



US 20080298241A1

(19) **United States**(12) **Patent Application Publication**
Ohana et al.(10) **Pub. No.: US 2008/0298241 A1**(43) **Pub. Date: Dec. 4, 2008**(54) **APPARATUS AND METHODS FOR
REDUCTION OF TRANSMISSION DELAY IN
A COMMUNICATION NETWORK****Related U.S. Application Data**


(60) Provisional application No. 60/941,020, filed on May 31, 2007.

(75) Inventors: **Yitshak Ohana**, Jerusalem (IL);
Philippe Klein, Jerusalem (IL); **Avi
Kliger**, Ramat Gan (IL); **Stephen
Palm**, Irvine, CA (US)**Publication Classification**(51) **Int. Cl.**
H04L 12/24 (2006.01)
H04L 12/56 (2006.01)
(52) **U.S. Cl.** **370/235; 370/390**
(57) **ABSTRACT**

Correspondence Address:

Weiss & Arons, LLP
1540 Route 202, Suite 8
Pomona, NY 10970 (US)

Apparatus and methods for reducing latency in coordinated networks. The apparatus and methods relate to a protocol that may be referred to as the Persistent Reservation Request ("p-RR"), which may be viewed as a type of RR (reservation request). p-RR's may reduce latency, on average, to one MAP cycle or less. A p-RR may be used to facilitate Ethernet audiovisual bridging. Apparatus and methods of the invention may be used in connection with coaxial cable based networks that serve as a backbone for a managed network, which may interface with a package switched network.

(73) Assignee: **Broadcomm Corporation**, Irvine,
CA (US)(21) Appl. No.: **12/041,368**(22) Filed: **Mar. 3, 2008**




FIG. 1




FIG. 2
PRIOR ART




FIG. 3




FIG. 4




FIG. 5







FIG. 6

*FIG. 7*

*FIG. 8*

*FIG. 9*




FIG. 10

APPARATUS AND METHODS FOR REDUCTION OF TRANSMISSION DELAY IN A COMMUNICATION NETWORK

CROSS REFERENCE TO RELATED APPLICATION

[0001] This is a nonprovisional of U.S. Application No. 60/941,020, filed on May 31, 2007, which is incorporated by reference herein in its entirety.

FIELD OF THE INVENTION

[0002] The present invention relates generally to communication networks and specifically to optimizing bandwidth utilization in a coordinated network based on a coaxial cable backbone.

BACKGROUND OF THE INVENTION

[0003] Many structures, including homes, have networks based on coaxial cable ("coax").

[0004] The Multimedia over Coax Alliance ("MoCA™"), provides at its website (www.mocalliance.org) an example of a specification (viz., that available under the trademark MoCA), which is hereby incorporated herein by reference in its entirety, for networking of digital information, including video information, through coaxial cable. The specification has been distributed to an open membership.

[0005] Technologies available under the trademark MoCA, and other specifications and related technologies (collectively, with MoCA, "the existing technologies"), often utilize unused bandwidth available on the coax. For example, coax has been installed in more than 70% of homes in the United States. Some homes have existing coax in one or more primary entertainment consumption locations such as family rooms, media rooms and master bedrooms. The existing technologies allow homeowners to utilize installed coax as a networking system for the acquisition and use of information with high quality of service ("QoS").

[0006] The existing technologies may provide high speed (270 mbps), high QoS, and the innate security of a shielded, wired connection combined with packet-level encryption. Coax is designed for carrying high bandwidth video. Today, it is regularly used to securely deliver millions of dollars of pay-per-view and video content on a daily basis.

[0007] Existing technologies provide throughput through the existing coaxial cables to the places where the video devices are located in a structure without affecting other service signals that may be present on the cable.

[0008] The existing technologies work with access technologies such as asymmetric digital subscriber lines ("ADSL"), very high speed digital subscriber lines ("VDSL"), and Fiber to the Home ("FTTH"), which provide signals that typically enter the structure on a twisted pair or on an optical fiber, operating in a frequency band from a few hundred kilohertz to 8.5 MHz for ADSL and 12 MHz for VDSL. As services reach such a structure via any type of digital subscriber line ("xDSL") or FTTH, they may be routed via the existing technologies and the coax to the video devices. Cable functionalities, such as video, voice and Internet access, may be provided to the structure, via coax, by cable operators, and use coax running within the structure to reach individual cable service consuming devices in the struc-

ture. Typically, functionalities of the existing technologies run along with cable functionalities, but on different frequencies.

[0009] The coax infrastructure inside the structure typically includes coax, splitters and outlets. Splitters typically have one input and two or more outputs and are designed to transmit signals in the forward direction (input to output), in the backward direction (output to input), and to isolate outputs from different splitters, thus preventing signals from flowing from one coax outlet to another. Isolation is useful in order to a) reduce interference from other devices and b) maximize power transfer from Point Of Entry ("POE") to outlets for best TV reception.

[0010] Elements of the existing technologies are specifically designed to propagate backward through splitters ("insertion") and from output to output ("isolation"). One outlet in a structure can be reached from another by a single "isolation jump" and a number of "insertion jumps." Typically isolation jumps have an attenuation of 5 to 40 dB and each insertion jump attenuates approximately 3 dB. MoCA™-identified technology has a dynamic range in excess of 55 dB while supporting 200 Mbps throughput. Therefore MoCA™-identified technology can work effectively through a significant number of splitters.

[0011] Networks based on the existing technologies are often coordinated networks, in which a processing unit serves as a network coordinator. The coordinator defines medium access plan ("MAP") cycles, prospectively assigns data transmission events to the cycles, and serially processes the cycles by executing or coordinating the events in each cycle. Coordinated network schemes, such as MoCA™-identified technology, may be used for transmission of streaming video and thus data throughput between outlets is desirable.

[0012] FIG. 1 shows known data flow 100 that may be implemented in a coordinated network. In data flow 100, network coordinator 102 grants only explicit reservation requests (i.e., a transmission opportunity is scheduled only to a granted reservation request made by the transmitter on a per frame or aggregated frame basis). Upon receiving reservation requests, network coordinator 102 grants and schedules transmission opportunities in the next MAP cycle. For example, network coordinator 102 grants transmission opportunities to requesting transmitter nodes such as transmitting node 104. Network coordinator 102 periodically allocates a reservation request opportunity ("RR Opportunity"), such as 106 in MAP Cycle i, to transmitting node 104. In response to RR Opportunity 106, transmitting node 104 transmits in MAP Cycle i+1 reservation request ("RR") 108 for the transmission of data frame k. Network coordinator 102 responds in MAP Cycle i+1 by transmitting to transmitting node 104 frame k transmission grant 110. In MAP Cycle i+2, transmitting node 104 transmits data frame k to receiving node 112 in frame k transmission 111. Data flow 100 includes subsequent exchanges between network coordinator 102 and transmitting node 104 in connection with data frame m: viz., RR Opportunity 114, RR 116, frame m transmission grant 118 and frame m transmission 119.

[0013] Data flow 100 has an average latency of three MAP cycles between the receipt of an Ethernet packet by transmitting node 104 and reception of the packet at receiving node 112. In existing technologies such as that identified by MoCA™ the nominal MAP cycle duration is 1 millisecond ("ms"), yielding a temporal latency of 3 ms. Such latency may limit the ability of the existing technologies to support time

sensitive applications such as Ethernet AV (audiovisual) Bridging, including class 5 AVB data transfer.

[0014] FIG. 2 shows typical communication network 200, which includes network segments 202, 204 and 206. Network segments 202 and 206 are Ethernet packet-switched segments and network segment 204 is a shared media packet mode network segment that requires a network coordinator. Packet-switched network segments 202 and 206 require architectural and message passing characteristics that are different from those required by packet mode network segment 204. For example, legacy switch 208 in packet switched segment 202 may not be configured for exchanging audiovisual data. As another example, Ethernet hub 210 in packet switched segment 204 may operate based on a half duplex protocol, whereas AV switch 212 in packet mode segment 206 may require a full duplex protocol. As a result of protocol mismatching, the quality of communication of audiovisual data from AV device 214 in segment 202 to device 216, which may be a display, in segment 204 may be degraded or be subject to latency delays.

[0015] It therefore would be desirable to provide systems and methods for reducing latency in coordinated networks.

[0016] It therefore also would be desirable to provide systems and methods for Ethernet AV bridging using shared media networks.

SUMMARY OF THE INVENTION

[0017] Systems and methods for reducing latency in coordinated networks, and for performing Ethernet AV bridging, are provided substantially as shown in and/or described in connection with at least one of the figures, and as set forth more completely in the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] The above and other features of the present invention, its nature and various advantages will be more apparent upon consideration of the following detailed description, taken in conjunction with the accompanying drawings, and in which:

[0019] FIG. 1 shows a diagram of a known data flow in a coordinated shared media communication network;

[0020] FIG. 2 shows a known arrangement of interfaced communication network segments;

[0021] FIG. 3 shows a diagram of a data flow in accordance with the principles of the invention;

[0022] FIG. 4 shows an arrangement of communication network segments that may be used in accordance with the principles of the invention;

[0023] FIG. 5 shows the arrangement of FIG. 4 in use in accordance with the principles of the invention;

[0024] FIG. 6 also shows the arrangement of FIG. 4 in use in accordance with the principles of the invention;

[0025] FIG. 7 shows a process in accordance with the principles of the invention;

[0026] FIG. 8 shows another process in accordance with the principles of the invention;

[0027] FIG. 9 shows yet another process in accordance with the principles of the invention; and

[0028] FIG. 10 shows a device in accordance with the principles of the invention.

DETAILED DESCRIPTION OF EMBODIMENTS

[0029] Apparatus and methods for reducing latency in coordinated networks are provided in accordance with the principles of the invention. The apparatus and methods relate to a protocol that may be referred to herein as the Persistent Reservation Request (“p-RR”), which may be viewed as a type of RR (reservation request). In some embodiments of the invention, p-RR can reduce latencies, on average, to one MAP cycle or less. In some embodiments of the invention, a persistent reservation request may be used to facilitate Ethernet audiovisual bridging. Apparatus and methods of the invention may be used in connection with networks that are described in one or more of MoCA® Specifications v10, February 2006; MoCA® Specifications v1.1 Extensions, May 2007; and IEEE802.1 AVB WG Specifications, all of which are incorporated by reference herein in their entireties.

[0030] Methods in accordance with the principles of the invention may include a method for managing transmission of data over a shared media communication network. The method may include receiving from a node a first reservation request. The reservation request may request allocation of network resources for transmitting the data. The method may further include establishing a data flow based on the request, wherein the node transmits at least two data frames before a second reservation request is received. The flow may have a duration that is greater than the duration of one medium access plan cycle. The reservation request may include a traffic specification (“TSPEC”). Some embodiments of the invention may include canceling the flow based on information included in the traffic specification and/or canceling the flow based on the execution of a link maintenance operation (“LMO”) process. Some embodiments of the invention may include managing jitter by allocating network resources during only a portion of one or more MAP cycles.

[0031] An exemplary device in accordance with the principles of the invention may be configured to manage transmission of data over a shared media communication network. The device may include a receiver module configured to receive a first reservation request from a node in communication with the network; a processor module configured to allocate network resources based on the request; and a transmitter module configured to inform the node about a grant of the request. The grant may authorize the node to transmit at least two data frames before the node transmits a second reservation request.

[0032] Some methods in accordance with the principles of the invention may include providing Ethernet AV bridging using coordinated shared media networks. Those methods may include transmitting multicast information from a first node to a second node. The first node may reside in a packet mode network and the second node may reside in a packet-switched network. The method may include receiving at the first node a registration request from the second node. The registration request may request receipt of the multicast information and may include a TSPEC. The method may further include transmitting, via the first node, a reservation request to a network coordinator residing in the packet mode network. The reservation request may conform to the traffic specification.

[0033] Some methods for Ethernet AV bridging in accordance with the principles of the invention may include using

a packet mode network to transmit data from a first packet-switched network segment to a second packet-switched network segment. The methods may include receiving the data from the first segment using a packet mode network ingress port and transmitting the data to the second segment using a packet mode network egress port.

[0034] FIGS. 3-10 show illustrative embodiments and features of the invention.

[0035] FIG. 3 shows illustrative p-RR data flow **300** that may be used to reduce latency in a coordinated shared media network. Data flow **300** may involve network coordinator **302**, transmitting node **304**, and receiving node **306**. In MAP Cycle *i*, network coordinator **302** transmits to transmitting node **304** RR Opportunity **308**. In MAP Cycle *i*+1, transmitting node **304** transmits to network coordinator **302** p-RR **210**, which is a persistent request for network allocation of network resources for subsequent transmission, during the lifetime of flow *f*, of data frames to receiving node **306**. Before the end of MAP Cycle *i*+1, network coordinator **302** grants p-RR **210** for flow *f* by transmitting to transmitting node **304** flow *f* transmission grant **312**. Based on grant **312**, transmitting node **304** may transmit to receiving node **306** data frames such as frames *k* (**314**), *l* (**316**), *m* (**318**) and *n* (**320**).

[0036] The aforementioned MoCA 1.1 specification includes protocol pQoS (parameterized QoS), which supports parameterized QoS (Quality of Service) flow transactions over a network. In pQoS flow, a transmitting node sends a TSPEC admission request to a receiving node, as described in MoCA 1.1 section 4.2. After the request is accepted, a flow is created and data transmission between the transmitting and receiving nodes occurs in accordance with the protocol described above in connection with FIG. 1.

[0037] In some embodiments of the invention, after a pQoS flow is created, the transmitting node may use the flow's TSPEC to generate a p-RR for the flow and send the p-RR to the network coordinator ("NC"). Table 1 shows illustrative information that may be included in the p-RR

TABLE 1

Illustrative p-RR Information	
p-RR Information	Characteristics
Leased time	Copied from the flow's TSPEC. Indicates length of time period during which the p-RR applies.
Packet duration	Derived from the TSPEC nominal packet size and calculated for the connection current PHY profile (bit loading).
Mean time	Indicates the required mean time-between-transmission-grants to meet latency requirements (multiple transmissions may be allocated by the network controller to a flow within a single MAP cycle)

[0038] In some embodiments, after the NC receives the p-RR, it can allocate to the transmitting node network resources for the flow based on the p-RR information, such as the leased time. In some embodiments, the p-RR may be allowed to persist (and is thus referred to as a "persistent grant" or "PG") until a new p-RR for the same flow is received. When a node transmits exclusively audiovisual bridge traffic, and is operating under a persistent grant, there is no need to poll it frequently (every or almost every MAP

CYCLE) as is usually done in MoCA 1.0 networks. This may reduce overhead and improve network throughput.

[0039] In response to the p-RR, the NC may allocate to the transmitting node a time-limited transmission opportunity. The time interval between consecutive transmission opportunities (viz., the "transmission ('Tx') service time"), could be significantly smaller than the MAP cycle duration to accommodate the requirements of specific time sensitive traffic moving over the network. In some embodiments, a maximum Tx service time may be 100 μ s when the MAP cycle duration is 1 ms.

[0040] If an LMO process (MoCA 1.0 section 3.7) causes a change in the connection profile associated with the transmitting and receiving nodes, the transmitting node may send an updated p-RR. If the data rate between the transmitting and receiving nodes changes, for example, as a result of an LMO process, the transmitting node may renew its p-RR. If the transmitting node fails to renew its p-RR, the NC may, after a predetermined number of MAP cycles, discontinue p-RR grants for the flow and send to the transmitting node a new p-RR opportunity. The NC may allocate a p-RR opportunity to the transmitting node after each LMO process between the transmitting and the receiving nodes. The allocation of a p-RR opportunity after each LMO process may be conditioned on the time remaining before the leased time expires.

[0041] FIG. 4 shows illustrative communication network **400** in which persistent reservation requests in accordance with the principles of the invention may be used to provide AV bridging, such as Ethernet AV bridging, between packet switched network segments **402** and **406** using packet mode network segment **404**. Segment **404** may include network coordinator **408** and AV bridges **410** and **412** that interface, respectively, with packet switches **414** and **416**. Bridges **410** and **412** may, together with network coordinator **408**, establish a data flow that is based on a p-RR such as that described above in connection with FIG. 3. Based on such a data flow, talker node **418** (which resides in packet switched network segment **406**) may stream AV data to listener node **420** (which resides in packet switched network segment **402**) via packet mode network segment **404** (which is managed by network coordinator **408**).

[0042] Packet switched network data streaming protocols (such as the stream reservation protocol ("SRP") defined by IEEE standard 802.1Qat) often require registration by a listener node. The registration identifies to a talker node that the listener node desires to receive a data stream. After registration, such protocols require that the talker node transmit a reservation to reserve network resources along a path from talker to listener. FIG. 5 shows an illustrative example, according to the principles of the invention, of how packet mode network segment **404** may bridge between talker node **418** and listener node **420**, which communicate under a streaming protocol such as that defined by IEEE 802.1 Qat. The network segments, and elements that make up the segments, are the same as those shown in FIG. 4.

[0043] FIG. 5 shows an illustrative registration process, with a packet mode network serving as a backbone, and bridging, between packet switching network segments. The process originates when listener node **420**, in packet switched network segment **402**, transmits registration **450** to AV Ethernet switch **414**. Switch **414** propagates registration **450** to packet mode bridge **410**, which then functions as a network segment ingress node, in packet mode network segment **404**. Bridge **410** propagates registration **450** to bridge **412**, which

then functions as a network segment egress mode for network segment **404**. Bridge **410** transmits registration notification **460** to network coordinator **408**, which propagates notification **460** to one or more ingress and egress bridges, including bridges **410** and **412**. Registration **450** may also be propagated to nodes that do not use registration **450**. In communication network **400**, those nodes include **420**, **422**, **424**, **426**, **428** and **430**.

[0044] From bridge **412**, registration **450** is propagated to talker **418**. It will be appreciated that inter-node communications regarding registration, notification, reservation and other information may require suitable modification of the communications, but for the sake of clarity, such modifications are ignored in FIGS. **5** and **6** and the same reference numeral is used to identify such a communication at more than one stage of propagation.

[0045] FIG. **6** shows an illustrative reservation process that corresponds to the registration process shown in FIG. **5**. Packet mode network segment **404** is shown serving as a backbone, and bridging, between packet switching network segments **402** and **406** in connection with the reservation process. In a packet switching network or network segment subject to a QoS protocol, a talker node transmits a reservation message to reserve downstream network resources. In the illustrative reservation process shown in FIG. **6**, talker node **418** (in packet switched network segment **406**) transmits reservation message **470** to reserve network resources for data transmission to listener **420** (in packet switched network segment **402**). Reservation message **470** is received by AV bridge **412**. AV bridge **412** then transmits data flow reservation **472** (for a data flow within packet mode network segment **404**) to network coordinator **408**, which may respond by transmitting data flow grant **474** to bridge **412**. The data flow in network segment **404** may be a parameterized quality of service ("pQoS") flow, which may have a protocol that supports persistent reservation requests (p-RR). AV bridge **412** may transmit a p-RR to network coordinator **408** in the same MAP cycle as the data flow is requested. AV bridge **412**, acting as ingress node, may propagate reservation message **470** to bridge **410**, which acts as egress node.

[0046] When bridge **412** receives reservation message **470** from switch **416** (in packet switched network segment **406**), bridge **412** may reject or accept reservation message **470**. Bridge **412** may do so based on the availability of bridge **412** resources. Bridge **412** may query network coordinator **408** regarding the availability of resources in packet mode network segment **404**. Network coordinator **408** may communicate reservation message **470** to bridge **410** along path **476**. Bridge **410** may then propagate reservation message **470**, via switch **414**, to listener node **420**.

[0047] The P802.1Qat SRP is a one-way declarative protocol with reservation messages propagated from the talker towards listeners. It typically contains no backward propagated acknowledgement or status report messages. If needed, a talker node could leverage higher layer applications for getting feedback from listeners. A talker can also simply initiate the reservation message and then wait enough time before starting data transmission to ensure that the stream data can be served appropriately.

[0048] A registration event may be initiated by a listener and passed by intermediate bridges, such as those of network segment **404**. The talker node may be triggered by the registration event, which indicates that the listener node desires to establish a stream (data flow). Upon the registration event, the

talker node may attempt to reserve the required resources in the local node and, for a shared media LAN, the talker may attempt to reserve resources on the LAN to which the talker node's egress port is attached. The talker may configure itself appropriately according to the local and LAN reservation message results. It then may record the result into the status information field of an updated reservation message and send the message toward the listener node.

[0049] SRP assumes that an admission control policy (which governs the disposition of a request by a node that resides in one network segment for network resources that reside in a different network segment) is implemented by a network segment egress port (e.g., switch **416**, shown in FIG. **6**) rather than an ingress port (e.g., bridge **412**, shown in FIG. **6**). When the egress port is connected to a point-to-point link (e.g., a packet switched network segment), the egress node can make admission control decisions according to its own policy. When the egress port is connected to a shared media LAN (such as packet mode network segment **404**), a segment resource manager for the LAN, such as a network controller (e.g., controller **408**), or some kind of distributed cooperation between all the ports on this LAN may be needed. The implementation of admission control policy of a local node or segment resource manager may be according to suitable known policies.

[0050] Reservation messages may be propagated over a subtree (e.g., a virtual LAN) by which the talker and listeners are connected. For each receiving bridge along the path from the talker to listeners, reservation messages can convey the result of resource reservation of the upstream bridges, and may trigger any necessary local or shared media LAN resource reservation operations.

[0051] In certain embodiments of the invention, a talker node, such as **418**, may refresh a reservation message on a regular basis. Each listener, such as **420**, may keep a timer which will timeout the registration request if there is no corresponding reservation message received during the timer period.

[0052] On receiving the first reservation message of a stream, some embodiments of the invention may allow a bridge to create a reservation record for the reservation message. The reservation record may contain information obtained from the reservation message. This information can include stream identifier, talker MAC address, traffic specification, upstream reservation status, reservation message inbound port and any other suitable information.

[0053] If a bridge, such as **412** (shown in FIG. **6**), receives a reservation message that carries a positive reservation status, the bridge may reserve the required resources for each egress port that corresponds to a listener that has registered for a data stream. Preferably, the bridge can configure its local forwarding resources. On shared media LANs, such as packet mode network segment **404**, the network coordinator may reserve resources according to the upstream reservation results, and then transmit the reservation message out of each registered egress port. Each reservation message carries the updated reservation status and per-hop resource details according to the local/LAN reservation result on the corresponding port.

[0054] Because SRP is a one way declarative protocol, a failed reservation in an intermediate bridge will not influence any reservation that has been made upstream. Relevant listeners can receive reservation messages with negative reservation status. For example, a listener could choose to either

withdraw the registration (e.g., by sending an update message corresponding to the registration), therefore releasing any unnecessary reserved resources, or keep the registration. By keeping the registration, the stream reservation can be made along the whole path when all necessary resources become available.

[0055] If the received reservation message carries a negative reservation status, the bridge can preferably configure its forwarding resources appropriately, update the per-hop resource details information in the reservation message, and then transmit the reservation message out of each registered egress port.

[0056] After a reservation record has been created, a bridge could later receive a reservation message that is inconsistent with the reservation record. While a discrepancy in the talker MAC address may be reported as an error and the reservation message rejected, a discrepancy in traffic specification, upstream reservation status, and reservation message inbound port may be accepted or rejected by the bridge based on predetermined policy.

[0057] FIG. 7 shows illustrative process 700, which may provide Ethernet AV bridging when a packet mode network segment node is a talker for a prospective multicast to nodes in a packet switched network. FIG. 7 illustrates process 700 using the example of a node in a network having protocols defined by a MoCA specification such as one of those identified above. At step 702, the talker MoCA node receives a multiple MAC address registration protocol (“MMRP”) registration message. At step 704, the node may transmit a MoCA pQoS flow request to its network coordinator. At step 706, the network coordinator approves the flow request. At step 708, the MoCA node may initiate a p-RR, such as that shown and described above. Alternatively, the node may initiate an RR in every MAP cycle. The p-RR and the RR may be consistent with a TSPEC that may be embedded in the MMRP.

[0058] FIG. 8 shows illustrative process 800, which may provide Ethernet AV bridging when a packet mode network segment node is a listener node for a prospective multicast originating in a packet switched network. FIG. 8 illustrates process 802 using the example of a node in a network having protocols defined by a MoCA specification such as one of those identified above. At step 802, the listener MoCA node receives, via a MoCA ingress port, a stream request protocol (“SRP”) resource reservation message. At step 804, the node may transmit a MoCA pQoS flow request to its network coordinator. The flow request may be for the establishment of a flow with the MoCA ingress port through which the data stream will flow. At step 806, the network coordinator approves the flow request. At step 808, the MoCA node may initiate a p-RR, such as that shown and described above. Alternatively, the node may initiate an RR in every MAP cycle. The p-RR and the RR may be consistent with a TSPEC that may be embedded in the SRP.

[0059] FIG. 9 shows illustrative method 900 for using a p-RR to couple a network coordinator flow request grant to the data flow itself. Process 900 may be implemented for reducing latency in a coordinated shared media network and/or for providing Ethernet AV bridging. In process 900 a p-RR is embedded as extra fields in pQoS Flow Request (see steps 704 and 804 in FIGS. 7 and 8, respectively). At step 902, a MoCA node that desires to transmit data (the “transmitting node”) transmits to its network coordinator a layer 2 management entity (“L2ME”) frame for creating (or updating) a

pQoS flow. At step 904, the network coordinator issues Wave 0, in which all nodes in the MoCA network segment are informed about the proposed pQoS flow, and collects resource allocation information from the MoCA network segment nodes. At step 906, the network coordinator issues Wave 1, in which the MoCA network segment nodes are informed about the coordinators disposition (e.g., grant or denial) of the request for a pQoS flow. At step 908, the network coordinator issues Wave 2, in which the transmitting node transmits to the network coordinator a p-RR (which may be in a form defined in the specification MoCA 1.0, section 3.9.1.1). At step 910, the network coordinator issues Wave 3, in which the transmitting node, and any other interested nodes, are informed that the requested data flow was completed. (Under Wave 3 protocols, the NC notifies other nodes regarding the occurrence of some events, such as the grant or denial of a p-RR. When a p-RR is granted, the NC may notify the nodes regarding the amount of bandwidth that remains available after the p-RR is granted.)

[0060] FIG. 10 shows an illustrative embodiment of device 1000 that may correspond to a node in a managed shared media network, such as package mode network segment 404 (shown in FIG. 4). Device 1000 may reduce latency in a coordinated shared media network and/or may provide Ethernet AV bridging. Device 1000 may include single or multi-chip module 1002, which can be one or more integrated circuits, and which may include logic configured to: request a data flow; grant a data flow request; transmit or receive an RR; transmit or receive a p-RR; transmit or receive a reservation message; or to perform any other suitable logical operations. Device 1000 may include one or more of the following components: I/O circuitry 1004, which may interface with coaxial cable, telephone lines, wireless devices, PHY layer hardware, a keypad/display control device or any other suitable media or devices; peripheral devices 1006, which may include counter timers, real-time timers, power-on reset generators or any other suitable peripheral devices; processor 1008, which may control process flow, and which may generate, receive, grant or deny requests such as requests for data flows and requests for network resource allocation; and memory 1010. Components 1002, 1004, 1006, 1008 and 1010 may be coupled together by a system bus or other interconnections 1012 and may be present on one or more circuit boards such as 1020. In some embodiments, the components may be integrated into a single chip.

[0061] For the sake of clarity, the foregoing description, including specific examples of parameters or parameter values, is sometimes specific to certain protocols such as those identified with the name MoCA™ and/or Ethernet protocols. However, this is not intended to be limiting and the invention may be suitably generalized to other protocols and/or other packet protocols. The use of terms that may be specific to a particular protocol such as that identified by the name MoCA™ or Ethernet to describe a particular feature or embodiment is not intended to limit the scope of that feature or embodiment to that protocol specifically; instead the terms are used generally and are each intended to include parallel and similar terms defined under other protocols.

[0062] It will be appreciated that software components of the present invention including programs and data may, if desired, be implemented in ROM (read only memory) form, including CD-ROMs, EPROMs and EEPROMs, or may be stored in any other suitable computer-readable medium such as but not limited to discs of various kinds, cards of various

kinds and RAMs. Components described herein as software may, alternatively, be implemented wholly or partly in hardware, if desired, using conventional techniques.

[0063] Thus, systems and methods for compensating for managing transmission of data over a shared media communication network and for facilitating Ethernet audiovisual bridging have been provided. Persons skilled in the art will appreciate that the present invention can be practiced using embodiments of the invention other than those described, which are presented for purposes of illustration rather than of limitation. The present invention is limited only by the claims which follow.

What is claimed is:

1. A method for managing transmission of data over a shared media communication network, the method comprising:

receiving from a node a first reservation request, the node being in communication with the network and the reservation request requesting allocation of network resources for transmitting the data; and

establishing a data flow based on the request;

wherein, in the data flow, the node transmits at least two data frames before a second reservation request is received.

2. The method of claim 1 wherein the flow has a duration that is greater than the duration of one medium access plan cycle.

3. The method of claim 1 wherein the reservation request comprises a traffic specification.

4. The method of claim 3 further comprising canceling the flow based on information included in the traffic specification.

5. The method of claim 1 further comprising canceling the flow based on the execution of a link maintenance operation process.

6. The method of claim 1 wherein the establishing comprises allocating network resources during a portion of at least one medium access plan cycle, the portion having a duration that is less than the duration of the medium access plan cycle.

7. The method of claim 1 wherein communication over the shared media network is performed in accordance with protocols defined by a specification available under the trademark MoCA.

8. A device for managing transmission of data over a shared media communication network, the device comprising:

a receiver module in communication with the network, the receiver module configured to receive a first reservation request from a node in communication with the network;

a processor module in communication with the receiver module, the processor module configured to allocate network resources based on the request; and

a transmitter module configured to inform the node of a grant of the request, the grant authorizing the node to transmit at least two data frames before the node transmits a second reservation request.

9. The device of claim 8 wherein the receiver and the transmitter communicate over the shared media network in accordance with protocols defined by a specification available under the trademark MoCA.

10. A method for transmitting multicast information from a first node to a second node, the first node residing in a packet mode network and the second node residing in a packet-switched network, the method comprising:

receiving at the first node a registration request from the second node, the registration request requesting receipt of the multicast information, the registration request including a traffic specification; and

transmitting, using the first node, a reservation request to a network coordinator residing in the packet mode network, the reservation request conforming to the traffic specification.

11. The method of claim 10 wherein the first node transmits at least two frames of multicast data before transmitting a second reservation request to the network coordinator.

12. The method of claim 10 further comprising, when the network coordinator generates medium access plan cycles and the reservation request corresponds to one of the cycles, transmitting at least one additional reservation request, each of the at least one additional reservation requests being transmitted in a respective subsequent cycle.

13. The method of claim 12 wherein communication over the packet mode network is performed in accordance with protocols defined by a specification available under the trademark MoCA.

14. A method for using a packet mode network to transmit data from a first packet-switched network segment to a second packet-switched network segment, the method comprising:

receiving the data from the first segment using a packet mode network ingress port; and

transmitting the data to the second segment using a packet mode network egress port.

15. The method of claim 14 further comprising, using the ingress port, transmitting to a network coordinator a data flow request for data flow from the ingress port to the egress port.

16. The method of claim 14 further comprising, using the ingress port:

transmitting to the network coordinator a reservation request for the allocation of network resources for transmitting the data to the egress port; and

transmitting at least two frames of the data before transmitting a second reservation request to the network coordinator.

17. The method of claim 14 further comprising, when the network coordinator generates medium access plan cycles:

transmitting to the network coordinator a reservation request for the allocation of network resources for transmitting the data to the egress port, the reservation request corresponding to one of the cycles; and

transmitting at least one additional reservation request, each of the at least one additional reservation requests being transmitted in a respective subsequent cycle.

18. The method of claim 17 wherein communication over the packet mode network is performed in accordance with protocols defined by a specification available under the trademark MoCA.

* * * * *