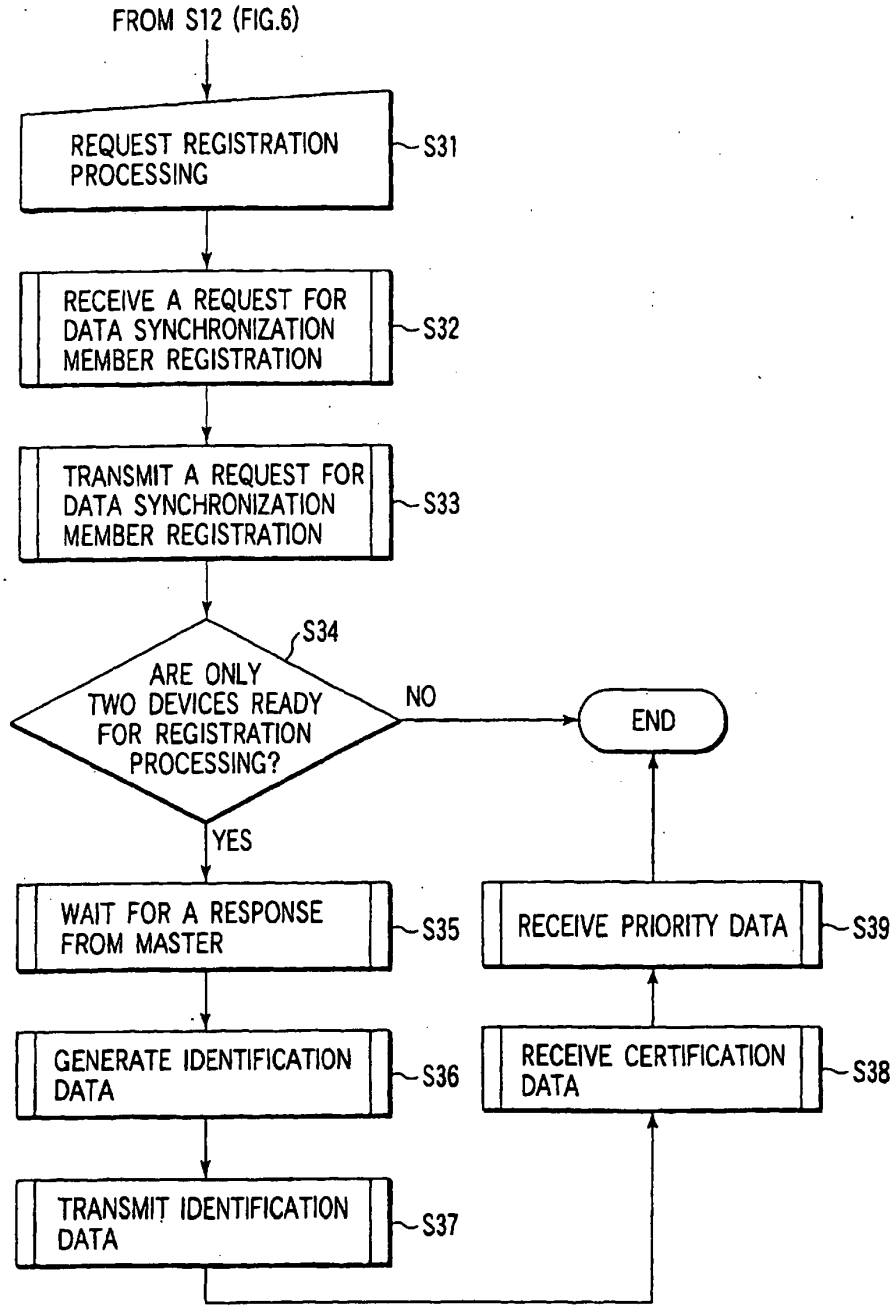

FROM S12 (FIG.6)
REQUEST REGISTRATION PROCESSING
S31
RECEIVE A REQUEST FOR DATA SYNCHRONIZATION MEMBER REGISTRATION
S32
TRANSMIT A REQUEST FOR DATA SYNCHRONIZATION MEMBER REGISTRATION
S33
S34
ARE ONLY TWO DEVICES READY FOR REGISTRATION PROCESSING?
NO
END
YES
WAIT FOR A RESPONSE FROM MASTER
S35
GENERATE IDENTIFICATION DATA
S36
TRANSMIT IDENTIFICATION DATA
S37
RECEIVE CERTIFICATION DATA
S38
RECEIVE PRIORITY DATA
S39

FIG. 7

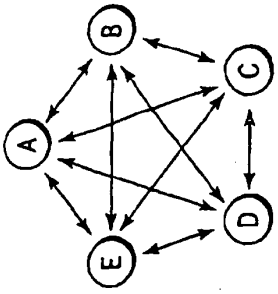


FIG. 8

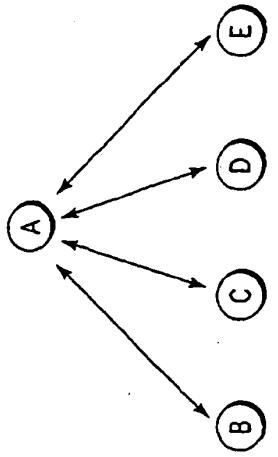


FIG. 9

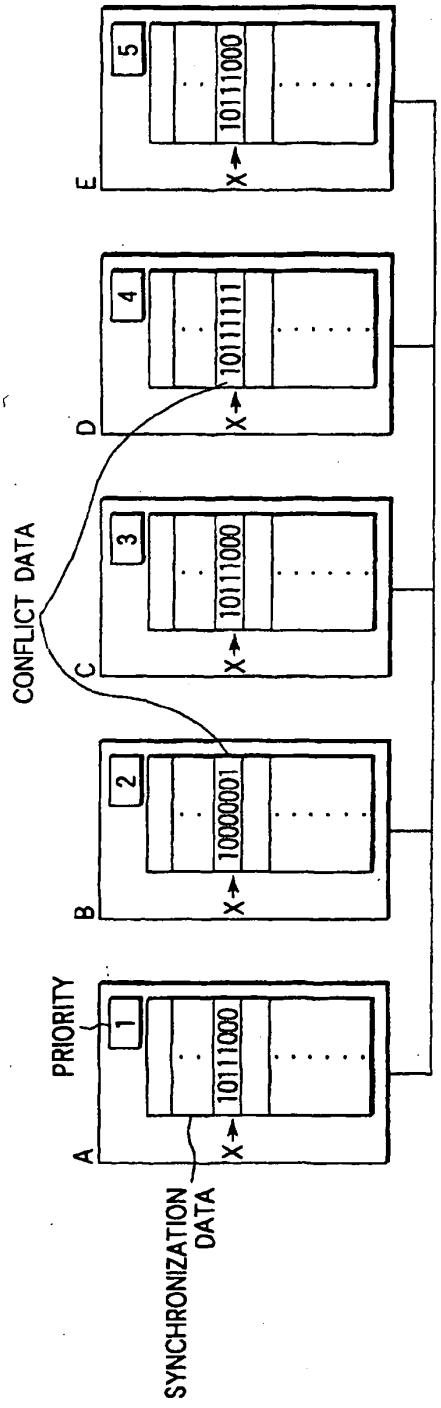


FIG. 10

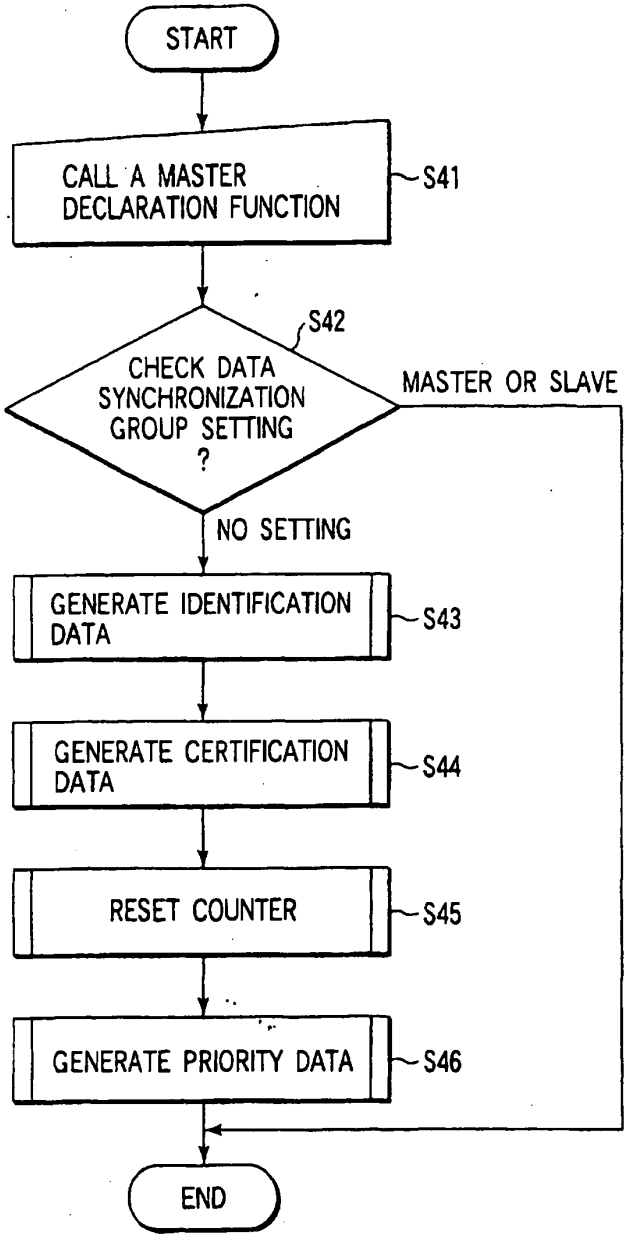

START
CALL A MASTER DECLARATION FUNCTION
S41
S42
CHECK DATA SYNCHRONIZATION GROUP SETTING ?
MASTER OR SLAVE
NO SETTING
GENERATE IDENTIFICATION DATA
S43
GENERATE CERTIFICATION DATA
S44
RESET COUNTER
S45
GENERATE PRIORITY DATA
S46
END

FIG. 11

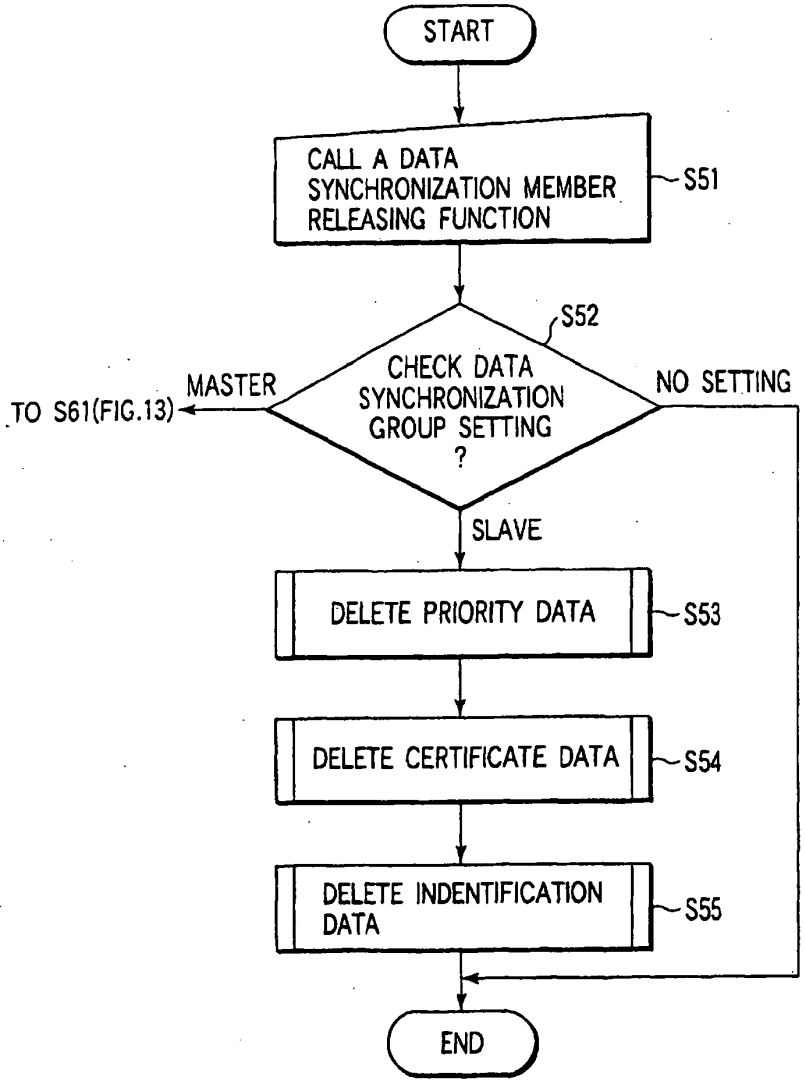

START
CALL A DATA SYNCHRONIZATION MEMBER RELEASING FUNCTION
S51
S52
CHECK DATA SYNCHRONIZATION GROUP SETTING ?
MASTER
TO S61(FIG.13)
NO SETTING
SLAVE
DELETE PRIORITY DATA
S53
DELETE CERTIFICATE DATA
S54
DELETE INDENTIFICATION DATA
S55
END

FIG. 12

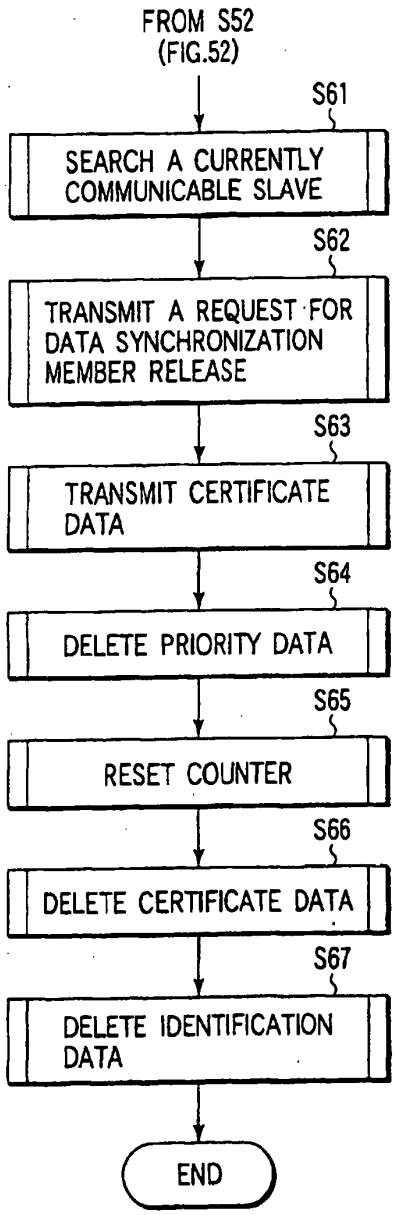

FROM S52
(FIG.52)
S61
SEARCH A CURRENTLY COMMUNICABLE SLAVE
S62
TRANSMIT A REQUEST FOR DATA SYNCHRONIZATION MEMBER RELEASE
S63
TRANSMIT CERTIFICATE DATA
S64
DELETE PRIORITY DATA
S65
RESET COUNTER
S66
DELETE CERTIFICATE DATA
S67
DELETE IDENTIFICATION DATA
END

FIG. 13

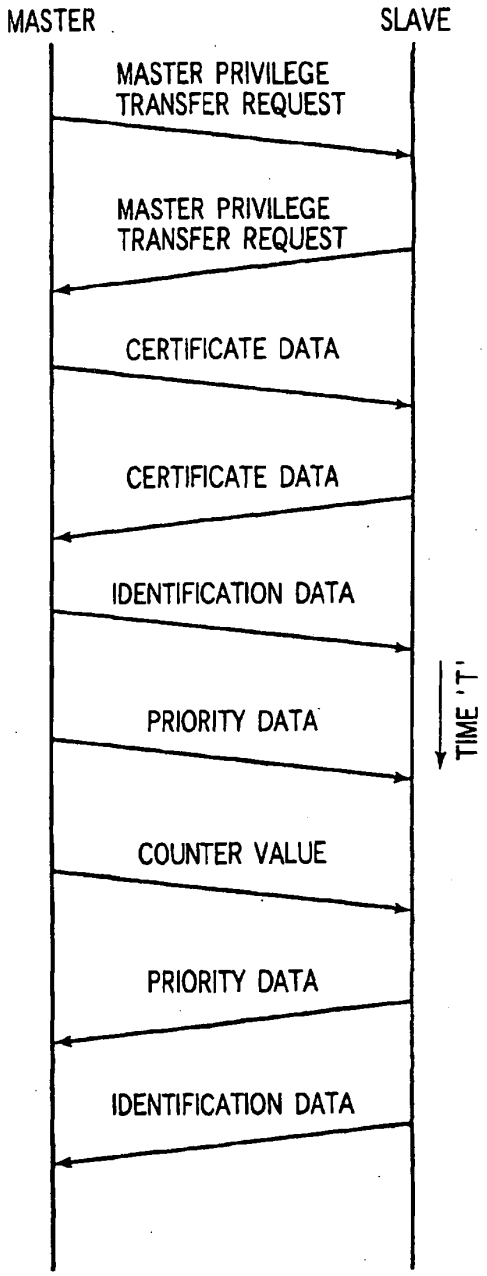

MASTER
SLAVE
MASTER PRIVILEGE TRANSFER REQUEST
MASTER PRIVILEGE TRANSFER REQUEST
CERTIFICATE DATA
CERTIFICATE DATA
IDENTIFICATION DATA
PRIORITY DATA
TIME 'T'
COUNTER VALUE
PRIORITY DATA
IDENTIFICATION DATA

FIG. 14

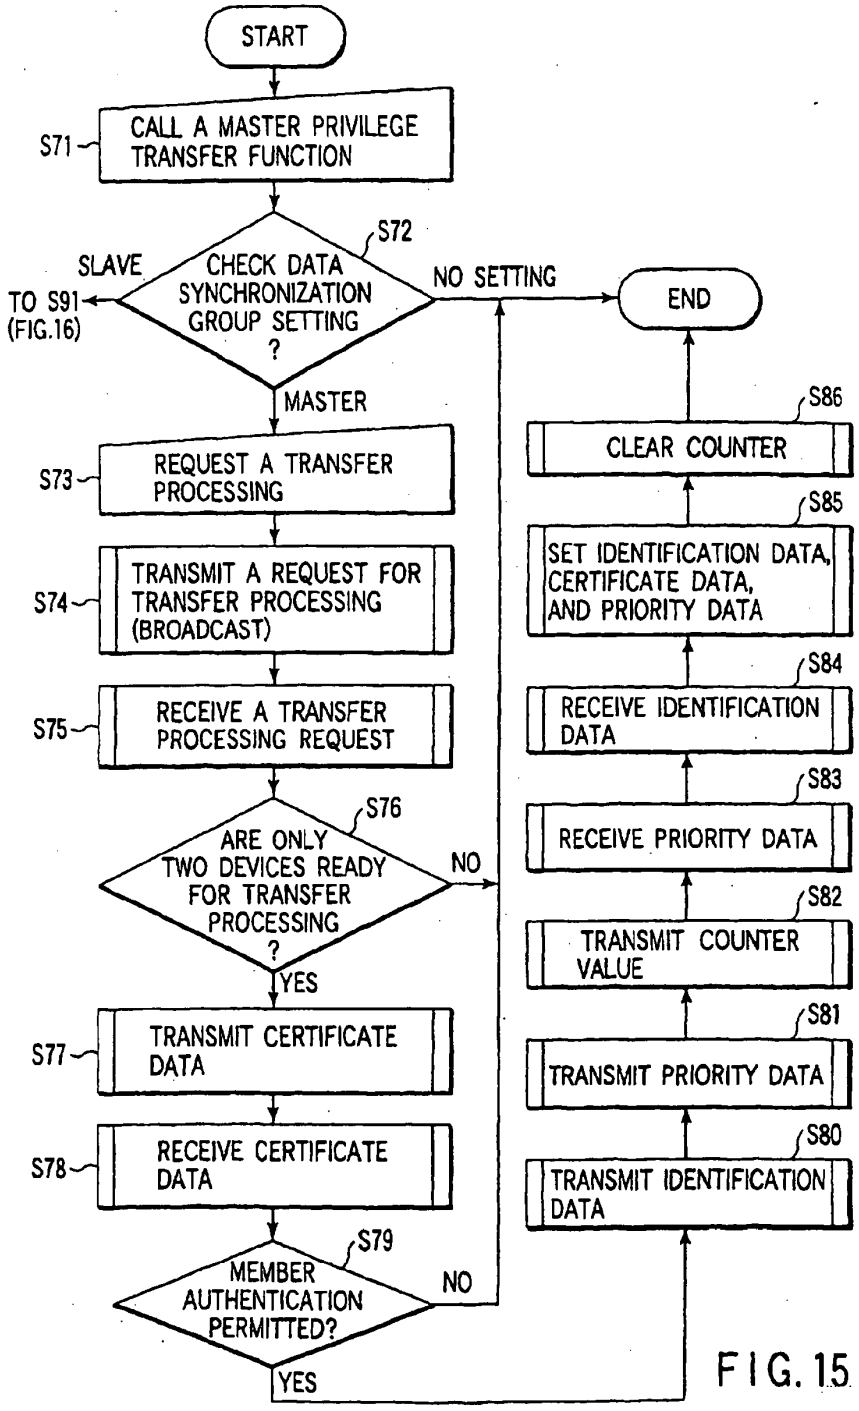

START
S71
CALL A MASTER PRIVILEGE TRANSFER FUNCTION
S72
CHECK DATA SYNCHRONIZATION GROUP SETTING ?
SLAVE
TO S91 (FIG.16)
NO SETTING
END
MASTER
S73
REQUEST A TRANSFER PROCESSING
S74
TRANSMIT A REQUEST FOR TRANSFER PROCESSING (BROADCAST)
S75
RECEIVE A TRANSFER PROCESSING REQUEST
S76
ARE ONLY TWO DEVICES READY FOR TRANSFER PROCESSING ?
NO
YES
S77
TRANSMIT CERTIFICATE DATA
S78
RECEIVE CERTIFICATE DATA
S79
MEMBER AUTHENTICATION PERMITTED?
NO
YES
S86
CLEAR COUNTER
S85
SET IDENTIFICATION DATA, CERTIFICATE DATA, AND PRIORITY DATA
S84
RECEIVE IDENTIFICATION DATA
S83
RECEIVE PRIORITY DATA
S82
TRANSMIT COUNTER VALUE
S81
TRANSMIT PRIORITY DATA
S80
TRANSMIT IDENTIFICATION DATA

FIG. 15

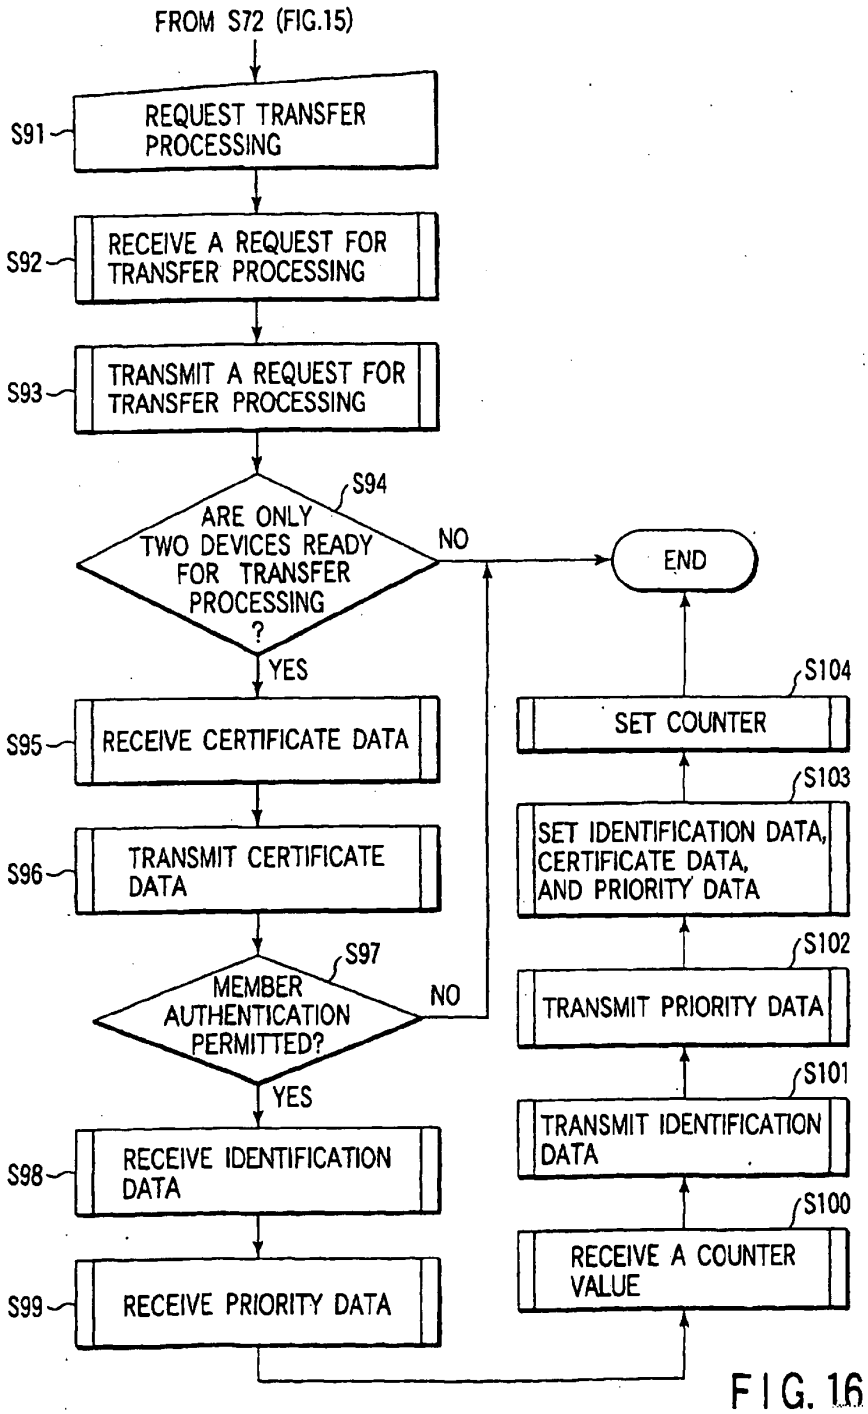

FROM S72 (FIG.15)
S91
REQUEST TRANSFER PROCESSING
S92
RECEIVE A REQUEST FOR TRANSFER PROCESSING
S93
TRANSMIT A REQUEST FOR TRANSFER PROCESSING
S94
ARE ONLY TWO DEVICES READY FOR TRANSFER PROCESSING ?
NO
END
YES
S95
RECEIVE CERTIFICATE DATA
S96
TRANSMIT CERTIFICATE DATA
S97
MEMBER AUTHENTICATION PERMITTED?
NO
YES
S98
RECEIVE IDENTIFICATION DATA
S99
RECEIVE PRIORITY DATA
S104
SET COUNTER
S103
SET IDENTIFICATION DATA, CERTIFICATE DATA, AND PRIORITY DATA
S102
TRANSMIT PRIORITY DATA
S101
TRANSMIT IDENTIFICATION DATA
S100
RECEIVE A COUNTER VALUE

FIG. 16

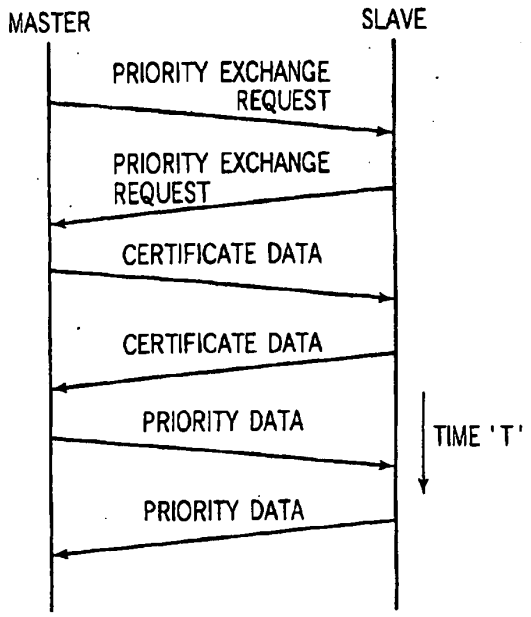

MASTER
SLAVE
PRIORITY EXCHANGE REQUEST
PRIORITY EXCHANGE REQUEST
CERTIFICATE DATA
CERTIFICATE DATA
PRIORITY DATA
PRIORITY DATA
TIME 'T'

FIG. 17

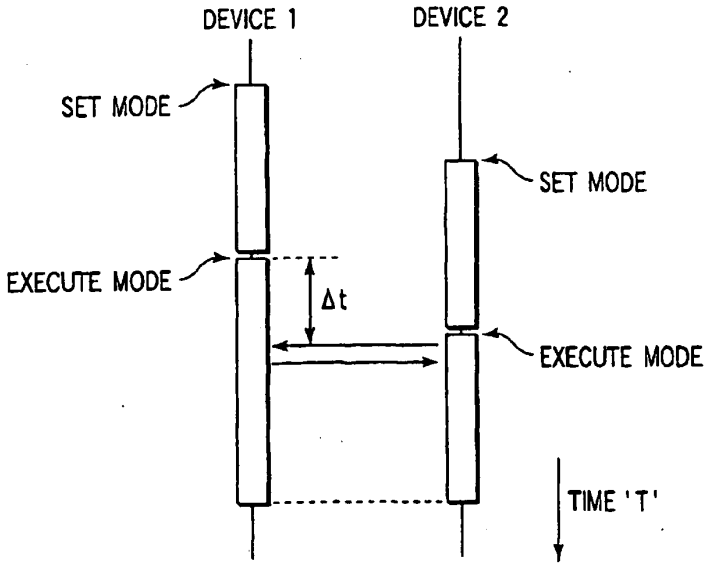

DEVICE 1
DEVICE 2
SET MODE
SET MODE
EXECUTE MODE
Δt
EXECUTE MODE
TIME 'T'

FIG. 19

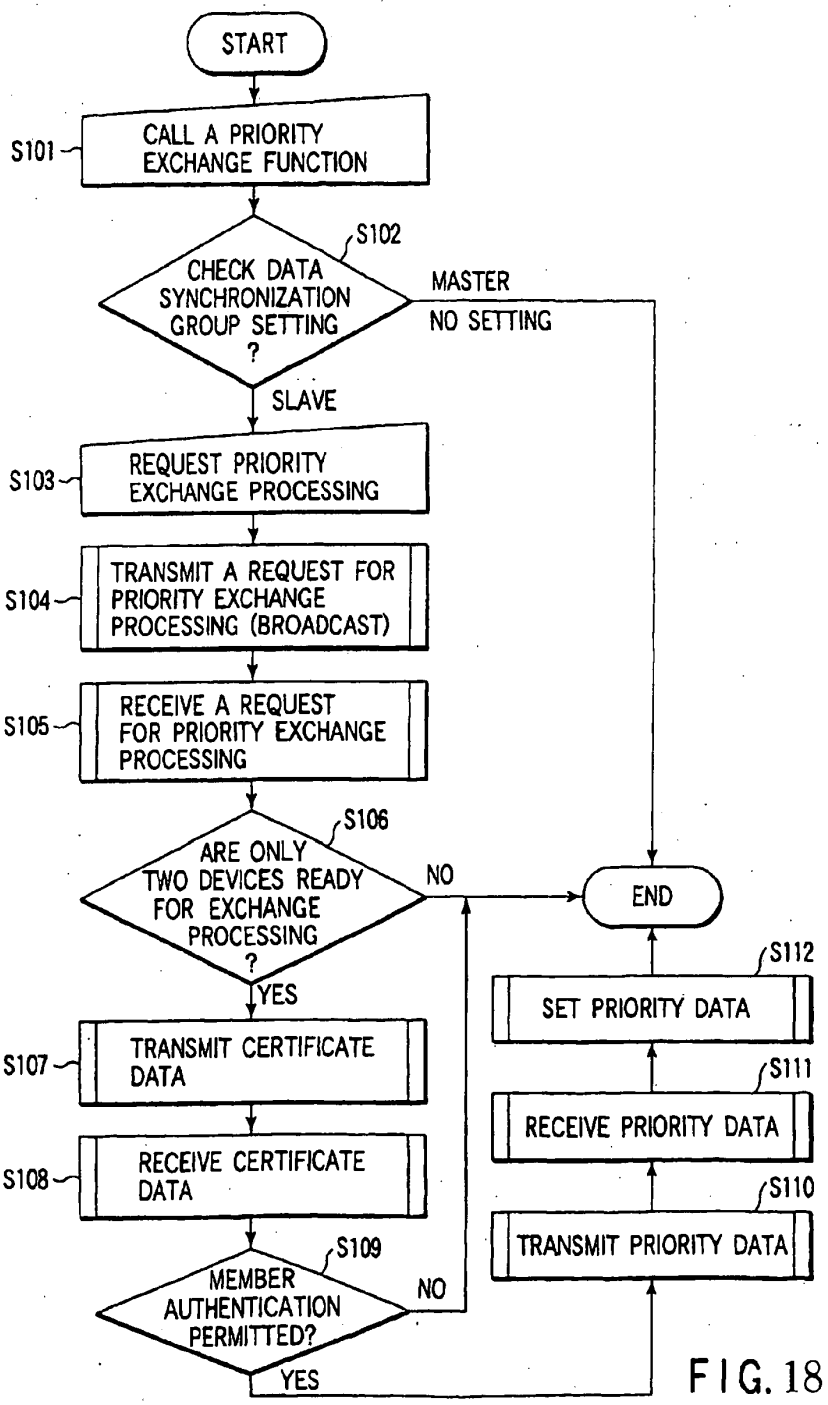

START
S101
CALL A PRIORITY EXCHANGE FUNCTION
S102
CHECK DATA SYNCHRONIZATION GROUP SETTING ?
MASTER
NO SETTING
SLAVE
S103
REQUEST PRIORITY EXCHANGE PROCESSING
S104
TRANSMIT A REQUEST FOR PRIORITY EXCHANGE PROCESSING (BROADCAST)
S105
RECEIVE A REQUEST FOR PRIORITY EXCHANGE PROCESSING
S106
ARE ONLY TWO DEVICES READY FOR EXCHANGE PROCESSING ?
NO
YES
END
S107
TRANSMIT CERTIFICATE DATA
S108
RECEIVE CERTIFICATE DATA
S109
MEMBER AUTHENTICATION PERMITTED?
NO
YES
S112
SET PRIORITY DATA
S111
RECEIVE PRIORITY DATA
S110
TRANSMIT PRIORITY DATA

FIG. 18

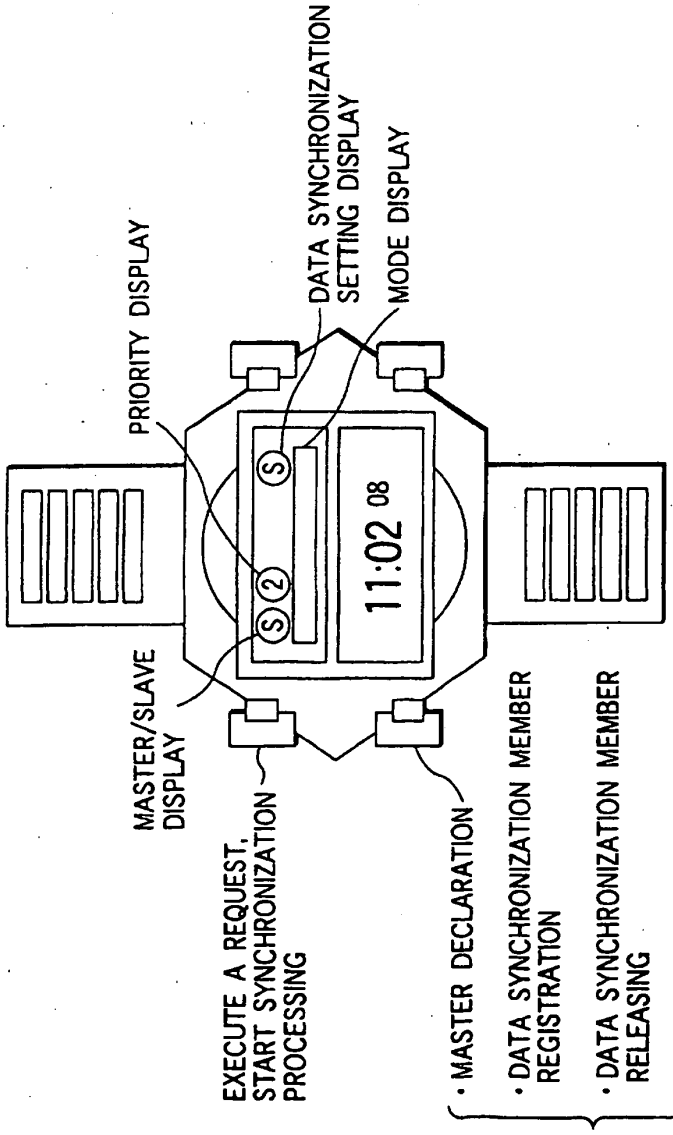

PRIORITY DISPLAY
DATA SYNCHRONIZATION SETTING DISPLAY
MODE DISPLAY
11:02 08
MASTER/SLAVE DISPLAY
EXECUTE A REQUEST, START SYNCHRONIZATION PROCESSING
• MASTER DECLARATION
• DATA SYNCHRONIZATION MEMBER REGISTRATION
• DATA SYNCHRONIZATION MEMBER RELEASING
• MASTER TRANSFERRING AND PRIORITY EXCHANGE

FIG. 20

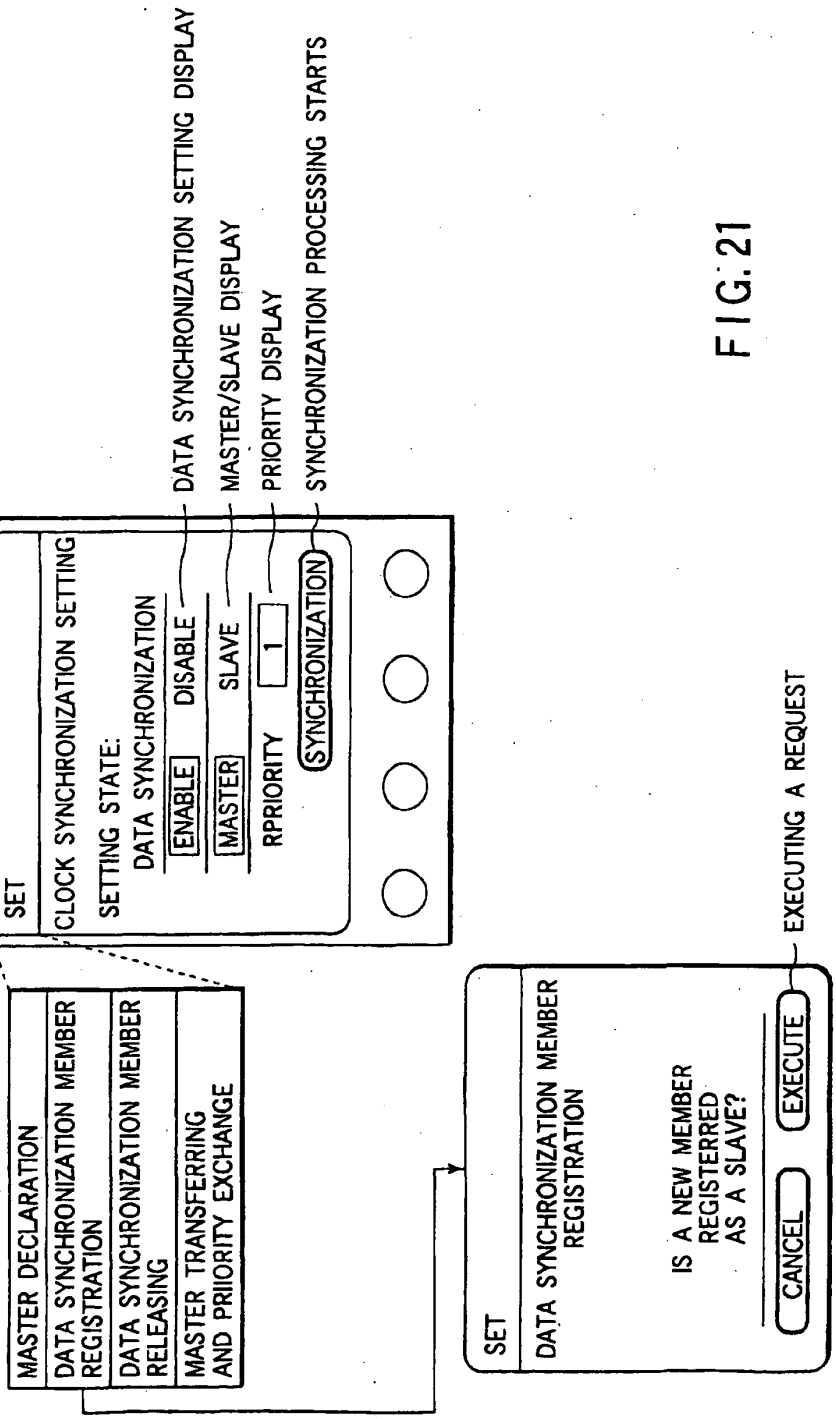

MASTER DECLARATION
DATA SYNCHRONIZATION MEMBER REGISTRATION
DATA SYNCHRONIZATION MEMBER RELEASING
MASTER TRANSFERRING AND PRIIORITY EXCHANGE
SET
CLOCK SYNCHRONIZATION SETTING
SETTING STATE:
DATA SYNCHRONIZATION
ENABLE
DISABLE
MASTER
SLAVE
RPRIORITY
1
SYNCHRONIZATION
DATA SYNCHRONIZATION SETTING DISPLAY
MASTER/SLAVE DISPLAY
PRIORITY DISPLAY
SYNCHRONIZATION PROCESSING STARTS
SET
DATA SYNCHRONIZATION MEMBER REGISTRATION
IS A NEW MEMBER REGISTERRED AS A SLAVE?
CANCEL
EXECUTE
EXECUTING A REQUEST

FIG. 21

PCT WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : G06F 17/30 | A1

(11) International Publication Number: WO 95/25313

(43) International Publication Date: 21 September 1995 (21.09.95)

(21) International Application Number: PCT/EP94/00833

(22) International Filing Date: 16 March 1994 (16.03.94)

(71) Applicant *(for all designated States except US)*: INTERNATIONAL BUSINESS MACHINES CORPORATION [US/US]; Old Orchard Road, Armonk, NY 10504 (US).

(72) Inventors; and

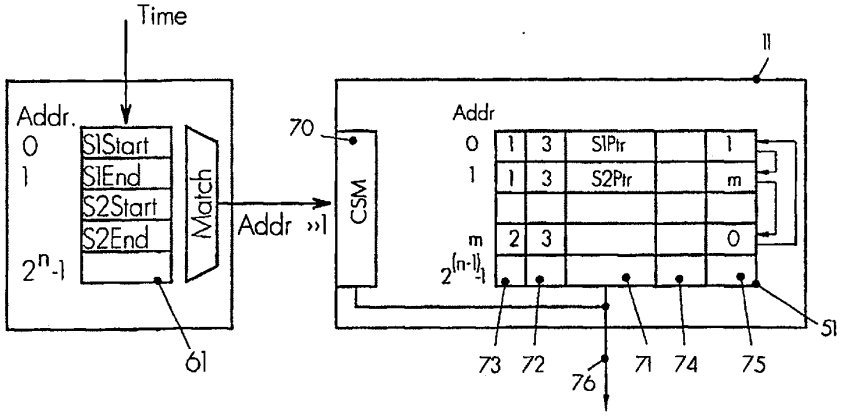
(75) Inventors/Applicants *(for US only)*: KAISERSWERTH, Matthias [DE/CH]; Leigrubenstrasse 2, CH-8805 Richterswil (CH). RÜTSCHE, Erich [CH/FR]; 1500, chemin des Adrets, F-06370 Mouans-Sartoux (FR).

(74) Agent: BARTH, Carl, O.; IBM Research Laboratory, Intellectual Property Dept., Säumerstrasse 4, CH-8803 Rüschlikon (CH).

(81) Designated States: JP, US, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published
With international search report.

(54) Title: METHOD AND APPARATUS FOR SYNCHRONIZATION AND SCHEDULING OF MULTIPLE DATA STREAMS AND REAL TIME TASKS



(57) Abstract

The present application concerns a method and an apparatus for the synchronization and the scheduling of multiple data streams and real time tasks. In accordance with the present invention, pointers pointing to the information streams to be synchronized and processed are kept in the table (51) of a selector (11). The time marks (S1Start etc.) of said information streams are stored in a table (61) of a timer. The system time is checked against said time marks and if a time mark has been reached, the table (51) of said selector (11) is updated. If the status of an information stream is such that its trigger condition is fulfilled, the pointer to said information stream is made available at the output of the selector (11). This pointer is then used to retrieve the respective information stream and to forward it to a playout device.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	**GB**	United Kingdom	**MR**	Mauritania
AU	Australia	**GE**	Georgia	**MW**	Malawi
BB	Barbados	**GN**	Guinea	**NE**	Niger
BE	Belgium	**GR**	Greece	**NL**	Netherlands
BF	Burkina Faso	**HU**	Hungary	**NO**	Norway
BG	Bulgaria	**IE**	Ireland	**NZ**	New Zealand
BJ	Benin	**IT**	Italy	**PL**	Poland
BR	Brazil	**JP**	Japan	**PT**	Portugal
BY	Belarus	**KE**	Kenya	**RO**	Romania
CA	Canada	**KG**	Kyrgystan	**RU**	Russian Federation
CF	Central African Republic	**KP**	Democratic People's Republic of Korea	**SD**	Sudan
CG	Congo			**SE**	Sweden
CH	Switzerland	**KR**	Republic of Korea	**SI**	Slovenia
CI	Côte d'Ivoire	**KZ**	Kazakhstan	**SK**	Slovakia
CM	Cameroon	**LI**	Liechtenstein	**SN**	Senegal
CN	China	**LK**	Sri Lanka	**TD**	Chad
CS	Czechoslovakia	**LU**	Luxembourg	**TG**	Togo
CZ	Czech Republic	**LV**	Latvia	**TJ**	Tajikistan
DE	Germany	**MC**	Monaco	**TT**	Trinidad and Tobago
DK	Denmark	**MD**	Republic of Moldova	**UA**	Ukraine
ES	Spain	**MG**	Madagascar	**US**	United States of America
FI	Finland	**ML**	Mali	**UZ**	Uzbekistan
FR	France	**MN**	Mongolia	**VN**	Viet Nam
GA	Gabon				

DESCRIPTION

Method and Apparatus for Synchronization and Scheduling of Multiple Data Streams and Real Time Tasks

TECHNICAL FIELD

The present invention concerns a method and apparatus for the synchronization and the scheduling of multiple data streams as well as for the scheduling of tasks in operating systems with hard real-time requirements.

BACKGROUND OF THE INVENTION

Multimedia information is composed of several objects that constitute a document. These objects could be of different data types, such as text, still images, video, and sound. In order for these objects to be meaningful to a user, there must be a logical and temporal structure of the objects constituting the documents' contents. To ensure this logical and temporal structure, e.g. when displaying a video sequence on a display device and providing a soundtrack for this video sequence via an adapter card to a speaker, the respective data streams have to be synchronized and scheduled before being sent to the output devices. The same principle applies to other related multimedia information which has to be synchronized in communicating with other devices or users (e.g. teleconferencing), before processing and presentation.

Synchronization and scheduling is a critical task for a wide spectrum of applications ranging from computer-integrated telephony, text with voice annotations, any kind of multi media systems, and interactive services (e.g. interactive television) to cooperative teleworking on documents including

foils and video, and distributed games. A further application, where information streams have to be synchronized, is known as medical imaging, for which high-resolution images must be distributed reliably and fast throughout a hospital.

To date, the synchronization is typically performed in software. However, if the scheduling intervals are too short, or if too many information streams must be synchronized, a software solution, which is very CPU intensive, is often the limiting factor as far as processing speed is concerned.

In real-time systems time-critical tasks have to be scheduled reliably and fast. In the present context, any system in which the time at which an output is produced is significant, is referred to as real-time system. The lag from the input time of a real-time system to its output time, i.e. the delay, must be sufficiently small for acceptable timeliness. With other words, a real-time system has to respond to an externally-generated input stimuli within a finite and specified period. In hard real-time systems, i.e. in those systems where it is absolutely imperative that responses occur within a specified deadline, scheduling of time-critical tasks is of particular interest. However, the present invention is not limited to hard real-time systems. It might be used in soft real-time systems, too. Such a soft real-time system would still function correctly if deadlines are occasionally missed, and a hardware solution in accordance with the present invention might be to costly under certain circumstances.

SUMMARY OF THE INVENTION

It is an object of the present invention to provide a method and apparatus for off-loading the synchronization of multiple information streams from the processor of the system handling these information streams.

It is an object of the present invention to provide a method and apparatus for fast and reliable synchronization and scheduling of multiple information streams.

It is another object of the present invention to provide a method and apparatus for scheduling tasks or processes in real-time systems.

The above objects have been accomplished by handling pointers pointing to the information streams to be synchronized and by maintaining a table with status information of each information stream. The time marks of said information streams are stored in a table and the system time is checked against said time marks. If a time mark has been reached the status information of the respective information stream is updated. If the status of an information stream is such that its trigger condition is fulfilled, the pointer to said information stream is made available for playout of this information stream.

DESCRIPTION OF THE DRAWINGS AND NOTATIONS USED

The invention is described in detail below with reference to the following drawings:

FIG. 1 shows a system, in accordance with the present invention, for playout of two information streams.

FIG. 2 illustrates the synchronization of multiple data streams.

FIG. 3 illustrates the synchronization and play out of two data streams when the respective trigger condition is fulfilled.

FIG. 4 illustrates how a priority value might be increased with advancing time.

FIG. 5 shows the two core components of an apparatus for synchronization and scheduling of data streams and real time tasks, in accordance with the present invention.

FIG. 6A shows a first exemplary embodiment of the timer.

FIG. 6B shows a second exemplary embodiment of the timer.

FIG. 7 shows an embodiment of the timer and selector, in accordance with the present invention.

FIG. 8 shows an example of an 8-bit timer where every second bit is used.

GENERAL DESCRIPTION

In multimedia systems multiple data streams must be synchronized and scheduled for playout to, for example, a speaker 15 and a video display 14, as illustrated in Figure 1. The synchronization of multiple data streams is illustrated in Figure 2. The three streams are synchronized at so-called synchronization points (SyncPoint) 22. A stream may be synchronized to one or multiple other streams or to time stamps 23 defined by an external clock. This synchronization is typically performed in software, as mentioned in the introductory portion of the present description. In the following, a mechanism is presented which performs synchronization and output scheduling of multiple data streams in hardware without central processing unit (CPU) interaction. The same mechanism can also be used for scheduling tasks in real-time systems.

A data stream is a sequence of data units. In Figure 3, the synchronization of two data units is illustrated. The SyncPoints are translated to time marks relative to the system which processes the received data streams. The SyncPoints might be translated to time marks by means of a clock or counter of said system. In the example given in Figure 3, a data unit from stream 1 cannot be played out before a time mark 30 (S1Start) and it must be played out before a time mark 31 (S1End). The same applies to the second data stream (stream 2). As illustrated in Figure 3, at least two different states (WAIT, READY) can be associated with a data stream. Before SxStart of a data stream x has been reached, the respective data unit is in the WAIT state. After SxStart and before SxEnd it is in the READY state. It is conceivable to assign further states, such as a DELAYED state, to these data streams. The DELAYED state might be associated to a data unit after SxEnd.

In order to be able to synchronize two data units for playout, trigger conditions are defined in accordance with the present invention. A typical trigger condition might be:

IF both data units are in the READY state
THEN they can be played out.

The trigger conditions are checked and as soon as the condition is fulfilled, the data units are played out. This concept is schematically illustrated in Figure 3, where means for checking the trigger condition 32 are shown.

In general, two or more date units can be related in time by trigger conditions, which can be an arbitrary logical combination of the different states of the respective data units. Fulfilling a given trigger condition then initiates the appropriate operation on the units.

The synchronization can be further enhanced with a priority value P to distinguish between data units (or tasks) which fulfill a trigger condition to select an older or more important stream (or task). This, for example, allows to prioritize streams that should not be delayed, whereas other data streams might be delayed. In Figure 4 an example is shown how the priority value P of a data unit (or task) can be increased in time. In this example, three different states (WAIT, READY, DELAYED) are associated with the data stream. The priority value P might also be increased or decreased by interaction of a processor or upon receipt of a feedback signal of a playout device.

A first embodiment of the present invention is schematically illustrated in Figure 1. According to this example, two data streams are received by means 18 for extraction of time marks. The first data stream originates from a storage disk 16, whereas the second data stream is sent via a network 17 to said means 18 for extraction of time marks. After extraction of the time marks, the data units of the two data streams are stored in buffers 19 and 20, and the respective time marks are provided to a timer 10. The buffer addresses are forwarded to a selector 11. In addition, a system clock or counter 13 is connected to said timer 10, and a processor, e.g. a CPU, might

be linked to the timer 10 and selector 11. Two playout devices, a display unit 14 and a loudspeaker 15, are connected via an interfacial playout device to the buffers 19 and 20.

The two core components of the present invention are the timer 10 and the selector 11. Their operation is described in connection with Figure 5.

The timer 10 holds a table 50 with entries concerning the time marks of each data unit received and stored in a buffer. These time marks are relative to a time signal either being generated by a system clock 13 or a counter. Each time mark reached is then output to the selector 11. The table 50 is written by an external processor. The selector comprises a table 51 where for each data unit in a buffer a pointer to said data unit, i.e. the respective buffer address, the state of said data unit, and the trigger condition are held. Data units to be synchronized are linked into a circular list via their table entries. For each time mark received from the timer 10, the sector 11 updates the status field of the concerned data unit and then checks the trigger condition of all data units to be synchronized with said data unit. Once the trigger condition is fulfilled, the pointer to the data unit (or units) is made available at the output of selector 11. The output of the pointer then triggers retrieval and playout of the data unit by means of the playout device 21. Table 51 of the selector 11 might be maintained and updated by an external processor, e.g. CPU 12.

The timer 10 reads the current system time and returns the address of a unit pointer in the selector 11. This function can be provided in several ways. First, a solution based on a content addressable memory (CAM) is presented which is suitable for a high resolution timer but relatively expensive. The second solution is less expensive and is suitable for lower resolution timers and for implementing timer priority as described below. The two solutions are shown in Figures 6A and 6B, respectively.

The CAM timer 60 is illustrated in Figure 6A. It comprises a CAM 61. The start and end time mark for each data unit to be scheduled are written into consecutive addresses of the CAM 61. The addresses then differ only in the least significant bit. Thus, for a CAM of 2^n addresses, n-1 bits determine the address of the scheduling unit pointer in the selector table 51. The width of a CAM entry is determined by the resolution of the time signal received at input 62. For a 32-bit time signal also a 32-bit wide CAM 61 is required.

For lower time resolution a timer 63 with counter 64, as shown in Figure 6B, might be employed. In this case the difference between the current time and the time mark is calculated and its two's complement is written to the timer table 65. The counter 64 increments each entry for each time tick at input 66. If the carry bit of an entry becomes set, its address is shifted right by one bit and output to the selector 11. As the counter 64 continues to increment the entries, the content of the entry (with its carry bit = 1) can then be used to reflect the time passed since reaching the start and/or end time. The counter value thus describes a time based priority. With this counter value a priority value P might be generated. The width of a timer entry can range from 8 to 32 bits depending on the timer resolution and the required timer differences. A free flag 67 might be provided in order to mark free table entries in table 65. This flag 67 is used by the counter to determine which entries it must increment.

For both schemes, a FIFO can hold free timer table entries to reduce the management load of the external processor. The processor gets a free entry from the FIFO and writes delayed entries, or entries which are freed, to the FIFO. The selector might directly write delayed entries to this FIFO.

It is obvious that the present invention is not limited to the above two implementations of timers since other timers are conceivable which are well suited, too.

The selector 11 comprises a table 51 and a comparator state machine (CSM) 70. The table 51 has $2^{(n-1)}$ entries (rows or columns). The exemplary implementation of the selector 11, given in Figure 7, shows a table 51 with $2^{(n-1)}$ rows. Each row has a pointer field 71 holding the pointer, i.e. the buffer address S1Ptr, to the data unit to be scheduled, a p-bit wide priority field 72, a k-bit wide status field 73, a 2-bit wide trigger field 74, and an (n-1)-bit wide synchronization field 75. The width of said status field 73 depends on the number of different states. In case of three different state (WAIT, READY, DELAYED) the status field 73 is 2-bits wide. The synchronization field 75 holds a pointer linking all those entries with which a data unit must be synchronized into a circular list. This circular list is schematically illustrated by means of arrows on the right hand side of table 51. The CSM 70 receives the address of a unit pointer from the timer 10 and updates the state of the status field 73 in this row. The CSM 70 then follows the linked list of synchronizations entries and computes whether the trigger condition is fulfilled by comparing the status fields 73 with the trigger fields 74. The trigger condition might also be checked by means of simple logical elements such as AND-gates, OR-gates etc.

In the following, it is assumed, without loss of generality, that the trigger condition is fulfilled when all entries in the respective trigger fields 74 are in the READY state. In this case, the trigger state fields are redundant and might be omitted.

Once all the entries linked by the synchronization field 75 are READY, the CSM 70 outputs the entries according to their priorities, if any. With $p=4$, 16 priority classes can be distinguished. The width w of the scheduling unit pointer is determined by the memory (buffer) architecture of the system. The width s of the synchronization field 75 is determined by the size of the selector table 51, with $s=n-1$.

In Figure 7 the problem shown in 3 is elaborated. The time marks 30 (S1Start) and 31 (S1End) of data unit S1 are written to the addresses 0 and 1

of the CAM 61. The address of the pointer to the data unit S1 is the CAM address shifted right by one bit. An external processor 12 writes the time marks to the timer table, i.e. the CAM 61 and the entries to the selector table 51. For a data unit to be scheduled which requires priority, the priority field 72 is set. In this example a 4-bit wide priority field 72 is used. For each time mark reached, the state, reflected by the status field 73, is updated. If the state is READY, the CSM 70 reads the synchronization field 75 and follows the linked list to determine whether the state of all other entries fulfills the trigger condition. If the trigger condition is reached, the unit pointer with the highest priority is made available at output 76 first. If multiple units have the same priority, or in case that no priority fields are provided, the unit pointers might either be output in parallel or in a consecutive manner, e.g. by order of their position in the selector table 51. It is obvious that the priority fields 72 can be omitted, too.

In many systems, multimedia as well as real-time, the scheduling of actions, data streams, tasks, processes, or the playout of data units is determined by a priority which increases with time. If a counter timer, as illustrated in Figure 6B, is used as the timer component 10, the entry in the timer table 65 also describes its priority. This priority can be used in the priority field 72 of the selector 11 to achieve time dependent scheduling. The four most significant bits of the timer entry are used to distinguish again 16 priorities. A different set of bits might be used if the priority must increase non-linearly. Figure 8 shows an example of an 8-bit timer where every second bit is used to get a priority mechanism that maps the 256 timer values to 16 non-linear priority values P.

According to the present invention, a synchronization point (SyncPoint) is set up by writing the time marks to the free entries in the timer table and setting up the associated entries in the selector table. For each stream this takes one write operation to the timer and to the selector table, plus the operations to set up the selector table entry. The actual scheduling takes place in hardware without any further intervention of an external processor.

For a solution in software a similar number of setup operations would be necessary but the search in a list of timers would take much longer. Typically timers involve interrupt processing which is very time consuming in particular in modern RISC processors with a large register file and instruction pipeline. In the present invention these operations are avoided and implemented in hardware. Additionally, the present invention can handle timers with much finer granularity than any software solution.

The present invention, which has been described by means of a multimedia related example, can be used in various multimedia communication subsystems, e.g. for multimedia multiparty teleconferencing, or multimedia protocol adapters. In ST-II and RTP protocol systems, for example, a playout time stamp can be extracted from the protocol header and used in the present scheduler (ST-II is the second version of an experimental Internet stream protocol, and RTP is a transport protocol for real-time applications). The present invention might also be used in multimedia bus systems which provide a time signal. Data units to be triggered at a specified time are written to the scheduler and the multimedia bus provides the timer input to the timer unit. The invention can also be used in a multimedia adapter such as the IBM ActionMedia II adapter (see IBM ActionMedia II Technical Reference, IBM 04G5144, 2nd Edition, June 1992, Mechanicsburg), or the MWave System (see "MWave Multimedia System", Technical Brief, Texas Instruments, 1992) to synchronize multiple data streams. The term ActionMedia is a trademark of the Intel Corporation and MWave is a trademark of the IBM Corporation.

As already mentioned in the introduction, the present invention might also be used in real-time systems for the scheduling of time-critical tasks or processes. For these kind of application there is no synchronization field in the table of the selector required. The CSM selects from the selector table the entry with the highest priority. The scheduling unit pointer, i.e. a pointer pointing to a task control block or to executable code, is then output. The timer priority mechanism described above can be used to implement

scheduling policies with an explicit deadline, e.g. the earliest-deadline-first or least-slack time scheduling algorithm as disclosed in the book "Real-Time Systems and their Programming Languages", by A. Burns and A. Wellings, Addison-Wesley Publishing Company, International Computer Science Series, 1990. An external processor calculates the difference to the time marks and writes it to the timer table. The entries are then used to derive the priority of a job.

CLAIMS

1. Apparatus for synchronization of at least two information streams which comprises:
   a) a timer (10) for holding a first table (50) with time marks for each one of said information streams, and for making each time mark reached available at the timer output,
   b) a selector (11) for holding a second table (51) with
      - pointer fields (71) carrying pointers to each of said information streams,
      - status fields (73) carrying information concerning the status of each of said information streams, the status field of an information stream being updated each time a time mark of said information stream is reached and made available at said timer output,
   c) means for checking whether a trigger condition of an information stream is fulfilled by watching the corresponding status field, and
   d) means for making the pointer to an information stream whose trigger condition is fulfilled available for playout of said information stream.

2. The apparatus of claim 1, wherein said second table (50) further comprises:
   - trigger fields (74) carrying information concerning the trigger conditions of each of said information streams.

3. The apparatus of claims 1 or 2, wherein said second table (51) further comprises:
   - synchronization fields (74) carrying, by means of a linked list, entries concerning information streams to be played out in a certain relationship to each other.

4. The apparatus of claims 1, 2, or 3, wherein said second table (51) further comprises:
   - priority fields (72) carrying priority values (P) such that the information stream with highest priority is played out first if the respective trigger condition has been fulfilled.

5. The apparatus of claim 4, further comprising means for increasing said priority values (P) with advancing time.

6. The apparatus of any of the preceding claims, further comprising:
   - means (18) for extracting said time marks from incoming information streams,
   - buffer means (19, 20) for storing said information streams,
   - playout means (21) for retrieval of said information streams stored in said buffer means (19, 20) and for playout to the respective playout devices (14, 15).

7. A multimedia system comprising an apparatus in accordance with any of the preceding claims.

8. Apparatus for scheduling at least two real time tasks or two processes which comprises:
   a) a timer (10) for holding a first table (50) with time marks for each one of said tasks or processes, and for making each time mark reached available at the timer output,
   b) a selector (11) for holding a second table (51) with
      - pointer fields (71) carrying pointers to task control blocks or executable code,
      - means for making a pointer to a task control block or an executable code available at an output of said selector each time a time mark of said task or process is reached in said timer.

9. The apparatus of claim 8, wherein the pointer to a task or process with highest priority is made available first.

10. Method for synchronization of at least two information streams, characterized by the following steps:
    a) checking whether a time mark of one of said information streams has been reached,
    b) updating the status information of said information stream whose time mark has been reached,
    c) determining whether the trigger condition of said information stream is fulfilled by watching said status information, and
    d) providing a pointer to said information stream for further processing if the trigger condition of said information stream is fulfilled.

11. The method of claim 10, wherein the playout of information streams, which have to be played out in a predefined relationship to each other, is scheduled by means of a linked list being organized such that each time when a status information of one of said information streams is updated, the status information of the other information streams are checked, and said information streams played out when all trigger conditions are fulfilled.

12. The method of claim 10, wherein a priority value (P) is assigned to each of said information streams, the information streams with highest priority being played out first.

13. The method of claim 12, wherein the priority value (P) of an information stream is increased with advancing time.

14. The method of claim 10, further comprising the following steps:
    - extracting and/or synchronizing the time marks of said information streams,
    - storing said information streams in buffers,

- making the addresses of said buffers available as pointers for later retrieval of said information streams,

15. The method of claim 14, wherein said information stream is retrieved from said buffer by means of said pointer, prior to being send to a playout device.

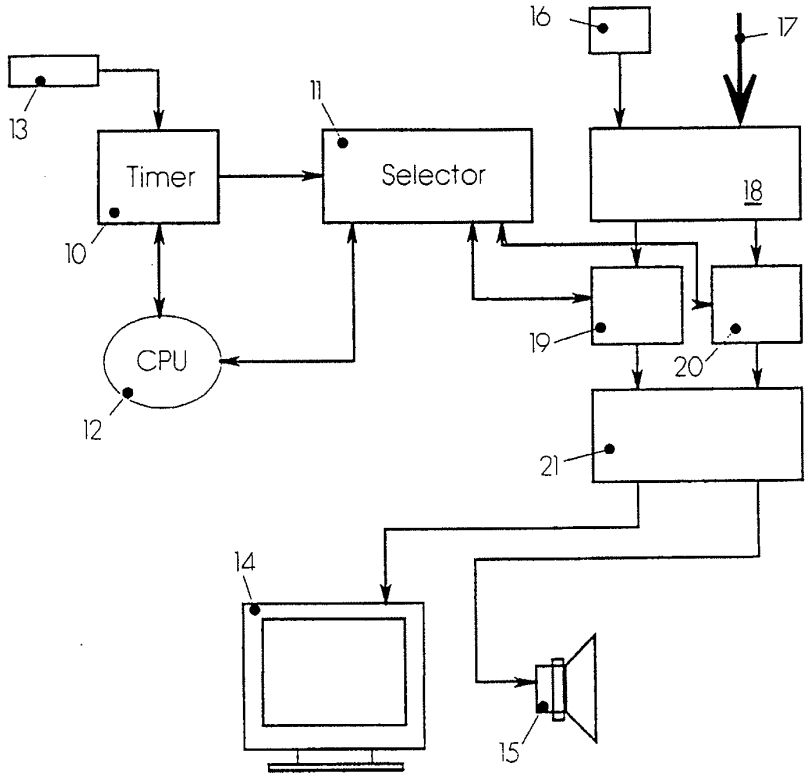

16
17
13
11
Timer
Selector
18
10
CPU
19
20
12
21
14
15

FIG. 1

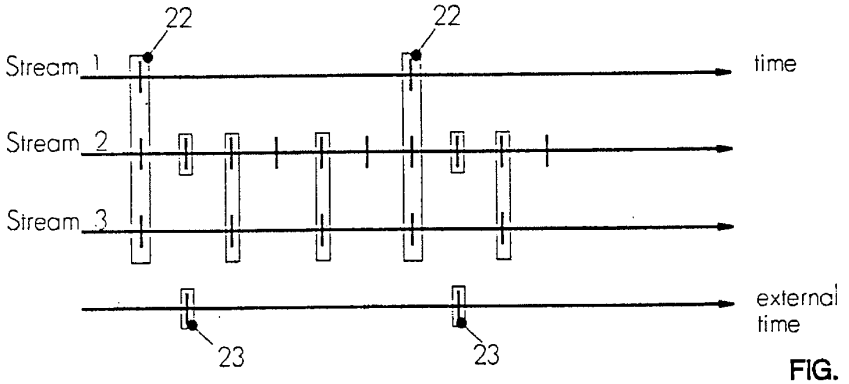

22
22
Stream 1
time
Stream 2
Stream 3
external time
23
23

FIG. 2

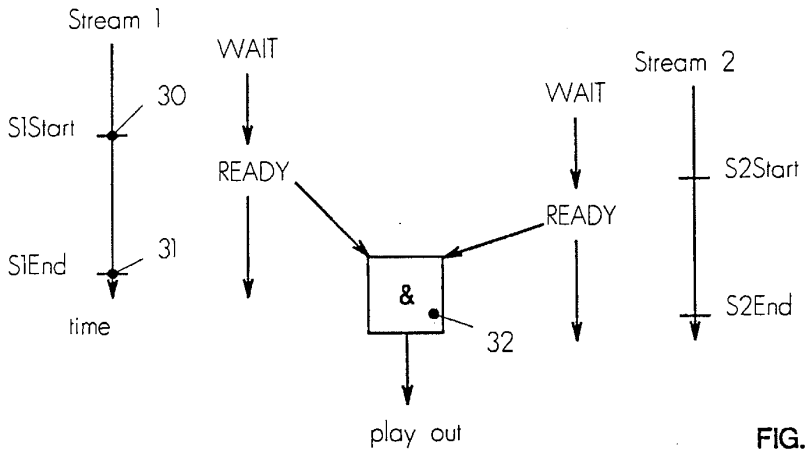


FIG. 3

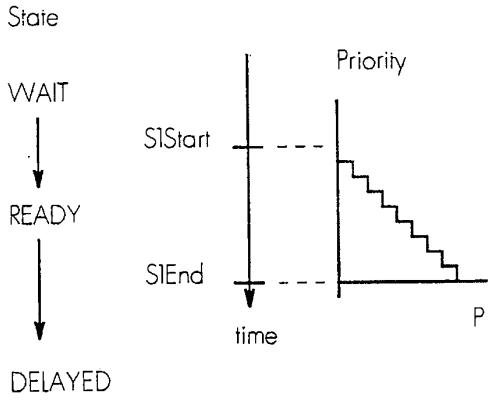


FIG. 4

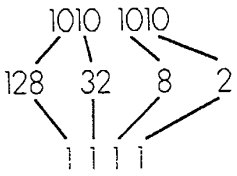


FIG. 8

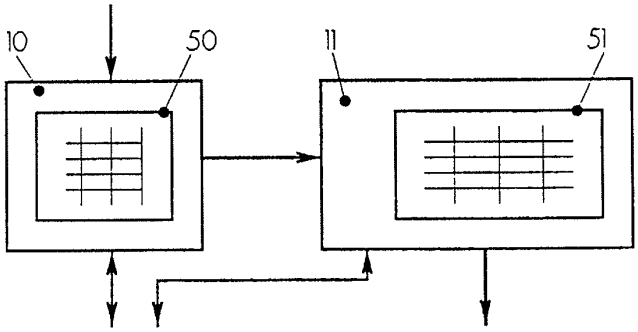


FIG. 5

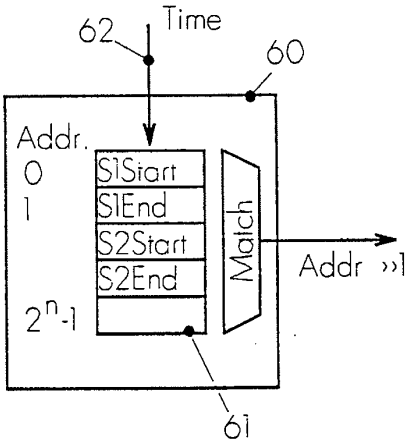


FIG. 6A

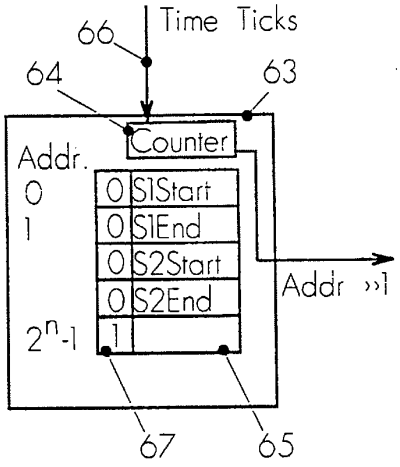


FIG. 6B

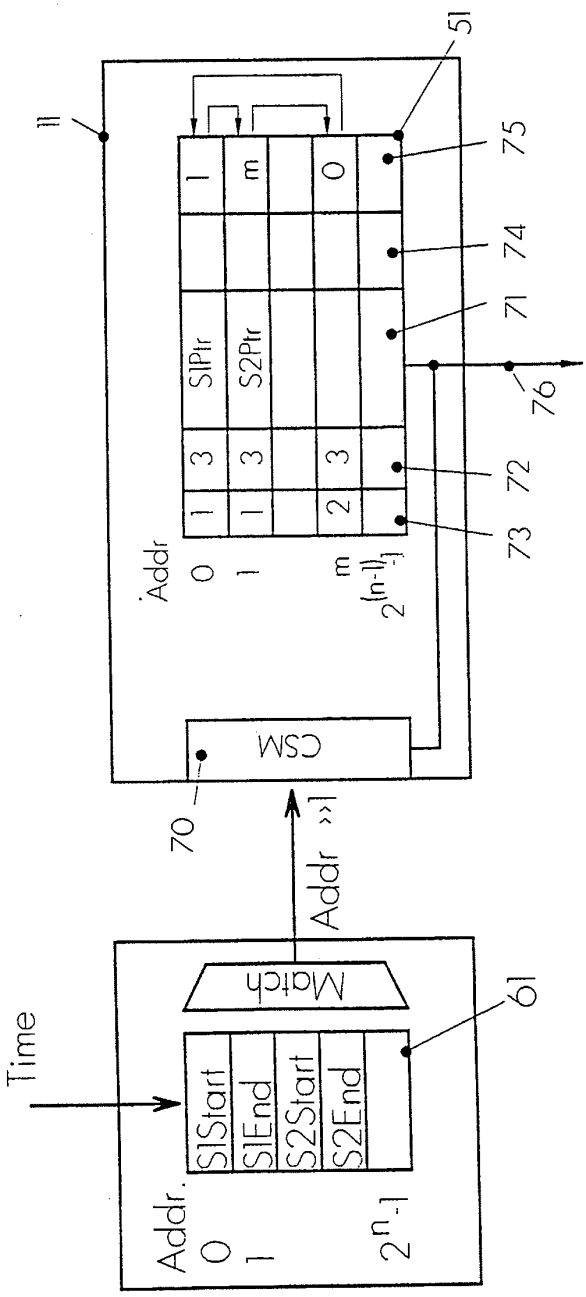

11
51
75
74
71
76
72
73
S1Ptr
S2Ptr
Addr
0
1
m
$2^{(n-1)}$
CSM
70
Addr »1
Match
61
Time
S1Start
S1End
S2Start
S2End
Addr.
0
1
2^{n-1}

FIG. 7

INTERNATIONAL SEARCH REPORT

Intern ıal Application No
PCT/EP 94/00833

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP,A,0 239 884 (I.B.M. CORPORATION) 7 October 1987 see the whole document ---	1,8,10
A	EP,A,0 550 197 (I.B.M. CORPORATION) 7 July 1993 see the whole document ---	1,8,10
A	COMPUTER COMMUNICATIONS, vol.16, no.3, March 1993, GUILDFORD GB pages 168 - 176 P.V. RANGAN ET AL. 'Techniques for multimedia synchronization in network file systems' see the whole document --- -/--	1,8,10

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
10 November 1994	01.12.94

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax (+31-70) 340-3016	Katerbau, R

Form PCT/ISA/210 (second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Inter. nal Application No
PCT/EP 94/00833

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	COMPUTER, vol.24, no.10, October 1991, LONG BEACH US pages 51 - 57 D.P. ANDERSON ET AL. 'A continuous media I/O server and its synchronization mechanism' see the whole document ---	1,8,10
A	BYTE, vol.18, no.2, February 1993, LONDON, GB pages 225 - 230 A.W. DAVIS ET AL. 'The Mac goes to the movies' see the whole document -----	1,8,10

1

Form PCT/ISA/210 (continuation of second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Information on patent family members

Intern. ial Application No
PCT/EP 94/00833

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
EP-A-0239884	07-10-87	US-A-	4893256	09-01-90
		DE-D-	3787553	04-11-93
		DE-T-	3787553	11-05-94
		JP-B-	6054469	20-07-94
		JP-A-	62239233	20-10-87
EP-A-0550197	07-07-93	US-A-	5333299	26-07-94
		CA-A-	2084574	01-07-93
		JP-A-	5265778	15-10-93

Form PCT/ISA/210 (patent family annex) (July 1992)

PCT WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : G06F 9/46

A1

(11) International Publication Number: **WO 99/61985**

(43) International Publication Date: 2 December 1999 (02.12.99)

(21) International Application Number: PCT/US99/06634

(22) International Filing Date: 26 March 1999 (26.03.99)

(30) Priority Data:
60/086,737 26 May 1998 (26.05.98) US

(71) Applicant: THE DOW CHEMICAL COMPANY [US/US]; 2030 Dow Center, Midland, MI 48674 (US).

(72) Inventors: WOODS, Randy, D.; 3788 Old Pine, Midland, MI 48642 (US). JACHIM, David, M.; Apartment L, 3063 Signature Boulevard, Ann Arbor, MI 48103 (US). DUPREE, Wayne, P.; 5301 Stony Creek, Midland, MI 48640 (US). VERNIERS, Gerrit, H.; Heirweg 88, B–9180 Moerbieke–Waas (BE). CHURCHILL, Stephen, G.; 2588 Waskevich, Midland, MI 48642 (US). FERNANDEZ, George, P.; 449 North 5 Mile, Midland, MI 48640 (US).

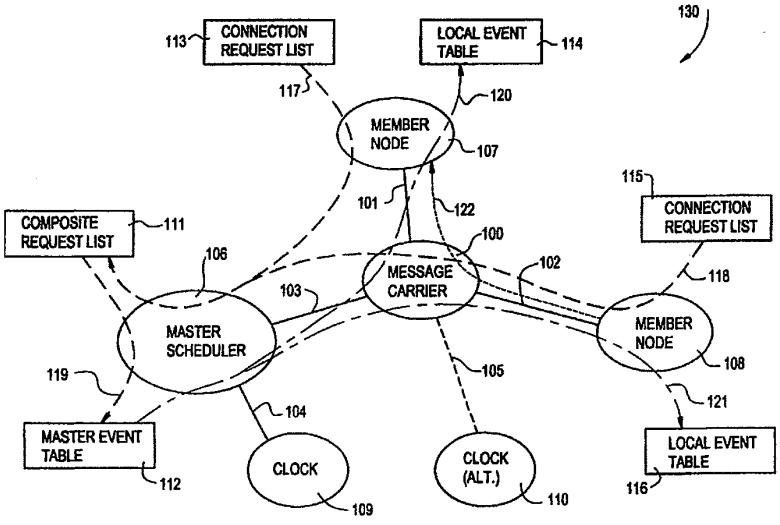
(74) Agent: SCHULTZ, Dale, H.; Patent Dept., P.O. Box 1967, Midland, MI 48641–1967 (US).

(81) Designated States: CA, CN, JP, MX, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published
With international search report.

(54) Title: DISTRIBUTED COMPUTING ENVIRONMENT USING REAL–TIME SCHEDULING LOGIC AND TIME DETERMINISTIC ARCHITECTURE

(57) Abstract



A system is presented for a plurality of computer nodes in mutual data communication to execute logic in a coordinated manner. The approach features aligning the definition of the time in the nodes so that each node has essentially the same definition of time,defining a schedule of communications for each computer node, with the schedule having future moments defined for initiating each communication, communicating each schedule to its respective node, and initiating each communication at the time defined in the schedule. The computer network has a message carrier supporting multiple communication paths; a plurality of computers coupled to the message carrier, a clock system for maintaining time synchronism among the computers; and a scheduler coupled to the message carrier for determining the sequence of communication between the computers. The message carrier establishes communication between selected computers based on the time–deterministic tasks being performed by the computers. A computer node in the plurality of computer nodes executes logic in coordination with the other nodes by defining a communication request list of each communication needed by the computer node with any other computer node respective to a future time period, sending the communication request list to the master scheduler, receiving back a communication schedule back from the master scheduler, defining a local task schedule from the communication schedule, and then executing the communication at the scheduled time.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	**ES**	Spain	**LS**	Lesotho	**SI**	Slovenia
AM	Armenia	**FI**	Finland	**LT**	Lithuania	**SK**	Slovakia
AT	Austria	**FR**	France	**LU**	Luxembourg	**SN**	Senegal
AU	Australia	**GA**	Gabon	**LV**	Latvia	**SZ**	Swaziland
AZ	Azerbaijan	**GB**	United Kingdom	**MC**	Monaco	**TD**	Chad
BA	Bosnia and Herzegovina	**GE**	Georgia	**MD**	Republic of Moldova	**TG**	Togo
BB	Barbados	**GH**	Ghana	**MG**	Madagascar	**TJ**	Tajikistan
BE	Belgium	**GN**	Guinea	**MK**	The former Yugoslav Republic of Macedonia	**TM**	Turkmenistan
BF	Burkina Faso	**GR**	Greece			**TR**	Turkey
BG	Bulgaria	**HU**	Hungary	**ML**	Mali	**TT**	Trinidad and Tobago
BJ	Benin	**IE**	Ireland	**MN**	Mongolia	**UA**	Ukraine
BR	Brazil	**IL**	Israel	**MR**	Mauritania	**UG**	Uganda
BY	Belarus	**IS**	Iceland	**MW**	Malawi	**US**	United States of America
CA	Canada	**IT**	Italy	**MX**	Mexico	**UZ**	Uzbekistan
CF	Central African Republic	**JP**	Japan	**NE**	Niger	**VN**	Viet Nam
CG	Congo	**KE**	Kenya	**NL**	Netherlands	**YU**	Yugoslavia
CH	Switzerland	**KG**	Kyrgyzstan	**NO**	Norway	**ZW**	Zimbabwe
CI	Côte d'Ivoire	**KP**	Democratic People's Republic of Korea	**NZ**	New Zealand		
CM	Cameroon			**PL**	Poland		
CN	China	**KR**	Republic of Korea	**PT**	Portugal		
CU	Cuba	**KZ**	Kazakstan	**RO**	Romania		
CZ	Czech Republic	**LC**	Saint Lucia	**RU**	Russian Federation		
DE	Germany	**LI**	Liechtenstein	**SD**	Sudan		
DK	Denmark	**LK**	Sri Lanka	**SE**	Sweden		
EE	Estonia	**LR**	Liberia	**SG**	Singapore		

DISTRIBUTED COMPUTING ENVIRONMENT USING REAL-TIME SCHEDULING LOGIC AND TIME DETERMINISTIC ARCHITECTURE

This invention is directed to the field of data processing control systems, particularly to integrated systems where a plurality of computers execute their logic in a coordinated manner within a network. The primary application of the system described in this specification is for computerized control of machines deployed on a large physical scale (for example a chemical manufacturing facility), although benefits to endeavors related to plural computer processors, parallel computer processors, and networked computer processors are also enabled.

When computational demands or geographical constraints require the physical distribution of computing facilities (for example such as for process control in a large, geographically dispersed chemical plant), implementation factors (such as product quality) may benefit from centralized coordination of the computing tasks. This invention describes a system and method for enabling a plurality of geographically dispersed computers to be networked for enabling time coordinated inter-activity. This resultant network enables task coordination and facilitates the implementation of other qualities beneficial to the task being performed by the system - qualities such as redundancy, security, and high throughput.

One embodiment of a system showing such substantive real-time networking to a set of physically distributed remote field units for use in a process control system is described the United States Patent 5,428,769 issued on Jun. 27, 1995 and entitled "Process control interface system having triply redundant remote field units" to Glaser, Robert S.; Hoy, Robert S.; Fernandez, G. Paul; and Grai, Timothy J. While the system is adequate for many process control situations, the networking methodology effectively addresses each system in each time period on the presumption that there is a need to do so; however, many of these addressing cycles are unnecessary if all variables for exchange are stable -- a more elegant scenario from the standpoint of network efficiency would provide for communications only when the status of either the inputs or the outputs require change.

Physical distribution of computing facilities combined with centralized coordination of the computing tasks is also a driving concept in high performance processors characterized by parallelism. Vector processors, as discussed in sections 7.1 and 7.2 ("Vector Processors") of Computer Architecture A Quantitative Approach by John L. Hennessy and David A.

Patterson (Morgan Kaufmann; San Mateo, California; 1990; pages 351-363) are one manifestation of such systems, and processor-memory pair approaches such as the Cm* design discussed from pages 11 to 18 of "Parallel Processing - The Cm* Experience" by E. F. Gehringer, D. P. Siewiorek, and Z. Segall (Digital Press via Editorial Inc.; Rockport, Massachusetts; 1987) provide a second example of these types of approaches.

Centralized coordination of computing tasks in a scheduled context for the future also draws on concepts from traditional project scheduling such as is discussed from pages 569 to 584 of the text Productions and Operations Management by J. William Gavett (Harcort, Brace, and World; New York; 1968).

Process control systems development is characterized by an ever-increasing incorporation and management of issues at large in the system being controlled, requiring that a comparably larger amount of logic on a comparably larger amount of data be executed in a time frame which is not permitted to expand. This can be achieved to some degree by hardware which processes logic more quickly and holds more data in an addressed physical store, but there are substantial costs in pursuing this solution. What is needed to resolve the ability of future systems to effectively expand their comprehensive treatment of the control situation is to truly provide a solution path for enabling a very large number of distributed computers to execute massive amounts of logic in a mutual manner which is time-synchronous across an entire network. The present invention provides such a method and system.

The invention provides a computer implemented method for enabling a plurality of computer nodes in mutual data communication to execute logic in a coordinated manner by

(a) aligning the definition of time between the computer nodes so that each computer node has essentially the same "understanding" of time;

(b) defining a schedule of communications for each computer node, where the schedule has a different time in the future defined for initiating each respective communication;

(c) communicating each schedule to its respective computer node; and

(d) initiating each communication from each node at the time defined in its schedule.

The described embodiment provides a computer network architecture, having a message carrier with a plurality of computers coupled to it; a clock system (accessible by all computers in the network) for maintaining time synchronism among the computers; and a scheduler for determining the sequence of communication between the computers by causing the message carrier to establish communication between selected computers based on time-deterministic tasks being performed by the computers.

The described embodiment further provides a global time system for aligning the sense of time in the set of computers to a globally defined reference.

The described embodiment further provides a computer implemented method for enabling a plurality of member nodes to execute logic in a coordinated manner.

The described embodiment further provides for concurrency in communications between computers via the message carrier.

The described embodiment further provides provides a computer implemented method for scheduling tasks for resource objects in a plurality of resource objects functioning as a collective group.

The described embodiment further provides a method by which different computers can align their definition of time by adjusting their oscillator's rate of oscillation.

The computer network architecture is enabled with a message carrier supporting multiple communication paths; a plurality of computers coupled to the message carrier, with each computer being programmed to perform time-deterministic tasks by cycling through a plurality of different operating states, including a communication state; a clock system accessible by all the computers for maintaining time synchronism among the computers; and a scheduler coupled to the message carrier for determining the sequence of communication between the computers by enabling the message carrier to establish communication

between selected computers based on the time-deterministic tasks being performed by the computers.

A computer node (element for computing with an independent central processing unit) in the plurality of computer nodes can execute logic in coordination with the other nodes by being connected to a message carrier along with the other computer nodes where there is also a provision for (a) master scheduling capability in one computer node to establish a master scheduler; (b) defining, in each computer node, a communication request list of each communication needed by the computer node with any other computer node respective to a future time period; (c) sending of the communication request list from each computer node to the master scheduler via the message carrier; (d) receiving of, in each computer node, a communication schedule from the master scheduler via the message carrier; and (e) defining, in each computer node, a local task schedule from the communication schedule.

The above approach enables a plurality of member nodes in a system to execute a large amount of logic in a coordinated manner on a group of distributed platforms.

In performing the above, the step of defining a reference moment time value is performed in one embodiment by the master scheduler.

In performing the above, the step of defining a communication request list in each member node in one embodiment limits the communications to a maximum number in the future time period and each communication has an duration time attribute and the value of the duration time attribute is less than a fixed maximum value in order to achieve efficient operation.

Use of a global time system in one embodiment enables a universal understanding of time across a large number of different computer nodes (member nodes) in an extended system.

The method for scheduling involves a process of scheduling tasks for resource (member) objects in a plurality of resource objects functioning as a collective group; and this computer implemented process (called an "apparent critical resource" process) involves the steps of:

defining a set of future tasks for the collective group;

defining each resource object needed for each future task in the set of future tasks;

defining each interaction needed between any two resource objects respective to the set of future tasks;

defining a use instance for each resource object in achieving each interaction where the use instance has an use duration attribute value;

defining an accessing resource object and an accessed resource object respective to each use instance;

determining a summed accessing time value for each accessing resource object as a sum of all use duration attribute values for all use instances respective to that resource object when it is the accessing resource object;

determining a summed accessed time value for each accessed resource object as a sum of all use duration attribute values for all interactions between the accessed resource object and the accessing resource object;

defining a set of sequential future time increments within a scheduling period in a schedule;

scheduling, into the earliest available time in the first future time increment, a future resource interaction event based on the first available future task in the set of future tasks between

(a) the accessing resource object having the highest summed accessing value respective to the set of future tasks and

(b) the accessed resource object having the highest summed accessed value respective to the set of future tasks (if a use instance exists between the accessed resource object and the accessing resource object -- the future resource interaction event in the schedule record has an initiation time equal to the earliest available time in the

schedule, an accessing resource object identifier attribute respective to the accessing resource object, an accessed resource object identifier attribute respective to the accessed resource object, and the aforementioned use duration attribute value)

removing from scheduling candidacy the use instance respective to the accessed resource object in the scheduled future resource interaction event and the accessing member object identifier attribute in the scheduled future resource interaction event to redefine the set of future tasks to those future tasks for which interactions have not been scheduled;

repeating the steps of determining a summed accessed value, determining a summed accessing value, scheduling a future communication task, and removing all use instances respective to scheduled resource objects until all use instances have been scheduled which can be scheduled into the first future time increments; and

recurrently executing the steps of determining a summed accessed value, determining a summed accessing value, scheduling a future resource interaction event, removing all resource objects respective to scheduled resource objects, and repeating such steps until all resource interaction events have been scheduled which can be scheduled into each future time increment and until all resource interaction events have been scheduled into the scheduling period.

In one embodiment, the message carrier has a plurality of channels for permitting concurrency in the communication events, the step of scheduling being respective to the plurality of channels as a group wherein the earliest available time is respective to any channel in the message carrier.

In some embodiments, each member node has a variable speed oscillator for adjusting a local clock to align with a global clock on the network. Time alignment logic and oscillator rate adjustment logic for execution by the computer circuitry are also provided to achieve rapid time alignment of the particular member node (computer node).

The message carrier used in the described embodiment is enabled with a network topology which can be manifested (in the preferred embodiment) with a crossbar switch, or, in an alternative embodiment, with a radio frequency communications system.

Additional features and advantages of the present invention will become more fully apparent from a reading of the detailed description of the preferred embodiment and the accompanying drawings.

Figure 1 presents a diagram of the data flows between the various elements in a distributed network.

Figure 2 depicts a distributed network of computing elements which are resident in physically different locations throughout the world. This figure includes a network of computing elements, a master scheduler, and a common source of time.

Figure 3 presents a diagram of a network in a bus structure.

Figure 4 shows a distributed computing environment where the network is implemented in a ring topology.

Figure 5 presents a diagram of a network in a switched star topology which is a preferred network architecture embodiment.

Figure 6 presents a diagram of data structures communicated between the computing elements and a master scheduler.

Figure 7 depicts a flowchart of the process implemented in the master scheduler in the scheduling of communications between two separate computing elements.

Figure 8 presents a simplified diagram of the scheduling logic (software in execution by electrical circuitry) implemented in the master scheduler in the scheduling of communications between two separate computing elements.

Figure 9 presents a detailed flowchart of the scheduling logic which is followed at the master scheduler.

Figures 10 presents a bar chart depicting the first scheduled communications between two separate computing elements after the first connection is scheduled.

Figure 11 presents a bar chart depicting the scheduled communications between computing elements after all connections are scheduled which can be scheduled in the first step.

Figure 12 presents a bar chart depicting the scheduled communications between computing elements after all connections have been scheduled.

Figure 13 provides a detailed composition of the Message Carrier, the Master Scheduler, the Clock, and the Member Nodes.

Figure 14 shows the relationships between the technologies which have been combined in the preferred embodiment, with emphasis on those most critical to the application.

Figure 15 presents a flowchart of the logic which executes in each of the computing elements in order to facilitate communication between the member nodes.

Figure 16 depicts a diagram of the data flows between the various elements in a distributed network where communications are facilitated by radio frequency.

As an introduction to this part of the specification, the system and method of the present invention implement centralized scheduling of cooperatively performed tasks, time deterministic event operations, and effective real-time synchronization and alignment of oscillators respective to the plurality of computers (also referenced as "member nodes") on the network. In this regard, the time deterministic network involves communicating between separate computing elements based on time rather than on some event (examples of events are the reception of a token or the detected availability of a particular bus). Centralized scheduling of cooperatively performed tasks, time deterministic event operations, and effective real-time synchronization and alignment of oscillators respective to the plurality of computers on the network requires that the computers share an essentially common view of the datalogical representation of an actual moment in real-time. This means, for instance

and assuming for this example (without limitation to the scope of the present invention) a system wide-precision requirement of one millionth of a second , that two different computers in the system essentially perceive that, when the moment of time in real-time defined by an undisputed source respective to the network is 08:00:0000000 hours, that the first of the two computers has a datalogical representation of the moment in real-time between 07:59:9999995 and 08:00:0000005 and that the second of the two computers also has a datalogical representation of the moment in real-time time between 07:59:9999995 and 08:00:0000005; as can be appreciated, the difference in the datalogical representation of the moment in real-time between the two computers is no more than 1 millionth of a second.

When the computers in a distributed computing network share an essentially common view of the datalogical representation of an actual moment in real-time, they can then be programmed to perform specific tasks based on time. Tasks which involve multiple computers can then be scheduled centrally and a basis is established for time-dependent activities to be invoked on geographically dispersed computers in a deterministic manner. In this regard, determinism in real-time operations characterizes the occurrence of a predefined event at a predefined moment of real-time.

The use of such time coordination in communications between computers also facilitates the employment of active redundancy. In active redundancy, there are two or more computers which are each executing essentially identical logic, exchanging the results from these executions, optionally arbitrating the independently derived results to come to a commonly understood mutual conclusion for further use, and acting to concurrently effect the result. This process proceeds most effectively when data received by each computer is shared with the other computer(s) executing the same logic prior to arbitrating the results, so there must be some mechanism by which the time to begin the arbitration process is mutually determined. In the present invention, the mechanism for providing time synchronization between the redundant computers and to schedule the tasks related to a communication instance is (1) to provide an essentially identical datalogical representation of the moment in real-time time among the computers; (2) to use a scheduling approach to define in each computer the moment in time when the communication instance in that computer will occur; and (3) to then to implement the appropriate sending, receiving, and linkage tasks in each system involved in the communication instance at the time defined for the communication instance. In this way, the time at which the referenced arbitration task will be performed is

pre-determined after the communications tasks have been scheduled (as opposed to the approach of waiting to determine the start time of the arbitration task after the communication tasks have actually executed).

The scheduling approach is implemented by execution of scheduling logic. The concept of logic relates to software (a computer program) whose purpose is achieved in execution by an enabling system of electrical circuitry (usually denoted by the term "computer"). The scheduling logic used for the scheduling of communications is based on analysis of demands for the resources (for example specific computers, data space within each computer, execution time in each computer, and specific circuits within each computer) which are being scheduled. This scheduling logic is used in the scheduling of network communications, but it has further applicability in the scheduling of any resources which require interaction among one another, and it is especially suited for real-time applications. The scheduling logic is further designed to be efficient respective to the resources it is responsible for scheduling and to be predictable as to the amount of time required to generate its output, a schedule, from the inputs (that is a collection of connection requests between any two computers on this network). In achieving this task, the scheduling logic schedules the available resources with the highest demand for the time that the schedule is being created. Reflecting the nature of (1) evaluating the resources with the highest demand at each point in the schedule and (2) giving precedence to the resources which have the highest demand, the particular scheduling logic method described in this specification is also referred to herein as Apparent Critical Resource, or ACR, scheduling logic. In this regard, in the preferred embodiment, the scheduling logic is not iterative in the sense that, once a particular task is scheduled, it is NOT removed from the schedule in order to pursue an alternative to its placement in the schedule. It is this quality which enables the preferred ACR scheduling logic type of scheduling logic to be well-suited for real-time applications.

In the application of scheduling logic to the scheduling of network communications, there are transmission resources and reception resources which are to be considered. Each time a network communication is scheduled, a transmission resource is associated with a reception resource.

In the scheduling of network communications, there are responsibilities which are handled by the master scheduler and those which are handled by the member nodes.

Figure 1 presents a data flow diagram 130 of the general elements and associated data flows between the various general elements in a distributed network and shows a message carrier 100, a first member node 107 and a second member node 108, a master scheduler 106 for coordinating communications between member nodes 107, 108 via data communication path 101 and data communication path 102, clock alternatives (as clock 109 and alternative clock 110) for achieving synchronization between member nodes 107, 108 (via data communication path 104, data communication path 105, data communication path 101, and data communication path 102), and the primary data tables and exchanges required to interact effectively in this distributed computing environment (via data communication path 118, data communication path 117, data communication path 119, data communication path 120, and data communication path 121).

The clock 109 is distributed to the member nodes through the message carrier from the master scheduler (alternatively, the message carrier 100 could distribute an alternative local clock 110 to the member nodes 107, 108 and to the master scheduler 106). This permits all nodes attached to the message carrier to share a common view of time. In this regard, clock 109 represents a global clock for use around the physical globe (for example in North America, in Europe, and in Asia) where the alternative clock 110 represents a clock which has been designated for the network referenced only by message carrier 100 and where a computer executing aspects of message carrier 100 arbitrarily defines time for its affiliated computers in the distributed network. Clock 109 or clock 110 periodically output a common clock signal via either linkage 104 or 105, respectively.

Each member node 107, 108 is responsible for receiving the common clock signal; effectively "synchronizing" (aligning the definition of time in the node to a close tolerance with the definition of time respective to the common clock signal to create a system time base operating in synchronism with an established clock frequency and anchor point) to that common clock signal; receiving an appropriate local event table 114, 116 from the master scheduler 106; formulating a respective task execution list based on the received local event table 116, 114 and any additional tasks which are not described in the local event table 114, 116; executing the task execution list; generating the request list of communications 113, 115; and transmitting the request list of communications to the master scheduler 106. Each member node 107, 108 generates a connection request list 113, 115 based on the tasks required of the node 107, 108, and forwards these requests at a prescribed time through the

message carrier 100 to the master scheduler 106. The master scheduler 106 combines the connection request lists 113, 115 from the member nodes into a composite request list 111, which it uses to generate a communications schedule represented in the master event table 112. This master event table 112 is then broken apart into sections applicable to each of the member nodes 107, 108 and sent through the message carrier 100 (also at prescribed times) to the member nodes 107, 108 for storage in local event tables 114, 116. The local event tables 114, 116 are used in the member nodes 107, 108 to control timing for applicable time dependent tasks, including communications with other member nodes.

Figure 2 depicts a distributed network 220 of computing elements which are resident in physically different locations throughout the world to show an instance of the situation abstracted in the data flow diagram 130 of Figure 1. This figure includes a network of computers 201-214 and two sources of time in clock 203 and clock 214 representing time sources respective to the data flow diagram clocks 109, 110; in this regard, Figure 2 shows clock 203 as the basis for time definition for all systems in North America and Eurasia while clock 206 defines a basis for time definition for the less globally dispersed domain of computers 206, 207, 208, 209, and 214 in Africa. It should be noted that clock 203, then, "maps" to clock 109 in the data flow diagram context of Figure 1 and that clock 214 "maps" to clock 110 in the data flow diagram context of Figure 1. Figure 2 thereby depicts two possible configurations for distributed computing environments (facilitating the associated message carriers 100, master schedulers 106, and clocks 104, 105). A distributed computing environment can be located in a geographically limited area or dispersed throughout the world. When spread across large geographical areas, the clock 203 is preferably a satellite timing distribution system such as the Global Positioning System.

Figure 3 presents a diagram of a network in a bus structure 330. Figure 3 shows the fundamental components of Figure 1 except that the message carrier 100 is a bus 300 connecting the member nodes 107, 108, to the master scheduler 106.

Figure 4 shows a distributed computing environment where the network is implemented in a ring topology 430. Figure 4 shows the fundamental components of Figure 1 except that the message carrier 100 is a ring enabled by ring segments 402, 403, 405, and 401 to effectively connect the member nodes 107, 108, to the master scheduler 106.

Figure 5 presents a diagram of a network in a switched star topology 530 which is a preferred network architecture embodiment. Figure 5 also shows the fundamental components of Figure 1 except that the message carrier 100 is explicitly a switched network implemented in a star topology which connects the member nodes 107, 108 and the master scheduler 106, and that this master scheduler 106 is embedded into the network switch 501. This diagram shows the preferred embodiment as it could be used in a large manufacturing process (a manufacturing process referencing the effecting of goods manufacture in an apparatus whose movable components are modified in real-time using control signals from a computer control system executing logic reading measurements from the apparatus and outputting the control signals) through a large number of input/output processing devices and computing platforms in a potentially geographically distributed arrangement. The star topology network 530 enables the reading of measurements and outputting of control signals respective to the apparatus for manufacture through I/O (input signal/output signal) interface 510, I/O interface 511, I/O interface 522, I/O interface 519, I/O interface 512, I/O interface 513, I/O interface 520, and I/O interface 521. I/O interface 510, I/O interface 511, I/O interface 512, and I/O interface 513 are all joined via ring breakout 508 and ring breakout 509 to control computer 504 and historian computer 505 through local network 515. In a similar manner, I/O interface 522, I/O interface 519, I/O interface 520, and I/O interface 521 are all joined via ring breakout 518 and ring breakout 517 to control computer 506 and historian computer 507 through local network 514. I/O interface 510, I/O interface 511, I/O interface 512, I/O interface 513, ring breakout 508, ring breakout 509, control computer 504, historian computer 505, and local network 515 operate with local logical integrity and high real-time datalogical cohesion. In a similar manner, I/O interface 522, I/O interface 519, I/O interface 520, I/O interface 521, ring breakout 518, ring breakout 517, control computer 506, historian computer 507, and local network 514 operate with local logical integrity and high real-time datalogical cohesion. Cohesion between

(a) I/O interface 510, I/O interface 511, I/O interface 512, I/O interface 513, ring breakout 508, ring breakout 509, control computer 504, historian computer 505, and local network 515

and

(b) I/O interface 522, I/O interface 519, I/O interface 520, I/O interface 521, ring breakout 518, ring breakout 517, control computer 506, historian computer 507, and local network 514

via messaging is enabled by network switch 501 which implements (1) communications for provision of an essentially identical datalogical representation of the moment in real-time time among the computers; (2) master scheduler 106 processes to define for control computer 506, historian computer 507, control computer 504, historian computer 505, user interface 503, network support 502, and network switch 501 each moment in time when each communication instance between any two of computer 506, historian computer 507, control computer 504, historian computer 505, and user interface 503 VIA THE USE OF network switch 501 will occur; and (3) appropriate sending, receiving, and linkage tasks respective to components involved in the communication instance at the time defined for the communication instance. Network support 502 is also coordinated as a node into switched star topology 530 as a single source of information for defining acceptable communication partners in switched star topology 530 and also for accumulating and providing access to diagnostic information in switched star topology 530. Network switch 501 references network support 502 for a listing of acceptable communication partners in switched star topology 530 and confirms communication instance requests against this listing in defining sending, receiving, and linkage tasks.

In one embodiment, I/O interface 522, I/O interface 519, I/O interface 520, and I/O interface 521 are facilitated by a process control system such as is described in United States Patent 5,555,424 (24Sederlund, et. al.) issued on Sept. 10, 1996 and entitled "Extended Harvard architecture computer memory system with programmable variable address increment" to Sederlund, Edward R.; Lindesmith, Robert J.; Root, Larry A.; Dupree, Wayne P.; and Thomas, Lowell V. This patent is expressly incorporated herein by reference in the present application for showing a manner of making and using the present invention.

In a more preferred embodiment, I/O interface 522, I/O interface 519, I/O interface 520, and I/O interface 521, are facilitated via a redundant process control computer system in a general process control system using two process control computers (such as the process control computer described in 24Sederlund, et.) as is described in United States Patent 5,583,757 (Baca, Jr., et. al.) issued on Dec. 10, 1996 and entitled "Method of input signal resolution for actively redundant process control computers" to Baca, Jr., Eloy; Dupree,

Wayne P.; Grinwis, Donald J.; Kanse, Johannes C.; Pelletier, Douglas P.; and Schulze, Oscar E. This patent is expressly incorporated herein by reference in the present application for showing a manner of making and using the present invention.

An embodiment of a system for achieving data access for the process control computer described in 24Sederlund, et al. is described in United States Patent 5,568,615 (15Sederlund, et. al.) issued on Oct. 22, 1996 and entitled "Stealth interface for process control computers" to Sederlund, Edward R.; Thomas, Nadene T.; Lindesmith, Robert J.; and Cowles, Russell W. This patent is expressly incorporated herein by reference in the present application for showing a manner of making and using the present invention.

An embodiment of a system providing a remote field unit for use with the process control computer described in 24Sederlund, et al. is described the United States Patent 5,428,769 (69Glaser, et. al.) issued on Jun. 27, 1995 and entitled "Process control interface system having triply redundant remote field units" to Glaser, Robert S.; Hoy, Robert S.; Fernandez, G. Paul; and Grai, Timothy J. This patent is expressly incorporated herein by reference in the present application for showing a manner of making and using the present invention.

An embodiment of a system providing an interface for reading electrical current in power distribution systems for use with the field unit described in 69Glaser, et. al. is described in United States Patent 5,151,866 (66Glaser, et. al.) issued on Sept. 29, 1992 and entitled "High speed power analyzer " to Glaser, R. Steven and Bade, Jeffrey M. This patent is expressly incorporated herein by reference in the present application for showing a manner of making and using the present invention.

An embodiment of a system providing a high speed gateway for use with the redundant process control computer system described in Baca, et al. and the system for achieving data access for the process control computer described in 24Sederlund, et al. is described in

(a) United States Patent 5,519,603 (Allbery, Jr., et. al.) issued on May 21, 1996 and entitled "Intelligent process control communication system and method having capability to time align corresponding data sets" to Allbery, Jr., James D.; Troisi, Peter A.; Johnson, Susan J.; Cullen, James H.; Butler, Richard L.; Ferreira, James P.; Ellison, Joseph ; Patel, Chiman L.; Uban, James E.; and Schultz, Dale H.;

(b) United States Patent 5,428,745 (45de Bruijn, et. al.) issued on Jun. 27, 1995 and entitled "Secure communication system for re-establishing time limited communication between first and second computers before communication time period expiration using new random number" to de Bruijn, Ronny P.; Verboven, Marc L. K.; van Weele, Leonardus A.; Vermeire, Roger R.; Schulze, Oscar E.; Schultz, Dale H.; and Bell, Brian G.; and

(c) United States Patent 5,561,770 (70de Bruijn, et. al.) issued on Oct. 1, 1996 and entitled "System and method for determining whether to transmit command to control computer by checking status of enable indicator associated with variable identified in the command" to de Bruijn, Ronny P.; van Weele, Leonardus A.; Verboven, Marc L. K.; Vermeire, Roger R.; Schulze, Oscar E.; Bell, Brian G.; and Schultz, Dale H.

These patents are expressly incorporated herein by reference in the present application for showing a manner of making and using the present invention.

An embodiment of a system providing human interfacing for use with the redundant process control computer system and interfaces described in the above patents is described in United States Patent 5,631,825 (van Weele, et. al.) issued on May 20, 1997 and entitled "Operator station for manufacturing process control system" to van Weele, Leonardus A.; de Bruijn, Ronny P.; Vermeire, Roger R.; Zemering, Christo; and Lenting, Ben. This patent is expressly incorporated herein by reference in the present application for showing a manner of making and using the present invention. While this human interface system could be used as one user interface in each or any one of I/O interface 522, I/O interface 519, I/O interface 520, and I/O interface 521, it should be noted that user interface 503 is preferably designed for interface via network switch 501 to the entire set of systems in a particular instance of switched star topology 530.

Embodiments of systems providing real-time interpretation of application program code executing in the above systems are described (a) in United States Patent 5,491,625 (Pressnall, et. al.) issued on Feb. 13, 1996 and entitled "Information display system for actively redundant computerized process control" to Pressnall, Dana W.; Polishak, Jeffery T.; Felix, Bradley K.; Durisin, Michael J.; and Ellison, Joseph. and (b) in United States Patent 5,408,603 (Van de Lavoir, et. al.) issued on Apr. 18, 1995 and entitled "Global process

control information system and method" to Van de Lavoir, Ronny; Follon, Marinus (Neerpelt, BE); and Ravenscroft, Ian. These patents are expressly incorporated herein by reference in the present application for showing a manner of making and using the present invention.

A large graphical overview system providing interface to humans is deployed in the preferred embodiment along with the Operator Station. This graphical overview system is described in United States Patent 5,726,668 (Clement) issued on Mar. 10, 1998 and entitled "Programmable graphics panel" to John L. Clement. This patent is expressly incorporated herein by reference in the present application for showing a manner of making and using the present invention.

In a conceptualized most preferred embodiment, I/O interface 510, I/O interface 511, I/O interface 522, I/O interface 519, I/O interface 512, I/O interface 513, I/O interface 520, I/O interface 521, ring breakout 508, ring breakout 509, control computer 504, historian computer 505, ring breakout 518, ring breakout 517, control computer 506, historian computer 507, I/O interface 522, I/O interface 519, I/O interface 520, I/O interface 521, ring breakout 518, ring breakout 517, control computer 506, network support 502, network switch 501, and historian computer 507 are implemented using a dedicated context cycling computer as is described in US patent application number 08/797,967 which was filed on February 12, 1997 by Wayne Dupree, Jeffery Lucas, Larry Root, Gerrit Verniers, and Stephen Churchill entitled "A Dedicated Context-Cycling Computer" with the system of 69Glaser, et. al. facilitating interface to field instrumentation. This application for patent is expressly incorporated herein by reference in the present application for showing the preferred contemplated manner of making and using the present invention.

The context cycling processor provides for a substantial number and variety of input and output circuits, with each of the specialized circuits having at least one dedicated register for retaining the process and configuration state of pertinent resources, attributes, or aspects associated with either the circuit or function while another context is active; with the combination of specialized circuit and dedicated register(s) establishing the electrical circuitry base for a context. Contexts also include affiliated logic specific for the circuitry. A plurality of contexts share common assets within the CPU.

As needed, a co-processor for use with the context cycling computer is described in United States Patent 5,655,133 (Dupree, et. al.) issued on Aug. 5, 1997 and entitled "Massively

multiplexed superscalar Harvard architecture computer" to Dupree, Wayne P.; Churchill, Stephen G.; Gallant, Jeffrey R.; Root, Larry A.; Bressette, William J.; Orr, III, Robert A.; Ramaswamy, Srikala; Lucas, Jeffrey A.; and Bleck, James. This patent is expressly incorporated herein by reference in the present application for showing the preferred contemplated manner of making and using the present invention. Network switch 501 also preferably uses a cross-bar (Xbar) switch 516 capable of implementing at least 99 parallel and simultaneous communication linkages.

The master scheduler 106 is implemented in network switch 501 to provide a common clock signal (as an access to either clock 109 or a manifestation of clock 110) and is further responsible for the distribution of the common clock signal; the collection of the requests from each of the member nodes 108, 107 (in this instance of network 530, control computer 506, historian computer 507, control computer 504, historian computer 505, user interface 503, network support 502, and network switch 501 are all instances of member nodes 107 and 108 in the data flow context of Figure 1); the scheduling of the communication tasks for each of the member nodes 108, 107; the transmission of the schedule to each of the member nodes 108, 107; and the implementing of the connections between the member nodes 108, 107.

The purpose of distributing a common clock signal is so that each member node 108, 107 on the network 530 has the same understanding of time. This aspect of having the same understanding of time is crucial in this invention since the scheduling of tasks is done based on time. Thus, the common clock signal is used to effectively synchronize and time-align all member nodes 108, 107 to a common time. On a pre-defined periodic basis, the master scheduler 106 in network switch 501 sends a message that is propagated to each of the member nodes 108, 107 to define the beginning of a time period.

The collection of the requests from each member node 108, 107 (any one of control computer 506, historian computer 507, control computer 504, historian computer 505, user interface 503, network support 502, and network switch 501) occurs once each scheduling period. The master scheduler 106 in network switch 501 sets aside the time to receive requests from each member node 108, 107 (any one of control computer 506, historian computer 507, control computer 504, historian computer 505, user interface 503, network support 502, and network switch 501) each scheduling period. Requests originate from the member nodes 108, 107 which will need the communications requested. The master

scheduler 106 in network switch 501 receives each of these requests and considers them in its scheduling logic. Each request contains information on the source, destination, and length of the communication. In addition, each request may contain a priority indicator as well as information needed to deliver its message to a particular application which is executing on the destination node. After the task of collecting the connection requests from all of the member nodes 108, 107 is completed, the result is collected into a composite request list 111. Note that it is also possible for entries in the composite request list 111 to originate inside of the master scheduler 106 in network switch 501 itself; it is also possible for particular connection requests to have a repetitive nature such that multiple connections in a set of future time periods are invoked from a single connection request.

The master scheduler 106 in network switch 501 then schedules the communications which are to occur in the future. The method which is used in order to perform the scheduling of the communications is not critical in the functioning of the time deterministic network, although the Apparent Critical Resource approach discussed in this specification is the preferred approach. Logic implements to assign times to the requested communications. The determination of when to stop the scheduling process during the scheduling period can occur in one of three ways: first, all of the communication requests which can be scheduled are scheduled and therefore, there is no more scheduling work remaining; second, there is no more time remaining in which to perform the scheduling (the processor must stop the scheduling process in order to continue on with its next task); and third, there have been more connections requested than could be scheduled in the available scheduling window of time increment (an exhaustion of bandwidth), in which case the excess requests are retained as a queue for the next scheduling attempt.

The resulting schedule is stored in a master event table 112. This master event table 112 contains the start time 621 of the scheduled communication in addition to some of the same information similar to that in the composite request list 111. Note that it is also possible that certain tasks are pre-loaded into the master event table 112 in order to bypass the scheduling process, such as is required to enforce certain timing requirements of a task.

The master scheduler 106 in network switch 501 then communicates the local event table 114, 116 (that portion of the schedule which pertains to a particular member node 107, 108) to each appropriate member node 107, 108. For each event appearing in the master event table 112 which includes an appropriate member node 108, 107 as either the source or the

destination, an entry exists in the local event table 114, 116. Each member node 108, 107 (any one of control computer 506, historian computer 507, control computer 504, historian computer 505, user interface 503, network support 502, and network switch 501) which is involved in any communication receives its own local event table 114, 116.

The master scheduler 106 in network switch 501 then implements the connections according to the times which are assigned and appear in the master event table 112. The processing of these connections when using the preferred embodiment of the non-blocking switch leads to the situation whereby there may be multiple connections which are implemented concurrently.

The Composite Request List 111 residing in the master scheduler 106 in network switch 501 may contain both continuously repetitive connections to be implemented without contingencies, or dynamically requested connections (implemented when requested, usually initiated by one of the nodes on the network). This Connection Request List contains, at a minimum, a Source specifier, a Destination specifier, and a Duration specifier (see Figure 6 showing record details 600). Additional parameters pertinent to scheduling optionally include such specifiers as Priority 617.

Figure 6 shows the composition of the primary tables and messages 600 used in the scheduling and task coordination process. The connection request list 115 has Source 614, Destination 615, Length 616, and possibly Priority 617 fields 610 in both the member nodes 108, 107 and when combined into the composite request list 111 in the master scheduler 106 in network switch 501. The scheduler logic 611 in master scheduler 106 uses the data in these fields 610 to construct a schedule in the form of a master event table 112 with fields 612 for the Source 618, Destination 619, Length 620 and Start Time 621 for each of the connections. Note that the scheduler logic 611 takes the Source 614, Destination 615, and Length 616 from the connection request list 111, while the Start Time 621 is a product of the scheduling process itself.

The local event tables 114, 116 are simply subsets of the master event table 112, broken into those portions applicable to the respective member node 108, 107 (any one of control computer 506, historian computer 507, control computer 504, historian computer 505, user interface 503, network support 502, and network switch 501) - no manipulation of the fields 613 is required. Virtual paths 601, 602, 603, 606, 607, and 608 show the flow of the

information from the connection request list fields 610 to the local event table fields 613. Virtual path 604 shows the input of priority 617 to the scheduler; and virtual path 609 shows the start time 621, 625 as communicated from scheduler logic 611.

This data is then used by the scheduler logic 611 to derive a schedule for servicing the requested connections. The outcome of the scheduling process of the scheduler logic 611 is recorded in the Master Event Table 112, which is in turn parsed up into respective Local Event Tables 116 for distribution to the network nodes (see Figure 1). The Event Tables contain Source 622, Destination 623, Duration 624, and Start Time 625 fields 613.

Within the network nodes, the Local Event Tables 114, 116 are merged together with the Local Task Lists and logic executes to ensure that all the necessary work can be completed and that all timing corresponds to (or doesn't conflict with) that specified by the scheduler logic 611.

Figure 7 summarizes the general scheduling process 710. Connection requests are collected during a defined request acceptance period. After a defined period of time, the collection of connection requests must be processed. Additional connection requests are not prohibited, but a cutoff point must be established. Figure 7 shows the basic steps involved in the master scheduler 106 in network switch 501 node which are required in scheduling the connection requests. The first three steps 701-703 provide for aligned time definition in all nodes . In step 704, connection requests are received from all member nodes 107, 108 (the nodes sending their requests as a communication at predefined moments in their local task schedules from prior scheduling operations). In step 705, the connection requesting process ends so that a schedule can be defined for the future as the scheduling process applies to those requests which were collected from the member nodes 108, 107 during a particular window of time increment which is about to conclude. Schedule connections requests step 706 can begin once the requests have been collected (more detail in Figure 8). Once the schedule is complete, the applicable portions of the schedule are extracted from the master event table 112 in the creation of the local event tables 114, 116, in step 707 which are then conveyed to the nodes 107, 108 in step 708 for later implementation in step 709.

For distributed computing environments in general, connections between nodes 107, 108 may be required at any time, meaning that connection requests may originate any time, and some service time must be tolerable in handling the connections. There are four processes,

then, that repetitively execute at a rate conformant to the desired response in the network 530: collecting connection requests, scheduling connection requests, distributing the resulting schedule to the network nodes, and communicating through the network 530 via the established requested connections (see Table 1).

Table 1 - Four Scheduling Steps

Request Connections
Schedule Connections
Convey Schedule
Communicate

Table 2 shows the timing relationship between these four processes which, as stated, are always in progress.

Table 2 - The Scheduling Pipeline

Request Connections			
Request Connections	Schedule Connections		
Request Connections	Schedule Connections	Convey Schedule A	
Request Connections	Schedule Connections	Convey Schedule B	Communicate A
Request Connections	Schedule Connections	Convey Schedule C	Communicate B
Request Connections	Schedule Connections	Convey Schedule D	Communicate C
Request Connections	Schedule Connections	Convey Schedule E	Communicate D
Request Connections	Schedule Connections	Convey Schedule F	Communicate E

Note in the rightmost column of Table 2 that communications are always taking place after an initial few iterations establish a basis for communications. For any particular communications in a window of time increment, the schedule for the communications being implemented was conveyed in a previous window of time increment, the requests were scheduled in a window of time increment prior to the conveyance of the schedule, and the requests were requested even prior to scheduling. Timing variations can exist in this pipeline, but this resembles the pipeline used in the preferred embodiment.

The logic used in the currently preferred embodiment has been designed with two objectives in mind: an efficient schedule must be the outcome, and the number of cycles required to generate that schedule must be bounded at a reasonable level. Efficient scheduling involves a simple process based on the premise that the efficiency of the resultant schedule is primarily a function of the efficiency to which the busiest resources have been utilized in the schedule. This amounts to making sure that, at all times in the scheduling process,

scheduling precedence is given to the resources with the highest remaining demands. In order to derive a schedule in an acceptable amount of time, the logic was designed to schedule each request once and only once (trial and error logics deemed unacceptable), and with a moderate handling requirement for each connection request.

The basic process in scheduling 806, as shown in Figure 8, is to

1) define a criteria for measuring resource demand;
2) sort the resources according to criticality;
3) service those connections which can be serviced in this step and in the order of resource criticality;
4) step forward to the completion of the previously serviced requests thereby freeing up the applicable resources;
5) repeat steps 2-4 until all requests have been serviced or until there's no opportunity to service more requests.

Any requests which weren't able to be serviced in this scheduling pass can then be resubmitted during the next scheduling pass.

The flowchart of Figure 8 outlines the process used in the Apparent Critical Resource scheduling logic. The first step is to define the criteria by which resource usage is to be measured 800. For the preferred embodiment, transmit and receive channels are the resources which are scheduled, and the usage duration for each of these resources is the criteria chosen.

Once the resources are evaluated for usage, they are sorted by usage 801. This allows the resources which are used most heavily to be scheduled in step 802 first. When all resources which can be scheduled have been scheduled (the existing requested connections may not require or permit the simultaneous use of all resources), the scheduling window of time increment is advanced in step 803 until applicable resources become available. If there are more connection requests to be serviced and more time in the schedule to service these requests, this process is repeated after step 804 resolves the question respective to the need for another iteration through steps 801 to 803). The sub process 806 of the scheduler logic proceeds to step 805 to exit upon resolution by step 804 that the scheduling instance has been completed.

Figure 9 presents a more detailed overview of the Apparent Critical Resource (ACR) preferred embodiment sub-process 950 of the scheduling logic. This diagram encompasses the steps mentioned in Figure 8 and some additional actions required for conditional operation and exception handling as the preferred implementation of the process shown in Figure 8. Step 901 represents the initiation point for the preferred ACR process. In step 902, the ACR table is constructed, and a determination is made in step 903 respecting the presence of unscheduled requests based on the values of the SAV in the ACR table (a rapid exit to step 918 is executed if the values indicates that no requests are in the ACR table). If requests exist, the ACR sub-process 950 proceeds to calculate window size in step 904, look for a specific request in step 905, and check on the existence of the specific request and parties for the communication in step 906. If the request and parties exist, the process proceeds to add the event to the event table in step 912 and modify the ACR path in step 913 before proceeding to the next request in step 914. If the high value is not zero in step 915, the process proceeds again to step 904. If the high value is zero in step 915, the definition of (future) time in the ACR is advanced in step 916, the bandwidth is checked in step 917 for sufficiency in time to enable further scheduling in this scheduling operation, and the process returns to step 914 for the next request. If in step 906, a receiver for the communication did not exist, an alternative receiver is defined in step 907 and evaluated in step 908. If step 908 does not define an alternative receiver, an alternative transmitter is selected in step 919 and evaluated in step 920; if acceptable (step 920), step 921 selects the receiver with the highest ACR value and returns the process to step 908; and, if not acceptable (step 920), the process advances to step 916. When an alternative is defined in step 908, the process proceeds to schedule via steps 909, 910 and 911 to step 912. The virtual datalogical process being executed by these steps can be best described by considering an example respective to the list of communication requests for a hypothetical network consisting of five nodes shown in Table 3:

Table 3: Initial Request List

Source	Destination	Length	Priority
1	2	15	1
1	4	23	1
1	5	2	1
2	1	24	1
3	1	16	1
3	1	12	1
3	2	20	1
3	3	8	1
3	4	19	1
3	4	16	1
3	5	4	1
3	5	9	1
4	1	13	1
4	5	9	1
5	1	21	1
5	1	8	1
5	2	8	1
5	3	1	1
5	4	18	1
Total		246	

Assume that the length is measured in some arbitrary units of time, perhaps in milliseconds, and that there are n different priority levels, 1 through n, where 1 is the most urgent and n is the least.

First, note that with a broadcast type network (where all transactions are required to share the bandwidth, such as a Carrier Sense Multiple Access with Collision Detection or CSMA/CD network), service time required to complete the above requests would require a theoretical minimum of 246 units of time.

The incorporation of a non-blocking switch permits the simultaneous operation of the 5 transmit resources and the 5 receive resources. In analyzing the benefit of implementing a non-blocking network, the scheduler logic of the example first assesses the demands on each of the resources by constructing a table such as Table 4:

Table 4: Initial Summed Access Values

X / R		1	2	3	4	5		SAV
1		-	15	-	23	2		40
2		24	-	-	-	-		24
3		28	20	8	35	13		104
4		13	-	-	-	9		22
5		29	8	1	18	-		56
SAV		94	43	9	76	24		246

The summed access values, or SAVs, are simply the row or column totals for the transmitter (X-in the rows of the table) and receiver (R-in the columns of the table) interaction times and reflect the total demands on the particular resource, as measured in the arbitrary units of length. Note now that the theoretical minimum time required to service the requests has dropped from 246 units of time to 104 by selecting a non-blocking switch solution. The minimum service time for these requests is probably constrained by transmitter 3 since transmitter 3 has the highest SAV and the other transmitters can typically be serviced while transmitter 3 is being serviced. While the theoretical lower limit may be 104 units of time, the actual minimum service time may be greater and may even be dependent upon the order in which the connections are serviced. The next objective, however, is determining the optimum scheduling method in order to maximize the capabilities of the existing network (minimizing the required service time).

Assume for this particular example that the resultant schedule, regardless of the number or composition of the requests, applies to a future window of time increment of time with a duration of 150 units of time. The example begin by looking at the connections to be made at

Time = 0 for that 150 unit schedule. The primary objective, judging by the SAVs, is to schedule some request involving Transmitter 3 and Receiver 1. Referring back to the initial Request List, the scheduler logic of the example finds a request involving Transmitter 3 and Receiver 1 which has a length of 16 units of time. Since the scheduler logic of the example is currently scheduling connections at Time = 0 and the scheduler logic of the example has 150 units of time, the scheduler logic of the example concludes that the 16 word request fits into the scheduling window of time increment , and records the transaction into the Event Table of Table 5.

Table 5: Intermediate Event Table

Source	Destination	Length	Priority	Start Time
3	1	16	1	0
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-

Pictorially, Table 5 may be shown as schedule chart 1001 in Figure 10. Transmitter scheduled element 1002 and receiver scheduled element 1003 denote the scheduled tasks.

The Request List must then somehow reflect the fact that this last request has been serviced. This modified Request List may now resemble that shown in Table 6.

Table 6: Intermediate Request List

Source	Destination	Length	Priority	Start Time
1	2	15	1	-
1	4	23	1	-
1	5	2	1	-
2	1	24	1	-
3	1	16	1	0
3	1	12	1	-
3	2	20	1	-
3	3	8	1	-
3	4	19	1	-
3	4	16	1	-
3	5	4	1	-
3	5	9	1	-
4	1	13	1	-
4	5	9	1	-
5	1	21	1	-
5	1	8	1	-
5	2	8	1	-
5	3	1	1	-
5	4	18	1	-
Total		246		

Since Transmitter 3 and Receiver 1 are now busy for Time = 0, their status must indicate this, so the updated SAV Table, following some reformatting, may now resemble Table 7:

Table 7: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	40	78
2	24	43
3	88	9
4	22	76
5	56	24

Note that the SAVs for Transmitter 3 and Receiver 1 have been adjusted to reflect the serviced request. This is because the scheduler logic continues to give scheduling precedence to the resources with the highest impact on the schedule. This happens at the expense of the other resource needs, and at some point the inefficiencies injected into the resources with low SAVs may cause those resources to have a higher impact on the remaining portion of the schedule than those resources which had the highest impact at the start. For this reason, the protocol is referred to as Apparent Critical Resource scheduling (scheduling precedence is always given to resources which appear to be on the critical path at the time of consideration), or ACR scheduling.

It is now evident that the scheduler logic of the example should try to schedule a request from Transmitter 5 to Receiver 4, the next two concurrent resources in regards to scheduling impact which remain available for scheduling. There is such a request, so the scheduler logic of the example executes so that the tables now resemble Tables 8, 9, and 10.

Table 8: Intermediate Request List

Source	Destination	Length	Priority	Start Time
1	2	15	1	-
1	4	23	1	-
1	5	2	1	-
2	1	24	1	-
3	1	16	1	0
3	1	12	1	-
3	2	20	1	-
3	3	8	1	-
3	4	19	1	-
3	4	16	1	-
3	5	4	1	-
3	5	9	1	-
4	1	13	1	-
4	5	9	1	-
5	1	21	1	-
5	1	8	1	-
5	2	8	1	-
5	3	1	1	-
5	4	18	1	0
Total		246		

Table 9: Intermediate Event Table

Source	Destination	Length	Priority	Start Time
3	1	16	1	0
5	4	18	1	0
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-

Table 10: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	40	78
2	24	43
3	88	9
4	22	58
5	38	24

Note that the Event Table is effectively a (chronologically) sorted list of the scheduled requests from the Request List. For this reason, intermediate stages are not shown for each scheduling step.

Repeating the process as before, the scheduler logic of the example schedules a 1-to-2 communication instance request (that is a connection from Transmitter 1 to Receiver 2), and the tables are modified to Tables 11 and 12 as follows:

Table 11: Intermediate Request List

Source	Destination	Length	Priority	Start Time
1	2	15	1	0
1	4	23	1	-
1	5	2	1	-
2	1	24	1	-
3	1	16	1	0
3	1	12	1	-
3	2	20	1	-
3	3	8	1	-
3	4	19	1	-
3	4	16	1	-
3	5	4	1	-
3	5	9	1	-
4	1	13	1	-
4	5	9	1	-
5	1	21	1	-
5	1	8	1	-
5	2	8	1	-
5	3	1	1	-
5	4	18	1	0
Total		246		

Table 12: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	25	78
2	24	28
3	88	9
4	22	58
5	38	24

The scheduler logic of the example next looks for a 2-to-5 communication instance request, but such a request was not made. Several alternatives exist in determining the next combination to be attempted, and the optimum solution is chosen based on the desired tradeoffs between bandwidth efficiency and processing costs. (With minimal bandwidth impact, a simple rule set can be chosen, and for the preferred embodiment, an alternative receiver is selected for the existing transmitter until all receiver choices have been exhausted, after which that particular transmitter is ignored and the next transmitter is selected, and the highest impact receiver is re-considered - see Figure 9.)

Using this logic, a 2 to 3 communication instance request is searched for but not found, then a 4 to 5 communication instance request is searched for, found, and selected. The SAV Table is updated, and it is determined that no other connections can be made at Time = 0. The Event Table now is shown in Table 13:

Table 13: Intermediate Event Table

Source	Destination	Length	Priority	Start Time
3	1	16	1	0
5	4	18	1	0
1	2	15	1	0
4	5	9	1	0
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-
-	-	-	-	-

Again, these four connections can be represented pictorially in schedule chart 1101 in Figure 11. Transmitter scheduled element 1002 and receiver scheduled element 1003 from Figure 10 have been denoted as scheduled elements 1102 and 1103, respectively, and have been joined by scheduled elements 1104 and 1105 (the 5-to-4 communication), 1106 and 1107 (the 1-to-2 communication), and 1108 and 1109 (the 4-to-5 communication).

The logic involved in Advancing Time amounts to advancing to the next chronological finish. In this case, the 4 to 5 communication instance connection is the next chronological finish,

so the example advances time to Time = 9, at which time Transmitter 4 and Receiver 5 are again enabled. The Request List and SAV Table are now as shown in Tables 14 and 15.

Table 14: Intermediate Request List

Source	Destination	Length	Priority	Start Time
1	2	15	1	0
1	4	23	1	-
1	5	2	1	-
2	1	24	1	-
3	1	16	1	0
3	1	12	1	-
3	2	20	1	-
3	3	8	1	-
3	4	19	1	-
3	4	16	1	-
3	5	4	1	-
3	5	9	1	-
4	1	13	1	-
4	5	9	1	0
5	1	21	1	-
5	1	8	1	-
5	2	8	1	-
5	3	1	1	-
5	4	18	1	0
Total		246		

Table 15: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	25	78
2	24	28
3	88	9
4	13	58
5	38	15

After searches for a 2-to-5 communication instance request, a 2-to-3 communication instance request, a 4-to-5 communication instance request and a 4-to-3 communication instance request, the scheduler logic of the example concludes that the example must again advance time - nothing could actually be scheduled at Time = 9. Now, the next finish is at Time = 15, at which time the 1-to-2 communication instance connection finishes and the SAV Table is now Table 16.

Table 16: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	25	78
2	24	28
3	88	9
4	13	58
5	38	15

The scheduler logic of the example now looks for a 1-to-2 communication instance request, but the only 1-to-2 communication instance has already been scheduled, so the example look for a 1-to-5 communication instance request, which the example finds and schedules. The new SAV Table is then Table 17.

Table 17: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	23	78
2	24	28
3	88	9
4	13	58
5	38	13

The scheduler logic of the example then looks for a 2-to-2 communication instance request, a 2-to-3 communication instance request, a 4-to-2 communication instance request and a 4-to-3 communication instance request before determining that it's time again to advance time. The example advances time from 15 to 16, at which time the SAV Table looks like Table 18.

Table 18: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	23	78
2	24	28
3	88	9
4	13	58
5	38	13

The scheduler logic of the example then schedules a 3-to-1 communication instance request, which turns out to be the only request which can be scheduled at Time = 16. The tables, after the scheduler logic of the example advances time to Time = 17, look like Tables 19, 20, and 21.

Table 19: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	23	66
2	24	28
3	76	9
4	13	58
5	38	13

Table 20: Intermediate Request List

Source	Destination	Length	Priority	Start Time
1	2	15	1	0
1	4	23	1	-
1	5	2	1	15
2	1	24	1	-
3	1	16	1	0
3	1	12	1	16
3	2	20	1	-
3	3	8	1	-
3	4	19	1	-
3	4	16	1	-
3	5	4	1	-
3	5	9	1	-
4	1	13	1	-
4	5	9	1	0
5	1	21	1	-
5	1	8	1	-
5	2	8	1	-
5	3	1	1	-
5	4	18	1	0
Total		246		

Nothing can be scheduled at Time = 17, so the example advance to Time = 18 where the 5-to-4 communication instance request finishes up.

Table 21: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	23	66
2	24	28
3	76	9
4	13	58
5	38	13

After unsuccessfully searching for a 5-to-4 communication instance request, a 5-to-2 communication instance request is found, a 1-to-4 communication instance request is found, and the tables now look like Tables 22 and 23:

Table 22: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	-	66
2	24	20
3	76	9
4	13	35
5	30	13

Table 23: Intermediate Request List

Source	Destination	Length	Priority	Start Time
1	2	15	1	0
1	4	23	1	18
1	5	2	1	15
2	1	24	1	-
3	1	16	1	0
3	1	12	1	16
3	2	20	1	-
3	3	8	1	-
3	4	19	1	-
3	4	16	1	-
3	5	4	1	-
3	5	9	1	-
4	1	13	1	-
4	5	9	1	0
5	1	21	1	-
5	1	8	1	-
5	2	8	1	18
5	3	1	1	-
5	4	18	1	0
Total		246		

The scheduler logic of the example now advances time to Time = 26 for which the scheduler logic of the example defines Table 24:

Table 24: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	-	66
2	24	20
3	76	9
4	13	35
5	30	13

The scheduler logic of the example now schedules a 5-to-3 communication instance, which is all that can be scheduled at this time. Time is then advanced to 27, then to 28, and the SAV Table looks like Table 25.

Table 25: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	-	66
2	24	20
3	76	8
4	13	35
5	29	13

Since there are no 3-to-1 communication instance requests, the scheduler logic of the example schedules a 3-to-2 communication instance request and a 5-to-1 communication instance request, and advances time to 31 and then to 48, at which time the SAV Table becomes Table 26.

Table 26: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	-	45
2	24	-
3	37	8
4	13	16
5	8	13

The scheduler logic of the example then schedules the 3-to-4 communication instance request and advances time to 49, giving Table 27 as a result.

Table 27: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	-	45
2	24	-
3	37	8
4	13	16
5	8	13

The scheduler logic of the example then books a 2-to-1 communication instance request and advances time to 67, giving Table 28.

Table 28: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	-	21
2	-	-
3	29	8
4	13	16
5	8	13

The scheduler logic of the example now books the 3-to-3 communication instance request and advances time to 73 as shown in Table 29.

Table 29: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	-	21
2	-	-
3	29	-
4	13	16
5	8	13

The scheduler logic of the example then books the 4-to-1 communication instance request, and advances time to 75, resulting in Table 30.

Table 30: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	-	8
2	-	-
3	29	-
4	-	16
5	8	13

This allows the remaining 3-to-4 communication instance request to be scheduled by the scheduler logic of the example, after which the scheduler logic of the example advances to Time = 86, resulting in Table 31 data.

Table 31: Intermediate Summed Access Value Table

Channel	Transmitter SAV	Receiver SAV
1	-	8
2	-	-
3	13	-
4	-	-
5	8	13

The scheduler logic of the example now schedules the remaining 5-to-1 communication instance request, advances time to 91, schedules the first 3-to-5 communication instance request, advances time to 94, then to 95, and schedules the final request, another 3-to-5 communication instance request. The resultant Event Table then looks like Table 32.

Table 32: Final Event Table

Source	Destination	Length	Priority	Start Time
3	1	16	1	0
5	4	18	1	0
1	2	15	1	0
4	5	9	1	0
1	5	2	1	15
3	1	12	1	16
5	2	8	1	18
1	4	23	1	18
5	3	1	1	26
3	2	20	1	28
5	1	21	1	28
3	4	19	1	48
2	1	24	1	49
3	3	8	1	67
4	1	13	1	73
3	4	16	1	75
5	1	8	1	86
3	5	4	1	91
3	5	9	1	95

Note that the entire list of request completes in 104 units of time - the theoretical minimum amount of time required to service these particular requests.

Pictorially, the final schedule is shown in as schedule chart 1201 in Figure 12. Transmitter scheduled element 1002 and receiver scheduled element 1003 from Figure 10 have been denoted as scheduled elements 1202 and 1203; scheduled elements 1104 and 1105, 1106 and 1107, and 1108 and 1109 from Figure 11 are denoted as scheduled elements 1204, 1205, 1206, 1207,1208, and 1209; and new elements 1216-1217 for the 1-to-4 communication instance, 1214-1215 for the 5-to-2 communication instance, 1212-1213 for the 3-to-1 communication instance, 1220-1221 for the 3-to-2 communication instance, 1222-1223 for the 5-to-1 communication instance, 1226-1227 for the 2-to-1 communication

instance, 1224-1225 for the 3-to-4 communication instance, 1228-1229 for the 3-to-3 communication instance, 1230-1231 for the 4-to-1 communication instance, 1232-1233 for the 3-to-4 communication instance, 1234-1235 for the 5-to-1 communication instance, and 1238-1239 for the 3-to-5 communication instance are shown. Three of the interactions are shown without drawing element identifiers (the 1-to-5 communication having a duration of 2, the 5-to-3 communication having a duration of 1, and the 3-to-5 communication having a duration of 4.

Note that transmitter 3, the critical resource, is booked solid until Time = 104, a criteria which must be met in order to reach the theoretical minimum scheduling time. Note also that receiver 1, the second most critical resource, was booked solid until Time = 94. While this logic doesn't always produce the theoretically optimum schedule, it is designed to be efficient with the resources having the highest demands.

If the scheduling process had applied to a 100 unit window of time increment of time rather than a 150 unit window of time increment , the very last request would have been precluded from being scheduled in this window of time increment , but would have been queued until the scheduling process was resumed for the next time window of time increment .

The above example showed a list of requests where all of the requests had priority levels of 1. In one embodiment, four different priority levels are supported. All of the highest level requests are scheduled according to the logic before any of the other priority requests are even considered for scheduling.

Once all of the guaranteed requests have been scheduled, connection requests from the other three priorities are scheduled around the guaranteed requests already in the schedule. When there are already connections in the schedule, the process of calculating the window of time increment size is a little more complex, but the logic is essentially the same. Though the ACR value (or SAV) is still the predominant factor in deciding what gets scheduled when, connection requests between the same two concurrent resources are selected based on priority. This means, for example, that the ACR logic may elect to schedule a low priority 4-to-5 communication instance request before it schedules a medium priority 4-to-2 communication instance request, but never before it schedules a medium priority 4-to-5 communication instance request.

Since there are requirements for certain connections to happen at a particular time rather than simply in some window of time increment of time, accommodations can be made for having these requests bypass scheduling. These fixed time connections are entered into the Event Table before scheduling begins and the scheduled connections must be worked in around the fixed connections. These fixed time connections obviously have the highest priority of any of the requests.

In the above example, advancing time was simply a matter of finding the next chronological finish of an already scheduled connection. However, the process of advancing time consumes cycles which might be better spent performing other functions. The network scheduler in the preferred embodiment uses a minimum time increment of 256 units (with a scheduling window of time increment of over 25,000 units rather than the 150 unit window of time increment described above) to reduce the processing requirements. While this does have an obvious negative impact on scheduling efficiencies, there is also a less obvious positive impact. On the average, more resources are candidates for scheduling at each time increment as a result of having processed more finishes between increments. With more resources contending simultaneously, those resources with higher critical impact values are less likely to be preempted by some other resource which had simply finished earlier. The end result of this compromise is that compute cycles are saved, but bandwidth efficiencies aren't substantially impacted.

In the above example, all requests were scheduled in their entirety - there was no need to break up the requests. Had any one of the requests been of a length longer than the defined window of time increment (150 units of time in the above example), however, it would have been segmented into at least two portions to enable scheduling. If a long message cannot be avoided by design, by breaking up the message and using a low priority attribute in this situation, long lower priority messages are less likely to preclude higher priority messages contending for the same resources.

In the preferred embodiment, requests greater than 1023 units are broken into 1000 unit "full" blocks and an optional "partial" block up to 1023 units in length.

In the preferred embodiment, the Itineraries (the particular messages used to convey the schedule) are themselves scheduled messages due to their varying, and potentially large,

size. The Agendas (the messages used to collect the connection requests), however, are small and handled as fixed time connections (they are not literally scheduled). The scheduling pipeline used in the preferred embodiment involves collecting the agendas and scheduling the Itineraries in the first window of time increment, where the second, third, and fourth window of time increments are used as described (for scheduling the messages, conveying the schedule, and implementing the connections, respectively).

Turning now to Figure 13, a more detailed representation of components respective to a member node 107, 108 is depicted. While Figure 13 reprises a number of elements of Figures 1 and 5, details of logic types are shown enabling concurrent transmission and receipt of messages respective to the other nodes. In this regard, Transmit Logic 1301 and Receive Logic 1302 are shown for the Master Scheduler 106, with companion Transmit Logic 1319 and Receive Logic 1320 on member node 107 and further companion Transmit Logic 1321 and Receive Logic 1322 on member node 108. Pathways 1311, 1305, and 1306 show vectored message paths outbound from message carrier 100 (with crossbar switch 516). Such messaging approaches are achievable with a context cycling computer such as has previously been discussed, enabling a particular node 107, 108, 106 to be engaged in highly interleaved transmission and receipt of messages. Further detail showing the local scheduler 1307, oscillator 1308, CPU 1309, and VCXO (Voltage Controlled Crystal Oscillator) 1309 is shown for node 107 along with the previously discussed local task list 1303 and local task schedule 1304 which have also been discussed in general.

Turning now to Figure 14, focus is directed to three technologies which, when combined, illustrate a unique and valuable abstracted characterization of considerations in networking distributed, real-time computers to further put the nature of the present invention into full perspective. In this regard key technologies relate to Tight Synchronization between computers; Time Deterministic operation, including network communications; and Centralized Scheduling of distributed tasks, particularly those involving network communications. This diagram shows the interrelationships 1400 between the various technologies implemented in the preferred embodiment, with emphasis on those technologies considered critical to the implementation. The fundamental technologies which, when combined, provide an architecture amenable to real time computing in a distributed computing environment are (1) Time Deterministic Event Control 1422; (2) Tight Synchronization 1423 between systems; and (3) Centralized Scheduling 1421 of the tasks critical to the operation of the real time systems in the distributed environment.

The relationships between the technologies are shown using arrows, where the arrowheads point from the enabling technology to the enabled technology, and where the term "enabling" is loosely used here to describe an improvement in the utility of that technology, or in the simplicity of applying the technology. Note that many of the technologies are mutually beneficial, as depicted with a double headed arrow.

The following table explains the relationships between the implementation characteristics as virtual linkages of Figure 14 between the features and goals of tight synchronization 1423, prioritization 1424, time deterministic event control 1422, security 1426, limit checking 1425, centralized scheduling 1421, physical circuit switching 1434, non-blocking tree topology 1429, reliability 1428, active redundancy 1427, and low cost 1435 are interrelated in an abstracted characterization.

Table 33 Interactions of key technical features

Connection In drawing	Quality
1401	Centralized scheduling assures that all network interactions occur at prescribed times. The first requirement for Time Determinism, by this description, is having the nodes on the network in tight synchronization. None of the network nodes can know the needs of the entire network, so in the absence of the central scheduler, the alternatives are to either remove all dynamics or to implement a solution tolerant of the variability. By incorporating centralized connection scheduling, prioritization, and limit checking, the (high priority) time critical interconnections can be time deterministic while the (low priority) high volume interconnections can exist through variability tolerant designs.
1402	In a distributed computing environment, deterministic networking can be achieved either with tightly synchronized systems, or by providing enough excess bandwidth to realize timely delivery in a worst case timing and bandwidth demand situation. Because such an analysis of worst case situation is difficult and because providing the excess bandwidth can be costly, synchronization is an enabling technology for Time

	Determinism. Additionally, the difficulty in achieving a desired level of synchronization is usually proportional to the extent that the communications are indeterminate.
1403	Active Redundancy enables Time Determinism primarily in regards to failure recovery. It is difficult to manage the possible failure scenarios in the absence of redundancy, yet it is also difficult to manage failure recovery in a Time Deterministic system with passive redundancy. Time Determinism facilitates the use of Active Redundancy in that one of the biggest challenges to active redundancy is in assuring that all sides operate identically. Furthermore, with a Time Deterministic environment, a redundant element of a system can be diagnosed as dysfunctional simply by noting that the appropriate actions aren't occurring at the prescribed times. The time deterministic network involves communicating based on time between separate computing elements. The fact that the communications are scheduled based on time facilitates the employment of active redundancy. In active redundancy, there are two or more computers which are each executing identical logics via its own sequence of program instructions, exchanging the results of this execution, and then arbitrating to come to a result. Due to the fact that the data received by each computer must be shared with the other computer(s) executing the same logic prior to arbitrating this data, there must be some mechanism by which the time to begin the arbitration process in determined. The mechanism which has been chosen is to provide time synchronization between the redundant computers and schedule the tasks related to the communications based on time. In this way, there is known the time at which the arbitration task can begin.
1404	One of the most difficult challenges in implementing active redundancy is in handling the variation between sides. All sides of a redundant system must be expected to match at defined checkpoints, and discrepancies must be resolved in order to retain the value in redundancy. All sides must know when a checkpoint has been reached, and rules must be defined to allow differences to be resolved. With centralized scheduling,

	variations in connection requests, either in content or in timing, can be reduced or even eliminated, thereby increasing the probability that all sides can maintain operational equivalence. On the other hand, centralized scheduling places a tremendous dependency on the availability of that scheduling element. Centralized scheduling becomes much more viable in the presence of redundancy.
1405	Centralized scheduling requires that all connection requests be assembled together for analysis. This permits limit checking to verify that (1) certain undesirable node-to-node combinations are prohibited and (2) traffic levels (possibly at a particular priority) between certain combinations of nodes doesn't exceed predefined levels.
1406	Limit checking enhances the security of the distributed computing system by assuring that certain restrictions aren't violated. These limit checks help against malicious attacks where certain systems shouldn't be allowed to interact with others, and against accidental complications such as overloading a system with legitimate low priority data or unintentional overloading of high priority data.
1407	Prioritization of the messages enhances the utility of the limit checking. Not only can the highest priority data be made time deterministic without excessively restricting the lower priority messages, but security measures can be put in place which take message priority into account.
1408	Limit checking improves the overall reliability of the system by preventing certain security risks and by retaining the time determinism. In the absence of limit checking, certain network faults could cripple the network, such as when a node were to accidentally request an excessive amount of traffic at the highest priority level.
1409	In the implementation of Limit Checking, the highest priority traffic is testing against throughput thresholds to assure that no errors on one channel preclude the highest priority messaging to take place between any nodes not including the node with errors.
1410	Non-blocking network architectures allow the network to realize throughputs higher than the throughput available to any particular node on the network. This translates to higher bandwidth at any particular

	price when compared to blocking architectures.
1411	Security can often be implemented more easily when a centralized security agent can be entrusted with the administration of the security policies. Specifically in the preferred embodiment, the central scheduler is given the security tables and the interconnecting systems can delegate the security task to the central authority, allowing the administration to be done in one place rather than in two. Furthermore, the same security policies don't have to be re-implemented in as many environments as can exist on the star.
1412	By introducing a centralized scheduler, the scheduling logic can be implemented so as to enforce prioritization on the received connection requests.
1413	Prioritization simplifies the task of implementing Time Determinism. This is because large but infrequent tasks such as program loading can be set to operate at a low priority, eliminating the analysis of this high impact operation would have on the simpler yet more critical task of peer to peer communications.
1415	On networks employing a broadcast protocol, all network nodes have the opportunity to view the data going to all other nodes (though encryption can make this viewing rather futile). Blocking protocols aren't always implemented in broadcast protocols, but non-blocking protocols preclude the possibility of broadcasting, thereby avoiding the obvious security risks.
1416	One of the obvious ways of achieving reliability is in the introduction of redundancy.
1417	High throughput, the primary objective of high bandwidth, is positively impacted in two ways through the introduction of prioritization. First, there is the perception of high throughput in that the data that has the most value is delivered in a more timely fashion by invoking prioritization. Secondly, higher volumes of data can be permitted knowing that the high priority data, which often must be delivered in a timely fashion, won't be sacrificed by permitting large volumes of low priority traffic to be requested.
1418	Active Redundancy requires that a failed element be capable of

	recovering from downtime. This often involves the exchange of historical data from the elements which did not fail. In order to minimize the impact of this exchange, prioritization can be introduced assuring that the recovering element can recover and that it can do so without a substantial detrimental impact on the healthy elements.
1419	Certain architectures and topologies are inherently at risk of network outages in the presence of system outages. Use of a non-blocking architecture and physical circuit switching, while not the only means to achieve a higher reliability, offers a level of isolation from network node faults.
1420	Centralized scheduling is an enabling technology for effectively using the bandwidth available to non-blocking networking architectures such as those implementing certain crossbar switches. In the absence of scheduling, fixed connection patterns must be implemented, thereby reducing the practical throughput available to the nodes.
1431	Physical Circuit Switching technology can be difficult to implement for several reasons, including challenges associated with determining what circuits should be put in place at any particular time. By knowing the needs ahead of time, the connection patterns for the circuit switch can be determined a priori and implemented as needed.
1432	Circuit switching of systems configured in a tree topology can be implemented with a separate circuit, or pair of circuits, for each system on the network. This allows each channel to have its own path for communications, permitting a non-blocking architecture.
1433	Because physical circuit switching can reduce the requirement for active circuitry, high line speeds can often lead to higher network bandwidths.
1436	The simplicity of the circuitry associated with physical circuit switching often permits a given set of capabilities to be implemented at a lower cost.

Turning now to Figure 15, the general process 1500 executed respective to scheduling in each member node 108, 107 (any one of control computer 506, historian computer 507, control computer 504, historian computer 505, user interface 503, network support 502, and network switch 501) is discussed. Each member node 108, 107 is responsible for receiving

the common clock signal; synchronizing to that common clock signal; receiving the local event table 114, 116 from the master scheduler 106 in network switch 501; formulating the task execution list based on the received local event table 114, 116 and merging in any additional tasks which are not described in the local event table 114, 116; executing the task execution list; generating the request list of communications; and transmitting the request list of communications to the master scheduler 106 in network switch 501.

Steps 1501 and 1502 - Since the network is so highly dependent on a global time base, the first step 1501 at a member node 108, 107 is to receive the common clock signal. The clock signal is communicated as a message from the master scheduler 106 in network switch 501 to each of the member nodes 108, 107. The member node 108, 107 receives this clock signal in step 1501, stamps the actual time of reception, and validates the integrity of this message. The validated message is used as an input to the process of synchronization and, in step 1502, the difference between the expected and actual times of time message receipt is optionally determined (optionally if the nature of the time message is a pulse instead of a data bearing message).

Steps 1503 and 1504 - Each member node 108, 107 must now synchronize to the common clock signal. The process of performing this synchronization can be broken up into two separate tasks: first, in step 1503, the adjustment of a voltage controlled oscillator (1308) to a frequency which matches the source of the common clock signal and second, in step 1504, the alignment of the time such that the start of each time period which is defined by the common clock signal occurs at the same time in the member nodes 108, 107 as it does in the master scheduler 106 in network switch 501.

Step 1505 - Once the synchronization of a member node 108, 107 to the master scheduler 106 in network switch 501 is complete, the member node 108, 107 must receive the local event table 114, 116 from the master scheduler 106 in network switch 501. The local event table 114, 116 contains the schedule for the communications which are involving this member node. Note that the reception of the local event table 114, 116 itself is scheduled by the master scheduler 106 in network switch 501. Therefore, the first time that the node is to receive the local event table 114, 116 it must be available to receive that message at any possible time that the local event table 114, 116 might be received.

Step 1506 - The next step for the member node 108, 107 is to formulate the task execution list. The task execution list is formulated from the local event table 114, 116 and the additional tasks which the member node 108, 107 must do which are not described in the local event table 114, 116. The local event table 114, 116 which is communicated from the master scheduler 106 in network switch 501 provides the timing associated with the transmission and reception tasks. In order to formulate the task execution list, the following sub-steps must be taken:

Substep 1506A: Adjust the start time received from the master scheduler 106 in network switch 501 to account for the propagation delay. The timing which is received from the master scheduler 106 in network switch 501 is relative to the master scheduler 106 in network switch 501. Although the member node 108, 107 is synchronized with the master scheduler 106 in network switch 501, there is still an inherent delay in the propagation through the media being communicated. Either this delay must be accounted for by the master scheduler 106 in network switch 501, the member node, or the communication window of time increment must be sufficiently large that the inherent delay is included as overhead in the time allotted to the connection. In the preferred embodiment, this delay is accounted for in the member node.

Substep 1506B: Schedule the tasks which are dependent on the communication tasks. There are tasks which must be scheduled prior to a transmission (for example, the transfer of the data to be transmitted from memory); there are also tasks which must be scheduled after a reception (for example, the arbitration of the data between the multiple, different sources in a redundant computing environment and the transfer of data to memory). The member node 108, 107 is responsible for the scheduling of these tasks.

Substep 1506C: Schedule the tasks which are neither described by the local event table 114, 116 nor are dependent on the communications. There are tasks which the member node 108, 107 must accomplish in addition to communication related tasks. These tasks may be for the operation of other I/O, localized display, etc.

Step 1507 - The next step for the member node 108, 107 is to execute the task execution list. Each task in the list contains the task which is to be executed, the start time for the task, and any parameters which are necessary in order to successfully complete the task.

Step 1508 - The member node 108, 107 can now formulate the communication request list, which is a list of the member nodes 108, 107 to which this member node 108, 107 is desiring to communicate. The information which is included in this communication request list is the source node identifier, destination node identifier, and the total length of the communication.

Step 1509 - This communication request list is communicated to the master scheduler 106 in network switch 501 for consideration in its scheduling. At this point, the member node 108, 107 has performed all of the necessary functions in a particular time period. The next step 1510 is to determine whether or not it is time to receive the next reference moment message. If it is, then the entire process repeats itself starting with the very first step 1501 of receiving the reference moment message. If it is not time to receive the next reference moment message, then the logic of the member node may forward to step 1505.

As one contemplated embodiment, Figure 16 presents the general data flow diagram 130 of Figure 1 as a data flow diagram 1600 where communications are effected via radio frequency means and transmissions 1601, 1602, and 1603 are effected via radio frequency transmitters and receivers.

The present invention is enabled through use of machines which are considered as a special purpose computers created in one embodiment by combining a computer with computer program code (alternatively termed, for example, software or logic, where logic denotes software in and for execution by an enabling system of electrical circuitry) directed to the particular function of process control so that, when connected to an electrical power source, powered by electricity, and activated to execute its computer program code (software), the special purpose computer executes the particular function related to the present invention pursuant to the instructions from the computer program code. The computer has electrically activated components denoted as a central processing unit, at least one physical memory connected to the central processing unit, an optional input keyboard connected to the central processing unit, an optional output display (which can include, without limitation, a printer, printing terminal where output is printed on paper, cathode ray tube monitor, and/or flat panel monitor) connected to the central processing unit, a computer clock pulse generator, and a connection to electrical power for energizing all of the electrically activated components of the computer. The central processor further has a control bus and specific computer circuits for either temporarily holding data (for example, a register or an

accumulator) or for executing a fundamental data processing operation (for example, an addition circuit); the specific computer circuits are connected in communication to the control bus through latching circuits which can individually be in either a conducting (communicating) or non-conducting (non-communicating) configuration; the collective configurations of all the latching circuits at a particular moment define a particular state of the central processor; and the state of the central processor is respectively modified and changed as the individual instructions of the computer program code are sequentially accessed by the central processing unit and output from the clock pulse generator is directed to the latches. The computer optionally also has computer program code in the form of a commercially available operating system which functions as a computer program code platform for enabling the execution of the special purpose computer program code directed to the particular function of the present invention. In overview, the present invention in this embodiment is implemented in the preferred embodiment by and through a plurality of computers which have been programmed to contribute to technology by performing the specifically useful purpose and practical application of the described system; each computer functions as a medium for realizing the functionality provided by the computer program code's functionality. In some embodiments, some communication linkages within an extended form of the special purpose computer may be implemented by electrical, optical, or radio-frequency signal communication interfaces and communication media which collectively use electrically-powered transmitter and receiver components which are directly or indirectly linked to at least one central processing unit.

The computer program code resident within a computer provides at least one component part of that machine. The computer program code may be in the form of, without limitation, inter-linked modules, routines and subroutines, program objects, and/or installed interactive processes. In some emergent computer designs, enabled by modern application specific integrated circuitry (ASIC) technology, computer program code may be at least partially expressed in hardwired circuitry reminiscent of traditional hardwired gate arrangements which execute simplified logical scenarios without need for a clock pulse. The Figures of this description show a number structural relationships between key modules and components within the computer program code and other aspects of the invention which enable the functionality of the machine of the present invention to be realized; in this regard, the computing elements in the system are constructed out of a combination of hardware and software.

The present invention has been described in an illustrative manner. In this regard, it is evident that those skilled in the art, once given the benefit of the foregoing disclosure, may now make modifications to the specific embodiment described herein without departing form the sprit of the present invention. Such modifications are to be considered within the scope of the present invention which is limited solely by the scope and spirit of the appended claims.

We claim:

1. A computer implemented method for enabling a plurality of computer nodes in mutual data communication to execute logic in a coordinated manner, comprising:

aligning real-time time between said plurality of computer nodes so that each computer node has essentially the same real-time time;

defining a schedule of communications for each computer node, said schedule having a different time in the future defined for initiating each respective communication;

communicating each schedule to its respective computer node; and

initiating each communication from each node at the time defined in its schedule.

2. A computer network architecture, comprising:

a message carrier supporting multiple communication paths;

a plurality of computers coupled to said message carrier, said computers being programmed to perform time-deterministic tasks by cycling through a plurality of different operating states, including a communication state;

a clock system accessible by all said computers for maintaining time synchronism among said computers; and

a scheduler coupled to said message carrier for determining the sequence of communication between said computers by causing said message carrier to establish communication between selected computers based on said time-deterministic tasks being performed by said computers.

3. A computer implemented method for enabling a computer node in a plurality of computer nodes to execute logic in coordination with the other nodes in said plurality of computer nodes, comprising:

connecting a message carrier to each computer node in said plurality of computer nodes;

providing master scheduling capability in one computer node to establish a master scheduler;

defining, in each computer node, a communication request list of each communication needed by said computer node with any other computer node respective to a future time period;

sending said communication request list from each computer node to said master scheduler via said message carrier;

receiving, in each computer node, a communication schedule from said master scheduler via said message carrier;

defining, in each computer node, a local task schedule from said communication schedule.

4. A computer implemented method for enabling a plurality of member nodes to execute logic in a coordinated manner, comprising:

connecting a message carrier to each member node in said plurality of member nodes;

providing a master scheduler in data communication with each said member node via said message carrier;

defining a reference moment time value for said plurality of member nodes;

communicating said reference moment time value to each member node and said master scheduler via said message carrier;

defining a communication request list in each said member node showing each communication needed by said member node with any other member node in a future time period;

communicating said communication request list from each member node to said master scheduler via said message carrier;

defining, in said master scheduler, scheduled communications between member nodes at communication event times respective to said reference moment time value and also respective to each said communication listed in each said received communication request list;

defining, in said master scheduler, a communication schedule for each member node respective to said scheduled communications ;

sending each said communication schedule from said master scheduler to each respective member node via said message carrier;

defining, in each member node, the difference in time between said reference moment time value and the time when said reference moment time value is communicated into said member node;

defining, in each member node, a local task schedule from said communication schedule and from said difference in time for the member node so that the member node will communicate to other member nodes via said message carrier at said scheduled communication event times;

executing said scheduled communications between said member nodes respective to each scheduled communication via said message carrier according to said respective local task schedules.

5. The method of Claim 4 wherein said step of defining a reference moment time value is performed by said master scheduler.

6. The method of Claim 4 wherein said step of defining a communication request list in each said member node limits said communications to a maximum number in said future time period.

7. The method of Claim 4 wherein each communication has a duration time attribute and the value of said duration time attribute is less than a fixed maximum value.

8. The method of Claim 4 wherein the step defining a reference moment time value for said plurality of member nodes is done by a global time system.

9. A computer implemented method for enabling a plurality of member nodes to execute logic in a coordinated manner, comprising:

connecting a message carrier to each member node in said plurality of member nodes, each node having a central processing unit;

providing a master scheduler in data communication with each said member node via said message carrier;

defining a reference moment time value as a moment in time for said plurality of member nodes;

communicating said reference moment time value to each member node and said master scheduler via said message carrier;

defining into a communication request list in each said member node, each access to any other member node needed by said member node in achieving a future task increment respective to a future time period;

communicating said communication request list from each said member node to said master scheduler via said message carrier;

defining, in said master scheduler, a communication event data object respective to each said access to any other member node needed by said member node in achieving a future task increment, said communication event data object having an accessing member identification attribute value, an accessed member identification attribute value, and an access duration attribute value;

collecting, in said master scheduler, all defined communication event data objects into a communication event set;

determining in said master scheduler a summed accessed value for each member node in said communication event set as a sum of all access duration attribute values in all communication event data objects having a accessed member identification attribute value respective to said member node;

determining in said master scheduler a summed accessing value for each member node in said communication event set as a sum of all access duration attribute values in all communication event data objects having an accessing member identification attribute value respective to said member node;

scheduling in said master scheduler and into the earliest available time in a schedule, a future communication task respective to said reference moment time value between the accessing member node having the highest said summed accessing value respective to said communication event set and the accessed member node respective to the accessed member identification attribute value in the communication event data object having the highest affiliated access duration attribute value for all communication event data objects respective to said accessing member node, said future communication task having a task initiation time equal to the earliest available time in the schedule, an accessing member node identifier attribute respective to the accessing member, and an accessed member node identifier attribute respective to the accessed member node;

removing, in said master scheduler, the communication event data object respective to the accessed member node identifier attribute in the scheduled future communication task and the accessed member node identifier attribute in the scheduled future communication task to redefine the communication event set to those communication event data objects which have not been scheduled; and

repeating, in said master scheduler, respective to the redefined communication event set, the steps of determining a summed accessed value, determining a summed accessing value, scheduling a future communication task, and removing all communication event data objects respective to scheduled member nodes until all communication event data objects have been scheduled into a communication schedule for each member node;

sending each said communication schedule from said master scheduler to each respective member node via said message carrier;

defining, in each member node, the difference in time between said reference moment time value and the time when said reference moment time value is communicated into said member node;

defining, in each member node, a local task schedule from said communication schedule and from said difference in time for the member node so that the member node will communicate to other member nodes via said message carrier at said scheduled communication event times; and

executing said scheduled communications between member nodes respective to each scheduled communication event via said message carrier according to said respective local task schedules.

10. A computer implemented method for scheduling communications between member objects in a plurality of member objects functioning as a collective group, comprising the steps of:

providing a computer;

defining in said computer, for each member object, a communication event data object respective to each access to any other member object needed by said member object in achieving a future task increment, said communication event data object having an accessing member identification attribute value, an accessed member identification attribute value, and an access duration attribute value;

collecting, in said computer, all defined communication event data objects into a communication event set;

determining, in said computer, a summed accessed value for each member object in said communication event set as a sum of all access duration attribute values in all communication event data objects having a accessed member identification attribute value respective to said member object;

determining, in said computer, a summed accessing value for each member object in said communication event set as a sum of all access duration attribute values in all communication event data objects having an accessing member identification attribute value respective to said member object;

defining a set of sequential future time increments within a scheduling period respective to a schedule;

scheduling, in said computer, into the earliest available time in the first of said future time increments, a future communication task based on the first available communication event data object in said communication event set between

the accessing member object having the highest said summed accessing value respective to said communication event set and

the accessed member object having the highest said summed accessed value respective to said communication event set if a communication event data object exists between said accessed member object and said accessing member object,

said future communication task having a task initiation time equal to the earliest available time in the schedule, an accessing member object identifier attribute respective to the accessing member object, an accessed member object identifier attribute respective to the accessed member object, and said access duration attribute value;

removing from scheduling candidacy, in said computer, the communication event data object respective to the accessed member object identifier attribute in the scheduled future communication task and the accessing member object identifier attribute in the scheduled future communication task to redefine the communication event set to those communication event data objects which have not been scheduled;

repeating, in said computer respective to the redefined communication event set, the steps of determining a summed accessed value, determining a summed accessing value, scheduling a future communication task, and removing all communication event data objects respective to scheduled member objects until all communication event data objects have been scheduled which can be scheduled into said first future time increment; and

recurrently executing, in said computer respective to the redefined communication event set and sequentially respective to each of the remaining of said plurality of future time increments, the steps of determining a summed accessed value, determining a summed accessing value, scheduling a future communication task, and removing all communication event data objects respective to scheduled member objects, and repeating such steps until all communication event data objects have been scheduled which can be scheduled into each said future time increment and until all communication event data objects have been scheduled into said scheduling period.

11. The method according to 10, wherein said defining is confined to future task increments which are needed within a predefined time period and said scheduling is respective to said predefined time period.

12. The method according to 11, further comprising the step of connecting a message carrier to each member object in said plurality of member objects, said message carrier having a plurality of channels for permitting concurrency in said communication events, said step of scheduling being respective to said plurality of channels as a group wherein said earliest available time is respective to any channel in said message carrier.

13. The method according to Claim 12 wherein each said member object is a member node being in data communication with said message carrier, each member node having a central processing unit.

14. A computer implemented method for scheduling tasks for resource objects in a plurality of resource objects functioning as a collective group, comprising the steps of:

providing a computer;

defining, in said computer, a set of future tasks for said collective group;

defining, in said computer, each resource object needed for each future task in said set of future tasks;

defining, in said computer, each interaction needed between any two resource objects respective to said set of future tasks;

defining, in said computer, a use instance for each said resource object in achieving each said interaction, said use instance having an use duration attribute value;

defining, in said computer, an accessing resource object and an accessed resource object respective to each said use instance;

firstly determining, in said computer, a summed accessing time value for each accessing resource object as a sum of all use duration attribute values for all said use instances respective to that resource object when it is the accessing resource object;

secondly determining, in said computer, a summed accessed time value for each accessed resource object as a sum of all use duration attribute values for all said interactions between said accessed resource object and said accessing resource object;

defining, in said computer, a set of sequential future time increments within a scheduling period respective to a schedule;

scheduling, in said computer, into the earliest available time in the first of said future time increments, a future resource interaction event based on the first available future task in said set of future tasks between

the accessing resource object having the highest said summed accessing value respective to said set of future tasks and

the accessed resource object having the highest said summed accessed value respective to said set of future tasks if a use instance exists between said accessed resource object and said accessing resource object,

said future resource interaction event having an initiation time equal to the earliest available time in the schedule, an accessing resource object identifier attribute respective to the accessing resource object, an accessed resource object identifier attribute respective to the accessed resource object, and said use duration attribute value;

removing from scheduling candidacy, in said computer, the use instance respective to the accessed resource object in the scheduled future resource interaction event and the accessing member object identifier attribute in the

scheduled future resource interaction event to redefine the set of future tasks to those future tasks for which interactions have not been scheduled;

repeating, in said computer, respective to the redefined future task set, the steps of determining a summed accessed value, determining a summed accessing value, scheduling a future communication task, and removing all use instances respective to scheduled resource objects until all use instances have been scheduled which can be scheduled into said first future time increment; and

recurrently executing, in said computer, respective to the redefined future task set and sequentially respective to each of the remaining of said plurality of future time increments, the steps of determining a summed accessed value, determining a summed accessing value, scheduling a future resource interaction event, removing all resource objects respective to scheduled resource objects, and repeating such steps until all resource interaction events have been scheduled which can be scheduled into each future time increment and until all resource interaction events have been scheduled into said scheduling period.

15. A computer implemented apparatus for enabling a plurality of computer nodes in mutual data communication to execute logic in a coordinated manner, comprising:

means for aligning real-time time between said plurality of computer nodes so that each computer node has essentially the same real-time time;

means for defining a schedule of communications for each computer node, said schedule having a different time in the future defined for initiating each respective communication;

means for communicating each schedule to its respective computer node;

means for initiating each communication from each node at the time defined in its schedule.

16. A computer-implemented apparatus having a message carrier, comprising:

a plurality of member nodes, each member node being in data communication with said message carrier, each member node having

a central processing unit,

time phase logic for receiving from said message carrier a reference moment message having a reference moment time value and for defining the difference in time between said reference moment time value and the time of receiving said reference moment message,

communication definition logic for identifying each internodal communication needed by said member node in a future time period and for establishing a communication request list respective to all said internodal communications,

communication request sending logic for communicating said communication request list via said message carrier to a master scheduler,

communication task receiving logic for receiving a communication schedule via said message carrier from said master scheduler,

task integration logic for defining a local task schedule from said communication schedule and said difference in time, and

communication implementation logic for communicating via said message carrier with other member nodes according to said local task schedule; and

a master scheduler in data communication with said plurality of member nodes via said message carrier, said master scheduler

having scheduler logic executed by a central processing unit, said scheduler logic having

reference moment definition logic for defining said reference moment time value as a moment in time,

reference moment communication logic for creating a reference moment message having said reference moment time value and for communicating said reference moment message via said message carrier to each member node,

communication request receiving logic for receiving each said communication request list from each said member node via said message carrier,

communication schedule definition logic for scheduling communications between member nodes at communication event times respective to said reference moment time value and also respective to each said internodal communication listed in each said received communication request list and for further defining a communication schedule for each member node respective to said scheduled communications, and

member node communication schedule response logic for sending each said communication schedule to each respective member node via said message carrier.

17. The apparatus of Claim 16, wherein each member node further comprises a variable speed oscillator.

18. The apparatus of Claim 16, wherein each member node further comprises time alignment logic and oscillator rate adjustment logic.

19. The apparatus of Claim 16, wherein said message carrier comprises a crossbar switch.

20. The apparatus of Claim 16, wherein said message carrier comprises a network topology.

21. The apparatus of Claim 16, wherein said message carrier comprises a radio frequency communications system.

FIG. 1

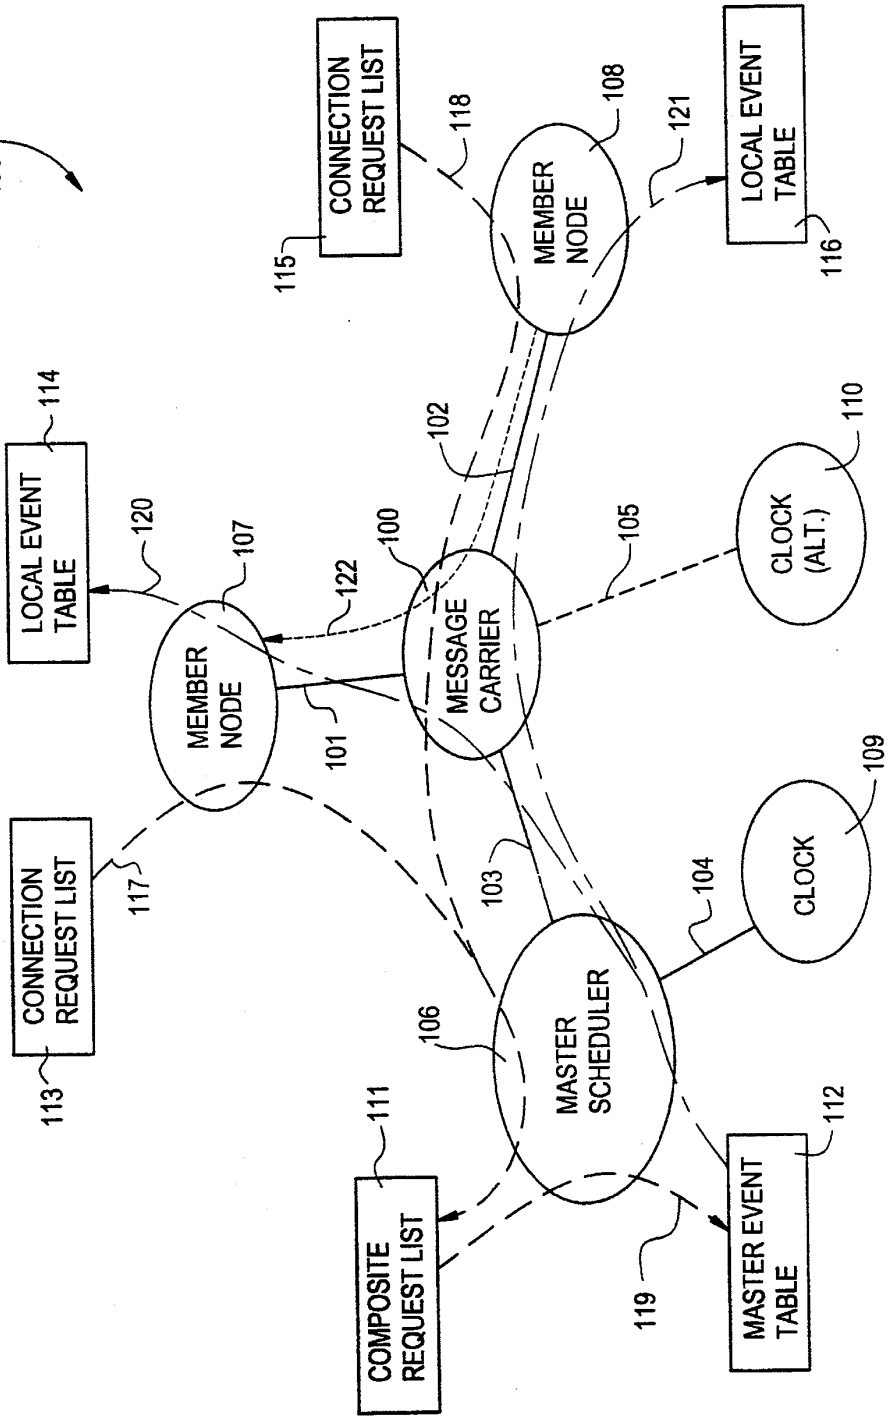

130
CONNECTION REQUEST LIST
115
118
MEMBER NODE
108
121
LOCAL EVENT TABLE
116
LOCAL EVENT TABLE
114
120
MEMBER NODE
107
102
100
122
MESSAGE CARRIER
105
CLOCK (ALT.)
110
101
CONNECTION REQUEST LIST
113
117
103
104
CLOCK
109
MASTER SCHEDULER
106
111
112
COMPOSITE REQUEST LIST
119
MASTER EVENT TABLE

FIG. 2

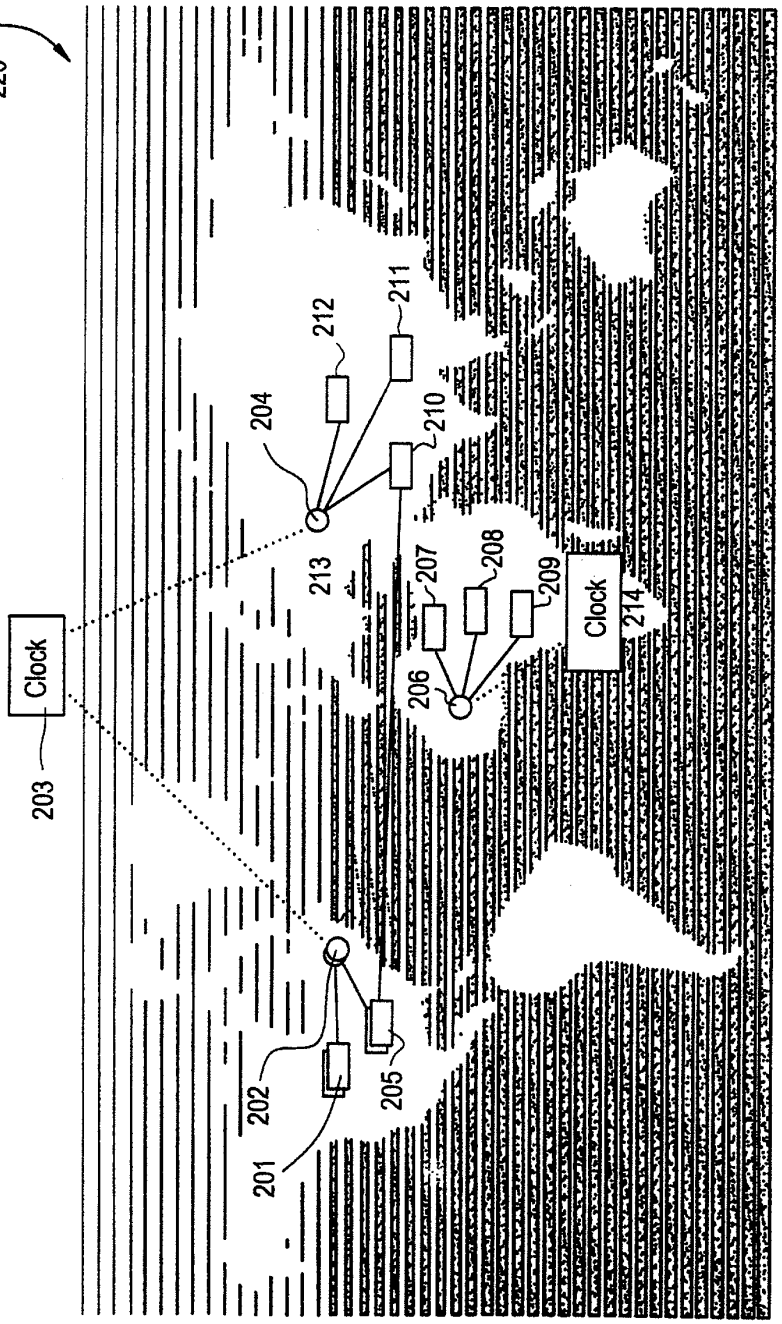

220
Clock
203
201
202
205
204
213
212
211
210
207
208
209
Clock
214
206

FIG. 3

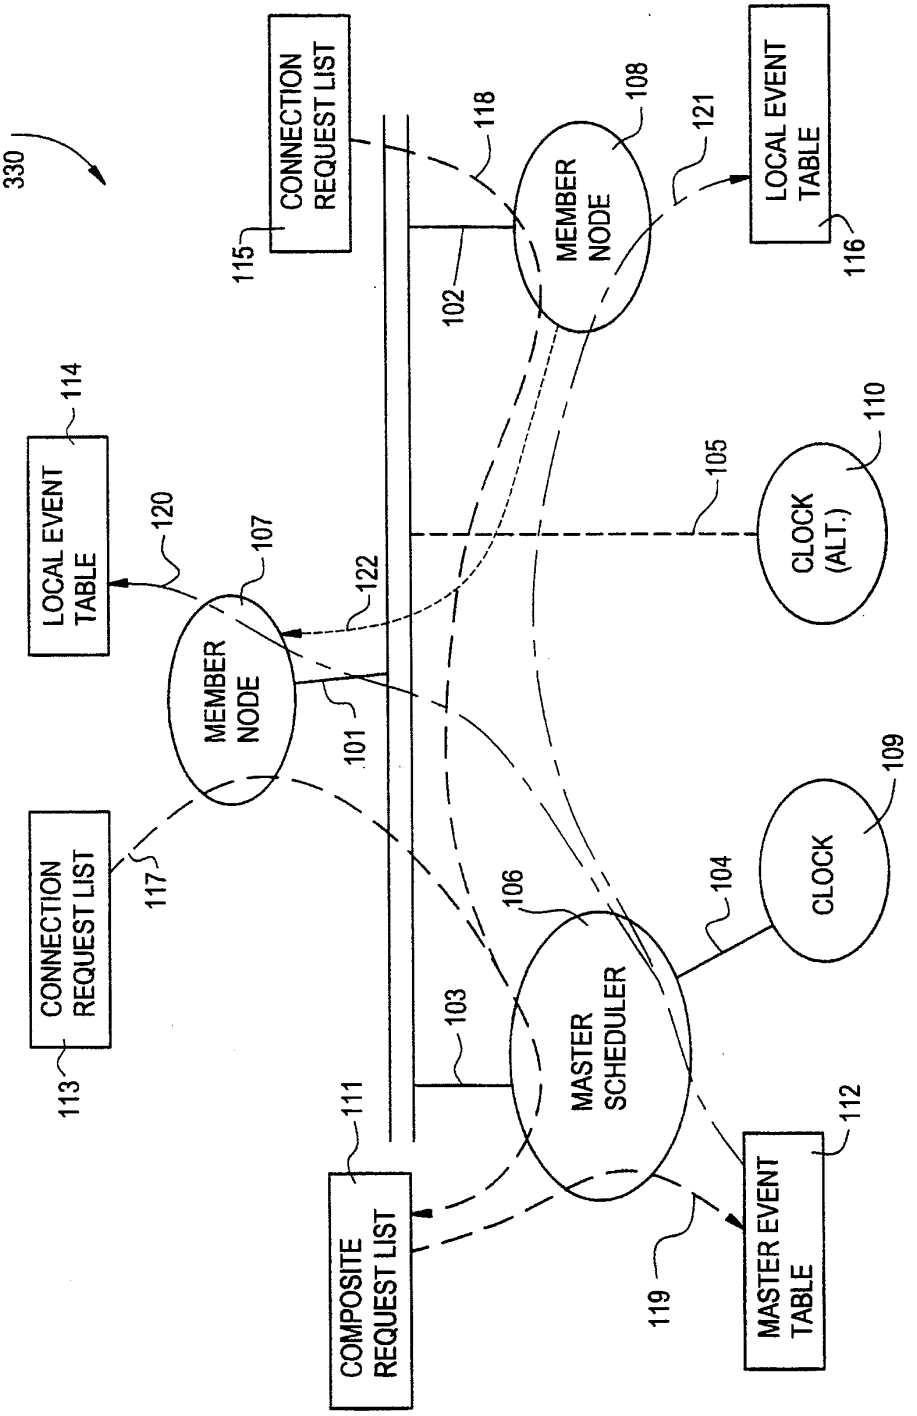

330
CONNECTION REQUEST LIST
115
118
MEMBER NODE
108
121
LOCAL EVENT TABLE
116
102
LOCAL EVENT TABLE
114
120
MEMBER NODE
107
122
105
CLOCK (ALT.)
110
101
CONNECTION REQUEST LIST
117
113
106
104
CLOCK
109
MASTER SCHEDULER
103
111
COMPOSITE REQUEST LIST
119
MASTER EVENT TABLE
112

FIG. 4

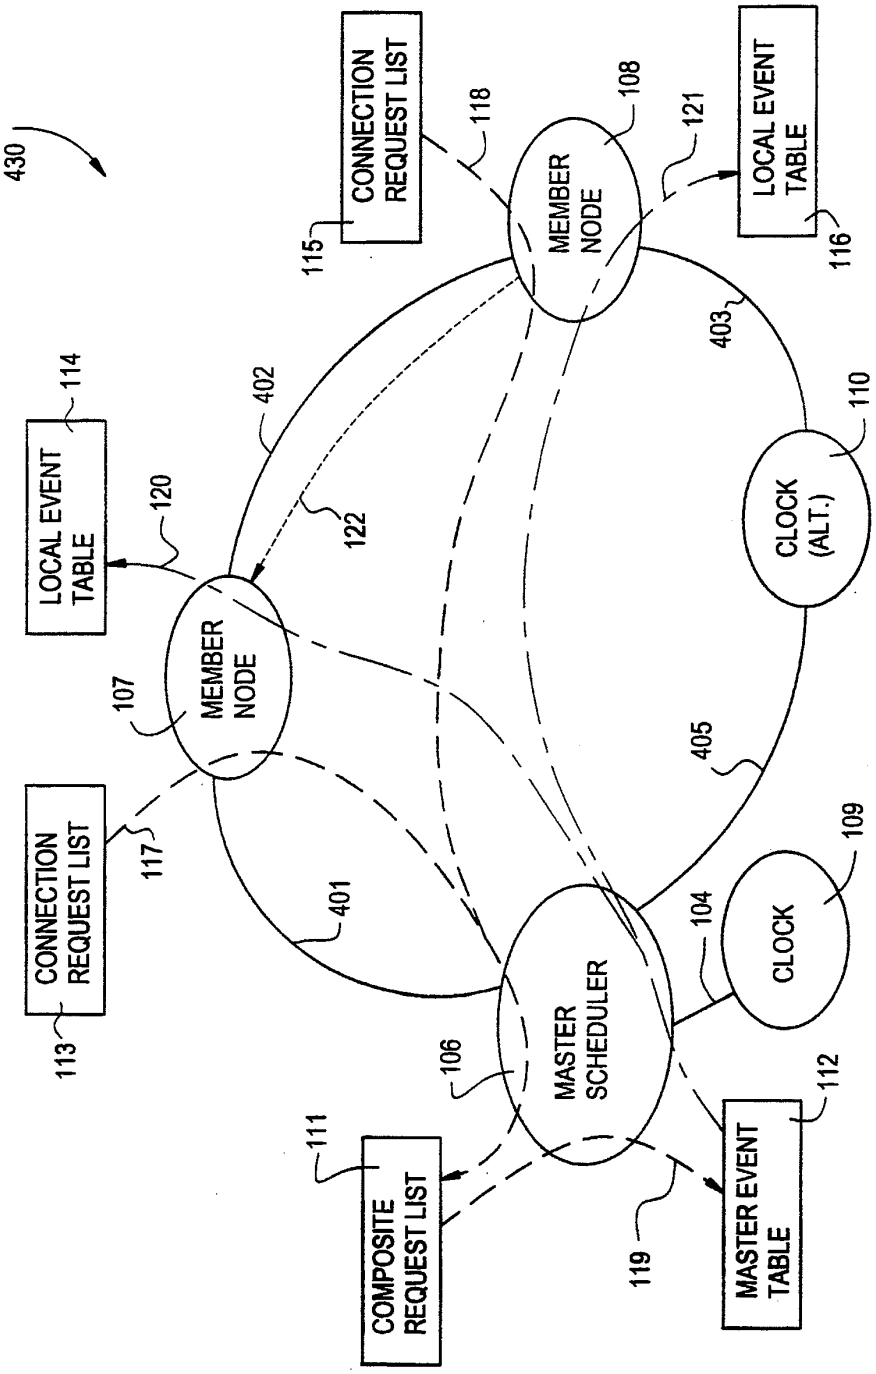

430
CONNECTION REQUEST LIST
115
118
MEMBER NODE
108
121
LOCAL EVENT TABLE
116
403
402
LOCAL EVENT TABLE
114
120
122
CLOCK (ALT.)
110
MEMBER NODE
107
405
CONNECTION REQUEST LIST
113
117
401
CLOCK
109
104
MASTER SCHEDULER
106
112
111
COMPOSITE REQUEST LIST
119
MASTER EVENT TABLE

FIG. 5

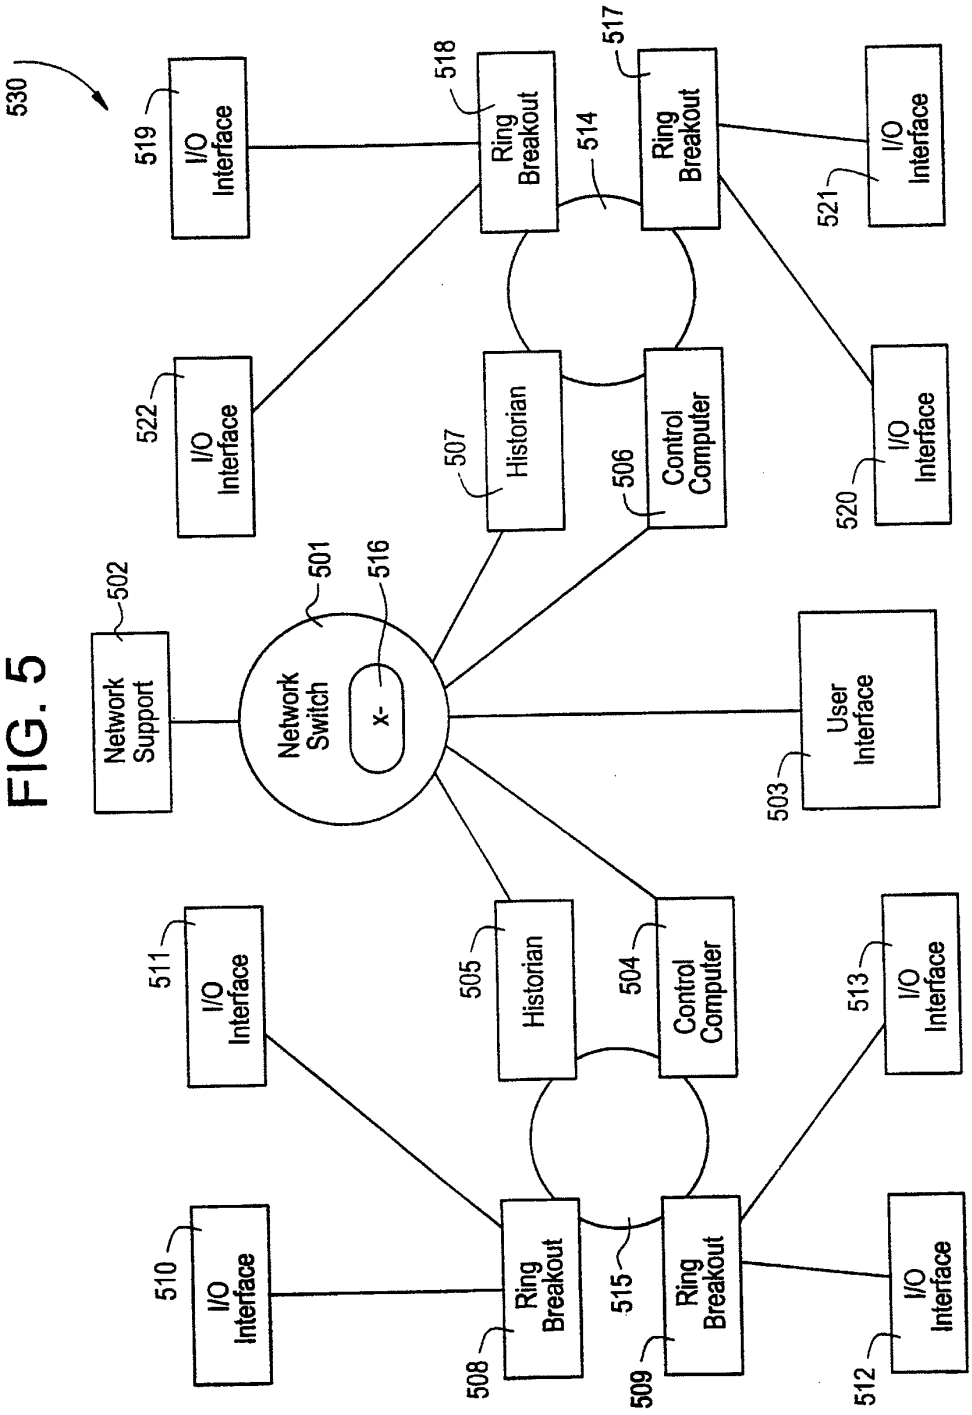

530
519
I/O Interface
518
Ring Breakout
517
Ring Breakout
514
I/O Interface
521
522
I/O Interface
507
Historian
506
Control Computer
I/O Interface
520
502
Network Support
501
Network Switch
516
x-
User Interface
503
511
I/O Interface
505
Historian
504
Control Computer
513
I/O Interface
510
I/O Interface
508
Ring Breakout
515
Ring Breakout
509
I/O Interface
512

FIG. 6

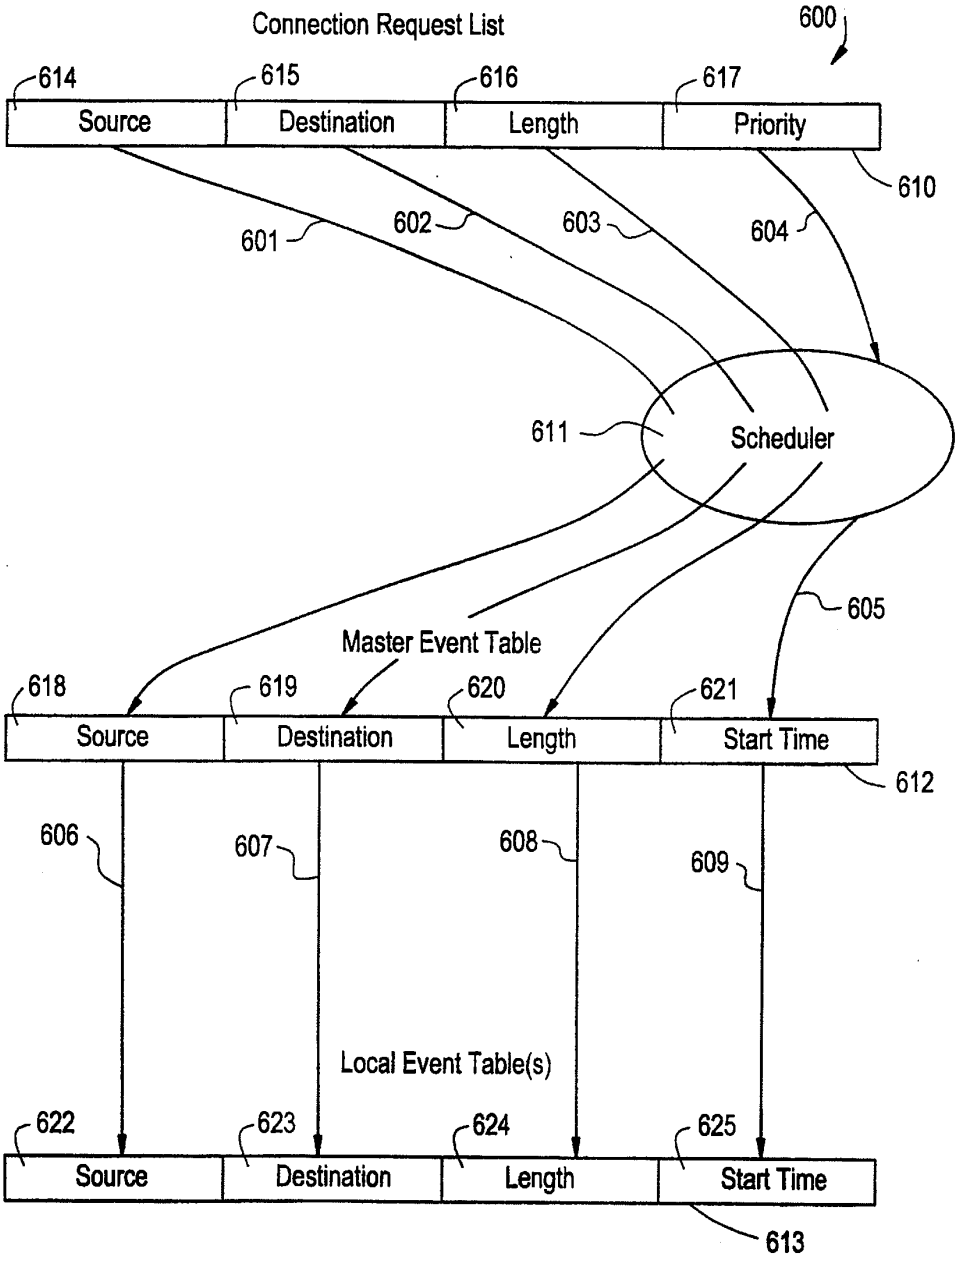

Connection Request List
600
614
615
616
617
Source
Destination
Length
Priority
610
601
602
603
604
611
Scheduler
605
Master Event Table
618
619
620
621
Source
Destination
Length
Start Time
612
606
607
608
609
Local Event Table(s)
622
623
624
625
Source
Destination
Length
Start Time
613

FIG. 7

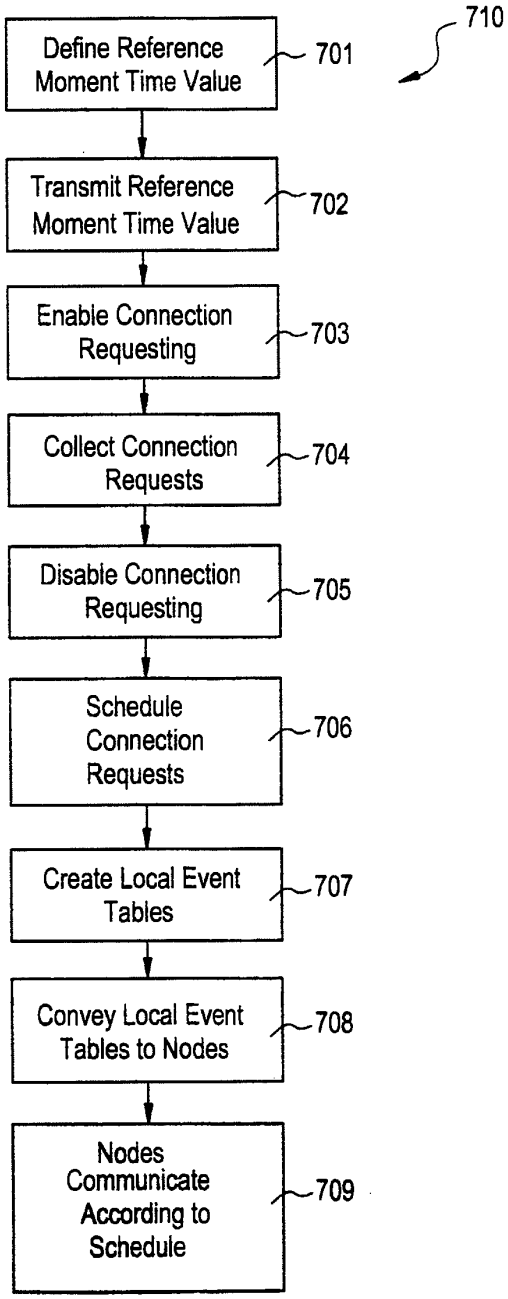

710
Define Reference Moment Time Value
701
Transmit Reference Moment Time Value
702
Enable Connection Requesting
703
Collect Connection Requests
704
Disable Connection Requesting
705
Schedule Connection Requests
706
Create Local Event Tables
707
Convey Local Event Tables to Nodes
708
Nodes Communicate According to Schedule
709

FIG. 8

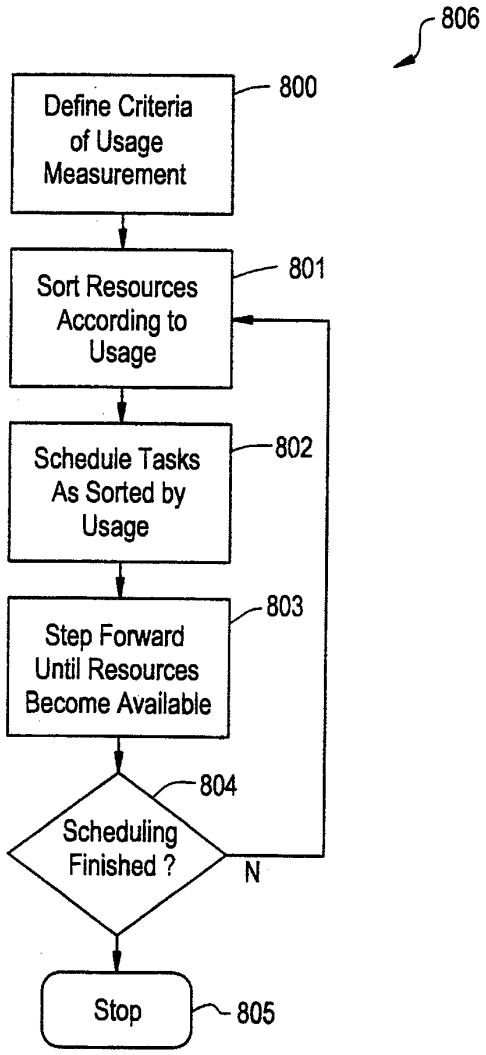

806
Define Criteria of Usage Measurement
800
Sort Resources According to Usage
801
Schedule Tasks As Sorted by Usage
802
Step Forward Until Resources Become Available
803
Scheduling Finished ?
804
N
Stop
805

FIG. 9

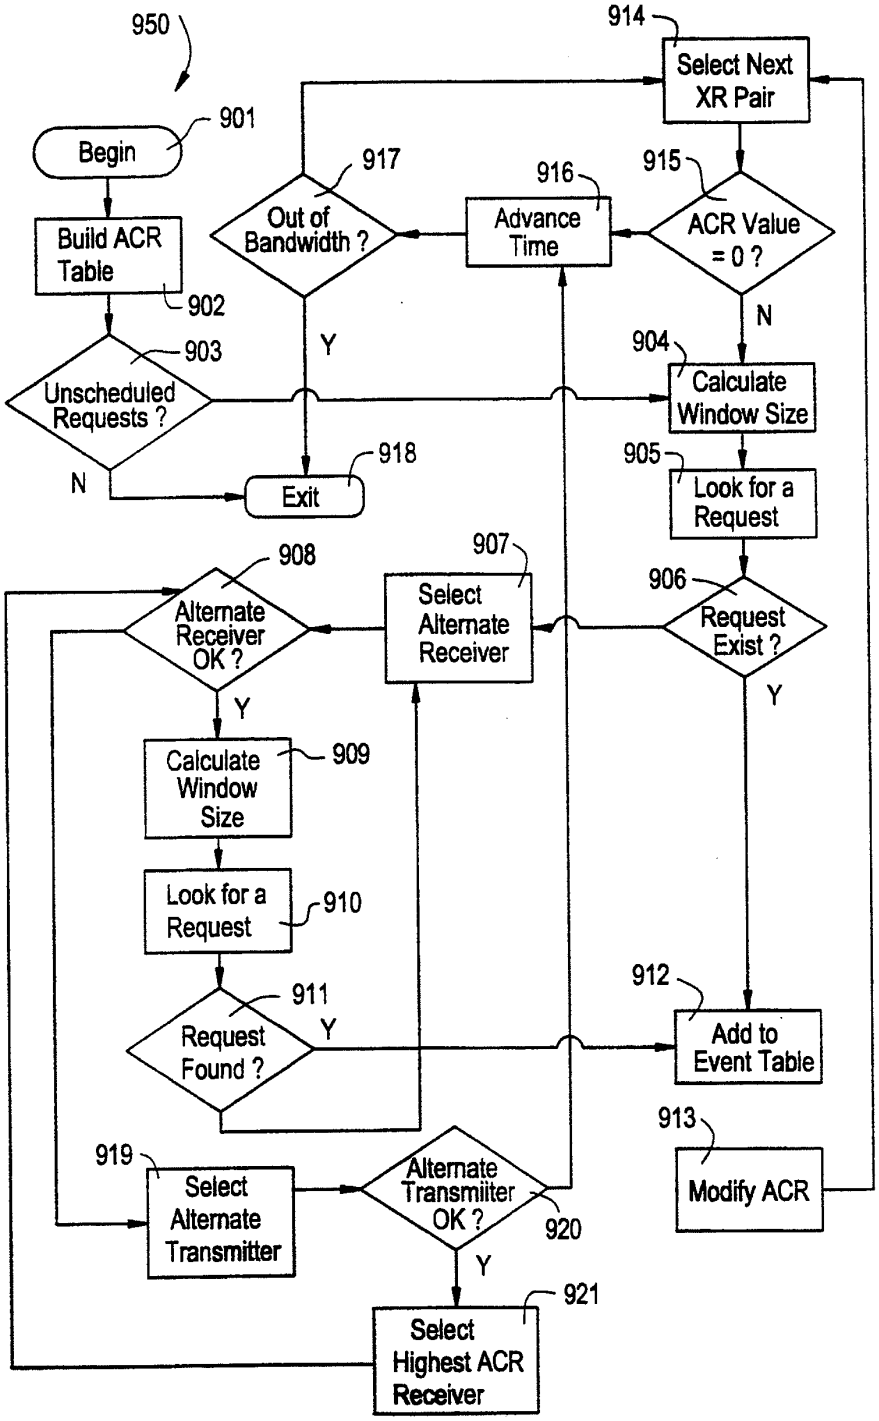

950
914
Select Next XR Pair
901
Begin
917
Out of Bandwidth ?
916
Advance Time
915
ACR Value = 0 ?
Build ACR Table
902
903
Unscheduled Requests ?
Y
N
904
Calculate Window Size
905
Look for a Request
918
Exit
N
908
907
Alternate Receiver OK ?
Select Alternate Receiver
906
Request Exist ?
Y
Y
Calculate Window Size
909
Look for a Request
910
911
Request Found ?
Y
912
Add to Event Table
913
Modify ACR
919
Select Alternate Transmitter
Alternate Transmiter OK ?
920
Y
921
Select Highest ACR Receiver

FIG. 10

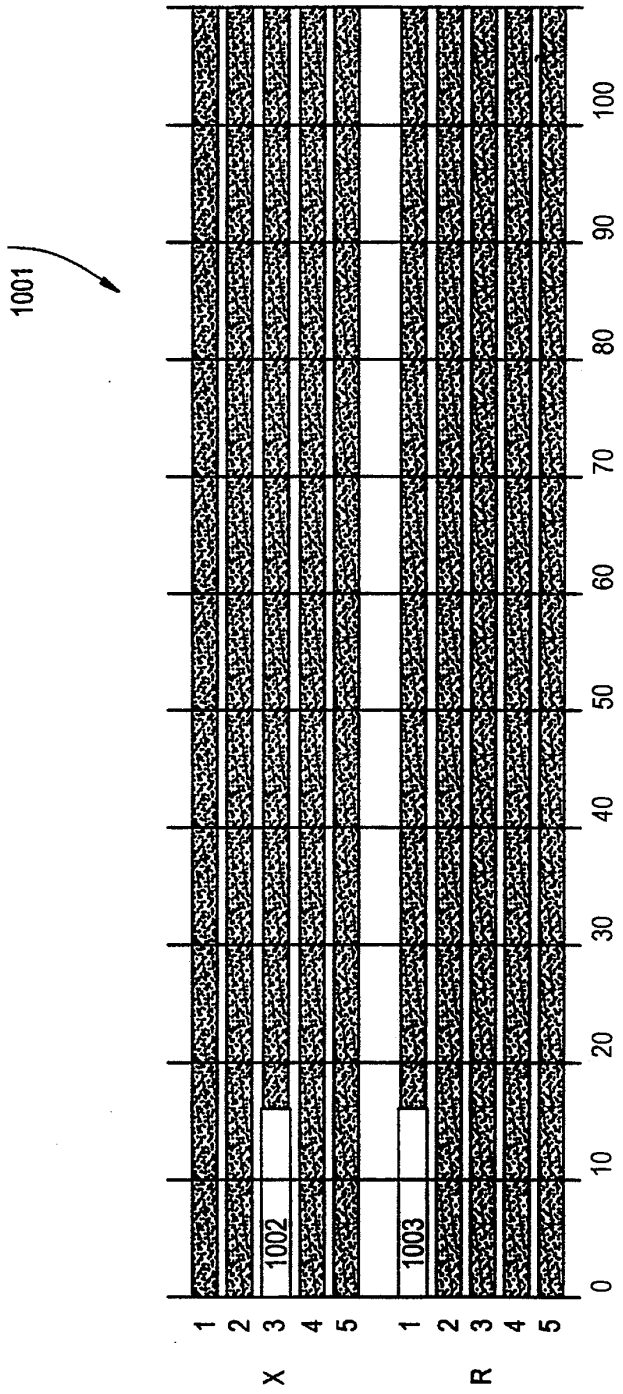


FIG. 11

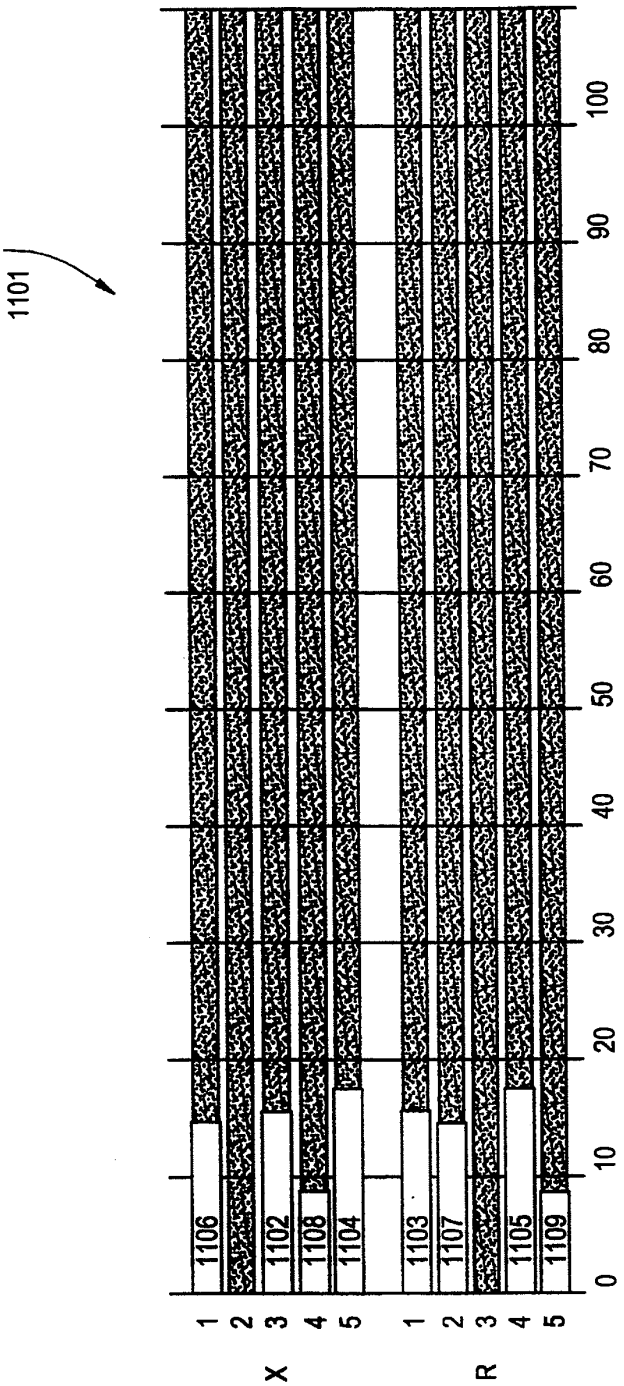

1101
1106
1102
1108
1104
1103
1107
1105
1109
0
10
20
30
40
50
60
70
80
90
100
1
2
3
4
5
X
1
2
3
4
5
R

FIG. 12

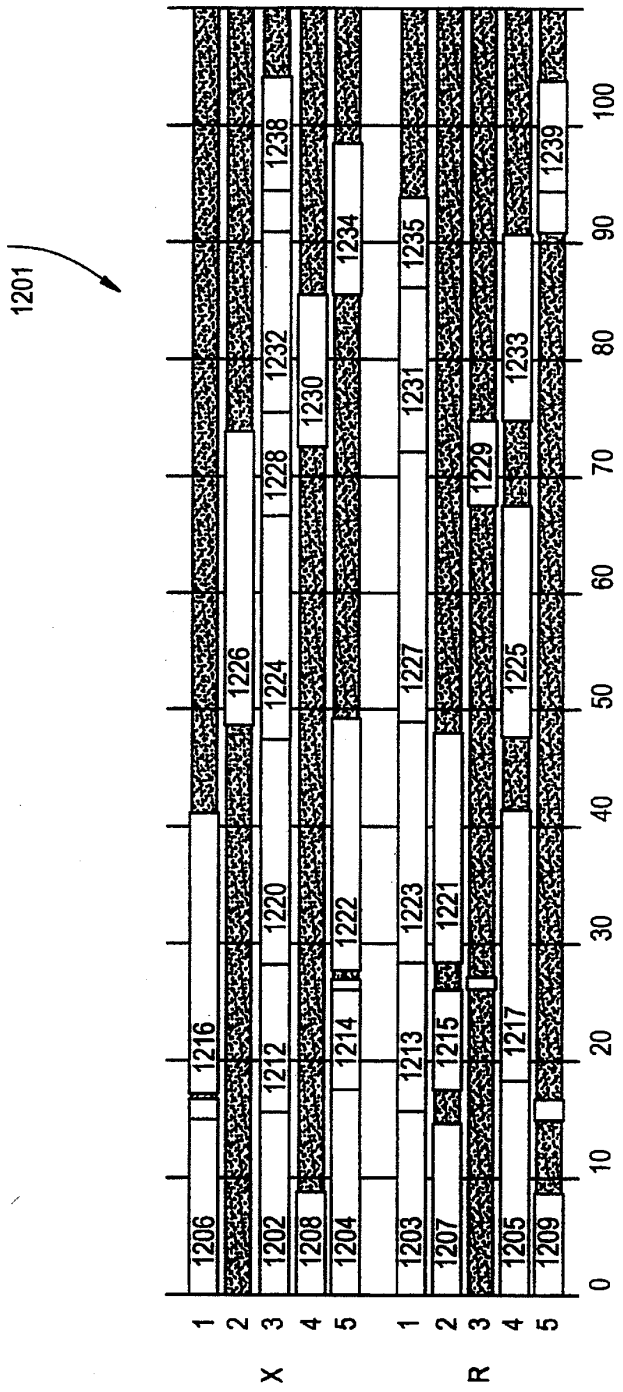

1201
1206
1216
1202
1212
1208
1220
1204
1214
1222
1226
1224
1228
1232
1230
1234
1238
1203
1213
1223
1227
1207
1215
1221
1231
1205
1217
1225
1229
1233
1209
1235
1239
0
10
20
30
40
50
60
70
80
90
100
1 2 3 4 5
X
1 2 3 4 5
R

FIG. 13

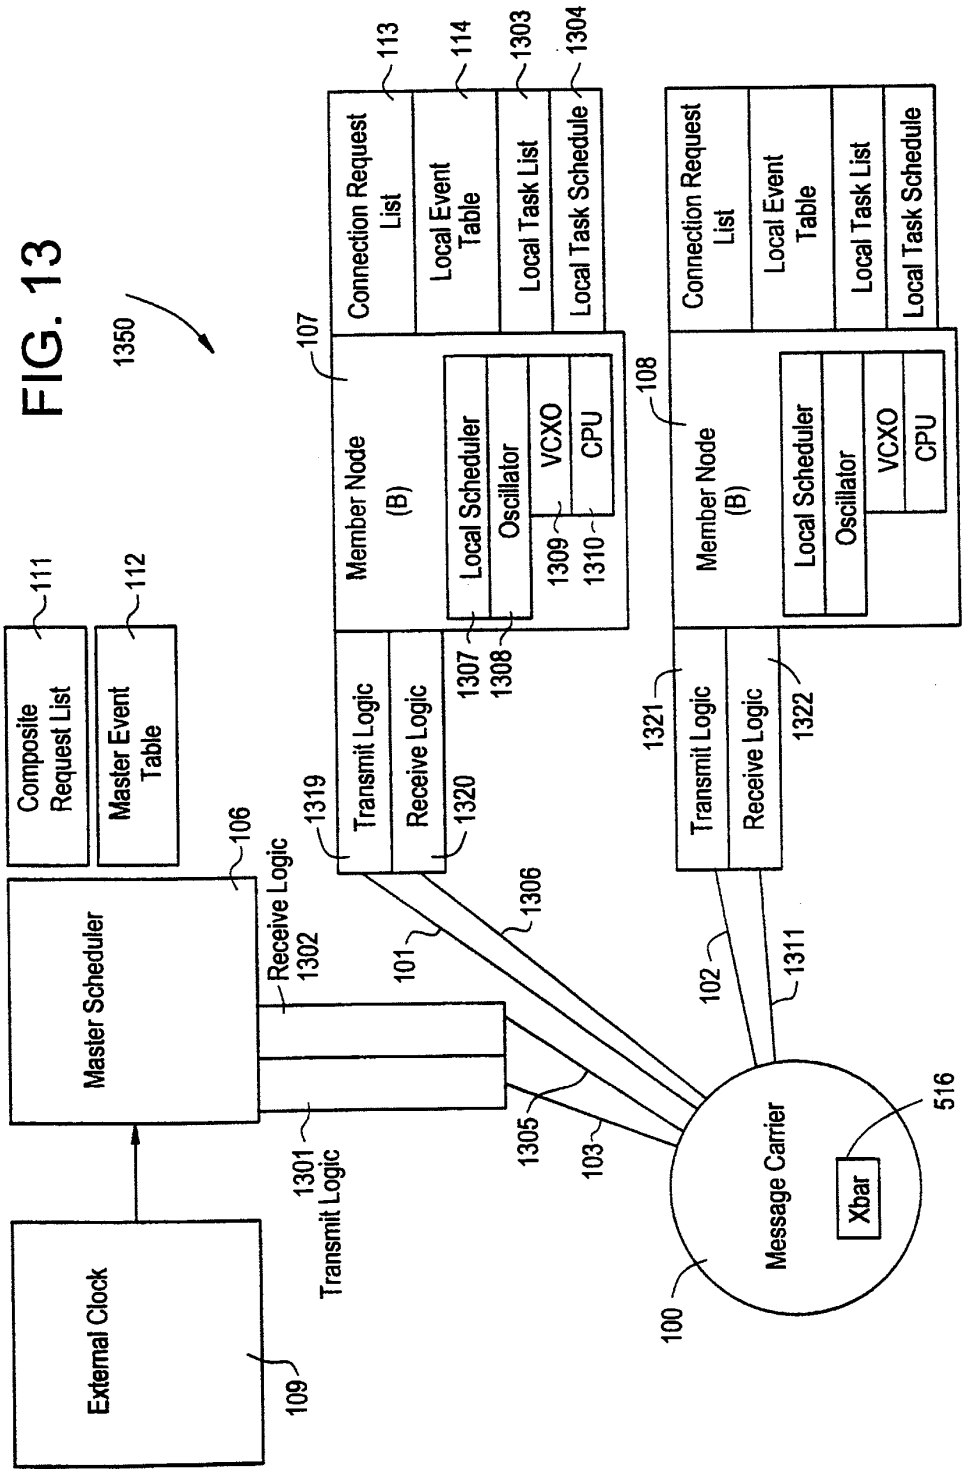
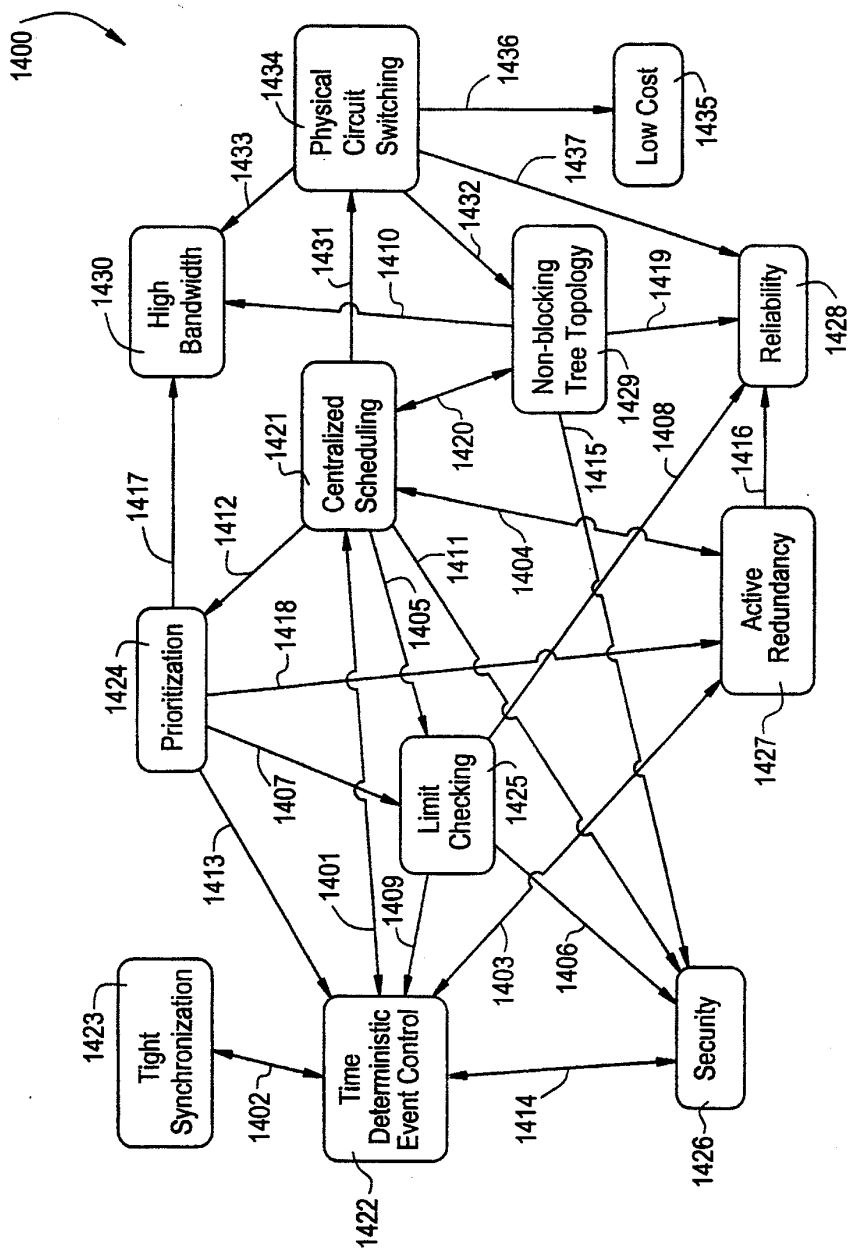

1350
External Clock
109
Master Scheduler
106
Composite Request List
111
Master Event Table
112
Receive Logic
1302
Transmit Logic
1301
101
1306
1305
103
Member Node (B)
107
Connection Request List
113
Local Event Table
114
Local Task List
1303
Local Task Schedule
1304
Local Scheduler
Oscillator
VCXO
CPU
1309
1310
1307
1308
Transmit Logic
1319
Receive Logic
1320
Member Node (B)
108
Connection Request List
Local Event Table
Local Task List
Local Task Schedule
Local Scheduler
Oscillator
VCXO
CPU
1321
Transmit Logic
Receive Logic
1322
102
1311
Message Carrier
Xbar
516
100

FIG. 14


1400
Physical Circuit Switching
1434
1436
Low Cost
1435
1437
1433
1432
High Bandwidth
1430
1431
1410
Non-blocking Tree Topology
1419
Reliability
1428
1429
Centralized Scheduling
1421
1420
1408
1415
1416
1417
1412
1411
1404
1405
Active Redundancy
Prioritization
1424
1418
1427
1407
Limit Checking
1425
1413
1401
1409
1403
1406
Tight Synchronization
1423
Time Deterministic Event Control
Security
1402
1414
1426
1422

15 / 16

FIG. 15

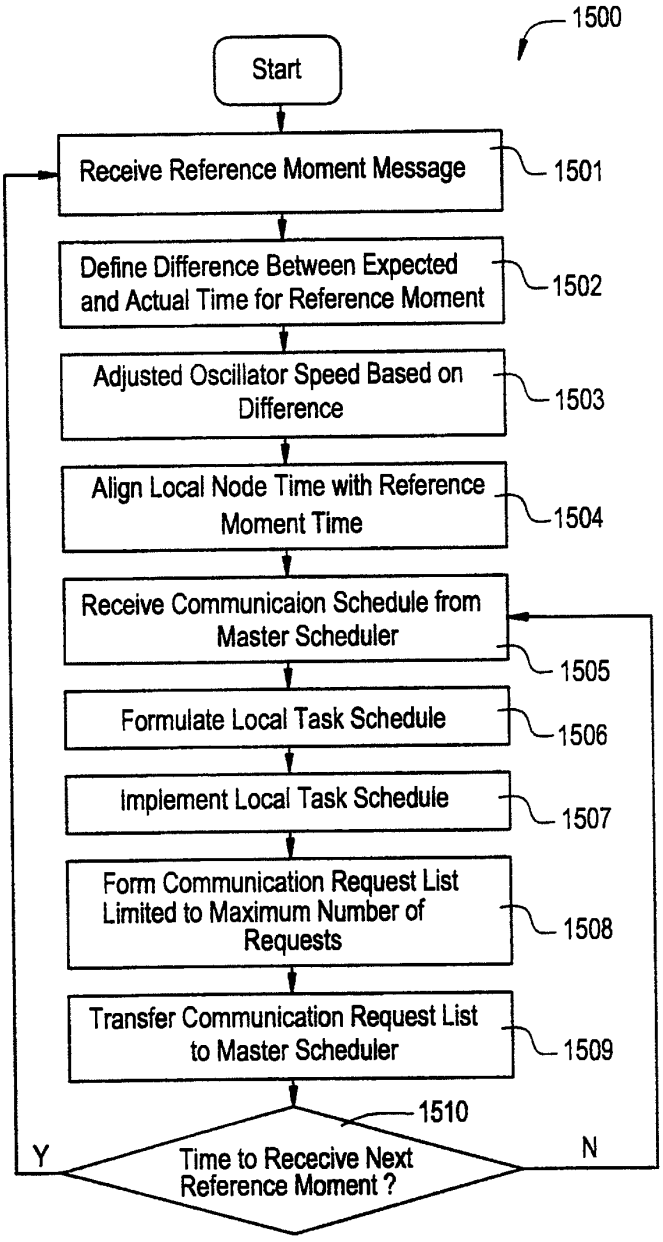
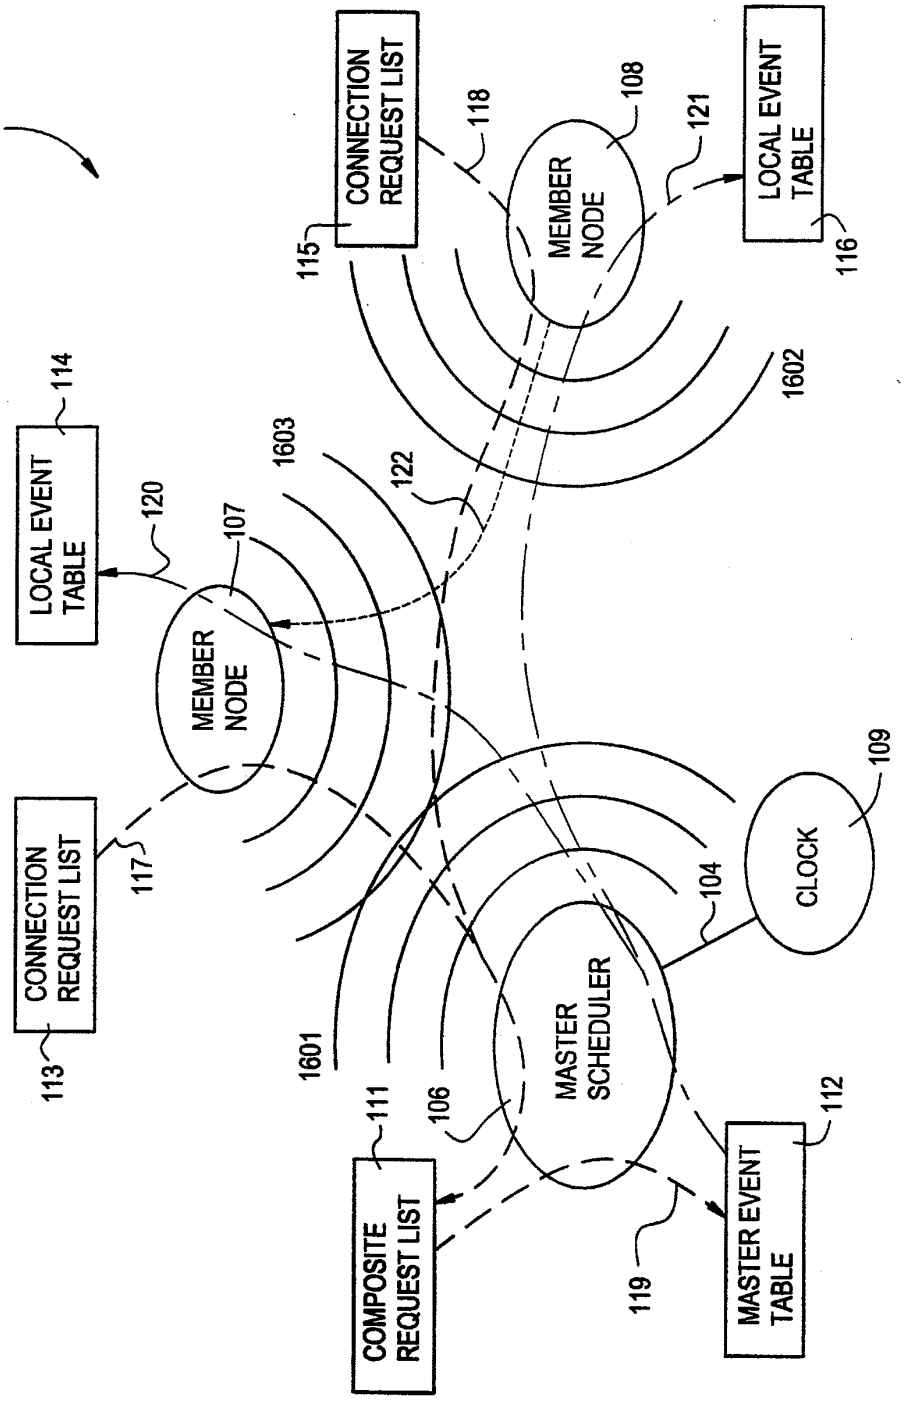

1500
Start
Receive Reference Moment Message
1501
Define Difference Between Expected and Actual Time for Reference Moment
1502
Adjusted Oscillator Speed Based on Difference
1503
Align Local Node Time with Reference Moment Time
1504
Receive Communicaion Schedule from Master Scheduler
1505
Formulate Local Task Schedule
1506
Implement Local Task Schedule
1507
Form Communication Request List Limited to Maximum Number of Requests
1508
Transfer Communication Request List to Master Scheduler
1509
1510
Time to Rececive Next Reference Moment ?
Y
N

FIG. 16


1600
CONNECTION REQUEST LIST
118
115
MEMBER NODE
108
121
LOCAL EVENT TABLE
116
1602
1603
122
LOCAL EVENT TABLE
114
120
107
MEMBER NODE
CONNECTION REQUEST LIST
117
113
1601
CLOCK
109
104
MASTER SCHEDULER
106
111
COMPOSITE REQUEST LIST
112
119
MASTER EVENT TABLE

INTERNATIONAL SEARCH REPORT

Int[illegible]ional Application No
PCT/US 99/06634

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate. of the relevant passages	Relevant to claim No.
A	SHIN K G ET AL: "REAL-TIME COMPUTING: A NEW DISCIPLINE OF COMPUTER SCIENCE AND ENGINEERING" PROCEEDINGS OF THE IEEE, vol. 82, no. 1, 1 January 1994 (1994-01-01), pages 6-23, XP000435878 ISSN: 0018-9219 page 8, paragraph II - page 17, paragraph VII ---	1-4,10, 11,14-16
A	EP 0 183 877 A (INMOS LTD) 11 June 1986 (1986-06-11) the whole document ---	1-4,10, 11,14-16
A	EP 0 250 011 A (INMOS LTD) 23 December 1987 (1987-12-23) the whole document ---	1-4,10, 11,14-16
	-/--	

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
27 July 1999	03/08/1999

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Michel, T

1

Form PCT/ISA/210 (second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Int ional Application No
PCT/US 99/06634

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication,where appropriate, of the relevant passages	Relevant to claim No.
A	RAMAKRISHNAN K V: "AN EVENT SCHEDULER FOR CONTROL APPLICATIONS" EDN ELECTRICAL DESIGN NEWS, vol. 41, no. 12, 6 June 1996 (1996-06-06), pages 155-158, 160, 16, XP000622016 ISSN: 0012-7515 the whole document	9-11

1

Form PCT/ISA/210 (continuation of second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

.nformation on patent family members

Int ional Application No
PCT/US 99/06634

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0183877 A	11-06-1986	NONE	
EP 0250011 A	23-12-1987	DE 3375844 A	07-04-1988
		DE 3376894 A	07-07-1988
		DE 3378083 A	27-10-1988
		DE 3382708 A	16-09-1993
		DE 3382708 T	10-02-1994
		EP 0111399 A	20-06-1984
		EP 0113178 A	11-07-1984
		EP 0110642 A	13-06-1984
		EP 0113516 A	18-07-1984
		EP 0110643 A	13-06-1984
		JP 1028967 B	07-06-1989
		JP 1547176 C	28-02-1990
		JP 59149543 A	27-08-1984
		JP 1789597 C	29-09-1993
		JP 4070652 B	11-11-1992
		JP 59133647 A	01-08-1984
		JP 1982715 C	25-10-1995
		JP 6101042 B	12-12-1994
		JP 59194256 A	05-11-1984
		JP 2111325 C	21-11-1996
		JP 8007739 B	29-01-1996
		JP 59133662 A	01-08-1984
		JP 2502960 B	29-05-1996
		JP 59133650 A	01-08-1984
		JP 2509453 B	19-06-1996
		JP 6342378 A	13-12-1994
		JP 2688320 B	10-12-1997
		JP 6260551 A	16-09-1994
		US 5452467 A	19-09-1995
		US 5506437 A	09-04-1996
		US 5491359 A	13-02-1996
		US 5031092 A	09-07-1991
		US 4680698 A	14-07-1987
		US 4704678 A	03-11-1987
		US 4724517 A	09-02-1988
		US 5243698 A	07-09-1993
		US 4819151 A	04-04-1989
		US 4967326 A	30-10-1990

Form PCT/ISA/210 (patent family annex) (July 1992)

Electronic Acknowledgement Receipt

EFS ID:	15720605
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A. J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 (13-0401)
Receipt Date:	08-MAY-2013
Filing Date:	
Time Stamp:	18:45:14
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	71301-US10-IDS.pdf	712398 d2e990aff3760e55cf683abec7b4f7f448a890a9	no	19

Warnings:	
Information:	

2	Foreign Reference	EP0251584A2-NorthernTelecom.pdf	1871427 e4060935063351b4642542a8f21ee0f379f46cdc	no	33
Warnings:					
Information:					
3	Foreign Reference	EP0672985A1-IBM.pdf	1715502 82829d8bae9a69fe31d924894be437476c07d679	no	31
Warnings:					
Information:					
4	Foreign Reference	EP1111527B1.pdf	2175253 54fb66a34a8c2836746d8102e61cd137b0659bdf	no	38
Warnings:					
Information:					
5	Foreign Reference	WO9525313.pdf	707368 11f89c485fe088f060c29d4ea21245e5a843160e	no	25
Warnings:					
Information:					
6	Foreign Reference	WO9961985.pdf	3653814 6c32d809f0d15c1a00bce9a3f87c4a0b7b334fd1	no	96
Warnings:					
Information:					
7	Non Patent Literature	AKYILDIZ-MultimediaGroup.pdf	2498395 69554162fcf4e51427d9c319e879a433693ddbfc	no	12
Warnings:					
Information:					
8	Non Patent Literature	Benslimane-MultimediaSynchronization.pdf	1209082 2ff64e8ed450d8cb5de0108808864985ed162ae5	no	8
Warnings:					
Information:					
9	Non Patent Literature	Intra-InterStreamSync-Biersack.pdf	731857 35794d83dd743d327e07825f217b68829a97479f	no	10
Warnings:					
Information:					
10	Non Patent Literature	MPEG2Tutorial-Bretl.pdf	1294244 911b6ab9efcf6e866f994de6b6a821bd9bd1a117	no	23
Warnings:					
Information:					

11	Non Patent Literature	Huang-SyncInfrastructure.pdf	2480357 b781148c0e2b6051001ce7913f481fbb079ed66d	no	11
Warnings:					
Information:					
12	Non Patent Literature	PCTUS04-023102-IntlSearchReport.pdf	239507 63ef304c9af50597a27e0d476723a348a893e612	no	6
Warnings:					
Information:					
13	Non Patent Literature	Ishibashi-GroupSync-LiveMedia.pdf	1692122 e5e93fef756c319896baa9cf95b6dffa1338616a	no	7
Warnings:					
Information:					
14	Non Patent Literature	Ishibashi-GroupSync-StoredMedia.pdf	1672500 08c674c017aee8409bf0d6327e7e0e3d23780544	no	9
Warnings:					
Information:					
15	Non Patent Literature	Joreference-synchronizedonetomany.pdf	542871 f6969e60ce4a5f4cd6d00b0627023766d6b5eafe	no	11
Warnings:					
Information:					
16	Non Patent Literature	Mills-NetworkTimeProtocol.pdf	6081281 f6207bf3fdc9206c3d5571514be05927e7a684ba	no	96
Warnings:					
Information:					
17	Non Patent Literature	Mills-PrecisionClockSync.pdf	1366650 53928c660ae7df7414888151a0eabdd85db75ca9	no	16
Warnings:					
Information:					
18	Non Patent Literature	Motorola2001SimplefiManual.pdf	945125 163341eb0c6d16fff34c3ee075d3e8b5afc30a61	no	111
Warnings:					
Information:					
19	Non Patent Literature	Park-GroupSyncMechanism.pdf	354862 0b1939edff956eb25c657d70adb9bb60df2b2ce5	no	5
Warnings:					
Information:					

20	Non Patent Literature	Prismiq-UserGuide.pdf	647554 ac49597896459778a52c6b4a843a0d3bfb81405b	no	44
Warnings:					
Information:					
21	Non Patent Literature	Rothermel-AdaptiveStreamSync.pdf	1009707 7fa20df4279b93306158108cc4100a36a685568b	no	22
Warnings:					
Information:					
22	Non Patent Literature	Schulzrinne-RTP-TransportProtocol.pdf	4486868 045bf0b66a252ccb18f334231c9ebe42f165c80b	no	89
Warnings:					
Information:					
23	Non Patent Literature	MPEG2Transport.pdf	317893 fdd15d79a37834e2d53ee663fcae10b090f507ce	no	5
Warnings:					
Information:					
24	Non Patent Literature	UPnP-reprint.pdf	5649097 c30d43df6be46fbed6e7f0168f57fdae1e66ed40	no	54
Warnings:					
Information:					
25	Non Patent Literature	AudioTronQuickStartGuide2.pdf	2102621 4b12d08d4da7e9d34220dd6edb49f5bb632eac17	no	24
Warnings:					
Information:					
26	Non Patent Literature	AudiotronReferenceManual-version3-70p.pdf	6675703 8996836378ac6cc677323bd9a88105b716386ff5	no	70
Warnings:					
Information:					
27	Non Patent Literature	Audiotronsetupguide-version1.pdf	4645843 070ec3b852e417efee9fc22b8842cf8a2e47d24c	no	38
Warnings:					
Information:					
28	Non Patent Literature	10816217-finalOA-13JUL2009.pdf	494964 fdb38a5dd85ff71a57798e8efe5c72ab2152327d	no	16
Warnings:					
Information:					

29	Non Patent Literature	10816217-finalOA-21OCT2011.pdf	640324 07f58dff07fb2e273c0f176593c25b69252c8d91	no	19
Warnings:					
Information:					
30	Non Patent Literature	10816217-finalOA-28JAN2011.pdf	672280 e87a3a10b4b4db76617f43bf151de7f860d9c66a	no	21
Warnings:					
Information:					
31	Non Patent Literature	10816217-finalOA-30JUN2008.pdf	838638 dd0e8f7fd8add62c9ccb2db62a82bc8439c9a8fc	no	30
Warnings:					
Information:					
32	Non Patent Literature	10816217-nonfinalOA-18JAN2008.pdf	1135142 59ce0684e9c6d36094203bf1817b88cc48971613	no	38
Warnings:					
Information:					
33	Non Patent Literature	10816217-nonfinalOA-21JUN2011.pdf	419738 c1ee6116de2b310276450c61abe974cfe7a54aee	no	13
Warnings:					
Information:					
34	Non Patent Literature	10816217-nonfinalOA-22JAN2009.pdf	577911 949b7a488f2fd167d23950d541cfd319f52d73d3	no	18
Warnings:					
Information:					
35	Non Patent Literature	10816217-nonfinalOA-25JUN2010.pdf	539926 6b0333e9cc2573392dad5b0110c331f45697660c	no	17
Warnings:					
Information:					
36	Non Patent Literature	10816217-NOA-27DEC2011.pdf	691357 16a70eceb799f369f8830e28a005b379434d801f	no	16
Warnings:					
Information:					
37	Non Patent Literature	13297000-nonfinalOA-29FEB2012.pdf	308664 a11567d38b79990e922273a4a24768f28ad12e7a	no	10
Warnings:					
Information:					

38	Non Patent Literature	13297000-FinalOA-13SEP2012.pdf	515124 304e3353e85741d7c9ed90b908213e6fc2166895	no	16

Warnings:	
Information:	
Total Files Size (in bytes):	64313969

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Nicholas A.J. Millington

Title: SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES

Appl. No.: TBD

Filing Date: 04/16/2013

Examiner: TBD

Art Unit: TBD

Confirmation Number:

PRELIMINARY AMENDMENT

Sir:

Please enter the following amendments prior to examination of the application.

Amendments to the Specification are presented on page 2;

Amendments to the Abstract are presented on page 3;

Amendments to the Claims are reflected in the listing of claims that begins on page 4 of this document.

Remarks/Arguments begin on page 8 of this document.

Amendments to the Specification:

On page 1 of the specification, please replace the first paragraph after the section heading "Incorporation By Reference" with the new paragraph as follows:

This application claims priority under 35 USC §120 of US Application No. 13/297,000, filed November 15, 2011, entitled "System And Method For Synchronizing Operations Among A Plurality Of Independently Clocked Digital Data Processing Devices", which is a continuation of Patent Application Serial No. 10/816,217, now US Patent 8,234,395, filed on April 1, 2004, entitled "System and Method for Synchronizing Channel Handoff as Among a Plurality of Independently Clocked Digital Devices", which is a non-provisional application claiming priority under 35 USC §119 to Provisional Patent Application Serial No. 60/490,768, filed on July 28, 2003, entitled "Method for Synchronizing Audio Playback Between Multiple Networked Devices," now expired, all of which are assigned to the assignee of the present application and are incorporated herein by reference.

Please amend the application as follows:

Amendments to the Abstract:

Please replace the Abstract with the new abstract as follows:

A plurality of zone players, including at least a first player and a second zone player, wherein the first zone player is designated as a master device of a synchrony group, a method comprising: determining, by the first zone player, that the second zone player should be designated as the master device of the synchrony group; instructing, by the first zone player, the second zone player to be designated as the master device of the synchrony group; and causing, at the first zone player, the first zone player to lose its designation as the master device of the synchrony group.

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1-20 (cancelled)

21. (new) In a network comprising a plurality of zone players, including at least a first player and a second zone player, wherein the first zone player is designated as a master device of a synchrony group, a method comprising:

determining, by the first zone player, that the second zone player should be designated as the master device of the synchrony group;

instructing, by the first zone player, the second zone player to be designated as the master device of the synchrony group; and

causing, at the first zone player, the first zone player to lose its designation as the master device of the synchrony group.

22. (new) The method of claim 21, wherein the first zone player determining that the second zone player should be designated as the master device of the synchrony group comprises the first zone player determining that it is overloaded.

23. (new) The method of claim 21, wherein the first zone player determining that the second zone player should be designated as the master device of the synchrony group comprises the first zone player receiving, from a user interface module, control information, wherein the control information indicates that the first zone player should disengage from the synchrony group.

24. (new) The method of claim 23, further comprising:

disengaging, by the first zone player, the first zone player from the synchrony group.

25. (new) The method of claim 21, further comprising:

transmitting, by the first zone player to the second zone player while the first zone player is designated as the master device of the synchrony group, a network address of an audio information channelization device, a multicast address that the audio information

channelization devices uses to broadcast audio information and playback timing information over the network.

26. (new) The method of claim 25, wherein during the transmitting, the second zone player is a slave device of the synchrony group.

27. (new) The method of claim 21, further comprising:

receiving, at the first zone player while the first zone player is designated as the master device of the synchrony group, control information from a user interface module, wherein the control information indicates that a third zone player should be added to the synchrony group; and

responsively transmitting, from the first zone player to the third zone player, slave control information, wherein the slave control information enables the third zone player to join the synchrony group as a slave device of the synchrony group.

28. (new) A tangible computer-readable memory having instructions stored thereon that when executed cause a first zone player that is designated as a master of a synchrony group to:

determine that a second zone player should be designated as the master device of the synchrony group;

instruct the second zone player to be designated as the master device of the synchrony group; and

cause the first zone player to lose its designation as the master device of the synchrony group.

29. (new) The computer-readable medium of claim 28, wherein the instructions that when executed cause the first zone player to determine that the second zone player should be designated as the master device of the synchrony group comprise instructions that when executed cause the first zone player to determine that it is overloaded.

30. (new) The computer-readable medium of claim 28, wherein the instructions that when executed cause the first zone player to determine that the second zone player should be designated as the master device of the synchrony group comprise instructions that when executed cause the first zone player to receive, from a user interface module, control information, wherein the control information indicates that the first zone player should disengage from the synchrony group.

31. (new) The computer-readable medium of claim 30, further comprising instructions that cause the first zone player to disengage from the synchrony group.

32. (new) The computer-readable medium of claim 28, further comprising instructions that when executed cause the first zone player, when it is designated as the master device of the synchrony group, to provide to the second zone player a network address of an audio information channelization device, and a multicast address that the audio information channelization devices uses to broadcast audio information and playback timing information over the network.

33. (new) The computer-readable medium of claim 32, wherein the second zone player is designated as a slave device of the synchrony group.

34. (new) The computer-readable medium of claim 28, further comprising instructions that when executed cause the first zone player to:

receive, while the first zone player is designated as the master device of the synchrony group, control information from a user interface module, wherein the control information indicates that a third zone player should be added to the synchrony group; and

responsively transmit, from the first zone player to the third zone player, slave control information, wherein the slave control information enables the third zone player to join the synchrony group as a slave device of the synchrony group.

35. (new) A system comprising:

a plurality of zone players, wherein (i) at least two of the plurality of zone players are members a synchrony group, (ii) one of the plurality of zone players is designated as a master device of the synchrony group, and (iii) at least one of the plurality of zone players is designated as a slave device of the synchrony group; and

a user interface module, wherein the user interface module is configured to transmit control information to the zone player currently designated as the master device of the synchrony group, and wherein the control information causes the zone player currently designated as the master device of the synchrony group to instruct another member of the synchrony group currently designated as a slave device to become designated as the master device of the synchrony group.

36. (new) The system of claim 35, wherein the control information further causes the zone player currently designated as the master device of the synchrony group to lose its status as the master device of the synchrony group.

37. (new) The system of claim 35, wherein the control information further causes the zone player currently designated as the master device of the synchrony group to disengage from the synchrony group.

38. (new) The system of claim 35, wherein the control information is transmitted via a network interface.

39. (new) The system of claim 35, wherein the zone player that is designated as the master device is configured to transmit to the members of the synchrony group a network address of an audio information channelization device.

40. (new) The system of claim 39, wherein the zone player that is designated as the master device is configured to transmit to the members of the synchrony group a multicast address that the audio information channelization devices uses to broadcast audio information and playback timing information over the network.

REMARKS

Applicants request that the above-presented amendments are entered prior to examination. The Examiner is invited to contact the undersigned if it would be helpful to advance the prosecution of the application.

Respectfully submitted,

Date: April 16, 2013

By: /Robert J. Irvine III/

TechLaw LLP
Telephone: (312) 834-4334

Robert J. Irvine III
Attorney for Applicant
Registration No. 41,865

Electronic Patent Application Fee Transmittal	
Application Number:	
Filing Date:	
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A.J. Millington
Filer:	Robert Johnson Irvine
Attorney Docket Number:	SNS-71301.US.10 / 13-0401
Filed as Large Entity	

Utility under 35 USC 111(a) Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Utility application filing	1011	1	280	280
Utility Search Fee	1111	1	600	600
Utility Examination Fee	1311	1	720	720
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**1600**

Electronic Acknowledgement Receipt

EFS ID:	15536226
Application Number:	13864248
International Application Number:	
Confirmation Number:	5095
Title of Invention:	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES
First Named Inventor/Applicant Name:	Nicholas A.J. Millington
Customer Number:	109473
Filer:	Robert Johnson Irvine
Filer Authorized By:	
Attorney Docket Number:	SNS-71301.US.10 / 13-0401
Receipt Date:	17-APR-2013
Filing Date:	
Time Stamp:	02:00:02
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Credit Card
Payment was successfully received in RAM	$1600
RAM confirmation Number	6994
Deposit Account	
Authorized User	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)

1	Application Data Sheet	ADS-71301_US_10.pdf	1068619 6977cd417df995a56db8c4c94f9b1f3089b594e8	no	6
Warnings:					
Information:					
2		71301_US-Specification.pdf	3285365 f36f9701ab8cd3989b81f0a21d17040f52bc7bca	yes	67

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Specification	1	58
	Claims	59	61
	Abstract	62	62
	Drawings-only black and white line drawings	63	67

Warnings:					
Information:					
3	Oath or Declaration filed	13297000-Dec.pdf	52966 b421f9358bd38188761d3667976f972dcfdfa526	no	2
Warnings:					
Information:					
4		71301US10_PA.pdf	94687 a7f8e6c7a8ceac2bc1f52d625830d137dbba3867	yes	8

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Preliminary Amendment	1	1
	Specification	2	2
	Abstract	3	3
	Claims	4	7
	Applicant Arguments/Remarks Made in an Amendment	8	8

Warnings:	
Information:	

5	Fee Worksheet (SB06)	fee-info.pdf	33255 6bc86d0efdde42fb2a18a0686f5ecce98d74eea3	no	2
Warnings:					
Information:					
			Total Files Size (in bytes):	4534892	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/AIA/14 (08-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	SNS-71301.US.10
	Application Number	
Title of Invention	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
Mr.	Nicholas	A. J.	Millington	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Santa Barbara	State/Province	CA	Country of Residence i	US

Mailing Address of Inventor:

Address 1	223 E. De La Guerra		
Address 2			
City	Santa Barbara	State/Province	CA
Postal Code	93101	Country i	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

☐ **An Address is being provided for the correspondence Information of this application.**

Customer Number	109473		
Email Address	bob@techlawllp.com	Add Email	Remove Email

Application Information:

Title of the Invention	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES		
Attorney Docket Number	SNS-71301.US.10	Small Entity Status Claimed	☐
Application Type	Nonprovisional		
Subject Matter	Utility		
Suggested Class (if any)		Sub Class (if any)	
Suggested Technology Center (if any)			
Total Number of Drawing Sheets (if any)	5	Suggested Figure for Publication (if any)	

PTO/AIA/14 (08-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	SNS-71301.US.10
	Application Number	
Title of Invention	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES	

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)
☐ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	109473		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

Prior Application Status	Pending		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
	Continuation of	13297000	2011-11-15
Prior Application Status	Pending		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
13297000	Continuation of	10816217	2004-04-01
Prior Application Status	Patented		Remove

Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)	Patent Number	Issue Date (YYYY-MM-DD)
10816217	non provisional of	60490768	2003-07-28	8234395	2012-07-31

Prior Application Status	Expired		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
60490768			

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button. Add

Foreign Priority Information:

PTO/AIA/14 (08-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	SNS-71301.US.10
	Application Number	
Title of Invention	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES	

This section allows for the applicant to claim benefit of foreign priority and to identify any prior foreign application for which priority is not claimed. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(a).

Remove

Application Number	Country i	Filing Date (YYYY-MM-DD)	Priority Claimed
			○ Yes ◉ No

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button. Add

Authorization to Permit Access:

☐ Authorization to Permit Access to the Instant Application by the Participating Offices

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date of filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

PTO/AIA/14 (08-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	SNS-71301.US.10
	Application Number	
Title of Invention	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES	

Applicant 1

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Remove

◉ Assignee	○ Legal Representative under 35 U.S.C. 117
○ Person to whom the inventor is obligated to assign.	○ Person who shows sufficient proprietary interest

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :

If the Assignee is an Organization check here. ☒

Organization Name	SONOS, INC.

Mailing Address Information:

Address 1	223 E De La Guerra St.		
Address 2			
City	Santa Barbara	**State/Province**	CA
Country [i]	US	Postal Code	93101
Phone Number		Fax Number	
Email Address			

Additional Applicant Data may be generated within this form by selecting the Add button. Add

Signature:

Remove

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications

Signature	/Robert J. Irvine III/			Date (YYYY-MM-DD)	2013-04-16
First Name	Robert	Last Name	Irvine	Registration Number	41865

Additional Signature may be generated within this form by selecting the Add button. Add

Application Data Sheet 37 CFR 1.76		Attorney Docket Number	SNS-71301.US.10
		Application Number	
Title of Invention	SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES		

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

ATTORNEY'S DOCKET NO. 11-0901

PATENTS

UNITED STATES PATENT APPLICATION

OF

NICHOLAS A. J. MILLINGTON

FOR

SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES

INCORPORATION BY REFERENCE

This application claims the benefit of Provisional Patent Application Serial No. 60/490,768, filed on July 28, 2003, entitled "Method for Synchronizing Audio Playback Between Multiple Networked Devices," and Non-provisional Patent Application Serial No. 10/816,217, filed on April 1, 2004, entitled "System and Method for Synchronizing Channel Handoff as Among a Plurality of Independently Clocked Digital Devices," both of which are assigned to the assignee of the present application and are incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates generally to the field of digital data processing devices, and more particularly to systems and methods for synchronizing operations among a plurality of independently-clocked digital data processing devices. The invention is embodied in a system for synchronizing operations among a plurality of devices, in relation to information that is provided by a common source. One embodiment of the invention enables synchronizing of audio playback as among two or more audio playback devices that receive audio information from a common information source, or channel.

More generally, the invention relates to the field of arrangements that synchronize output generated by a number of output generators, including audio output, video output, combinations of audio and video, as well as other types of output as will be appreciated by those skilled in the art, provided by a common channel. Generally, the invention will find utility in connection with any type of information for which synchrony among independently-clocked devices is desired.

BACKGROUND OF THE INVENTION

There are a number of circumstances under which it is desirable to maintain synchrony of operations among a plurality of independently-clocked digital data processing devices in relation to, for example, information that is provided thereto by a common source. For example, systems are

being developed in which one audio information source can distribute audio information in digital form to a number of audio playback devices for playback. The audio playback devices receive the digital information and convert it to analog form for playback. The audio playback devices may be located in the same room or they may be distributed in different rooms in a residence such as a house or an apartment, in different offices in an office building, or the like. For example, in a system installed in a residence, one audio playback device may be located in a living room, while another audio playback device is be located in a kitchen, and yet other audio playback devices may be located in various bedrooms of a house. In such an arrangement, the audio information that is distributed to various audio playback devices may relate to the same audio program, or the information may relate to different audio programs. If the audio information source provides audio information relating to the same audio program to two or more audio playback devices at the same time, the audio playback devices will generally contemporaneously play the same program. For example, if the audio information source provides audio information to audio playback devices located in the living room and kitchen in a house at the same time, they will generally contemporaneously play the same program.

One problem that can arise is to ensure that, if two or more audio playback devices are contemporaneously attempting to play back the same audio program, they do so simultaneously. Small differences in the audio playback devices' start times and/or playback speeds can be perceived by a listener as an echo effect, and larger differences can be very annoying. Differences can arise because for a number of reasons, including delays in the transfer of audio information over the network. Such delays can differ as among the various audio playback devices for a variety of reasons, including where they are connected into the network, message traffic and other reasons as will be apparent to those skilled in the art.

Another problem arises from the following. When an audio playback device converts the digital audio information from digital to analog form, it does so using a clock that provides timing information. Generally, the audio playback devices that are being developed have independent clocks, and, if they are not clocking at precisely the same rate, the audio playback provided by the various devices can get out of synchronization.

SUMMARY OF THE INVENTION

The invention provides a new and improved system and method for synchronizing operations among a number of digital data processing devices that are regulated by independent clocking devices. Generally, the invention will find utility in connection with any type of information for which synchrony among devices connected to a network is desired. The invention is described in connection with a plurality of audio playback devices that receive digital audio information that is to be played back in synchrony, but it will be appreciated that the invention can find usefulness in connection with any kind of information for which coordination among devices that have independent clocking devices would find utility.

In brief summary, the invention provides, in one aspect, a system for maintaining synchrony of operations among a plurality of devices that have independent clocking arrangements. The system includes a task distribution device that distributes tasks to a synchrony group comprising a plurality of devices that are to perform the tasks distributed by the task distribution device in synchrony. The task distribution device distributes each task to the members of the synchrony group over a network. Each task is associated with a time stamp that indicates a time, relative to a clock maintained by the task distribution device, at which the members of the synchrony group are to execute the task. Each member of the synchrony group periodically obtains from the task distribution device an indication of the current time indicated by its clock, determines a time differential between the task distribution device's clock and its respective clock and determines therefrom a time at which, according to its respective clock, the time stamp indicates that it is to execute the task.

In one embodiment, the tasks that are distributed include audio information for an audio track that is to be played by all of the devices comprising the synchrony group synchronously. The audio track is divided into a series of frames, each of which is associated with a time stamp indicating the time, relative to the clock maintained by an audio information channel device, which, in that embodiment, serves as the task distribution device, at which the members of the synchrony group are to play the respective frame. Each member of the synchrony group, using a very accurate

protocol, periodically obtains the time indicated by the audio information channel device, and determines a differential between the time as indicated by its local clock and the audio information channel device's clock. The member uses the differential and the time as indicated by the time stamp to determine the time, relative to its local clock, at which it is to play the respective frame. The members of the synchrony group do this for all of the frames, and accordingly are able to play the frames in synchrony.

BRIEF DESCRIPTION OF THE DRAWINGS

This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

FIG. 1 schematically depicts an illustrative networked audio system, constructed in accordance with the invention;

FIG. 2 schematically depicts a functional block diagram of a synchrony group utilizing a plurality of zone players formed within the networked audio system depicted in FIG. 1;

FIG. 2A schematically depicts two synchrony groups, illustrating how a member of one synchrony group can provide audio information to the members of another synchrony group;

FIG. 3 depicts an functional block diagram of a zone player for use in the networked audio system depicted in FIG. 1; and

FIG. 4 is useful in understanding an digital audio information framing methodology useful in the network audio system depicted in FIG. 1.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

FIG. 1 depicts an illustrative network audio system 10 constructed in accordance with the invention. With reference to FIG. 1, the network audio system 10 includes a plurality of zone players 11(1) through 11(N) (generally identified by reference numeral 11(n)) interconnected by a local network 12, all of which operate under control of one or more user interface modules generally identified by reference numeral 13. One or more of the zone players 11(n) may also be connected to one or more audio information sources, which will generally be identified herein by reference numeral 14(n)(s), and/or one or more audio reproduction devices, which will generally be identified by reference numeral 15(n)(r). In the reference numeral 14(n)(s), index "n" refers to the index "n" of the zone player 11(n) to which the audio information source is connected, and the index "s" ($s=1,...,S_n$) refers to the "s-th" audio information source connected to that "n-th" zone player 11(n). Thus, if, for example, a zone player 11(n) is connected to four audio information sources 14(n)(1) through 14(n)(4), the audio information sources may be generally identified by reference numeral 14(n)(s), with $S_n=4$. It will be appreciated that the number of audio information sources S_n may vary as among the various zone players 11(n), and some zone players may not have any audio information sources connected thereto. Similarly, in the reference numeral 15(n)(r), index "n" refers to the index "n" of the zone player 11(n) to which the audio reproduction device is connected, and the index "r" ($r=1,...,R_n$) refers to the "r-th" audio information source connected to that "n-th" zone player 11(n). In addition to the audio information sources 14(n)(s), the network audio system 10 may include one or more audio information sources 16(1) through 16(M) connected through appropriate network interface devices (not separately shown) to the local network 12. Furthermore, the local network may include one or more network interface devices (also not separately shown) that are configured to connect the local network 12 to other networks, including a wide area network such as the Internet, the public switched telephony network (PSTN) or other networks as will be apparent to those skilled in the art, over which connections to audio information sources may be established.

The zone players 11(n) associated with system 10 may be distributed throughout an establishment such as residence, an office complex, a hotel, a conference hall, an amphitheater or auditorium, or other types of establishments as will be apparent to those skilled in the art or the like. For example, if the zone players 11(n) and their associated audio information source(s) and/or audio reproduction device(s) are distributed throughout a residence, one, such as zone player 11(1) and its associated audio information source(s) and audio reproduction device(s) may be located in a living room, another may be located in a kitchen, another may be located in a dining room, and yet others may be located in respective bedrooms, to selectively provide entertainment in those rooms. On the other hand, if the zone players 11(n) and their associated audio information source(s) and/or audio reproduction device(s) are distributed throughout an office complex, one may, for example, be provided in each office to selectively provide entertainment to the employees in the respective offices. Similarly, if the zone players 11(n) and associated audio information source(s) and/or audio reproduction device(s) are used in a hotel, they may be distributed throughout the rooms to provide entertainment to the guests. Similar arrangements may be used with zone players 11(n) and associated audio information source(s) and/or audio reproduction device(s) used in an amphitheater or auditorium. Other arrangements in other types of environments will be apparent to those skilled in the art. In each case, the zone players 11(n) can be used to selectively provide entertainment in the respective locations, as will be described below.

The audio information sources 14(n)(s) and 16(m) may be any of a number of types of conventional sources of audio information, including, for example, compact disc ("CD") players, AM and/or FM radio receivers, analog or digital tape cassette players, analog record turntables and the like. In addition, the audio information sources 14(n)(s) and 16(m) may comprise digital audio files stored locally on, for example, personal computers (PCs), personal digital assistants (PDAs), or similar devices capable of storing digital information in volatile or non-volatile form. As noted above, the local network 12 may also have an interface (not shown) to a wide area network, over which the network audio system 10 can obtain audio information. Moreover, one or more of the audio information sources 14(n)(s) may also comprise an interface to a wide area network such as the Internet, the public switched telephony network (PSTN) or any other source of audio

information. In addition, one or more of the audio information sources 14(n)(s) and 16(m) may comprise interfaces to radio services delivered over, for example, satellite. Audio information obtained over the wide area network may comprise, for example, streaming digital audio information such as Internet radio, digital audio files stored on servers, and other types of audio information and sources as will be appreciated by those skilled in the art. Other arrangements and other types of audio information sources will be apparent to those skilled in the art.

Generally, the audio information sources 14(n)(s) and 16(m) provide audio information associated with audio programs to the zone players for playback. A zone player that receives audio information from an audio information source 14(n)(s) that is connected thereto can provide playback and/or forward the audio information, along with playback timing information, over the local network 12 to other zone players for playback. Similarly, each audio information source 16(m) that is not directly connected to a zone player can transmit audio information over the network 12 to any zone player 11(n) for playback. In addition, as will be explained in detail below, the respective zone player 11(n) can transmit the audio information that it receives either from an audio information source 14(n)(s) connected thereto, or from an audio information source 16(m), to selected ones of the other zone players 11(n'), 11(n''),... (n not equal to n', n'',...) for playback by those other zone players. The other zone players 11(n'), 11(n''),... to which the zone player 11(n) transmits the audio information for playback may be selected by a user using the user interface module 13. In that operation, the zone player 11(n) will transmit the audio information to the selected zone players 11(n'), 11(n''),... over the network 12. As will be described below in greater detail, the zone players 11(n), 11(n'), 11(n''),... operate such that the zone players 11(n'), 11(n''),... synchronize their playback of the audio program with the playback by the zone player 11(n), so that the zone players 11(n), 11(n'), 11(n'') provide the same audio program at the same time.

Users, using user interface module 13, may also enable different groupings or sets of zone players to provide audio playback of different audio programs synchronously. For example, a user, using a user interface module 13, may enable zone players 11(1) and 11(2) to play one audio program, audio information for which may be provided by, for example, one audio information source 14(1)(1). The same or a different user may, using the same or a different user interface

module 13, enable zone players 11(4) and 11(5) to contemporaneously play another audio program, audio information for which may be provided by a second audio information source, such as audio information source 14(5)(2). Further, a user may enable zone player 11(3) to contemporaneously play yet another audio program, audio information for which may be provided by yet another audio information source, such as audio information source 16(1). As yet another possibility, a user may contemporaneously enable zone player 11(1) to provide audio information from an audio information source connected thereto, such as audio information source 14(1)(2), to another zone player, such as zone player 11(6) for playback.

In the following, the term "synchrony group" will be used to refer to a set of one or more zone players that are to play the same audio program synchronously. Thus, in the above example, zone players 11(1) and 11(2) comprise one synchrony group, zone player 11(3) comprises a second synchrony group, zone players 11(4) and 11(5) comprise a third synchrony group, and zone player 11(6) comprises yet a fourth synchrony group. Thus, while zone players 11(1) and 11(2) are playing the same audio program, they will play the audio program synchronously. Similarly, while zone players 11(4) and 11(5) are playing the same audio program, they will play the audio program synchronously. On the other hand, zone players that are playing different audio programs may do so with unrelated timings. That is, for example, the timing with which zone players 11(1) and 11(2) play their audio program may have no relationship to the timing with which zone player 11(3), zone players 11(4) and 11(5), and zone player 11(6) play their audio programs. It will be appreciated that, since "synchrony group" is used to refer to sets of zone players that are playing the same audio program synchronously, zone player 11(1) will not be part of zone player 11(6)'s synchrony group, even though zone player 11(1) is providing the audio information for the audio program to zone player 11(6).

In the network audio system 10, the synchrony groups are not fixed. Users can enable them to be established and modified dynamically. Continuing with the above example, a user may enable the zone player 11(1) to begin providing playback of the audio program provided thereto by audio information source 14(1)(1), and subsequently enable zone player 11(2) to join the synchrony group. Similarly, a user may enable the zone player 11(5) to begin providing playback of the audio program

provided thereto by audio information source 14(5)(2), and subsequently enable zone player 11(4) to join that synchrony group. In addition, a user may enable a zone player to leave a synchrony group and possibly join another synchrony group. For example, a user may enable the zone player 11(2) to leave the synchrony group with zone player 11(1), and join the synchrony group with zone player 11(6). As another possibility, the user may enable the zone player 11(1) to leave the synchrony group with zone player 11(2) and join the synchrony group with zone player 11(6). In connection with the last possibility, the zone player 11(1) can continue providing audio information from the audio information source 14(1)(1) to the zone player 11(2) for playback thereby.

A user, using the user interface module 13, can enable a zone player 11(n) that is currently not a member of a synchrony group to join a synchrony group, after which it will be enabled to play the audio program that is currently being played by that synchrony group. Similarly, a user, also using the user interface module 13, can enable a zone player 11(n) that is currently a member of one synchrony group, to disengage from that synchrony group and join another synchrony group, after which that zone player will be playing the audio program associated with the other synchrony group. For example, if a zone player 11(6) is currently not a member of any synchrony group, it, under control of the user interface module 13, can become a member of a synchrony group, after which it will play the audio program being played by the other members of the synchrony group, in synchrony with the other members of the synchrony group. In becoming a member of the synchrony group, zone player 11(6) can notify the zone player that is the master device for the synchrony group that it wishes to become a member of its synchrony group, after which that zone player will also transmit audio information associated with the audio program, as well as timing information, to the zone player 11(6). As the zone player 11(6) receives the audio information and the timing information from the master device, it will play the audio information with the timing indicated by the timing information, which will enable the zone player 11(6) to play the audio program in synchrony with the other zone player(s) in the synchrony group.

Similarly, if a user, using the user interface module 13, enables a zone player 11(n) associated with a synchrony group to disengage from that synchrony group, and if the zone player 11(n) is not the master device of the synchrony group, the zone player 11(n) can notify the master

device, after which the master device can terminate transmission of the audio information and timing information to the zone player 11(n). If the user also enables the zone player 11(n) to begin playing another audio program using audio information from an audio information source 14(n)(s) connected thereto, it will acquire the audio information from the audio information source 14(n)(s) and initiate playback thereof. If the user enables another zone player 11(n') to join the synchrony group associated with zone player 11(n), operations in connection therewith can proceed as described immediately above.

As yet another possibility, if a user, using the user interface module 13, enables a zone player 11(n) associated with a synchrony group to disengage from that synchrony group and join another synchrony group, and if the zone player is not the master device of the synchrony group from which it is disengaging, the zone player 11(n) can notify the master device of the synchrony group from which it is disengaging, after which that zone player will terminate transmission of audio information and timing information to the zone player 11(n) that is disengaging. Contemporaneously, the zone player 11(n) can notify the master device of the synchrony group that it (that is, zone player 11(n)) is joining, after which the master device can begin transmission of audio information and timing information to that zone player 11(n). The zone player 11(n) can thereafter begin playback of the audio program defined by the audio information, in accordance with the timing information so that the zone player 11(n) will play the audio program in synchrony with the master device.

As yet another possibility, a user, using the user interface module 13, may enable a zone player 11(n) that is not associated with a synchrony group, to begin playing an audio program using audio information provided to it by an audio information source 14(n)(s) connected thereto. In that case, the user, also using the user interface module 13 or a user interface device that is specific to the audio information source 14(n)(s), can enable the audio information source 14(n)(s) to provide audio information to the zone player 11(n). After the zone player 11(n) has begun playback, or contemporaneously therewith, the user, using the user interface module 13, can enable other zone players 11(n'), 11(n''),... to join zone player 11(n)'s synchrony group and enable that zone player

11(n) to transmit audio information and timing information thereto as described above, to facilitate synchronous playback of the audio program by the other zone players 11(n'), 11(n")....

A user can use the user interface module 13 to control other aspects of the network audio system 10, including but not limited to the selection of the audio information source 14(n)(s) that a particular zone player 11(n) is to utilize, the volume of the audio playback, and so forth. In addition, a user may use the user interface module 13 to turn audio information source(s) 14(n)(s) on and off and to enable them to provide audio information to the respective zone players 11(n).

Operations performed by the various devices associated with a synchrony group will be described in connection with FIG. 2, which schematically depicts a functional block diagram of a synchrony group in the network audio system 10 described above in connection with FIG. 1. With reference to FIG. 2, a synchrony group 20 includes a master device 21 and zero or more slave devices 22(1) through 22 (G) (generally identified by reference numeral 22(g)), all of which synchronously play an audio program provided by an audio information channel device 23. Each of the master device 21, slave devices 22(g) and audio information channel device 23 utilizes a zone player 11(n) depicted in FIG. 1, although it will be clear from the description below that a zone player may be utilized both for the audio information channel device for the synchrony group 20, and the master device 21 or a slave device 22(g) of the synchrony group 20. As will be described below in more detail, the audio information channel device 23 obtains the audio information for the audio program from an audio information source, adds playback timing information, and transmits the combined audio and playback timing information to the master device 21 and slave devices 22(g) over the network 12 for playback. The playback timing information that is provided with the audio information, together with clock timing information provided by the audio information channel device 23 to the various devices 21 and 22(g) as will be described below, enables the master device 21 and slave devices 22(g) of the synchrony group 20 to play the audio information simultaneously.

The master device 21 and the slave devices 22(g) receive the audio and playback timing information, as well as the clock timing information, that are provided by the audio information channel device 23, and play back the audio program defined by the audio information. The master

device 21 is also the member of the synchrony group 20 that communicates with the user interface module 13 and that controls the operations of the slave devices 22(g) in the synchrony group 20. In addition, the master device 21 controls the operations of the audio information channel device 23 that provides the audio and playback timing information for the synchrony group 20. Generally, the initial master device 21 for the synchrony group will be the first zone player 11(n) that a user wishes to play an audio program. However, as will be described below, the zone player 11(n) that operates as the master device 21 can be migrated from one zone player 11(n) to another zone player 11(n'), which preferably will be a zone player that is currently operating as a slave device 22(g) in the synchrony group.

In addition, under certain circumstances, as will be described below, the zone player 11(n) that operates as the audio information channel device 23 can be migrated from one zone player to another zone player, which also will preferably will be a zone player that is currently operating as a member of the synchrony group 20. It will be appreciated that the zone player that operates as the master device 21 can be migrated to another zone player independently of the migration of the audio information channel device 23. For example, if one zone player 11(n) is operating as both the master device 21 and the audio information channel device 23 for a synchrony group 20, the master device 21 can be migrated to another zone player 11(n') while the zone player 11(n) is still operating as the audio information channel device 23. Similarly, if one zone player 11(n) is operating as both the master device 21 and the audio information channel device 23 for a synchrony group 20, the audio information channel device 23 can be migrated to another zone player 11(n') while the zone player 11(n) is still operating as the master device 21. In addition, if one zone player 11(n) is operating as both the master device 21 and the audio information channel device 23 for a synchrony group 20, the master device 21 can be migrated to another zone player 11(n') and the audio information channel device can be migrated to a third zone player 11(n").

The master device 21 receives control information from the user interface module 13 for controlling the synchrony group 20 and provides status information indicating the operational status of the synchrony group to the user interface module 13. Generally, the control information from the user interface module 13 enables the master device 21 to, in turn, enable the audio information

channel device 23 to provide audio and playback timing information to the synchrony group to enable the devices 21 and 22(g) that are members of the synchrony group 20 to play the audio program synchronously. In addition, the control information from the user interface module 13 enables the master device 21 to, in turn, enable other zone players to join the synchrony group as slave devices 22(g) and to enable slave devices 22(g) to disengage from the synchrony group. Control information from the user interface module 13 can also enable the zone player 11(n) that is currently operating as the master device 21 to disengage from the synchrony group, but prior to doing so that zone player will enable the master device 21 to transfer from that zone player 11(n) to another zone player11(n'), preferably to a zone player 11(n') that is currently a slave device 22(g) in the synchrony group 20. The control information from the user interface module 13 can also enable the master device 21 to adjust its playback volume and to enable individual ones of the various slave devices 22(g) to adjust their playback volumes. In addition, the control information from the user interface module 13 can enable the synchrony group 20 to terminate playing of a current track of the audio program and skip to the next track, and to re-order tracks in a play list of tracks defining the audio program that is to be played by the synchrony group 20.

The status information that the master device 21 may provide to the user interface module 13 can include such information as a name or other identifier for the track of the audio work that is currently being played, the names or other identifiers for upcoming tracks, the identifier of the zone player 11(n) that is currently operating as the master device 21, and identifiers of the zone players that are currently operating as slave devices 22(g). In one embodiment, the user interface module 13 includes a display (not separately shown) that can display the status information to the user.

It will be appreciated that the zone player 11(n) that is operating as the audio information channel device 23 for one synchrony group may also comprise the master device 21 or any of the slave devices 22(g) in another synchrony group. This may occur if, for example, the audio information source that is to provide the audio information that is to be played by the one synchrony group is connected to a zone player also being utilized as the master device or a slave device for the other synchrony group. This will be schematically illustrated below in connection with FIG. 2A. Since, as noted above, the zone player 11(n) that is operating as the audio information channel

device 23 for the synchrony group 20 may also be operating as a master device 21 or slave device 22(g) for another synchrony group, it can also be connected to one or more audio reproduction devices 15(n)(r), although that is not depicted in FIG. 2. Since the master device 21 and slave devices 22(g) are all to provide playback of the audio program, they will be connected to respective audio reproduction devices 15(n)(r). Furthermore, it will be appreciated that one or more of the zone players 11(n) that operate as the master device 21 and slave devices 22(g) in synchrony group 20 may also operate as an audio information channel device for that synchrony group or for another synchrony group and so they may be connected to one or more audio information sources 14(n)(s), although that is also not depicted in FIG. 2. In addition, it will be appreciated that a zone player 11(n) can also operate as a audio information channel device 23 for multiple synchrony groups.

If the audio information channel device 23 does not utilize the same zone player as the master device 21, the master device 21 controls the audio information channel device by exchanging control information over the network 12 with the audio information channel device 23. The control information is represented in FIG. 2 by the arrow labeled CHAN_DEV_CTRL_INFO. The control information that the master device 21 provides to the audio information channel device 23 will generally depend on the nature of the audio information source that is to provide the audio information for the audio program that is to be played and the operation to be enabled by the control information. If, for example, the audio information source is a conventional compact disc, tape, or record player, broadcast radio receiver, or the like, which is connected to a zone player 11(n), the master device 21 may merely enable the zone player serving as the audio information channel device 23 to receive the audio information for the program from the audio information source. It will be appreciated that, if the audio information is not in digital form, the audio information channel device 23 will convert it to digital form and provide the digitized audio information, along with the playback timing information, to the master device 21 and slave devices 22(g).

On the other hand, if the audio information source is, for example, a digital data storage device, such as may be on a personal computer or similar device, the master device 21 can provide a play list to the audio information channel device 23 that identifies one or more files containing the audio information for the audio program. In that case, the audio information channel device 23 can

retrieve the files from the digital data storage device and provide them, along with the playback timing information, to the master device 21 and the slave devices 22(g). It will be appreciated that, in this case, the audio information source may be directly connected to the audio information channel device 23, as, for example, an audio information source 14(n)(s), or it may comprise an audio information source 16(m) connected to the network 12. As a further alternative, if the audio information source is a source available over the wide area network, the master device 21 can provide a play list comprising a list of web addresses identifying the files containing the audio information for the audio program that is to be played, and in that connection the audio information channel device 23 can initiate a retrieval of the files over the wide area network. As yet another alternative, if the audio information source is a source of streaming audio received over the wide area network, the master device 21 can provide a network address from which the streaming audio can be received. Other arrangements by which the master device 21 can control the audio information channel device 23 will be apparent to those skilled in the art.

The master device 21 can also provide control information to the synchrony group's audio information channel device 23 to enable a migration from one zone player 11(n) to another zone player 11(n'). This may occur if, for example, the audio information source is one of audio information sources 16 or a source accessible over the wide area network via the network 12. The master device 21 can enable migration of the audio information channel device 23 for several reasons, including, for example, to reduce the loading of the zone player 11(n), to improve latency of message transmission in the network 12, and other reasons as will be appreciated by those skilled in the art.

As noted above, the audio information channel device 23 provides audio and playback timing information for the synchrony group to enable the master device 21 and slave devices 22(g) to play the audio program synchronously. Details of the audio and playback timing information will be described in detail below in connection with FIGS. 3 and 4, but, in brief, the audio information channel device 23 transmits the audio and playback timing information in messages over the network 12 using a multi-cast message transmission methodology. In that methodology, the audio information channel device 23 will transmit the audio and playback timing information in a series

of messages, with each message being received by all of the zone players 11(n) comprising the synchrony group 20, that is, by the master device 21 and the slave devices 22(g). Each of the messages includes a multi-cast address, which the master device 21 and slave devices 22(g) will monitor and, when they detect a message with that address, they will receive and use the contents of the message. The audio and playback timing information is represented in FIG. 2 by the arrow labeled "AUD+PBTIME_INF0," which has a single tail, representing a source for the information at the audio information channel device 23, and multiple arrowheads representing the destinations of the information, with one arrowhead extending to the master device 21 and other arrowheads extending to each of the slave devices 22(g) in the synchrony group 20. The audio information channel device 23 may make use of any convenient multi-cast message transmission methodology in transmitting the audio and playback timing information to the synchrony group 20. As will be described in detail in connection with FIG. 4, the audio and playback timing information is in the form of a series of frames, with each frame having a time stamp. The time stamp indicates a time, relative to the time indicated by a clock maintained by the audio information channel device 23, at which the frame is to be played. Depending on the size or sizes of the messages used in the selected multi-cast message transmission methodology and the size or sizes of the frames, a message may contain one frame, or multiple frames, or, alternatively, a frame may extend across several messages.

The audio information channel device 23 also provides clock time information to the master device 21 and each of the slave devices 22(g) individually over network 12 using a highly accurate clock time information transmission methodology. The distribution of the clock time information is represented in FIG. 2 by the arrows labeled "AICD_CLK_INF (M)" (in the case of the clock time information provided to the master device 21) and "AICD_CLK_INF (S_1)" through "AICD_CLK_INF (S_G)" (in the case of audio information channel device clock information provided to the slave devices 22(g)). In one embodiment, the master device 21 and slave devices 22(g) make use of the well-known SNTP (Simple Network Time Protocol) to obtain current clock time information from the audio information channel device 23. The SNTP makes use of a unicast message transfer methodology, in which one device, such as the audio information channel device 23, provides clock time information to a specific other device, such as the master device 21 or a

slave device 22(g), using the other device's network, or unicast, address. Each of the master device 21 and slave devices 22(g) will periodically initiate SNTP transactions with the audio information channel device 23 to obtain the clock time information from the audio information channel device 23. As will be described below in more detail, the master device 21 and each slave device 22(g) make use of the clock time information to determine the time differential between the time indicated by the audio information channel device's clock and the time indicated by its respective clock, and use that time differential value, along with the playback time information associated with the audio information and the respective device's local time as indicated by its clock to determine when the various frames are to be played. This enables the master device 21 and the slave devices 22(g) in the synchrony group 20 to play the respective frames simultaneously.

As noted above, the control information provided by the user to the master device 21 through the user interface module 13 can also enable the master device 21 to, in turn, enable another zone player 11(n') to join the synchrony group as a new slave device 22(g). In that operation, the user interface module 13 will provide control information, including the identification of the zone player 11(n') that is to join the synchrony group to the master device 21. After it receives the identification of the zone player 11(n') that is to join the synchrony group, the master device 21 will exchange control information, which is represented in FIG. 2 by the arrows labeled SLV_DEV_CTRL_INF (S_1) through SLV_DEV_CTRL_INF (S_G) group slave control information, over the network 12 with the zone player 11(n') that is identified in the control information from the user interface module 13. The control information that the master device 21 provides to the new zone player 11(n') includes the network address of the zone player 11(n) that is operating as the audio information channel device 23 for the synchrony group, as well as the multi-cast address that the audio information channel device 23 is using to broadcast the audio and playback timing information over the network. The zone player that is to operate as the new slave device 22(g') uses the multi-cast address to begin receiving the multi-cast messages that contain the audio information for the audio program being played by the synchrony group.

It will be appreciated that, if the zone player 11(n) that is operating as the master device 21 for the synchrony group 20 is also operating the audio information channel device 23, and if there

are no slave devices 22(g) in the synchrony group 20, the audio information channel device 23 may not be transmitting audio and playback timing information over the network. In that case, if the new slave device 22(g') is the first slave device in the synchrony group, the zone player 11(n) that is operating as both the master device 21 and audio information channel device 23, can begin transmitting the audio and playback timing information over the network 12 when the slave device 22(g') is added to the synchrony group 20. The zone player 11(n) can maintain a count of the number of slave devices 22(g) in the synchrony group 20 as they join and disengage, and, if the number drops to zero, it can stop transmitting the audio and playback timing information over the network 12 to reduce the message traffic over the network 12.

The new slave device 22(g') added to the synchrony group 20 uses the network address of the audio information channel device 23 for several purposes. In particular, the new slave device 22(g') will, like the master device 21 (assuming the zone player 11(n) operating as the master device 21 is not also the audio information channel device 23), engage in SNTP transactions with the audio information channel device 23 to obtain the clock timing information from the audio information channel device 23. In addition, the new slave device 22(g') can notify the audio information channel device 23 that it is a new slave device 22(g') for the synchrony group 20 and provide the audio information channel device 23 with its network address. As will be described below, in one embodiment, particularly in connection with audio information obtained from a source, such as a digital data storage device, which can provide audio information at a rate that is faster than the rate at which it will be played, the audio information channel device 23 will buffer audio and timing information and broadcast it over the network 12 to the synchrony group 20 generally at a rate at which it is provided by the source. Accordingly, when a new slave device 22(g') joins the synchrony group 20, the playback timing information may indicate that the audio information that is currently being broadcast by the audio information channel device 23 using the multi-cast methodology is to be played back some time in the future. To reduce the delay with which the new slave device 22(g') will begin playback, the audio information channel device 23 can also retransmit previously transmitted audio and timing information that it had buffered to the new slave device 22(g') using the unicast network address of the slave device 22(g').

The master device 21 can also use the slave device control information exchanged with the slave devices 22(g) for other purposes. For example, the master device 21 can use the slave device control information to initiate a migration of the master from its zone player 11(n) to another zone player 11(n'). This may occur for any of a number of reasons, including, for example, that the master device 21 is terminating playback by it of the audio program and is leaving the synchrony group 20, but one or more of the other devices in the synchrony group is to continue playing the audio program. The master device 21 may also want to initiate a migration if it is overloaded, which can occur if, for example, the zone player 11(n) that is the master device 21 for its synchrony group is also operating as an audio information channel device 23 for another synchrony group.

The user can also use the user interface module 13 to adjust playback volume by the individual zone players 11(n) comprising the synchrony group. In that operation, the user interface module 13 provides information identifying the particular device whose volume is to be adjusted, and the level at which the volume is to be set to the master device 21. If the device whose volume is to be adjusted is the master device 21, the master device 21 can adjust its volume according to the information that it receives from the user interface module 13. On the other hand, if the device whose volume is to be adjusted is a slave device 22(g), the master device 21 can provide group slave control information to the respective slave device 22(g), to enable it to adjust its volume.

The user can also use the user interface module 13 to enable a synchrony group 20 to cancel playing of the track in an audio program that is currently being played, and to proceed immediately to the next track. This may occur, for example, if the tracks for the program is in the form of a series of digital audio information files, and the user wishes to cancel playback of the track that is defined by one of the files. In that case, when the master device 21 receives the command to cancel playback of the current track, it will provide channel device control information to the audio information channel device 23 so indicating. In response, the audio information channel device 23 inserts control information into the audio and playback timing information, which will be referred to as a "resynchronize" command. In addition, the audio information channel device 23 will begin transmitting audio information for the next track, with timing information to enable it to be played

immediately. The resynchronize command can also enable playback of a track to be cancelled before it has been played. Details of these operations will be described below.

As noted above, there may be multiple synchrony groups in the network audio system 10, and further that, for example, a zone player 11(n) may operate both as a master device 21 or a slave device 22(g) in one synchrony group, and as the audio information channel device 23 providing audio and playback timing information and clock timing information for another synchrony group. An illustrative arrangement of this will be described in connection with FIG. 2A. With reference to FIG. 2A, that FIG. depicts elements of two synchrony groups, identified by reference numerals 20(1) and 20(2), respectively. For clarity, FIG. 2A does not show a number of elements, the presence of which would be evident from FIGS. 1 and 2 as described above. For example, FIG. 2A does not depict the audio information sources from which audio information is obtained for the synchrony groups or the audio reproduction devices that are used to produce sound for the master and slave devices, which are depicted in both FIGS. 1 and 2. In addition, FIG. 2A does not depict arrows that represent control information provided by the respective master devices to the slave devices in the respective synchrony groups, or to the audio information channel devices that provide audio and timing information for the respective synchrony groups, which are depicted in FIG. 2. In addition, FIG. 2A does not depict the arrows that represent the clock timing information provided by the audio information channel devices to the respective members of the respective synchrony groups, which are also depicted in FIG. 2. As will be noted below, however, FIG. 2A does depict arrows representing the audio and playback timing information provided by the respective audio information channel devices for the respective synchrony groups 20(1), 20(2), to the master and slave devices comprising the respective synchrony groups 20(1), 20(2).

Each synchrony group 20(1), 20(2) comprises elements of a number of zone players. A functional block diagram of a zone player will be described below in connection with FIG. 3. Synchrony group 20(1) includes a master device 21(1) and "K" slave devices 22(1)(1) through 22(K)(1) (the index "1" in reference numeral 21(1) and the last index in reference numeral 22(1)(1) through 21(K)(1) corresponds to the index of the synchrony group 20(1) to which they belong) utilize zone players 11(1) and 11(K+1) respectively. Similarly, synchrony group 20(2) includes a

master device 21(2) and "L" slave devices 22(1)(2) through 22(L)(2) that utilize zone players 11(K+2) through 11(K+L+2). In the illustrative arrangement depicted in FIG. 2A, both synchrony groups 20(1) and 20(2) are controlled by the user interface module 13, which can provide control information to, and receive status information from, the master devices 21(1) and 21(2) independently. It will be appreciated that separate user interface modules may be provided to provide control information to, and receive status information from, the respective master devices 21(1), 21(2).

As noted above, the slave device 22(1)(2) in synchrony group 20(2) utilizes zone player 11(K+3). In the illustrative arrangement depicted in FIG. 2A, the audio information channel device 23(1) that provides audio and playback timing information to the master and slave devices 21(1), 22(1)(1),..., 22(K)(1) of synchrony group 20(1) also utilizes zone player 11(K+3). As noted above, this may occur if, for example, the audio information source that is to provide audio information to be played by the synchrony group 20(1) is connected to the zone player 11(K+3). Thus, when the master device 21(1) of synchrony group 20(1) exchanges channel device control information with the audio information channel device 23(1), it is effectively exchanging channel device control information with the zone player 11(K+3). Similarly, when the master and slave devices 21(1), 22(1)(1),..., 22(K)(1) of synchrony group 20(1) receive audio and playback timing information, as well as clock timing information, from the audio information channel device 23(1), they are effectively receiving the information from the zone player 11(K+3). FIG. 2A depicts a multi-headed arrow representing audio and playback timing information transmitted by the zone player 11(K+3), as audio information channel device 23(1), to the master and slave devices 21(1), 22(1)(1),...,11(K)(1) comprising synchrony group 20(1).

On the other hand, in the illustrative arrangement depicted in FIG. 2A, the synchrony group 20(2) utilizes a zone player 11(K+L+3) as its audio information channel device 23(2). As with synchrony group 20(1), when the master device 21(2) of synchrony group 20(2) exchanges channel device control information with the audio information channel device 23(2), it is effectively exchanging channel device control information with the zone player 11(K+L+3). Similarly, when the master and slave devices 21(2), 22(1)(2),..., 22(L)(2) of synchrony group 20(2) receive audio and

playback timing information, as well as clock timing information, from the audio information channel device 23(2), they are effectively receiving the information from the zone player 11(K+L+3). FIG. 2A depicts a multi-headed arrow representing audio and playback timing information transmitted by the zone player 11(K+3) as audio information channel device 23(2) to the master and slave devices 21(2), 22(1)(2),...,22(L)(2) comprising synchrony group 20(2).

In the illustrative arrangement depicted in FIG. 2A, zone player 11(K+L+3), which is the audio information channel device 23(2) for synchrony group 20(2), is not shown as being either a master or a slave device in another synchrony group. However, it will be appreciated that zone player 11(K+L+3) could also be utilized as the master device or a slave device for another synchrony group. Indeed, it will be appreciated that the zone player that is utilized as the audio information channel device for synchrony group 20(2) may also be a zone player that is utilized as the master device 21(1) or a slave device 22(1)(1),..., 22(K)(1) in the synchrony group 20(1).

A zone player 11(n) that is utilized as a member of one synchrony group may also be utilized as the audio information channel device for another synchrony group if the audio information source that is to supply the audio information that is to be played by the other synchrony group is connected to that zone player 11(n). A zone player 11(n) may also be utilized as the audio information channel device for the other synchrony group if, for example, the audio information source is an audio information source 16(m) (FIG. 1) that is connected to the network 12 or an audio information source that is available over a wide area network such as the Internet. The latter may occur if, for example, the zone player 11(n) has sufficient processing power to operate as the audio information channel device and it is in an optimal location in the network 12, relative to the zone players comprising the other synchrony group (that is the synchrony group for which it is operating as audio information channel device) for providing the audio and playback timing information to the members of the other synchrony group. Other circumstances under which the zone player 11(n) that is utilized as a member of one synchrony group may also be utilized as the audio information channel device for another synchrony group will be apparent to those skilled in the art.

As was noted above, the master device 21 for a synchrony group 20 may be migrated from one zone player 11(n) to another zone player 11(n'). As was further noted above, the audio information channel device 23 for a synchrony group 20 may be migrated from one zone player 11(n) to another zone player 11(n'). It will be appreciated that the latter may occur if, for example, the audio information source that provides the audio program for the synchrony group is not connected to the zone player 11(n) that is operating as the audio information channel device 23, but instead is one of the audio information sources 16(m) connected to the network 12 or a source available over a wide area network such as the Internet. Operations performed during a migration of an audio information channel device 23 from one zone player 11(n) to another zone player 11(n') will generally depend on the nature of the audio information that is being channeled by the audio information channel device 23. For example, if the audio information source provides streaming audio, the zone player 11(n) that is currently operating as the audio information channel device 23 for the synchrony group 20, can provide the following information to the other zone player 11(n') that is to become the audio information channel device 23 for the synchrony group 20:

(a) the identification of the source of the streaming audio information,

(b) the time stamp associated with the frame that the zone player 11(n) is currently forming, and

(c) the identifications of the zone players that are operating as the master device 21 and slave devices 22(g) comprising the synchrony group 20.

After the zone player 11(n') receives the information from the zone player 11(n), it will begin receiving the streaming audio from the streaming audio information source identified by the zone player 11(n), assemble the streaming audio information into frames, associate each frame with a time stamp, and transmit the resulting audio and playback timing information over the network 12. The zone player 11(n') will perform these operations in the same manner as described above, except that, instead of using the time indicated by its digital to analog converter clock 34 directly in generating the time stamps for the frames, the initial time stamp will be related to the value of the time stamp that is provided by the zone player 11(n) (reference item (b) above), with the rate at which the time

stamps are incremented corresponding to the rate at which its (that is, the zone player 11(n')'s) clock increments. In addition, the zone player 11(n') will notify the zone players that are operating as the master device 21 and slave devices 22(g) of the synchrony group 20 that it is the new audio information channel device 23 for the synchrony group 20, and provide the multi-cast address that it will be using to multi-cast the audio and playback timing information, as well as its unicast network address. After the members of the synchrony group 20 receive the notification from the zone player 11(n') indicating that it is the new audio information channel device 23 for the synchrony group 20, they will receive the audio and playback timing information from the zone player 11(n') instead of the zone player 11(n), using the multi-cast address provided by the zone player 11(n'). In addition, they can utilize the zone player 11(n')'s unicast network address to obtain current time information therefrom. It will be appreciated that the zone player 11(n') will determine its current time in relation to the time stamp that is provided by the zone player 11(n) (reference item (b) above) or the current time information that it received from the zone player 11(n) using the SNTP protocol as described above.

Generally similar operations can be performed in connection with migrating the audio information channel device from one zone player 11(n) to another zone player 11(n') if the audio information is from one or more audio information files, such as may be the case if the audio information comprises MP3 or WAV files that are available from sources such as sources 16(m) connected to the network 12 or over from sources available over a wide area network such as the Internet, except for differences to accommodate the fact that the audio information is in files. In that case, the zone player 11(n) that is currently operating as the audio information channel device 23 for the synchrony group 20 can provide the following information to the zone player 11(n') that is to become the audio information channel device 23 for the synchrony group 20:

(d) a list of the audio information files containing the audio information that is to be played;

(e) the identification of the file for which the zone player 11(n) is currently providing audio and playback timing information, along with the offset into the file for which the current item of

audio and playback timing information is being generated and the time stamp that the zone player 11(n) is associating with that frame, and

(f) the identifications of the zone players that comprise the master device 21 and slave devices 22(g) comprising the synchrony group 20.

After the zone player 11(n') receives the information from the zone player 11(n), it will begin retrieving audio information from the file identified in item (e), starting at the identified offset. In addition, the zone player 11(n') can assemble the retrieved audio information into frames, associate each frame with a time stamp and transmit the resulting audio and playback timing information over the network 12. The zone player 11(n') will perform these operations in the same manner as described above, except that, instead of using the time indicated by its digital to analog converter clock 34 directly in generating the time stamps for the frames, the value of the initial time stamp will be related to the time stamp that is provided by the zone player 11(n) (reference item (e) above), with the rate at which the time stamps are incremented corresponding to the rate at which its (that is, the zone player 11(n')'s) clock increments. In addition, the zone player 11(n') will notify the zone players that are operating as the master device 21 and slave devices 22(g) of the synchrony group 20 that it is the new audio information channel device 23 for the synchrony group 20, and provide the multi-cast address that it will be using to multi-cast the audio and playback timing information, as well as its unicast network address. After the members of the synchrony group 20 receive the notification from the zone player 11(n') indicating that it is the new audio information channel device 23 for the synchrony group 20, they will receive the audio and playback timing information from the zone player 11(n') instead of the zone player 11(n), using the multi-cast address provided by the zone player 11(n'). In addition, they can utilize the zone player 11(n')'s unicast network address to obtain current time information therefrom. It will be appreciated that the zone player 11(n') will determine its current time in relation to the time stamp that is provided by the zone player 11(n) (reference item (b) above) or the current time information that it received from the zone player 11(n) using the SNTP protocol as described above. The zone player 11(n') will process successive audio information files in the list that it receives from the zone player 11(n) (reference item (d)).

Operations performed by the zone players 11(n) and 11(n') in connection with migration of the audio information channel device 23 for other types of audio information will be apparent to those skilled in the art. In any case, preferably, the zone player 11(n) will continue operating as an audio information channel device 23 for the synchrony group 20 for at least a brief time after it notifies the zone player 11(n') that it is to become audio information channel device for the synchrony group, so that the zone player 11(n') will have time to notify the zone players in the synchrony group 20 that it is the new audio information channel device 23 for the synchrony group.

Before proceeding further in describing operations performed by the network audio system 10, it would be helpful to provide a detailed description of a zone player 11(n) constructed in accordance with the invention. FIG. 3 depicts a functional block diagram of a zone player 11(n) constructed in accordance with the invention. All of the zone players in the network audio system 10 may have similar construction. With reference to FIG. 3, the zone player 11(n) includes an audio information source interface 30, an audio information buffer 31, a playback scheduler 32, a digital to analog converter 33, an audio amplifier 35, an audio reproduction device interface 36, a network communications manager 40, and a network interface 41, all of which operate under the control of a control module 42. The zone player 11(n) also has a device clock 43 that provides timing signals that control the general operations of the zone player 11(n). In addition, the zone player 11(n) includes a user interface module interface 44 that can receive control signals from the user interface module 13 (FIGS. 1 and 2) for controlling operations of the zone player 11(n), and provide status information to the user interface module 13.

Generally, the audio information buffer 31 buffers audio information, in digital form, along with playback timing information. If the zone player 11(n) is operating as the audio information channel device 23 (FIG. 2) for a synchrony group 20, the information that is buffered in the audio information buffer 31 will include the audio and playback timing information that will be provided to the devices 21 and 22(g) in the synchrony group 20. If the zone player 11(n) is operating as the master device 21 or a slave device 22(g) for a synchrony group, the information that is buffered in the audio information buffer 31 will include the audio and playback timing information that the zone player 11(n) is to play.

The audio information buffer 31 can receive audio and playback timing information from two sources, namely, the audio information source interface 30 and the network communications manager 40. In particular, if the zone player 11(n) is operating as the audio information channel device 23 for a synchrony group 20, and if the audio information source is a source 14(n)(s) connected to the zone player 11(n), the audio information buffer 31 will receive and buffer audio and playback timing information from the audio information source interface 30. On the other hand, if the zone player 11(n) is operating as the audio information channel device 23 for a synchrony group 20, and if the audio information source is a source 16(m) connected to the network 12, or a source available over the wide area network, the audio information buffer 31 will receive and buffer audio and playback timing information from the network communications manager 40. It will be appreciated that, if the zone player 11(n) is not a member of the synchrony group, the zone player 11(n) will not play this buffered audio and playback timing information.

On yet another hand, if the zone player 11(n) is operating as the master device 21 or a slave device 22(g) in a synchrony group, and if the zone player 11(n) is not also the audio information channel device 23 providing audio and playback timing information for the synchrony group 20, the audio information buffer 31 will receive and buffer audio and playback timing information from the network communications manager 40.

The audio information source interface 30 connects to the audio information source(s) 14(n)(s) associated with the zone player 11(n). While the zone player 11(n) is operating as audio information channel device 23 for a synchrony group 20, and if the audio information is to be provided by a source 14(n)(s) connected to the zone player 11(n), the audio information source interface 30 will selectively receive audio information from one of the audio information source(s) 14(n)(s) to which the zone player is connected and store the audio information in the audio information buffer 21. If the audio information from the selected audio information source 14(n)(s) is in analog form, the audio information source interface 30 will convert it to digital form. The selection of the audio information source 14(n)(s) from which the audio information source interface 30 receives audio information is under control of the control module 42, which, in turn, receives control information from the user interface module through the user interface module interface 44.

The audio information source interface 30 adds playback timing information to the digital audio information and buffers the combined audio and playback timing information in the audio information buffer 21.

More specifically, as noted above, the audio information source interface 30 receives audio information from an audio information source 14(n)(s), converts it to digital form if necessary, and buffers it along with playback timing information in the audio information buffer 21. In addition, the audio information source interface 30 will also provide formatting and scheduling information for the digital audio information, whether as received from the selected audio information source 14(n)(s) or as converted from an analog audio information source. As will be made clear below, the formatting and scheduling information will control not only playback by the zone player 11(n) itself, but will also enable other zone players 11(n'), 11(n''),... that may be in a synchrony group for which the zone player 11(n) is the master device, to play the audio program associated with the audio information in synchrony with the zone player 11(n).

In one particular embodiment, the audio information source interface 30 divides the audio information associated with an audio work into a series of frames, with each frame comprising digital audio information for a predetermined period of time. As used herein, an audio track may comprise any unit of audio information that is to be played without interruption. On the other hand, an audio program may comprise a series of one or more audio tracks that are to be played in succession. It will be appreciated that the tracks comprising the audio program may also be played without interruption, or alternatively playback between tracks may be interrupted by a selected time interval. FIG. 4 schematically depicts an illustrative framing strategy used in connection with one embodiment of the invention for a digital audio stream comprising an audio work. More specifically, FIG. 4 depicts a framed digital audio stream 50 comprising a sequence of frames 51(1) through 51(F) (generally identified by reference numeral 51(f)). Each frame 51(f), in turn, comprises a series of audio samples 52(f)(1) through 52(f)(S) (generally identified by reference numeral 52(f)(s)) of the audio track. Preferably all of the frames will have the same number "S" of audio samples, although it will be appreciated from the following that that is primarily for convenience. On the other hand, it will be appreciated that, the number of audio samples may differ from "S"; this

may particularly be the case if the frame 51(f) contains the last audio samples for the digital audio stream for a particular audio work. In that case, the last frame 51(F) will preferably contain samples 52(F)(1) through 52(F)(x), where "x" is less than "S." Generally, it is desirable that the number of samples be consistent among all frames 51(f), and in that case padding, which will not be played, can be added to the last frame 51(F).

Associated with each frame 51(f) is a header 55(f) that includes a number of fields for storing other information that is useful in controlling playback of the audio samples in the respective frame 51(f). In particular, the header 55(f) associated with a frame 51(f) includes a frame sequence number field 56, an encoding type field 57, a sampling rate information field 58, a time stamp field 60, an end of track flag 61, and a length flag field 62. The header 55(f) may also include fields (not shown) for storing other information that is useful in controlling playback. Generally, the frame sequence number field 56 receives a sequence number "f" that identifies the relative position of the frame 51(f) in the sequence of frames 51(1)..51(f)...51(F) containing the digital audio stream 50. The encoding type field 57 receives a value that identifies the type of encoding and/or compression that has been used in generating the digital audio stream. Conventional encoding or compression schemes include, for example, the well-known MP3 and WAV encoding and/or compression schemes, although it will be appreciated that other schemes may be provided for as well. The sampling rate information field 58 receives sampling rate information that indicates the sampling rate for the audio samples 52(f)(s). As will be apparent to those skilled in the art, the sampling rate determines the rate at which the zone player 11(n) is to play the audio samples 52(f)(s) in the frame, and, as will be described below, determines the period of the digital to analog converter clock 34.

The condition of the end of work flag 61 indicates whether the frame 51(f) contains the last digital audio samples for the audio track associated with the framed digital audio work 50. If the frame 51(f) does not contain the audio samples that are associated with the end of the digital audio stream 50 for a respective audio work, the end of work flag will be clear. On the other hand, if the frame 51(f) does contain the audio samples that are associated with the end of the digital audio stream 50 for a respective audio work, the end of work flag 61 will be set. In addition, since the number of valid audio samples 52(F)(s) in the frame 51(F), that is, the samples that are not padding,

may be less than "S," the default number of audio samples in a frame 51(f), the length flag field 62 will contain a value that identifies the number of audio samples in 52(F)(s) in the last frame 51(F) of the audio work 50. If, as noted above, the frames have a consistent number "S" of samples, the samples 52(F)(x+1) through 52(F)(S) will contain padding, which will not be played.

The time stamp field 60 stores a time stamp that identifies the time at which the zone player 11(n) is to play the respective frame. More specifically, for each frame of a framed digital audio stream 50 that is buffered in the audio information buffer 21, the audio information source interface 30, using timing information from the digital to analog converter clock 34, will determine a time at which the zone player 11(n) is to play the respective frame, and stores a time stamp identifying the playback time in the time stamp field 60. The time stamp associated with each frame will later be used by the playback scheduler 32 to determine when the portion of the digital audio stream stored in the frame is to be coupled to the digital to analog converter 33 to initiate play back. It will be appreciated that the time stamps that are associated with frames in sequential frames 51(1), 51(2),...,51(F), will be such that they will be played back in order, and without an interruption between the sequential frames comprising the digital audio stream 50. It will further be appreciated that, after a time stamp has been determined for the first frame, stored in frame 51(1), of a digital audio stream 50, the audio information source interface 30 can determine time stamps for the subsequent frame 51(2), 51(3),....,51(F) in relation to the number of samples "S" in the respective frames and the sample rate. The time stamps will also preferably be such that frames will be played back after some slight time delay after they have been buffered in the audio information buffer 21; the purpose for the time delay will be made clear below.

Returning to FIG. 3, in addition to dividing the digital audio information into frames, the audio information source interface 30 also aggregates and/or divides the frames 51(f) as necessary into packets, each of which will be of a length that would fit into a message for transmission over the network, and associates each packet with a packet sequence number. For example, if a packet will accommodate multiple frames 51(f), 51(f+1),... 51(f+y-1), it will aggregate them into a packet and associate them with a packet number, for example p(x). If the entire frames 51(f) and 51(f+y-1) was accommodated in packet p(x), where "x" is the sequence number, which will occur if the size

of a packet is an exact multiple of the frame size, the next packet, p(x+1) will begin with frame 51(f+y) and will include frames 51(f+y),..., 51(f+2y-1). Subsequent packets p(x+2),.... will be formed in a similar manner. On the other hand, if the packet length will not accommodate an exact multiple of the frame size, the last frame in the packet will be continued at the beginning of the next packet.

If the audio information source interface 30 is aware of track boundaries, which may be the case if the tracks are divided into files, the packets will reflect the track boundaries, that is, the packets will not contain frames from two tracks. Thus, if the last frames associated with a track are insufficient to fill a packet, the packet will contain padding from the last frame associated with the track to the end of the packet, and the next packet will begin with the first frames associated with the next track.

In one embodiment, the audio information source interface 30 stores the packets in the audio information buffer 31 in a ring buffer. As is conventional, a ring buffer includes a series of storage locations in the buffer. Each entry will be sufficient to store one packet. Four pointers are used in connection with the ring buffer, a first pointer pointing to the beginning of the ring buffer, a second pointer pointing to the end of the ring buffer, an third "write" pointer pointing to the entry into which a packet will be written and a fourth "read" pointer pointing to the entry from which packet will be read for use in playback. When a packet is read from the ring buffer for playback, it will be read from the entry pointed to by the read pointer. After the packet has been read, the read pointer will be advanced. If the read pointer points beyond the end of the ring buffer, as indicated by the end pointer, it will be reset to point to the entry pointed to by the beginning pointer, and the operations can be repeated.

On the other hand, when the audio information source interface 30 stores a packet in the ring buffer, first determine whether the entry pointed to by the write pointer points to the same entry as the entry pointed to by the read pointer. If the write pointer points to the same entry as the entry pointed to by the read pointer, the entry contains at least a portion of a packet that has not yet been read for playback, and the audio information source interface 30 will delay storage of the packet until

the entire packet has been read and the read pointer advanced. After the read pointer has been advanced, the audio information source interface 30 can store the packet in the entry pointed to by the write pointer. After the packet has been stored, the audio information source interface 30 will advance the write pointer. If the write pointer points beyond the end of the ring buffer, as indicated by the end pointer, it will be reset to point to the entry pointed to by the beginning pointer, and the operations can be repeated.

As noted above, the zone player 11(n) can operate both as an audio information channel device 23 and a member of the synchrony group 20 of which it is a member. In that case, the audio information buffer 31 can contain one ring buffer. On the other hand, the zone player 11(n) can operate as an audio information channel device 23 for one synchrony group 20(1) (FIG. 2A) and a member of another synchrony group 20(2). In that case, the audio information buffer 31 would maintain two ring buffers, one for the audio and timing information associated with synchrony group 20(1), and the other for the audio and timing information associated with synchrony group 20(2). It will be appreciated that, in the latter case, the zone player 11(n) will only use the audio and timing information that is associated with synchrony group 20(2) for playback.

The playback scheduler 32 schedules playback of the audio information that is buffered in the audio information buffer 31 that is to be played by the zone player 11(n). Accordingly, under control of the playback scheduler 32, the digital audio information that is buffered in the audio information buffer 21 that is to be played by the zone player 11(n) is transferred to the digital to analog converter 33 for playback. As noted above, if the zone player 11(n) is operating as an audio information channel device 23 for a synchrony group 20 for which it is not a member, the playback scheduler 32 will not schedule the digital audio information that is to be played by that synchrony group 20 for playback. The playback scheduler 32 only schedules the digital audio information, if any, that is buffered in the audio information buffer 31 that is associated with a synchrony group for which the zone player 11(n) is a member, whether as master device 21 or a slave device 22(g).

Essentially, the playback scheduler 32 makes use of the read pointer associated with the circular buffer that contains the audio and playback timing information that is to be played by the

zone player 11(n). The playback scheduler 32 retrieves the packet information from the entry of the ring buffer pointed to by the read pointer, and then advances the ring pointer as described above. The playback scheduler 32 determines the boundaries of the frames in the packet and uses the time stamps in the time stamp fields 60 associated with the respective frame 51(f), along with timing information provided by the zone player 11(n)'s digital to analog converter clock 34, to determine when the respective frame is to be transferred to the digital to analog converter 33. Generally, when the time stamp associated with a buffered digital audio information frame corresponds to the current time as indicated by the digital to analog converter clock 34, the playback scheduler 32 will enable the respective frame to be transferred to the digital to analog converter 33.

The digital to analog converter 33, also under control of the digital to analog converter clock 34, converts the buffered digital audio information to analog form, and provides the analog audio information to the audio amplifier 35 for amplification. The amplified analog information, in turn, is provided to the audio reproduction devices 15(n)(r) through the audio reproduction device interface 36. The audio reproduction devices 15(n)(r) transform the analog audio information signal to sound thereby to provide the audio program to a listener. The amount by which the audio amplifier 35 amplifies the analog signal is controlled by the control module 42, in response to volume control information provided by the user through the user interface module 13.

The network communications manager 40 controls network communications over the network 12, and the network interface 41 transmits and receives message packets over the network 12. The network communications manager 40 generates and receives messages to facilitate the transfer of the various types of information described above in connection with FIG. 2, including the channel device control information, slave device control information, audio and playback timing information and the audio information channel device's clock timing information. In connection with the channel device control information and the slave device control information, the network communications manager 40 will generate messages for transfer over the network 12 in response to control information from the control module 42. Similarly, when the network communications manager 40 receives messages containing channel device control information and slave device

control information, the network communications manager will provide the information to the control module 42 for processing.

With regards to the audio information channel device's clock timing information, as noted above, the master device 21 and slave devices 22(g) of the synchrony group 20 obtain the clock timing information from the audio information channel device 23 using the well-known SNTP. If the zone player 11(n) is operating as the audio information channel device 23 for a synchrony group, during the SNTP operation, it will provide its current time, particularly a current time as indicated by its digital to analog converter clock 34. On the other hand, if the zone player 11(n) is operating as the master device 21 or slave device 22(g) of a synchrony group 20, it will receive the clock timing information from the audio information channel device 23. After the respective device 21, 22(g) has obtained the audio information channel device's clock timing information, it will generate a differential time value ΔT representing the difference between the time T indicated by its digital to analog converter clock 34 and the current time information from the audio information channel device 23. The differential time value will be used to update the time stamps for the frames of the digital audio stream 50 (FIG. 4) that are received from the audio information channel device.

With regards to the audio and playback timing information, operations performed by the network communications manager 40 will depend on whether

(i) the audio and playback timing information has been buffered in the audio information buffer 31 for transmission, as audio information channel device 23, over the network 12 to the master device 21 and/or slave devices 22(g) of a synchrony group, or

(ii) the audio and playback timing information has been received from the network 12 to be played by the zone player 11(n) as either the master device 21 for a synchrony group or a slave device in a synchrony group.

It will be appreciated that the network communications manager 40 may be engaged in both (i) and (ii) contemporaneously, since the zone player 11(n) may operate both as the audio information channel device 23(1) for a synchrony group 20(1) (reference FIG. 2A) of which it is not a member,

and a member of another synchrony group 20(2) for which another zone player 11(n') is the audio information channel device 23(2). With reference to item (i) above, after a packet that is to be transmitted has been buffered in the respective ring buffer, the network communications manager 40 retrieves the packet, packages it into a message and enables the network interface 41 to transmit the message over the network 12. If the control module 42 receives control information from the user interface module 13 (if the master device 21 is also the audio information channel device 23 for the synchrony group 20) or from the master device (if the master device 21 is not the audio information channel device 23 for the synchrony group 20) that would require the transmission of a "resynchronize" command as described above, the control module 42 of the audio information channel device 23 enables the network communications manager 40 to insert the command into a message containing the audio and playback timing information. Details of the operations performed in connection with the "resynchronize" command will be described below. As noted above, the "resynchronize" command is used if the user enables a synchrony group to terminate the playback of a track that is currently being played, or cancel playback of a track whose playback has not begun.

On the other hand, with reference to item (ii) above, if network interface 41 receives a message containing a packet containing frames of audio and playback timing information that the zone player 11(n) is to play either as a master device 21 or a slave device for a synchrony group 20, the network interface 41 provides the audio and playback timing information to the network communications manager 40. The network communications manager 40 will determine whether the packet contains a resynchronize command and, if so, notify the control module 42, which will enable operations to be performed as described below. In any case, the network communications manager 40 will normally buffer the various frames comprising the audio and playback timing information in the audio information buffer 31, and in that operation will generally operate as described above in connection with the audio information source interface 30. Before buffering them, however, the network communications manager 40 will update their time stamps using the time differential value described above. It will be appreciated that the network communications manager 40 will perform similar operations whether the messages that contain the packets were multi-cast messages or unicast messages as described above

The updating of the time stamps by the master device 21 and the slave devices 22(g) in the synchrony group 20 will ensure that they all play the audio information synchronously. In particular, after the network communications manager 40 has received a frame 51(f) from the network interface 41, it will also obtain, from the digital to analog converter clock 34, the zone player 11(n)'s current time as indicated by its digital to analog converter clock 34. The network communications manager 40 will determine a time differential value that is the difference between the slave device's current clock time, as indicated by its digital to analog converter 34, and the audio information channel device's time as indicated by the audio information channel device's clock timing information. Accordingly, if the master or slave device's current time has a value T_S and the audio information channel device's current time, as indicated by the clock timing information, has a value T_C, the time differential value $\Delta T=T_S-T_C$. If the current time of the master or slave device in the synchrony group 20, as indicated by its digital to analog converter clock 34, is ahead of the audio information channel device's clock time as indicated by the clock timing information received during the SNTP operation, the time differential value will have a positive value. On the other hand, if the master or slave device's current time is behind the audio information channel device's clock time, the time differential value ΔT will have a negative value. If the zone player 11(n) obtains clock timing information from the audio information channel device 23 periodically while it is a member of the synchrony group 20, the network communications manager 40 can generate an updated value for the time differential value ΔT when it receives the clock timing information from the audio information channel device 23, and will subsequently use the updated time differential value.

The network communications manager 40 uses the time differential value ΔT that it generates from the audio information channel device timing information and zone player 11(n)'s current time to update the time stamps that will be associated with the digital audio information frames that the zone player 11(n) receives from the audio information channel device. For each digital audio information frame that is received from the audio information channel device, instead of storing the time stamp that is associated with the frame as received in the message in the audio information buffer 21, the network communications manager 40 will store the updated time stamp with the digital audio information frame. The updated time stamp is generated in a manner so that, when the

zone player 11(n), as a member of the synchrony group plays back the digital audio information frame, it will do so in synchrony with other devices in the synchrony group.

More specifically, after the zone player 11(n)'s network interface 41 receives a message containing a packet that, in turn, contains one or more frames 51(f), it will provide the packet to the network communications manager 40. For each frame 51(f) in the packet that the network communications manager 40 receives from the network interface 41, the network communications manager 40 will add the time differential value ΔT to the frame's time stamp, to generate the updated time stamp for the frame 51(f), and store the frame 51(f), along with the header 55(f) with the updated time stamp in the audio information buffer 31. Thus, for example, if a frame's time stamp has a time value T_F, the network communications manager 40 will generate an updated time stamp T^U_F having a time value $T^U_F = T_F + \Delta T$. Since time value T^U_F according to the slave device's digital to analog converter clock 34 is simultaneous to the time value T_F according to the audio information channel device's digital to analog converter clock 34, the zone player 11(n) device will play the digital audio information frame at the time determined by the audio information channel device 23. Since all of the members of the synchrony group 20 will perform the same operations, generating the updated time stamps T^U_F for the various frames 51(f) in relation to their respective differential time values, all of the zone players 11(n) in the synchrony group 20 will play them synchronously. The network communications manager 40 will generate updated time stamps T^U_F for all of the time stamps 60 in the packet, and then store the packet in the audio information buffer 31.

It will be appreciated that, before storing a packet in the audio information buffer 21, the network communications manager 40 can compare the updated time stamps T^U_F associated with the frames in the packet to the slave device's current time as indicated by its digital to analog converter clock 34. If the network communications manager 40 determines the time indicated by the updated time stamps of frames 51(f) in the packet are earlier than the zone player's current time, it can discard the packet instead of storing it in the audio information buffer 21, since the zone player 11(n) will not play them. That is, if the updated time stamp has a time value T^U_F that identifies a time that is earlier than the zone player's current time T_S as indicated by the zone player's digital to analog converter clock 34, the network communications manager 40 can discard the packet.

If the zone player 11(n) is operating as the master device 21 of a synchrony group 20, when the user, through the user interface module 13, notifies the zone player 11(n) that another zone player 11(n') is to join the synchrony group 20 as a slave device 22(g), the control module 42 of the zone player 11(n) enables the network communications manager 40 to engage in an exchange of messages, described above in connection with FIG. 2, to enable the other zone player 11(n') to join the synchrony group 20 as a slave device. As noted above, during the message exchange, the messages generated by the network communications manager 40 of the zone player 11(n) will provide the network communications manager of the zone player 11(n') that is to join the synchrony group 20 with information such as the multi-cast address being used by the audio information channel device 23 that is providing the audio program to the synchrony group 20, as well as a unicast network address for the audio information channel device 23. After receiving that information, the network communications manager and network interface of the zone player 11(n') that is to join the synchrony group 20 can begin receiving the multi-cast messages containing the audio program for the synchrony group, engage in SNTP transactions with the audio information channel device 23 to obtain the latter's current time, and also enable the audio information channel device 23 to send the zone player 11(n') frames 51(f) that it had previously broadcast using the unicast message transmission methodology as described above.

On the other hand, if the network communications manager 40 and network interface 41 of the zone player 11(n) receive a message over the network 12 indicating that it is to become a slave device 22(g) of a synchrony group for which another zone player 11(n') is the master device, the network communications manager 40 for zone player 11(n) will provide a notification to the control module 42 of zone player 11(n). Thereafter, the control module 42 of zone player 11(n) can enable the network communications manager 40 of zone player 11(n) to perform the operations described above to enable it to join the synchrony group 20.

As noted above, the user, using user interface module 13, can enable the synchrony group to terminate playback of a track of an audio program that is currently being played. After playback of a track that is currently being played has been terminated, playback will continue in a conventional manner with the next track that has been buffered in the audio information buffer 31.

It will be appreciated that that could be the next track that is on the original play list, or a previous track. In addition, the user can enable the synchrony group 20 to cancel playback of a track that it has not yet begun to play, but for which buffering of packets has begun in the synchrony group 20. Both of these operations make use of the "resynchronize" command that the master device 21 of the synchrony group 20 can enable the audio information channel device 23 to include in the multi-cast message stream that it transmits to the synchrony group 20. Generally, in response to receipt of the resynchronize command, the members of the synchrony group 20 flush the ring buffer containing the packets that they are to play in the future. In addition, if the members of the synchrony group provide separate buffers for their respective digital to analog converters 33, the members will also flush those buffers as well. After the audio information channel device transmits a packet containing the resynchronize command:

(i) in the case of the use of the resynchronize command to terminate playing of a track that is currently being played, the audio information channel device 23 will begin multi-casting packets for the next track, to begin play immediately, and will continue through the play list in the manner described above; and

(ii) in the case of the use of the resynchronize command to cancel play of a track for which buffering has begun, but which is to be played in the future, the audio information channel device 23 will begin multi-casting packets for the track after the track that has been cancelled, to be played beginning at the time the cancelled track was to begin play, and will also continue through the play list in the manner described above.

It will be appreciated that,

(a) in the first case (item (i) directly above), the resynchronize command can enable the read pointer to be set to point to the entry in the circular buffer into which the first packet for the next track will be written, which will correspond to the entry to which the write pointer points, but

(b) in the second case (item (ii) directly above), the resynchronize command can enable the write pointer for the circular buffer to be set to point to the entry that contains the first packet for the track whose play is being cancelled.

It will further be appreciated that, if a track is cancelled for which buffering has not begun, the resynchronize command will generally not be necessary, since the audio information channel device 23 for the synchrony group 20 need only delete that track from the play list.

Operations performed in connection with use of the resynchronize command to cancel playback of a track that is currently being played will be described in connection with Packet Sequence A below, and operations performed in connection with the use of the resynchronize command to cancel playback of a track that is has not yet begun to play, but for which buffering of packets has begun, will be described in connection with Packet Sequence B below.

Packet Sequence A

(A1.0) [packet 57]
(A1.1 [continuation of frame 99]
(A1.2) [frame 100, time = 0:00:01, type = mp3 audio]
(A1.3) [frame 101, time = 0:00:02, type = mp3 audio]
(A1.4) [frame 102, time = 0:00:03, type = mp3 audio]

(A2.0) [packet 58]
(A2.1) [continuation of frame 102]
(A2.2) [frame 103, time = 0:00:04, type = mp3 audio]
(A2.3) [frame 104, time = 0:00:05, type = mp3 audio]
(A2.4) [frame 105, time = 0:00:06, type = mp3 audio]

(A3.0) [packet 59]
(A3.1) [continuation of frame 105]
(A3.2) [frame 106, time = 0:00:07, type = mp3 audio]

(A3.3) [frame 107, time = 0:00:08, type = mp3 audio]
(A3.4) [frame 108, time = 0:00:09, type = mp3 audio]

(A4.0) [packet 60]
(A4.1) [continuation of frame 108]
(A4.2) [frame 109, time = 0:00:10, type = mp3 audio]
(A4.3) [Resynchronize command]
(A4.4) [Padding, if necessary]

(A5.0) [packet 61]
(A5.1) [frame 1, time = 0:00:07, type = mp3 audio]
(A5.2) [frame 2, time = 0:00:08, type = mp3 audio]
(A5.3) [frame 3, time = 0:00:09, type = mp3 audio]
(A5.4) [frame 4, time = 0.00.10, type = mp3 audio]

(A6.0) [packet 62]
(A6.1) [continuation of frame 4]
(A6.2) [frame 5, time = 0:00:11, type = mp3 audio]
(A6.3) [frame 6, time = 0:00:12, type = mp3 audio]
(A6.4) [frame 7, time = 0:00:13, type = mp3 audio]

Packet Sequence A comprises a sequence of six packets, identified by packet 57 through packet 62, that the audio information channel device 23 multi-casts in respective messages to the members of a synchrony group 20. It will be appreciated that the series of messages that the audio information channel device 23 may multi-cast to the synchrony group 20 may include a messages prior to the packet 57, and may also include messages after packet 62. Each packet comprises a packet header, which is symbolized by lines (A1.0), (A2.0),...(A6.0) in Packet Sequence A, and will generally also include information associated with at least a portion of a frame. In the packets represented in Packet Sequence A, each packet includes information associated with a plurality of

frames. Depending on the lengths of the packets, each packet may contain information associated with a portion of a frame, an entire frame, or multiple frames. In the illustration represented by Packet Sequence A, it is assumed that each packet may contain information associated with multiple frames. In addition, it is assumed that a packet does not necessarily contain information associated with an integral number of frames; in that case, a packet may contain information associated with a portion of a frame, and the next packet will contain the information associated with the rest of a frame.

The frames and associated header playback timing information contained in the various packets are symbolized by lines (A1.1), (A1.2),...,(A1.4), (A2.1),....(A6.4) of Packet Sequence A. Thus, for example, line (A1.2) of packet 57 represents the one-hundredth frame, that is, frame 51(100) (reference FIG. 4), of the track whose audio information is being transmitted in the sequence of packets that includes packet 57. The frame 51(100) is to be played at an illustrative time, according to the audio information channel device's digital to analog converter clock, of "time = 0:00:01," and the frame is encoded and/or compressed using the well-known MP3 encoding and compression methodology. In that case, the legend"time = 0:00:01" represents the time stamp that would be included in field 60 (FIG. 4) of the header associated with the frame 50(100) as multi-cast by the audio information channel device for the synchrony group. It will be appreciated that the playback time and encoding/compression methodology will be referred in the header 55(100) that is associated with the frame 51(100). It will also be appreciated that the header may also contain additional information as described above.

Similarly, line (A1.3) of packet 57 represents the one-hundred and first frame, that is, frame 51(101), of the track whose audio information is being transmitted in the sequence of packets that includes packet 57. The frame 51(101) is to be played at an illustrative time, according to the audio information channel device's digital to analog converter clock, of "0:00:02," and the frame is also encoded and/or compressed using the MP3 encoding and compression methodology. Line (A1.4) of packet 57 represents similar information, although it will be appreciated that, depending on the length of packet 57, the line may not represent the information for an entire frame 51(102) and/or its associated header. If the length of packet 57 is not sufficient to accommodate the information for

the entire frame 51(102) and/or associated header, the information will continue in packet 58, as represented by line (A2.1) in Packet Sequence A. Similarly, if the length of packet 56 was not sufficient to contain the information for an entire frame 51(99) preceding frame 51(100), packet 57 (lines (A1.0) through 1.4) may contain any information from frame 51(99) that packet 56 was unable to accommodate.

As noted above, when the master device 21 or a slave device 22(g) in the synchrony group 20 receives the packet 57, its respective network communications manager 40 will update the time stamps associated with the various frames 51(f) as described above before buffering the respective frames in the respective audio information buffer 31.

Packets 58 and 59 contain information that is organized along the lines described above in connection with packet 57.

Packet 60 also contains, as represented by lines (A4.1) and (A4.2), information that is organized along the lines of the information represented by lines (Ax.1) and (Ax.2) ("x" equals an integer) described above in connection with packets 57 through 59. On the other hand, packet 60 contains a resynchronize command, as represented by line (A4.3). Packet 60 also may contain padding, as represented by line 4.4, following the resynchronize command. As noted above, the master device 21 of a synchrony group 20 will enable the audio information channel device 23 that is providing audio information to the synchrony group 20 to multi-cast a message containing the resynchronize command when it receives notification from the user interface module 13 that the user wishes to cancel playback of a track that is currently being played. In the example depicted in Packet Sequence A, as will be described below, the audio information channel device 23 receives notification from the master device 21 that the user wishes to cancel playback of a track at a time corresponding to "time = 0:00:07" according to its digital to analog converter clock 34, and, in line (A4.3) of packet 60 it will provide the resynchronize command, followed by padding, if necessary.

As will be apparent from examining lines (A3.1) through (A3.4) of packet 59 and lines (A4.1) and (A4.2) of packet 60, although the audio information channel device 23 has received the notification from the synchrony group's master device 21 to multi-cast the resynchronize command

at a time corresponding to "time = 0:00:07" according to the clock time indicated by its digital to analog converter clock 34, it (that is, the audio information channel device 23) has already multi-cast messages containing frames that are to be played at that time and subsequently. That is, the audio information channel device 23 has, multi-cast in packet 59, frames 51(106) through 51(108) that contain time stamps "time = 0:00:07," "time = 0:00:08" and "time = 0:00:09," respectively, and, in packet 60, in addition to the continuation of frame 51(108), frame 51(109) that contains time stamp "time = 0:00:10." (It will be appreciated that the times indicated by the illustrative time stamps are for illustration purposes only, and that in an actual embodiment the time stamps may have different values and differentials.)

As noted above, the audio information channel device 23 multi-casts a message containing a packet that, in turn, contains the resynchronize command when it receives the notification from the master device 21 to do so. In the example depicted in Packet Sequence A, the packet will be multi-cast when the audio information channel device's digital to analog converter clock time corresponds to "0:00:07." Subsequently, two things happen. In one, aspect, when the master device 21 and slave devices 22(g) receive the packet that contains the resynchronize command, they will stop playing the audio program that they are playing.

In addition, the audio information channel device 23 will begin transmitting frames containing audio information for the next track, including therewith time stamps immediately following the digital to analog converter clock time at which the packet including the resynchronize command was transmitted. Accordingly, and with further reference to Packet Sequence A, the audio information channel device 23 will multi-cast a message containing packet 61. As indicated above, packet 61 contains, as represented in lines (A5.1) through (A5.3), frames 51(1) through 51(3), which are the first three frames of the next track of the audio program that is to be played. It is also compressed and encoded using the MP3 encoding and compression scheme, and it is accompanied by time stamps "time = 0:00:07," "time = 0:00:08" and "time = 0:00:10." As noted above, the time stamp "time = 0:00:07" corresponds to the clock time at which the audio information channel device 23 multi-casts the resynchronize command, and, when the master device 21 and slave devices 22(g) receive these frames, they would be expected to begin playing them very shortly, if not immediately

after the audio information channel device 23 multi-casts the message containing the packet that, in turn, contains the resynchronize command. Packet 61 also includes at least a portion of the next frame, that is, frame 51(4), for that track. In addition, Packet Sequence A depicted above further includes a subsequent packet, namely, packet 62, that contains any necessary continuation of frame 51(4), as well as three subsequent frames. If any additional packets are required for the track, as well as for subsequent tracks, they can be multi-cast in a similar manner.

As further noted above, the resynchronize command can also be used to cancel playing of one or more tracks for which playback has begun. This will be illustrated in connection with Packet Sequence B:

Packet Sequence B

(B1.0) [packet 157]
(B1.1) [continuation of frame 99]
(B1.2) [frame 100, time = 0:00:01, type = mp3 audio]
(B1.3) [frame 101, time = 0:00:02, type = mp3 audio]
(B1.4) [frame 102, time = 0:00:03, type = mp3 audio]

(B2.0) [packet 158]
(B2.1) [continuation of frame 102]
(B2.2) [frame 103, time = 0:00:04, type = mp3 audio]
(B2.3) [frame 104, time = 0:00:05, type = mp3 audio]
(B2.4) [frame 105, time = 0:00:06, type = mp3 audio]

(B3.0) [packet 159]
(B3.1) [continuation of frame 105]
(B3.2) [frame 106, time = 0:00:07, type = mp3 audio]
(B3.3) [track boundary notification]
(B3.4) [Padding, if necessary]

(B4.0) [packet 160]
(B4.1) [frame 1, time = 0:00:08, type = mp3 audio]
(B4.2) [frame 2, time = 0:00:09, type = mp3 audio]
(B4.3) [frame 3, time = 0:00:10, type = mp3 audio]

(B5.0) [packet 161]
(B5.1) [continuation of frame 3]
(B5.2) [frame 4, time = 0:00:11, type = mp3 audio]
(B5.3) [Resynchronize, after packet 159]
(B5.4) [Padding, if necessary]

(B6.0) [packet 162]
(B6.1) [frame 1, time = 0:00:08, type = mp3 audio]
(B6.2) [frame 2, time = 0:00:09, type = mp3 audio]
(B6.3) [frame 3, time = 0:00:10, type = mp3 audio]
(B6.4) [frame 4, time = 0:00:11, type = mp3 audio]

(B7.0) [packet 163]
(B7.1) [continuation of frame 4]
(B7.2) [frame 5, time = 0:00:12, type = mp3 audio]
(B7.3) [frame 6, time = 0:00:13, type = mp3 audio]
(B7.4) [frame 7, time = 0:00:14, type = mp3 audio]

Packet Sequence B comprises a series of seven packets, identified by packet 157 through 163, that the audio information channel device 23 multi-casts to the members of a synchrony group 20. As with Packet Sequence A, it will be appreciated that the series of packets that the audio information channel device 23 may multi-cast to the synchrony group 20 may include packets prior to the packet 157, and may also include packets after packet 162. Each packet comprises a packet

header, which is symbolized by lines (B1.0), (B2.0),...(B7.0) in Packet Sequence B. As in Packet Sequence A, each packet will also generally include information associated with at least a portion of a frame 51(f)along with its associated frame 55(f). As in the packets represented in Packet Sequence A, each packet includes information associated with a plurality of frames. Depending on the lengths of the packets, each packet may contain information associated with a portion of a frame, an entire frame, or multiple frames. Further, as with Packet Sequence A, it is assumed that each packet may contain information associated with multiple frames. In addition, it is assumed that a packet does not necessarily contain information associated with an integral number of frames; in that case, a packet may contain information associated with a portion of a frame, and the next packet will contain the information associated with the rest of a frame.

The structures of the packets represented by Packet Sequence B are similar to those described above in connection with Packet Sequence A, and will not be repeated here. Generally, Packet Sequence B illustratively contains a sequence of packets that represent at least portions of three tracks that may have been selected from, for example, a play list. In particular, packets 157 through 159 represent frames from a portion of one track, packets 160 and 161 represent frames from a second track and packets 162 and 163 represent frames from a third track. The play list indicated that the first, second and third tracks were to be played in that order. With particular reference to Packet Sequence B, it should be noted that line (B3.3) indicates that packet 159 includes an indication that that packet contains the last frame for the track, and line (B3.4) provides for padding to the end of the packet. The first frame of the next track begins in packet 160.

In connection with the use of the resynchronize command to cancel playback of a track, at least a portion of which the audio information channel device 23 has multi-cast to the members of the synchrony group, packet 161, in line (B5.3) represents a resynchronize command that indicates that resynchronization is to occur after packet 159, that is, immediately after the packet that contains the last frame of the first of the three tracks represented by the packets in Packet Sequence B. It should be noted that the resynchronize command is in the packet 161, while the resynchronization is to occur at packet 160, that is, the synchrony group is to not play the track starting with packet 160, but instead is to begin playing the track frames for which begin with the next packet, that is,

packet 162. As with Packet Sequence A, in Packet Sequence B the audio information channel device 23, in packet 162 and 163, multi-casts frames whose time stamps indicate that they are to be played when the frames that were multi-cast in packets 160 and 161 were to be played. By use of the resynchronize command and specifying a packet in this manner, the audio information channel device can cancel playback of a track for which playback has not yet begun.

It will be appreciated that the resynchronize command is generally not necessary for cancelling play back of a track that the audio information channel device 23 has not started multi-casting to the synchrony group 20, since the audio information channel device 23 itself can re-order the play list to accommodate the cancellation.

The invention provides a number of advantages. In particular, the invention provides a network audio system in which a number of devices share information can reproduce audio information synchronously, notwithstanding the fact that packets, which may contain digital audio information, transmitted over the network to the various zone players connected thereto may have differing delays and the zone players operate with independent clocks. Moreover, although the invention has been described in connection with audio information, it will be appreciated that the invention will find utility in connection with any type of isochronous information for which synchrony among devices is desired. The system is such that synchrony groups are created and destroyed dynamically, and in such a manner as to avoid requiring a dedicated device as the master device.

It will be appreciated that a number of changes and modifications may be made to the network audio system 10 as described above. For example, although the invention has been described as providing that the audio information channel device 23 provides digital audio information to the members synchrony group 20 that has been encoded using particular types of encoding and compression methodologies, it will be appreciated that the audio information channel device 23 can provide digital audio information to various members of the synchrony group 20 that have been encoded and compressed using different types of encoding and compression methodologies, and, moreover, for which different sampling rates have been used. For example, the

audio information channel device 23 may provide digital audio information to the master device 21 and slave devices 22(1) through 22(g_1) using the MP3 methodology at a specified sampling rate, the digital audio information for the same program to slave devices 22(g_1+1) through 22(g_2) using the WAV methodology at one specified sampling rate, and to slave devices 22(g_2+1) through 22(G) using the WAV methodology at another specified sampling rate. In that case, the audio information channel device 23 can specify the particular encoding and compression methodology that has been used in the encoding type field 57 associated with each frame and the sampling rate in the sampling rate field 58. Moreover, since the encoding and compression type and sampling rate are specified for each frame, the encoding and compression type and sampling rate can be changed from frame to frame. The audio information channel device 23 may use different multi-cast addresses for the different encoding and compression types and sampling rates, but it will be appreciated that that would not be required.

It will be appreciated that two advantages of providing that the encoding and compression methodology and the sampling rate is provided on a frame-by-frame basis, instead of on, for example, a track-by-track basis, is that that would facilitate a slave device joining the synchrony group 20 at a frame mid-track, without requiring, for example, the master device 21 or the audio information channel device 23 to notify it of the encoding and compression methodology or the sampling rate.

Another modification is that, instead of the network communications manager 40 of a member of a synchrony group 20 generating the updated time stamp T^u_F for a digital audio information frame by adding the time differential value ΔT to the time stamp T_F associated with a frame, the network communications manager 40 may instead generate the updated time stamp T^U_F by subtracting the differential time value ΔT from the member's current time T_S as indicated by the member's digital to analog converter clock 34 at the time at which the digital audio information is received. It will be appreciated, however, that there may be variable time delays in processing of messages by the slave device's network communications manager 40, and so it may be preferable to generate the time differential value ΔT using the time stamp T_F provided by the audio information channel device 23.

In addition, instead of the network communications manager 40 of a member of a synchrony group generating an updated time stamp to reflect the difference between the times indicated by the member's digital to analog converter clock and the audio information channel device's digital to analog converter clock, the network communications manager 40 can generate the time differential value ΔT and provide it to the member's playback scheduler 32. In that case, the member's network communications manager 40 can store each digital audio information frame along with the time stamp T_F as received from the master device in the audio information buffer 21. The playback scheduler 32 can utilize the time differential value ΔT, and the time stamps T_F associated with the digital audio information frames, to determine when the respective digital audio information frames are to be played. In determining when a digital audio information frame is to be played, the playback scheduler can add the time differential value to the time stamp T_F associated with the digital audio frame, and enable the digital audio frame to be coupled to the digital to analog converter 33 when the time indicated by the sum corresponds to the current time as indicated by the slave device's digital to analog converter clock 34. Alternatively, when the member's digital to analog converter clock 34 updates its current time T_S, the playback scheduler can generate an updated current time T'_S by subtracting the differential time value ΔT from the current time T_S, and using the updated current time T'_S to determine when to play a digital audio information frame.

As described above, the members of a synchrony group 20 periodically obtain the audio information channel device's current time value and uses the current time value that it receives from the audio information channel device to periodically update the time differential value ΔT that it uses in updating the time stamps associated with the various frames. It will be appreciated that, if the digital to analog converter clock(s) associated with the member(s) of a synchrony group 20 are ensured to have the same rate as the digital to analog converter clock, a member need only obtain the current time value from the audio information channel device once, at the beginning of playback.

As another alternative, if the zone players are provided with digital to analog converter clock 34 whose time and rate can be set by an element such as the network communications manager 40, when a zone player 11(n) is operating as a member of a synchrony group 20, its network communications manager 40 can use the various types of timing information that it receives from

the audio information channel device 23, including the current time information and the playback timing information indicated by the time stamps that are associated with the various frames 51(f) comprising the audio and playback timing information that it receives, to adjust the synchrony group member's digital to analog converter clock's time value and/or the clock rate that it uses for playback. If the clock's time value is to be adjusted, when the synchrony group member's network communications manager 40 initially receives the current time information from the audio information channel device 23 for the synchrony group 20, the network communications manager 40 can set the synchrony group member's digital to analog converter clock 34 to the current time value as indicated by the audio information channel device's current time information. The network communications manager 40 can set the clock 34 to the current time value indicated by the audio information channel device's current time information once, or periodically as it receives the current time information.

Alternatively or in addition, the synchrony group member's network communications manager 40 can use one or both of the current time information and/or the playback timing information in the time stamps associated with the respective frames 51(f) to adjust the clock rate of the clock 34 that it uses for playback. For example, when the synchrony group member's network communications manager 40 receives a frame $51(f_X)$ having a time stamp having a time value T_{f_X}, it can generate the updated time value $T^U_{f_X} = T_{f_X} + \Delta T$ as described above, and store the frame with the time stamp with the updated time value in the audio information buffer 30. In addition, since both the number of samples in a frame and the sampling rate, which determines the rate at which the frame is to be played, are known to the network communications manager 40, it can use that information, along with the updated time value $T^U{}_{F_X}$ that is to be used for frame $51(f_X)$ to generate an expected updated time value $T^E{}_{f_{X+1}}$ that is expected for the updated time stamp of the next frame $51(f_{X+1})$. After the synchrony group member's network communications manager 40 receives the next frame $51(f_{X+1})$, it can generate the updated time value $T^U{}_{f_{X+1}}$ and compare that

value to the expected updated time value $T^E_{f_{X+1}}$. If the two time values do not correspond, or if the difference between them is above a selected threshold level, the clock that is used by the audio information channel device 23 to generate the time stamps is advancing at a different rate than the synchrony group member's digital to analog converter clock 34, and so the network communications manager 40 can adjust the rate of the digital to analog converter clock 34 to approach that of the clock used by the audio information channel device 23 so that the differential time value ΔT is constant. On the other hand, if the two time values do correspond, then the time differential value ΔT is constant, or the difference is below a threshold level, and the network communications manager 40 need not change the clock rate of the digital to analog converter clock 34. It will be appreciated that, if the clock rate is to be adjusted, the rate adjustment can be fixed, or it can vary based on, for example, the difference between the updated time value $T^U_{f_{X+1}}$ and the expected updated time value $T^E_{f_{X+1}}$.

It will also be appreciated that, if no rate adjustment is performed for one frame $51(f_{X+1})$, the synchrony group member's network communications manager 40 can generate an expected updated time value $T^E_{f_{X+2}}$ that is expected for the updated time stamp of the next frame $51(f_{X+2})$ using the updated time value $T^U_{F_X}$ determined for frame $51(f_X)$, along with the number of samples in a frame and the sampling rate, and compare the expected updated time value $T^E_{f_{X+2}}$ to the updated time value $T^U_{f_{X+2}}$ that it generates when it receives frame $51(f_{X+2})$. At that point, if the network communications manager 41 determines that two time values do not correspond, or if the difference between them is above a selected threshold level, it can adjust the rate of the digital to analog converter clock 34. Similar operations can be performed if no rate adjustment is performed for several successive frames $51(f_{X+1})$, $51(f_{X+2})$,.... This will accommodate the possibility that the rate differential between the clock 34 and the clock used by the audio information channel device 23 in

generating the time stamps have rates that differ by an amount sufficiently small that it cannot be detected using time stamps of two or more successive frames.

Instead or in addition to adjusting the clock rate as described above, the synchrony group member's network communications manager 40 can perform similar operations in connection with adjusting the clock rate in connection with the current time information that it receives from the audio information channel device 23.

Furthermore, although the network audio system 10 has been described such that the master device 21 of a synchrony group 20 can, in response to control information provided thereto by a user through the user interface module 13, provide a notification to a zone player 11(n) that it is to become a member of its synchrony group 20 as a slave device 22(g), it will be appreciated that the user interface module 13 can provide the notification directly to the zone player 11(n) that is to become a member of the synchrony group 20. In that case, the zone player 11(n) can notify the master device 21 that it is to become a slave device 22(g) in the synchrony group 20, after which the master device 21 can provide information regarding the synchrony group 20, including the multi-cast and unicast addresses of the audio information channel device and other information as described above.

Similarly, although the network audio system 10 has been described such that the master device 21 of a synchrony group 20 can, in response to control information provided thereto by a user through the user interface module 13, provide a command to a slave device 22(g) to enable the slave device 22(g) to adjust its volume, it will be appreciated that the user interface module 13 can provide control information directly to the slave device 22(g) to enable the slave device 22(g) to adjust its volume.

In addition, although the network audio system 10 has been described such that each frames 51(f) is associated with a frame sequence number (reference field 56, FIG. 4), it will be appreciated that, if the packets described above in connection with Packet Sequence A and Packet Sequence B are provided with packet sequence numbers, the frame sequence numbers need not be provided, since the packet sequence numbers can suffice for defining the frame sequencing.

Furthermore, although the network audio system 10 has been described such that the zone players 11(n) are provided with an audio amplifier 35 for amplifying the analog signal provided by the respective digital to analog converters 33, it will be appreciated that a zone player may be provided that does not itself include an audio amplifier. In that case, the analog signal may be coupled to an external amplifier for amplification as necessary before being provided to the audio reproduction device(s) 15(n)(r). It will be appreciated that a single zone player 11(n) may be provided with multiple audio amplifiers and audio reproduction device interfaces, and, if necessary, multiple digital to analog converters 33, to provide audio programs for corresponding numbers of synchrony groups.

Similarly, although the zone players 11(n) have been described such that they may be connected to one or more audio information sources, it will be appreciated that an audio information source may form part of and be integrated into a zone player 11(n). For example, a zone player may include a compact disk player, cassette tape player, broadcast radio receiver, or the like, that has been integrated into it. In addition, as noted above, an individual zone player 11(n) may be connected to multiple audio information sources and may contemporaneously operate as the audio information channel device 23 for multiple synchrony groups .

In addition, although FIG. 1 shows the network audio system 10 as including one user interface module 13, it will be appreciated that the system 10 may include a plurality of user interface modules. Each user interface module be useful for controlling all of the zone players as described above, or alternatively one or more of the user interface modules may be useful for controlling selected subsets of the zone players.

Moreover, it will be appreciated that, although the invention has been described in connection with audio information, it will be appreciated that the invention will find utility in connection with any type of information for which synchrony among devices connected to a network is desired.

As noted above, while a zone player 11(n) is operating as audio information channel device 23 for a synchrony group 20, when the zone player 11(n)'s audio information source interface 30 or network communications manager 40 stores digital audio information frames based on audio

information from an audio information source 14(n)(s) in the audio information buffer 31, it will provide time stamps for the respective frames to schedule them for playback after some time delay after they have been buffered in the audio information buffer 31. The delay is provided so that, for other zone players 11(n'), 11(n''),... that are operating as members of a synchrony group, there will be sufficient time for the audio and playback timing information to be transferred over the network 12 to those other zone players 11(n'), 11(n''),... so that it can be processed and played by them at the appropriate time as described above. The time period that is selected for the time delay may be fixed or variable, and in either case may be based on a number of factors. If the time period selected for the time delay is fixed, it may be based on, for example, factors such as an estimate of the maximum latency in the network 12, the estimated maximum loading of the various components comprising the zone players 11(n), and other estimates as will be appreciated by those skilled in the art.

The time delay may be the same for audio information from all types of audio information sources, and may be constant over the entire period that the synchrony group 20 is playing an audio work. Alternatively, different time delays may be utilized based on various criteria. For example, if the audio information is to be played independently of information associated with other types of media, the time delay may be selected to be relatively long, on the order of a significant fraction of a second, or longer. On the other hand, if the audio information is to be played contemporaneously with, for example, video information, which may be supplied by, for example, a video disk, video tape cassette, over cable, satellite, or broadcast television, which may not be buffered or which may be displayed independently of the network audio system 10, it may be undesirable to provide for such a lengthy delay, since the time delay of the audio playback, in relation to the video display, may be noticeable. In that case, the zone player 11(n) may provide for a much shorter time delay. In one embodiment, the time delay provided for audio information to be played concurrently with video information is selected to be generally on the order of fifty milliseconds, which would barely, if at all, be perceptible to someone viewing the video. Other desirable time delays for information from other types of sources will be apparent to those skilled in the art.

As yet a further possibility, the zone player 11(n), when operating as an audio information channel device 23 for a synchrony group 20, can dynamically determine the time delay based on a

number of conditions in the network audio system 10, including, for example, the message transfer latency in network 12, the loading of microprocessors or other components that are used in the various zone players 11(n'), 11(n''),... that may comprise a synchrony group 20, as well as other factors. For example, if the audio information channel device 23 determines that the latency in the network 12 has increased beyond a selected threshold, the audio information channel device 23 can adjust the delay to increase the likelihood that the members of the synchrony group 20 will be able to receive the packets and process the frames so that they will be able to play them at the appropriate times. Similarly, if the audio information channel device 23 is notified that a member of the synchrony group 20 to which it provides audio information requires additional time to receive and process the frames that it transmits, the audio information channel device 23 can adjust the delay accordingly. It will be appreciated that, to reduce or minimize possible discontinuities in the audio playback by the members of the synchrony group, the audio information channel device 23 can, instead of adjusting the time delay during a particular audio track, adjust the time delay between tracks, during silent periods of a track or otherwise as will be appreciated by those skilled in the art. In addition, the audio information channel device 23 can use conventional audio compression methodologies to facilitate a speeding up and/or slowing down of playback of an audio track while it is in the process of providing additional time delay. Generally, the members of the synchrony group 20 can provide notifications to the audio information channel device 23 if they determine that they will need an additional time delay, and the audio information channel device 23 can adjust the time delay in accordance with the notifications from the members of the synchrony group 20.

It will be appreciated that a system in accordance with the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof, any portion of which may be controlled by a suitable program. Any program may in whole or in part comprise part of or be stored on the system in a conventional manner, or it may in whole or in part be provided in to the system over a network or other mechanism for transferring information in a conventional manner. In addition, it will be appreciated that the system may be operated and/or otherwise controlled by means of information provided by an operator using operator input elements (not shown) which may be connected directly to the system or which may

transfer the information to the system over a network or other mechanism for transferring information in a conventional manner.

The foregoing description has been limited to a specific embodiment of this invention. It will be apparent, however, that various variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention. It is the object of the appended claims to cover these and such other variations and modifications as come within the true spirit and scope of the invention.

What is claimed as new and desired to be secured by Letters Patent of the United States is:

1. A method for providing synchronized audio play back amongst a plurality of audio devices, the method comprising:

retrieving, by a computing device, audio information from any of a locally stored digital audio file, a remotely stored digital audio file via a wide area network, and a streaming digital audio service;

transmitting a plurality of frames to a plurality of audio devices over a local network from the computing device, each frame of the plurality of frames to include the retrieved audio information and timing information inserted at the computing device, wherein the timing information is to be used to play the audio information by the plurality of audio devices in synchrony.

2. The method of claim 1, wherein the computing device comprises a personal computer.

3. The method of claim 2, further comprising transmitting the plurality of frames to the plurality of audio devices responsive to receiving a user command via a user interface of the personal computer.

4. The method of claim 1, wherein the computing device comprises a mobile computing device.

5. The method of claim 4, further comprising transmitting the plurality of frames to the plurality of audio devices responsive to receiving a user command via a user interface of the mobile computing device.

6. The method of claim 1, wherein the computing device is configured to play the audio information in synchrony with the plurality of audio devices.

7. The method of claim 1, wherein the computing device is configured to play different audio information than the plurality of audio devices and at the same time as the plurality of audio devices play the audio information.

8. The method of claim 1, wherein the locally stored digital audio file is stored on a personal computer or a mobile computing device configured to store digital information.

9. The method of claim 1, further comprising wirelessly transmitting the plurality of frames to the plurality of audio devices.

10. The method of claim 1, wherein the computing device comprises a user interface to receive a user command to transmit the plurality of frames.

11. The method of claim 10, wherein the user interface is to enable a user to at least one of select an audio track, display status information, and display an upcoming audio playback selection.

12. The method of claim 10, wherein the user interface is to enable a user to select one or more of the plurality of audio devices for audio playback.

13. A non-transitory computer readable medium having stored therein instructions executable by a processor to perform a method comprising:

retrieving, by a computing device, audio information from any of a locally stored digital audio file, a remotely stored digital audio file via a wide area network, and a streaming digital audio service;

transmitting a plurality of frames to a plurality of audio devices over a local network from the computing device, each frame of the plurality of frames to include the retrieved audio information and timing information inserted at the computing device, wherein the timing information is to be used to play the audio information by the plurality of audio devices in synchrony.

14. A method for providing synchronized audio amongst a plurality of audio devices, the method comprising:

receiving, by an audio device, a plurality of frames over a local network from a computing device, each frame of the plurality of frames to include audio information and timing information, wherein the timing information is to be used to play the audio information by the audio device in synchrony with another receiving audio device; and

converting, by the audio device, the audio information from a digital form into an analog signal to be output based on the timing information via an amplifier and a speaker driver.

15. The method of claim 15, wherein the computing device comprises a mobile computing device.

16. The method of claim 15, wherein the computing device comprises a personal computer.

17. The method of claim 15, wherein the audio information is from any of a locally stored digital audio file, a remotely stored digital audio file via a wide area network, and a streaming digital audio service.

18. The method of claim 15, wherein the local network comprises a wired network.

19. The method of claim 15, wherein the local network comprises a wireless network.

20. A non-transitory computer readable medium having stored therein instructions executable by a processor to perform a method comprising:

receiving, by an audio device, a plurality of frames over a local network from a computing device, each frame of the plurality of frames to include audio information and timing information, wherein the timing information is to be used to play the audio information by the audio device in synchrony with another receiving audio device; and

converting, by the audio device, the audio information from a digital form into an analog signal to be output based on the timing information via an amplifier and a speaker driver.

ABSTRACT OF THE DISCLOSURE

A system is described for maintaining synchrony of operations among a plurality of devices that have independent clocking arrangements. The system includes a task distribution device that distributes tasks to a synchrony group comprising a plurality of devices that are to perform the tasks distributed by the task distribution device in synchrony. The task distribution device distributes each task to the members of the synchrony group over a network. Each task is associated with a time stamp that indicates a time, relative to a clock maintained by the task distribution device, at which the members of the synchrony group are to execute the task. Each member of the synchrony group periodically obtains from the task distribution device an indication of the current time indicated by its clock, determines a time differential between the task distribution device's clock and its respective clock and determines therefrom a time at which, according to its respective clock, the time stamp indicates that it is to execute the task.

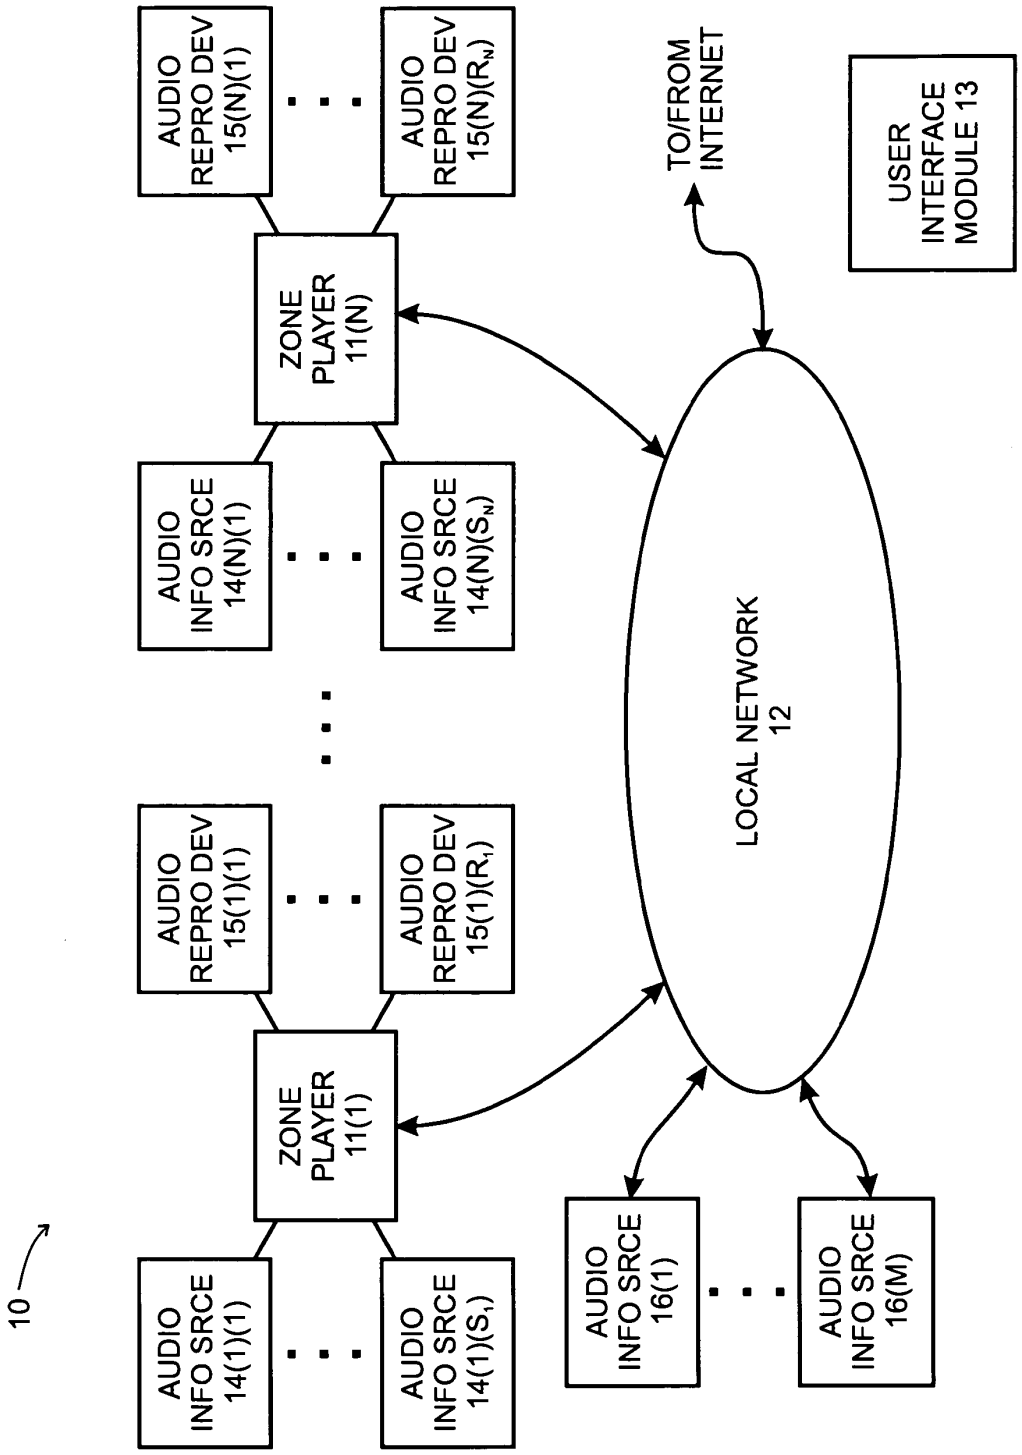

10
AUDIO INFO SRCE 14(1)(1)
AUDIO INFO SRCE 14(1)(S_1)
AUDIO REPRO DEV 15(1)(1)
AUDIO REPRO DEV 15(1)(R_1)
ZONE PLAYER 11(1)
AUDIO INFO SRCE 14(N)(1)
AUDIO INFO SRCE 14(N)(S_N)
AUDIO REPRO DEV 15(N)(1)
AUDIO REPRO DEV 15(N)(R_N)
ZONE PLAYER 11(N)
LOCAL NETWORK 12
TO/FROM INTERNET
AUDIO INFO SRCE 16(1)
AUDIO INFO SRCE 16(M)
USER INTERFACE MODULE 13

FIG. 1

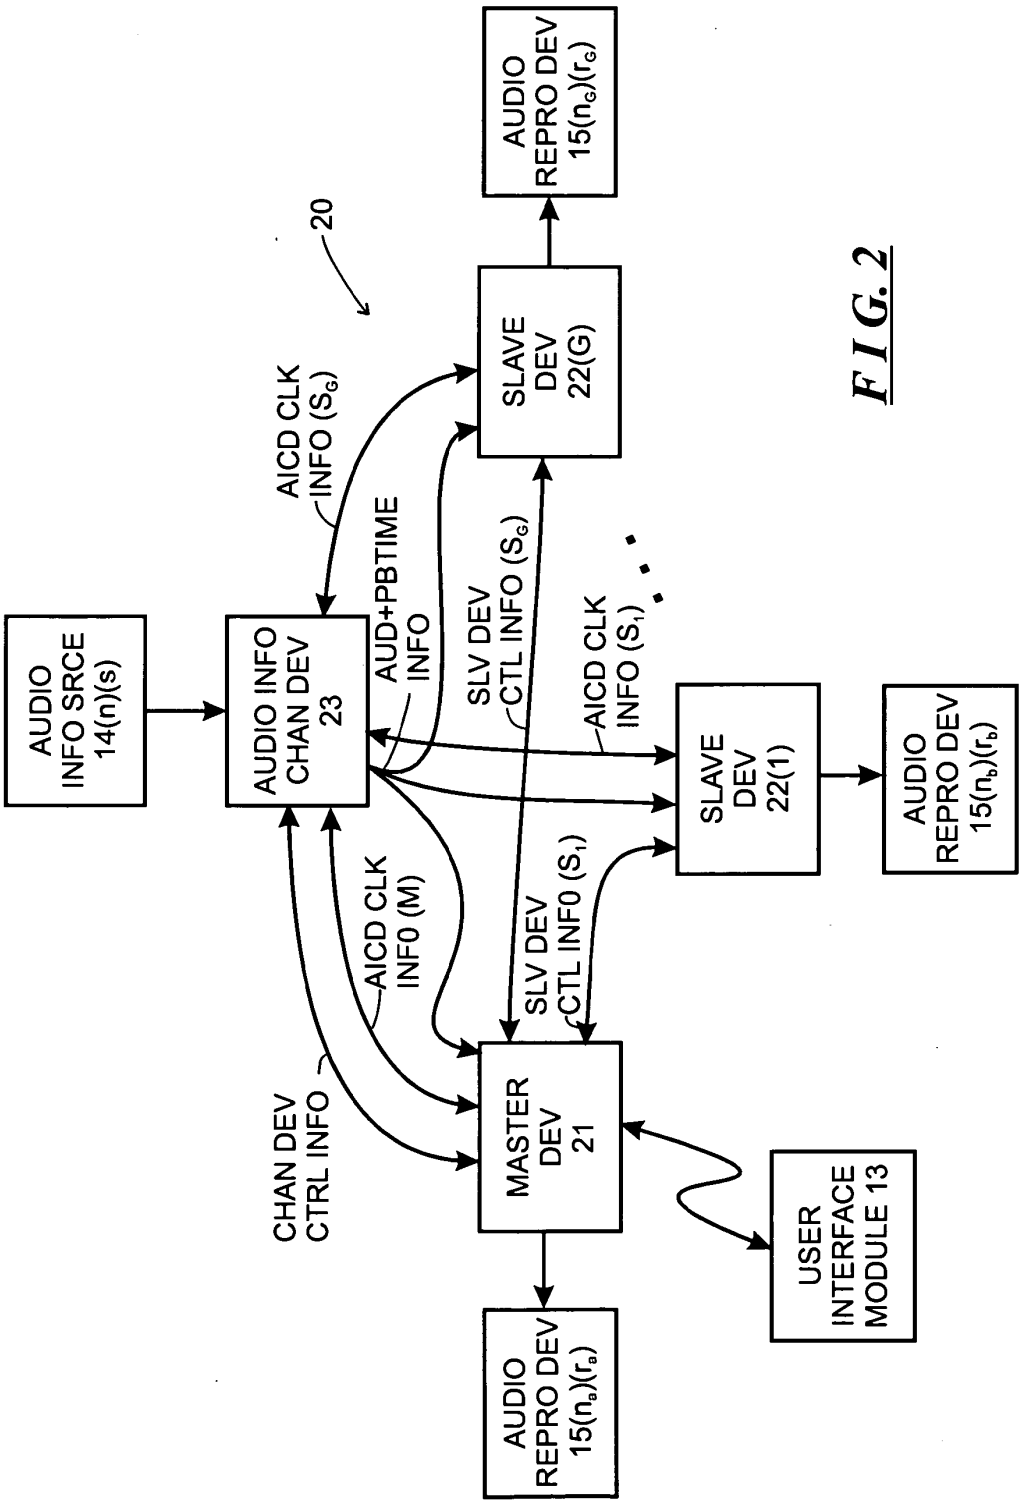

20
AUDIO INFO SRCE 14(n)(s)
AUDIO INFO CHAN DEV 23
AUD+PBTIME INFO
AICD CLK INFO (S_G)
SLV DEV CTL INFO (S_G)
AICD CLK INFO (S_1)
SLAVE DEV 22(G)
AUDIO REPRO DEV 15(n_G)(r_G)
SLAVE DEV 22(1)
AUDIO REPRO DEV 15(n_b)(r_b)
AICD CLK INFO (M)
SLV DEV CTL INFO (S_1)
CHAN DEV CTRL INFO
MASTER DEV 21
USER INTERFACE MODULE 13
AUDIO REPRO DEV 15(n_a)(r_a)

FIG. 2

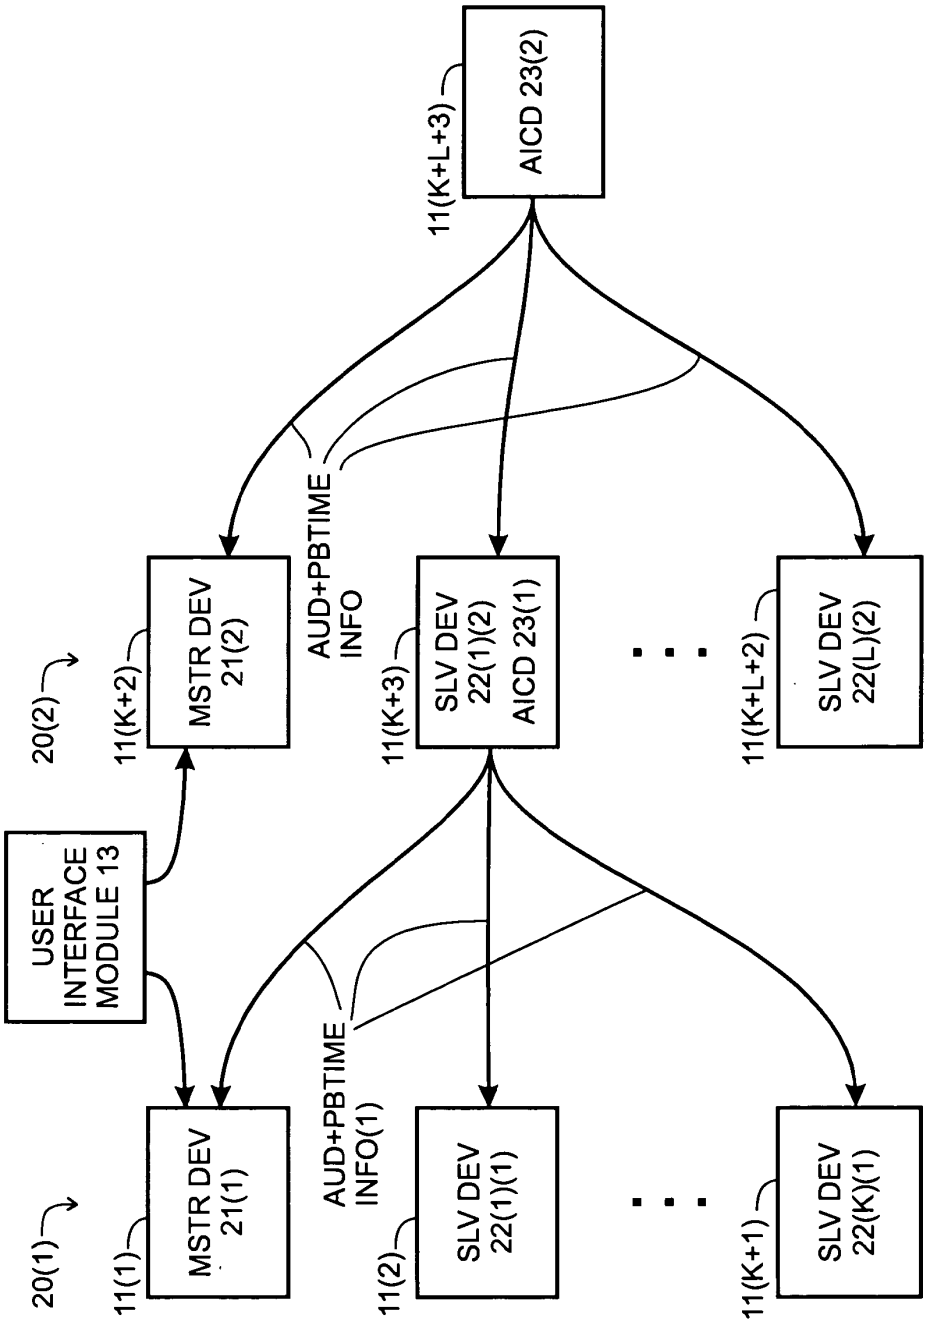

AICD 23(2)
11(K+L+3)
AUD+PBTIME INFO
MSTR DEV 21(2)
11(K+2)
20(2)
SLV DEV 22(1)(2) AICD 23(1)
11(K+3)
SLV DEV 22(L)(2)
11(K+L+2)
USER INTERFACE MODULE 13
AUD+PBTIME INFO(1)
MSTR DEV 21(1)
11(1)
20(1)
SLV DEV 22(1)(1)
11(2)
SLV DEV 22(K)(1)
11(K+1)

FIG. 2A

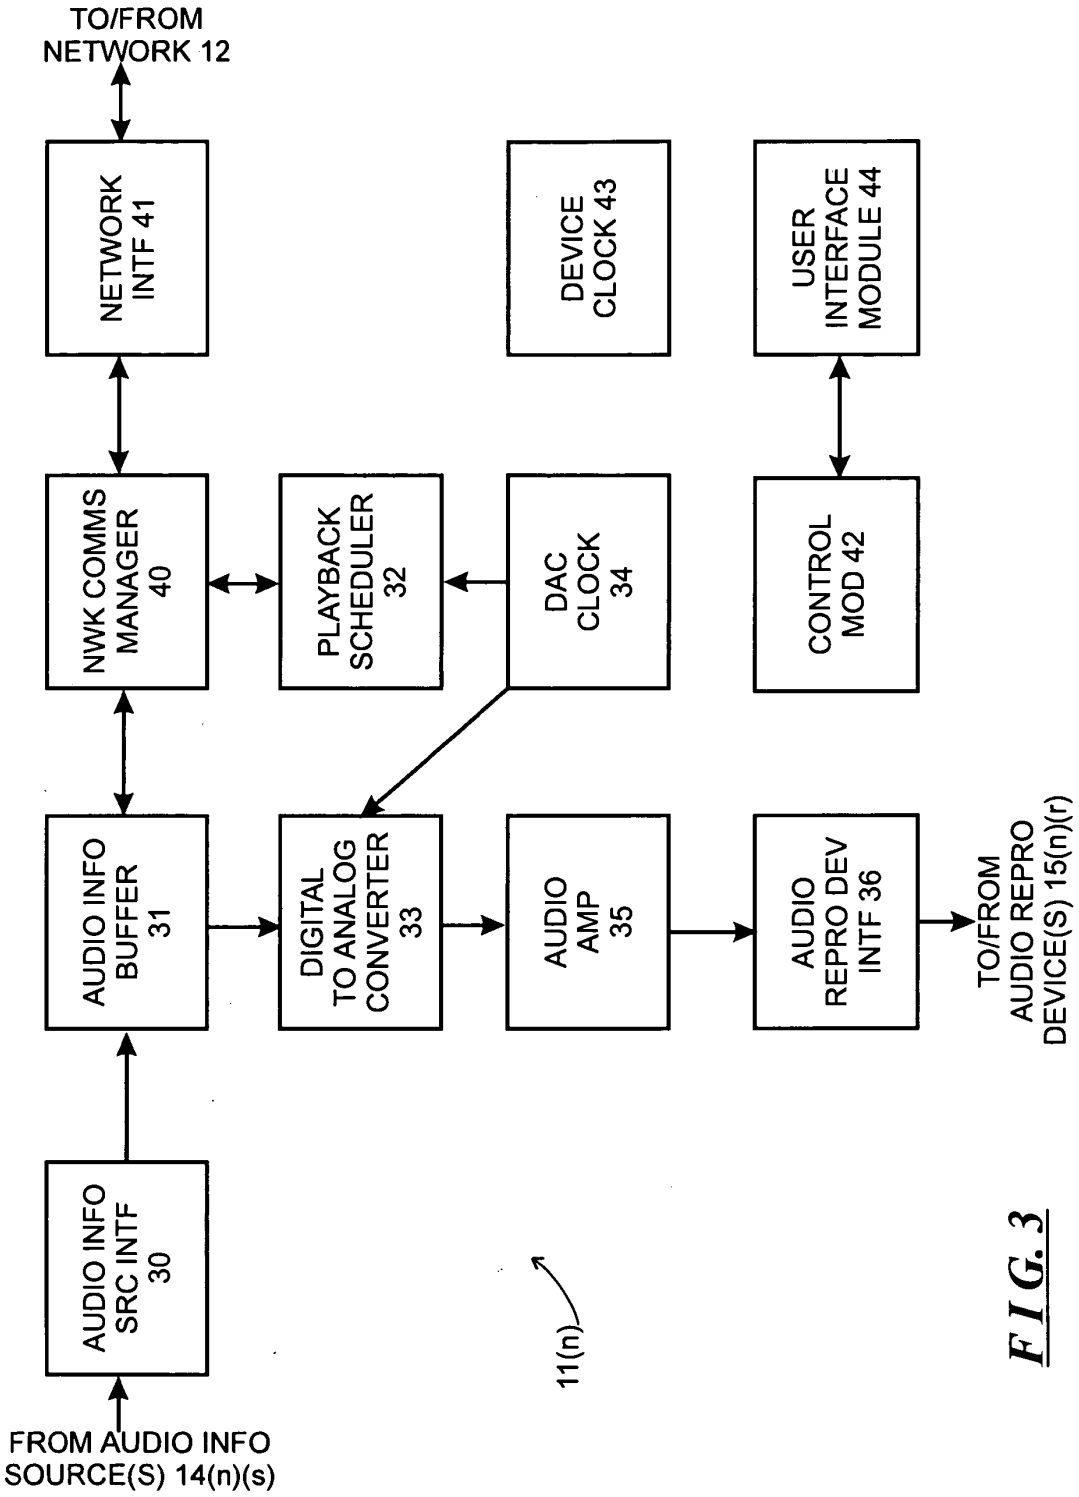

TO/FROM NETWORK 12
NETWORK INTF 41
DEVICE CLOCK 43
USER INTERFACE MODULE 44
NWK COMMS MANAGER 40
PLAYBACK SCHEDULER 32
DAC CLOCK 34
CONTROL MOD 42
AUDIO INFO BUFFER 31
DIGITAL TO ANALOG CONVERTER 33
AUDIO AMP 35
AUDIO REPRO DEV INTF 36
TO/FROM AUDIO REPRO DEVICE(S) 15(n)(r)
AUDIO INFO SRC INTF 30
11(n)
FROM AUDIO INFO SOURCE(S) 14(n)(s)

FIG. 3

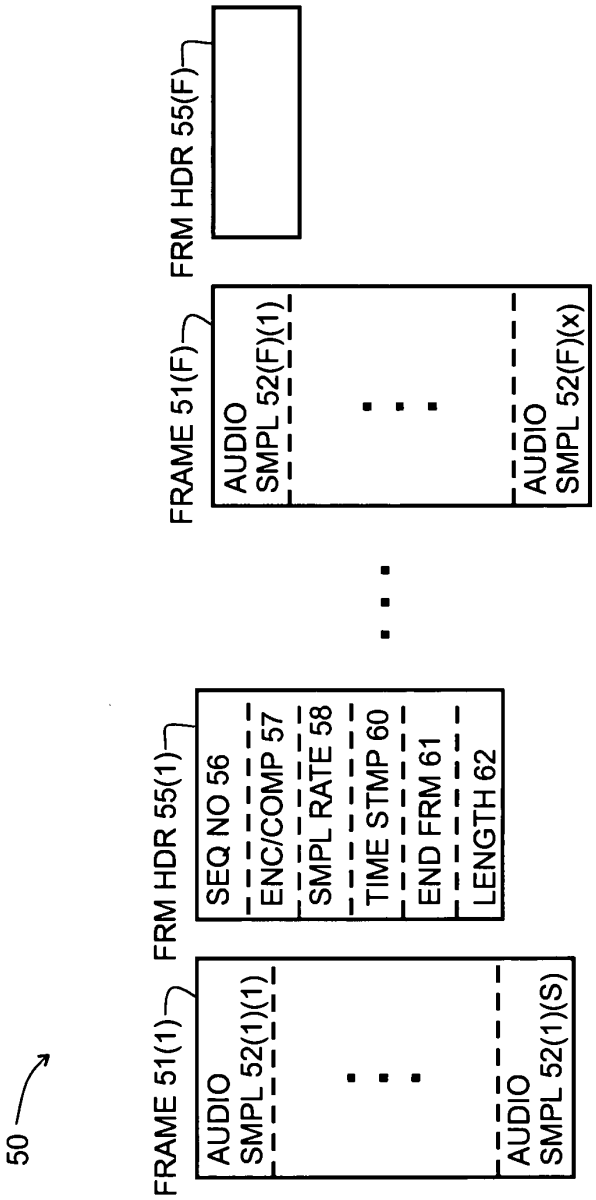

50
FRAME 51(1)
AUDIO SMPL 52(1)(1)
AUDIO SMPL 52(1)(S)
FRM HDR 55(1)
SEQ NO 56
ENC/COMP 57
SMPL RATE 58
TIME STMP 60
END FRM 61
LENGTH 62
FRAME 51(F)
AUDIO SMPL 52(F)(1)
AUDIO SMPL 52(F)(x)
FRM HDR 55(F)

FIG. 4

PATENTS

Attorney Docket No. RIN-001

DECLARATION AND POWER OF ATTORNEY

As a below-named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled "System and Method For Synchronizing Operations Among A Plurality Of Independently Clocked Digital Data Processing Devices," the specification of which is filed herewith and identified by Attorney Docket No. RIN-001.

I hereby state that I have reviewed and understand the contents of the above-identified application specification, including the claims.

I acknowledge the duty to disclose information that is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, section 1.56(a).

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment or both under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint Richard A. Jordan, Reg. No. 27,807 my attorney, with full power of substitution, delegation and revocation, to prosecute this application, to make alterations and amendments therein, to receive the patent and to transact all business in the Patent and Trademark Office connected therewith. Please direct all telephone calls to Richard A. Jordan at (781) 431-1357. Please address all correspondence to Richard A. Jordan, at P. O. Box 81363, Wellesley Hills, MA 02481-0004.

Nicholas A. J. Millington

3/31/04
Date

Residence: 1221 Salsipuedes Street
Santa Barbara, CA 93103

Attorney Docket No. RIN-001

Citizenship U. S.

Post Office Address: 1221 Salsipuedes Street
Santa Barbara, CA 93103

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 13/864,248	Filing Date 04/17/2013	☐ To be Mailed

ENTITY: ☒ LARGE ☐ SMALL ☐ MICRO

APPLICATION AS FILED – PART I

	(Column 1)	(Column 2)		
FOR	NUMBER FILED	NUMBER EXTRA	RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	

APPLICATION AS AMENDED – PART II

AMENDMENT	04/17/2013	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 20	Minus	** 20	= 0	X $80 =	0
	Independent (37 CFR 1.16(h))	* 3	Minus	***3	= 0	X $420 =	0
	☐ Application Size Fee (37 CFR 1.16(s))						
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
						TOTAL ADD'L FEE	0

AMENDMENT		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))						
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
						TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

LIE
/TAMMY L. ACREE/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.