

EXHIBIT D-5

Invalidity of 8,112,104 in view of Apple Prior Art Systems

On information and belief, Apple prior art systems including, utilizing, or otherwise comprising speech synthesis and speech recognition technologies (“Apple Prior Art Systems”) were sold, on–sale, in public use, or otherwise publicly available prior to the earliest priority date of U.S. Patent No. 8,112,104 (“’104 patent”) and qualify as prior art to the ’104 patent at least under pre-AIA 35 U.S.C. §§ 102(a) and/or (b). Apple Prior Art Systems anticipate or, alone or with other references, renders obvious one or more of the Asserted Claims of the ’104 patent.

The chart below illustrates how exemplary Apple Prior Art Systems render the Asserted Claims of the ’104 patent invalid as anticipated or obvious over the Apple Prior Art Systems alone or in view of additional secondary references, including but not limited to the references listed below and prior art identified in Defendant’s Invalidity Contentions. One exemplary Apple Prior Art System identified in this claim chart includes, for example, one or more of the PowerBook G3, iMac, and Power Mac G4 with Apple Macintosh Speech Manager and MacSpeech iListen installed, which were sold, on–sale, in public use, or otherwise publicly available at least as early as January 2003. Documentation describing those features (e.g., “User Guides”) was also available at that time. The PowerBook G3 system is evidenced by, for example, Inside Macintosh Sound_1994 and iListen_15_New_speechrecogniti, which published at least as early as the earliest claimed priority date of the ’104 patent, as well as the additional information cited herein, including at <https://support.apple.com/en-us/112450>.

The citations and descriptions provided herein refer to the Apple Macintosh Speech Manager and MacSpeech iListen system as a representative system demonstrating the relevant prior art features and functionality of the Apple Prior Art Systems, and is not intended to limit the invalidity contentions to this specific implementation. On the contrary, as demonstrated herein, components of the Apple Prior Art Systems could be used interchangeably with compatible systems. Accordingly, the chart identifies other examples of exemplary components that could be substituted for or used interchangeably with the Apple Macintosh Speech Manager and MacSpeech iListen device, including but not limited to IBM ViaVoice. Moreover, and as discussed in greater detail in the following chart, the Apple Prior Art systems may comprise one or more of the following components which could be used interchangeably and include substantially similar structures and functionality:

- Communication Systems: Apple computers available before the priority date of the ’104 patent, including but not limited to: PowerBook G3, iBook G3 (Clamshell), iBook G3 (Dual USB), iBook G4, PowerBook G4 Titanium, PowerBook G4 Aluminum, iMac, iMac (Flat Panel), Power Mac G4;
- Speech Synthesis Technologies compatible with the Communication Systems described above, including but not limited to: Apple Macintosh Speech Manager, MacinTalk, MacinTalk Pro, MacinTalk 2, MacinTalk 3; and
- Speech Recognition Technologies compatible with the Communication Systems described above, including but not limited to: MacSpeech iListen, IBM ViaVoice, IBM ViaVoice Mobility Suite, MacSpeech iDictate.

All Communication Systems are wireless devices (e.g., laptops, desktops, etc.) that are capable of wireless connections with external networks and devices. For example, the Communication Systems utilize or are capable of utilizing components such as USB wireless adapters, PCMCIA wireless adapters, and/or Apple AirPort to enable wireless communications, including with wireless networks and/or other devices such as microphones, speakers, mice, and keyboards. On information and belief, other configurations, models, or embodiments may also satisfy one or more claim limitations, and the inclusion of the examples identified in this chart do not preclude the existence or applicability of alternative implementations. Defendant reserves the right to supplement, amend, or otherwise update these contentions as additional information becomes available, including but not limited to information obtained through discovery.

To the extent IngenioSpec contends the Apple Prior Art Systems do not disclose one or more limitations of the claims, it would have been obvious to combine the Apple Prior Art Systems with the knowledge of a person of ordinary skill in the art (“POSITA”) and/or with one or more of the following references to render the Asserted Claims the ’104 patent invalid:

- US7398209B2 (“Kennewick ’209”) was filed on June 3, 2003, and published on March 4, 2004, and is therefore prior art under at least 35 U.S.C. § 102(e).
- US20100115149A1 (“Ewer ’149”) was filed on December 2, 2002, and published on May 6, 2010, and is therefore prior art under at least 35 U.S.C. § 102 (e).
- Jabra FreeSpeak BT200 was released in 2002, and is therefore prior art under at least 35 U.S.C. § 102(a).

For any claim limitation, Defendant may rely on excerpts cited for any other limitation and/or additional excerpts not set forth fully herein to the extent necessary to provide further explanation for a reference’s disclosure of a limitation. Where an excerpt refers to or discusses a figure or figure items, that figure and any additional descriptions of that figure should be understood to be incorporated by reference as if set forth fully herein.

With respect to the proposed modifications to the Apple Prior Art System, such modification would have been obvious to try, an obvious combination of prior art elements according to known methods to yield predictable results, a simple substitution of one known element for another to obtain predictable results, a use of known techniques to improve a similar devices or method in the same way, an application of a known technique to a known device or method ready for improvement to yield predictable results, a variation of a known work in one field of endeavor for use in either the same field or a different one based on design incentives or other market forces with variations that are predictable to a POSITA, and/or obvious in view of teachings, suggestions, and motivations in the prior art that would have led a POSITA to modify or combine the prior art references.

To the extent that these Invalidity Contentions rely on or otherwise embody particular interpretations of terms or phrases in the asserted claims, Defendant is not proposing any such interpretations as proper interpretations or constructions of those terms or phrases. Various positions put forth in this document are predicated on IngenioSpec’s apparent interpretation of its claims as evidenced by its Infringement Contentions. Those positions are not intended to and do not necessarily reflect Defendant’s interpretation of the true and proper scope of the asserted claims. Defendant further incorporates by reference all reservations of rights set forth in the Invalidity Contentions, as if fully set forth herein.

Defendant’s arguments throughout these invalidity contentions are based on IngenioSpec’s apparent interpretation of the claims in its infringement contentions. Defendant does not agree that such interpretation is proper and reserve the right to argue that the mapped actions do not infringe any claims of the ’104 patent. In the alternative, should the Court or factfinder construe and/or apply the claims in a manner consistent with IngenioSpec’s interpretations in its infringement contentions for the ’104 patent, the claims are invalid for at least the reasons set forth herein.

Claims	Prior Art
1[pre]. A wireless communication system, comprising:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick ’209, Ewer ’149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless communication system.” The Apple Prior Art Systems are “wireless communication system[s].” For example, the Communications Systems included in the Apple Prior Art Systems are wireless devices (e.g., laptops, desktops, etc.) that are capable of wireless connections with external devices. For example, the Communication Systems utilize or

Claims	Prior Art
	are capable of utilizing components such as USB wireless adapters, PCMCIA wireless adapters, and/or Apple AirPort to enable wireless communications, including with wireless networks and/or other devices such as headsets, microphones, speakers, mice, and/or keyboards. The following citations show examples of how the Apple Prior Art Systems disclose this limitation: https://support.apple.com/en-us/112502 (iMac Flat Panel Technical Specifications): “Wireless networking – Airport ready. Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b DSSS compliant” “Headphone jack” “Built-in microphone” “Total of five USB 1.1 ports” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink . . . , AOL” https://support.apple.com/en-us/112501 (iMac Technical Specifications): “Support for optional 11-Mbps AirPort Card; IEEE 802.11 DSSS compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply)” “Support for external USB audio devices” “ Microsoft Internet Explorer, Microsoft Outlook Express, Netscape Communicator” https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications): “Wireless networking – AirPort ready” “Ready for wireless networking: Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply).” “Support for external third-party USB digital audio peripherals such as microphones, speakers, and MIDI devices” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink” Inside Macintosh Sound_1994 at 1-10 Some recent models of Macintosh computers contain built-in sound hardware that extends the Apple Sound Chip’s features. For example, Macintosh computers with built-in microphones include the


Claims

Prior Art

Enhanced Apple Sound Chip (EASC). Some Macintosh computers contain DSP chips that provide very high-quality sound (16-bit stereo sound, at rates up to 44 kHz). There are also NuBus expansion cards available from third-party developers that provide other audio DAC hardware.

Kennewick '209 also discloses a speech unit in communication with a main unit, which is an example of the claimed "wireless communication system." The following citations show examples of where Kennewick '209 discloses this limitation.

Kennewick '209 at FIG. 1




Kennewick '209 at 10:33–48

The one or more speech units 128 and the main unit 98 may communicate over a communication link. The communication link can include a wired or wireless link. According to one embodiment, the communication link comprises an RF link. The transceiver 130 on the speech unit communicates coded speech data bi-directionally over the communication link with the transceiver 126 on the main unit 98. According to another embodiment, RF link may use any standard local area wireless data protocols including the IEEE 802.11, Bluetooth or other standards. Alternatively, an infrared data link conforming to any suitable standard such as IrDA or other infrared standards can be used. In an alternative embodiment, wires connect the speech unit 128 and the main unit 98, eliminating the need for one speech


Claims	Prior Art
	coder 138. Other wired or wireless analog or digital transmission techniques can be used. Kennewick '209 at 10:65–11:10 The agents 106 return results of questions as responses to users. The response may be created using the results of information queries, the system personality 108 and the user preferences or other data in the user profile 110. The agents 106 generally present these results using the speech unit 128. The agents 106 create a response string, which is sent to the text to speech engine 124. The text to speech engine 124 creates the required utterances, which may be encoded and compressed by the speech coder 122. Once coded, the utterances are transmitted from the main unit 98 by the transceiver 126 to the transceiver 130 on the speech unit 128. The utterance is then decoded and decompressed by the speech coder 138 and output by the speaker 136. Ewer '149 similarly discloses communications between a headset and communication device such as a phone. The following citations show examples of where Ewer '149 discloses this limitation. Ewer '149 at Abstract A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on processor-based hosts using digital signaling are disclosed. The signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The signaling module may include a connection status detector and a status signal generating processor. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed. Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the

Claims	Prior Art
	term “connection” utilized herein generally refers to both wired and wireless connections. Jabra FreeSpeak similarly discloses communications between a headset and communication device such as a phone. The following citations show examples of where Jabra FreeSpeak discloses this limitation. Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak™— cord-free comfort and hands-free convenience in a Bluetooth® headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators. Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones.*

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 2 Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02  • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide Jabra FreeSpeak Datasheet at Page 2 About Bluetooth Bluetooth is a radio technology developed to deliver short-range wireless mobility. Bluetooth technology eliminates wires between devices, such as mobile phones and headsets, and enables communication between devices that are in close proximity to one another. Bluetooth facilitates fast, secure transmissions of both voice and data. More information is available at www.bluetooth.com.  A POSITA would have been motivated to combine any combination of the Apple Prior Art Systems with the teachings of Kennewick ‘209, Ewer ‘149, and Jabra FreeSpeak to meet the claim limitations. Kennewick ‘209 is configured to utilize a wireless speech unit (e.g., a headset) in combination with a main device (e.g., a communication device). Ewer ‘149 similarly uses a headset in combination with a communication device. And Jabra FreeSpeak is one such headset used in combination with a communication device. Apple Prior Art Systems are configured to work with microphones on various devices, including computers and PDAs. And as discussed above the Communication Systems were wireless communication systems that are compatible with external audio devices. A POSITA would have understood that a wireless headset such as the ear piece disclosed Kennewick ‘209, the headset disclosed in Ewer ‘149, or the Jabra FreeSpeak would be compatible with Communication Systems to further provide a wireless microphone and speaker to a user. Further, as demonstrated by the Jabra FreeSpeak BT200, such wireless headset connectivity already existed and could be implemented in devices that did not have wireless capabilities. Jabra FreeSpeak Datasheet at Pages 1-2. A POSITA would thus be motivated to modify Apple Prior Art Systems with the teachings of Kennewick ‘209 and Ewer ‘149.

Claims	Prior Art
	A POSITA would have had a reasonable expectation of success in combining Apple Prior Art Systems with the teachings of Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak. Apple Prior Art Systems include and provide connection to headsets, microphones, and/or speakers. <i>See, e.g.,</i> Inside Macintosh Sound_1994 at 1-9 (“A Macintosh computer produces sound when the Sound Manager sends some data through a sound channel to the available audio hardware, usually at the request of an application. The audio hardware is a digital-to-analog converter (DAC) that translates digital sound data into analog audio signals. Those signals are then sent to the internal speaker, to a sound output connector (to which the user can connect headphones, external speakers, or sound amplification equipment), or to other sound output hardware.”); 2-23 (“Some Macintosh computers play <i>only</i> the left channel of stereo sounds out the internal speaker. Other machines (for example, the Macintosh SE/30 and Macintosh IIsi) mix both channels together before sending a signal to the internal speaker. You can use the Gestalt function to determine if a particular machine mixes both left and right channels to the internal speaker. All Macintosh computers except the Macintosh SE and the Macintosh Plus, however, play stereo signals out the headphone jack”). Further, Jabra FreeSpeak BT200 discloses an adapter designed to provide wireless connectivity capabilities to devices not otherwise designed for such communications. Jabra FreeSpeak BT200 brochure, Page 1–2. Thus, a POSITA would merely need to incorporate the wireless connectivity of the headset as disclosed by Kennewick ‘209 (using, for example, the teachings of the Jabra FreeSpeak BT200) into the wired headsets disclosed in and supported by Apple Prior Art Systems. Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak are thus easily combined based on familiar components, and it would have been straightforward to modify the wired headset to create a single wireless headset. To the extent that Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[1a] a wireless headset; and	Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak expressly or inherently disclose “a wireless headset.” For the reasons discussed above and in 1[pre], the Communications Systems included in the Apple Prior Art Systems are wireless devices (e.g., laptops, desktops, etc.) that connect with external devices, including wireless headsets, e.g., via USB wireless adapters, PCMCIA wireless adapters, and/or Apple AirPort. <i>See</i> above claim 1[pre].

Claims	Prior Art
	https://support.apple.com/en-us/112502 (iMac Flat Panel Technical Specifications): “Wireless networking – Airport ready. Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b DSSS compliant” “Headphone jack” “Built-in microphone” “Total of five USB 1.1 ports” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink . . . , AOL” https://support.apple.com/en-us/112501 (iMac Technical Specifications): “Support for optional 11-Mbps AirPort Card; IEEE 802.11 DSSS compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply)” “Support for external USB audio devices” “Microsoft Internet Explorer, Microsoft Outlook Express, Netscape Communicator” https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications): “Wireless networking – AirPort ready” “Ready for wireless networking: Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply).” “Support for external third-party USB digital audio peripherals such as microphones, speakers, and MIDI devices” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink” Inside Macintosh Sound_1994 at 1-10 Some recent models of Macintosh computers contain built-in sound hardware that extends the Apple Sound Chip’s features. For example, Macintosh computers with built-in microphones include the Enhanced Apple Sound Chip (EASC). Some Macintosh computers contain DSP chips that provide very high-quality sound (16-bit stereo sound, at rates up to 44 kHz). There are also NuBus expansion cards available from third-party developers that provide other audio DAC hardware.

Claims	Prior Art
	As discussed below, wireless headsets were known to a POSITA and indeed commercially available. A POSITA would have known and found it obvious to connect a wireless headset to the Apple Prior Art systems using known methods, such as those expressly enabled by the Apple Prior Art Systems. A POSITA would have reasonably expected success in doing so because it would have been nothing more than the combination of well-known components (the wireless connection capabilities of the Communication Systems with known wireless headsets) to achieve the predictable and desired result of enabling wireless audio communication. For example, Kennewick '209 discloses a speech unit which is an example of the claimed “wireless headset.” The following citations show examples of where Kennewick '209 discloses this limitation. Kennewick '209 at FIG. 1  FIG. 1 Kennewick '209 at 10:33–48 The one or more speech units 128 and the main unit 98 may communicate over a communication link. The communication link can include a wired or wireless link. According to one embodiment, the communication link comprises an RF link. The transceiver 130 on the speech unit communicates coded speech data bi-directionally over the communication link with the transceiver 126 on the main unit 98. According to another embodiment, RF link may use any standard local area wireless data protocols

Claims	Prior Art
	including the IEEE 802.11, Bluetooth or other standards. Alternatively, an infrared data link conforming to any suitable standard such as IrDA or other infrared standards can be used. In an alternative embodiment, wires connect the speech unit 128 and the main unit 98, eliminating the need for one speech coder 138. Other wired or wireless analog or digital transmission techniques can be used. Kennewick '209 at 10:65–11:10 The agents 106 return results of questions as responses to users. The response may be created using the results of information queries, the system personality 108 and the user preferences or other data in the user profile 110. The agents 106 generally present these results using the speech unit 128. The agents 106 create a response string, which is sent to the text to speech engine 124. The text to speech engine 124 creates the required utterances, which may be encoded and compressed by the speech coder 122. Once coded, the utterances are transmitted from the main unit 98 by the transceiver 126 to the transceiver 130 on the speech unit 128. The utterance is then decoded and decompressed by the speech coder 138 and output by the speaker 136. Kennewick '209 at 15:25–31 According to an embodiment of the invention, the microphone array, filters and speech coder 138 are physically separated from the main unit 98 into a speech unit 128, and connected using a wireless link. Since bandwidth on a wireless connection is at a premium, the speech coder dynamically adapts the digitization rate and compression of the captured speech. Kennewick '209 at 16:17–23 In an alternative embodiment, the array microphone can be replaced by a set of physically distributed microphones or a headset worn by the user. The distributed microphones can be placed in different parts of a room or in different rooms. The distributed microphones can create a three-dimensional array to improve signal to noise ratio. The headset may use a wireless or wired connection. Ewer '149 similarly discloses a wireless headset. The following citations show examples of where Ewer '149 discloses this limitation. Ewer '149 at Abstract A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on processor-based hosts using digital signaling are disclosed. The signaling system includes a signaling module for communicating with the

Claims	Prior Art
	host and a headset selectively in communication with the signaling module for use with the application software. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The signaling module may include a connection status detector and a status signal generating processor. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed. Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the term “connection” utilized herein generally refers to both wired and wireless connections. Jabra FreeSpeak is also a “wireless headset.” The following citations show examples of where Jabra FreeSpeak discloses this limitation. Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak™— cord-free comfort and hands-free convenience in a Bluetooth® headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators.

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones.* Jabra FreeSpeak Datasheet at Page 2 Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02 • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide  Jabra FreeSpeak Datasheet at Page 2 About Bluetooth Bluetooth is a radio technology developed to deliver short-range wireless mobility. Bluetooth technology eliminates wires between devices, such as mobile phones and headsets, and enables communication between devices that are in close proximity to one another. Bluetooth facilitates fast, secure transmissions of both voice and data. More information is available at www.bluetooth.com.  A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak to meet the claim limitations with reasonable expectation of success for the reasons described above herein and with respect to claim 1[pre]. Specifically, a POSITA would have

Claims	Prior Art
	been motivated to implement a wireless headset as disclosed by these references/systems in the Apple Prior Art Systems to facilitate the desired and expected result of enabling wireless audio communication within the system. To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[1b] a wireless communication device that supports wireless communication, including messaging, for a first user, said wireless communication device capable of wirelessly communicating with the wireless headset,	Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “a wireless communication device that supports wireless communication, including messaging, for a first user, said wireless communication device capable of wirelessly communicating with the wireless headset.” For the reasons discussed above and in 1[pre], [1a], the Communications Systems included in the Apple Prior Art Systems are wireless devices (e.g., laptops, desktops, etc.) that support wireless communication, including wireless networking. The Communications Systems also supported wireless communications for messaging, including email, for a user. And, as discussed in 1[pre], [1a], the Communications Systems, alone or as modified as discussed above were capable of wirelessly communicating with the wireless headset. See above claim 1[pre], [1a]. https://support.apple.com/en-us/112502 (iMac Flat Panel Technical Specifications): “Wireless networking – Airport ready. Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b DSSS compliant” “Headphone jack” “Built-in microphone” “Total of five USB 1.1 ports” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink . . . , AOL” https://support.apple.com/en-us/112501 (iMac Technical Specifications): “Support for optional 11-Mbps AirPort Card; IEEE 802.11 DSSS compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply)”


Claims	Prior Art
	“Support for external USB audio devices” “Microsoft Internet Explorer, Microsoft Outlook Express, Netscape Communicator” https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications): “Wireless networking – AirPort ready” “Ready for wireless networking: Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply).” “Support for external third-party USB digital audio peripherals such as microphones, speakers, and MIDI devices” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink” Inside Macintosh Sound, 1994 at 1-10 Some recent models of Macintosh computers contain built-in sound hardware that extends the Apple Sound Chip’s features. For example, Macintosh computers with built-in microphones include the Enhanced Apple Sound Chip (EASC). Some Macintosh computers contain DSP chips that provide very high-quality sound (16-bit stereo sound, at rates up to 44 kHz). There are also NuBus expansion cards available from third-party developers that provide other audio DAC hardware. A POSITA would have understood that the Communications Systems included in the Apple Prior Art Systems are wireless communication systems configured to facilitate communication between users over one or more networks, including by sending text and email messages, as was common and conventional at the time. As discussed below, wireless headsets were known to a POSITA and indeed commercially available. A POSITA would have known and found it obvious to connect a wireless headset to the Apple Prior Art systems using known methods, such as those expressly enabled by the Apple Prior Art Systems. A POSITA would have reasonably expected success in doing so because it would have been nothing more than the combination of well-known components (the wireless connection capabilities of the Communication Systems with known wireless headsets) to achieve the predictable and desired result of enabling wireless audio communication.

Claims

Prior Art

Kennewick '209 also discloses a main unit, which is an example of the claimed “wireless communication device” that wirelessly communicates with the speech unit (e.g., the wireless headset) configured to support wireless messaging such as email. The following citations show examples of where Kennewick '209 discloses this limitation.

Kennewick '209 at FIG. 1



Kennewick '209 at 5:47–6:6

In addition to the ability to process user questions and commands, the invention may provide local or remote control functions for the system or for other devices. Users may initiate commands locally or remotely. Several techniques for executing remote operations are possible, for example, the use of a telephone or other audio connection. Other remote command techniques may be used. The processing of these commands is performed in a nearly identical manner to a question. One difference being that the result of the command is generally an action rather than a response. Examples of control operations supported by the invention include, finding and playing music, film, games or other entertainment media, control of the appliances or other electronic or electro-mechanical devices, security systems control, audio, text and video conferencing (including Voice over IP conferencing), operating integrated messaging environments and other control operations. By using an optional analog or digital telephone interface or a data network interface, the invention can provide remote control capability of the computer

Claims	Prior Art
	or and auxiliary devices connected to the computer. For example, if the computer is connected via an interface to one or more auxiliary devices, a remote user can use this capability to remotely control or program the operation of the auxiliary device. For example, this enables recording of a television program, checking the status of a home security system, checking voice mail messages, hearing and responding to email messages, and much more. Kennewick '209 at 7:47–49 automatic telephone dialing, reading and sending emails and pages by voice and other communications control functions; Kennewick '209 at 9:66–10:7 The system 90 may include a main unit 98 and a speech unit 128. Alternatively, the system 98 may only comprise of the main unit 98, the speech unit 128 being a completely separate system. The event manager 100 may mediate interactions between other components of the main unit 98. The event manager 100 provides a multi-threaded environment allowing the system 98 to operate on multiple commands or questions from multiple user sessions without conflict and in an efficient manner, maintaining real-time response capabilities. Kennewick '209 at 10:21–32 The main user interface for the system 90, according to an embodiment of the invention, is through one or more speech units 128. The speech unit 128 includes one or more microphones, for example array microphone 134, to receive the utterances of the user. The speech received at the microphone 134 may be processed by filter 132 and passed to the speech coder 138 for encoding and compression. In one embodiment, a transceiver module 130 transmits the coded speech to the main unit 98. Coded speech received from the main unit 98 is detected by the transceiver 130, then decoded and decompressed by the speech coder 138 and annunciated by the speaker 136. Kennewick '209 at 10:33–48 The one or more speech units 128 and the main unit 98 may communicate over a communication link. The communication link can include a wired or wireless link. According to one embodiment, the communication link comprises an RF link. The transceiver 130 on the speech unit communicates coded speech data bi-directionally over the communication link with the transceiver 126 on the main unit 98. According to another embodiment, RF link may use any standard local area wireless data protocols


Claims	Prior Art
	including the IEEE 802.11, Bluetooth or other standards. Alternatively, an infrared data link conforming to any suitable standard such as IrDA or other infrared standards can be used. In an alternative embodiment, wires connect the speech unit 128 and the main unit 98, eliminating the need for one speech coder 138. Other wired or wireless analog or digital transmission techniques can be used. Kennewick '209 at 10:65–11:10 The agents 106 return results of questions as responses to users. The response may be created using the results of information queries, the system personality 108 and the user preferences or other data in the user profile 110. The agents 106 generally present these results using the speech unit 128. The agents 106 create a response string, which is sent to the text to speech engine 124. The text to speech engine 124 creates the required utterances, which may be encoded and compressed by the speech coder 122. Once coded, the utterances are transmitted from the main unit 98 by the transceiver 126 to the transceiver 130 on the speech unit 128. The utterance is then decoded and decompressed by the speech coder 138 and output by the speaker 136. Ewer '149 similarly discloses a telephony (e.g., “wireless communication device”) in wireless communication with the headset which supports messaging. The following citations show examples of where Ewer '149 discloses this limitation. Ewer '149 at Abstract A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on processor-based hosts using digital signaling are disclosed. The signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The signaling module may include a connection status detector and a status signal generating processor. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed. Ewer '149 at [0011] A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on personal computers or other

Claims	Prior Art
	processor-based hosts using digital signaling are disclosed. Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the term “connection” utilized herein generally refers to both wired and wireless connections. Ewer '149 at [0028] The host device 104 is generally any suitable processor-based device such as a personal computer (PC), a personal digital assistant (PDA), a digital music player (e.g., MP3 player), a video player (e.g., DVD player), a video game player, and a processor-based telephone. The host device 104 executes application software such as a telephony application software that uses the headset 110, for example, for receiving the user's voice as input and/or for outputting sounds to the user as output. When the user disconnects the headset 110 at the connector 108, the headset adapter 106 transmits a digital flag signal indicating a change in the headset connection status to the host 104 running the application software. Similarly, when the user reconnects the headset 110 at the connector 108, the headset adapter 106 transmits the flag signal to indicate a change in the headset connection status to the application software running on the host 104. Thus, as is evident, the adapter 106 functions at least in part as a headset connection status signaling module to the host 104. Jabra FreeSpeak is also a “wireless headset.” The following citations show examples of where Jabra FreeSpeak discloses this limitation.

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak™— cord-free comfort and hands-free convenience in a Bluetooth™ headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators. Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones*. Jabra FreeSpeak Datasheet at Page 2 Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02  • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide


Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 2  A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209 and/or Ewer '149 to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim 1[pre]. Specifically, a POSITA would have been motivated to implement a wireless headset as disclosed by Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to the Apple Prior Art Systems, and a POSITA would have a reasonable expectation of success in implementing a wireless headset to the Apple Prior Art Systems based on the disclosure of Jabra FreeSpeak BT200. To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, and Ewer '149 those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[1c] wherein said wireless communication device is configured to perform the method of: receiving an incoming text message from a second user;	Apple Prior Art Systems expressly or inherently disclose “wherein said wireless communication device is configured to perform the method of: receiving an incoming text message from a second user.” Apple Prior Art Systems, including the Speech Recognition Technologies, run on or are combinable with devices such as computers and PDAs, including the Communication Systems, which are configured to receive text messages such as emails. Apple Prior Art Systems, including the Speech Recognition Technologies, run on or are combinable with devices such as computers and PDAs, including the Communication Systems, which are configured to receive text messages such as emails. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. <i>See above claim [1b]</i>

Claims	Prior Art
	https://support.apple.com/en-us/112502 (iMac Flat Panel Technical Specifications): “Wireless networking – AirPort ready. Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b DSSS compliant” “Headphone jack” “Built-in microphone” “Total of five USB 1.1 ports” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink . . ., AOL” https://support.apple.com/en-us/112501 (iMac Technical Specifications): “Support for optional 11-Mbps AirPort Card; IEEE 802.11 DSSS compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply)” “Support for external USB audio devices” “Microsoft Internet Explorer, Microsoft Outlook Express, Netscape Communicator” https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications): “Wireless networking – AirPort ready” “Ready for wireless networking: Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply).” “Support for external third-party USB digital audio peripherals such as microphones, speakers, and MIDI devices” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink” IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Mobility Suite gives you voice access to your PIM applications. It allows you to use simple, natural voice commands such as "Go to Contacts" and "Read John Smith's home phone number". When you "Go to Mail", you can have all of your mail read to you, or you can read only selected messages. In the "Calendar" function, you can check your calendar for your next appointment, or add new ones. You can easily keep track of your busy schedule, all by using voice commands. In addition, this IBM software gives you the added convenience of being able to listen to rather than read information

Claims	Prior Art
	and messages. IBM ViaVoice Mobility Suite v.4 at Page 2  The screenshot shows a table titled "IBM Embedded ViaVoice Mobility Suite Features". It has two columns: "Inbox (Mail) Commands" and "Sample Commands". Under "Inbox (Mail) Commands", it says "Reads incoming mail; gives you the opportunity to reply and send responses to others." Under "Sample Commands", it lists: "Read mail", "Read new mail", "Next message", "Create reply", and "Send reply". IBM Mobility Suite Introduction at 3 Welcome to IBM Embedded ViaVoice Mobility Suite, a productivity application that controls your Personal Digital Assistant (PDA) by voice. In response to your commands, IBM Embedded ViaVoice Mobility Suite will help manage your appointments, contacts, tasks, and mail. IBM Mobility Suite Introduction at 5 • To hear important information, for example, say Read mail or Read my very next appointment. • To hear prompts, messages, or requested information again, say Repeat that. IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in


Claims	Prior Art
	the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message. • read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. A POSITA would have understood that the Apple Prior Art Systems, including the Communication Systems running or in combination with the Speech Recognition Technologies, are configured to incoming text messages and emails from another user. Because the Speech Recognition Systems expressly contemplate reading received messages, a POSITA would have understood that for this functionality to operate, the Communication System would first have received a message to be read aloud, thus receiving an incoming text message from a second user, as claimed. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[1d] converting the incoming text message to an incoming audio message;	Apple Prior Art Systems expressly or inherently disclose “converting the incoming text message to an incoming audio message.” Apple Prior Art Systems, including the Speech Recognition Systems as running on or in combination with the Communication Systems, disclose synthesizing speech using various functions within Mac computers such as a “read” command, which performs text-to-audio conversion. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. <i>See above claim [1b]</i>

Claims	Prior Art
	https://support.apple.com/en-us/112502 (iMac Flat Panel Technical Specifications): “Wireless networking – Airport ready. Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b DSSS compliant” “Headphone jack” “Built-in microphone” “Total of five USB 1.1 ports” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink . . . , AOL” https://support.apple.com/en-us/112501 (iMac Technical Specifications): “Support for optional 11-Mbps AirPort Card; IEEE 802.11 DSSS compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply)” “Support for external USB audio devices” “Microsoft Internet Explorer, Microsoft Outlook Express, Netscape Communicator” https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications): “Wireless networking – AirPort ready” “Ready for wireless networking: Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply).” “Support for external third-party USB digital audio peripherals such as microphones, speakers, and MIDI devices” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink” IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Mobility Suite gives you voice access to your PIM applications. It allows you to use simple, natural voice commands such as "Go to Contacts" and "Read John Smith's home phone number". When you "Go to Mail", you can have all of your mail read to you, or you can read only selected messages. In the "Calendar" function, you can check your calendar for your next appointment, or add new ones. You can easily keep track of your busy schedule, all by using voice commands. In addition, this IBM software gives you the added convenience of being able to listen to rather than read information


Claims	Prior Art				
	and messages. IBM ViaVoice Mobility Suite v.4 at 1 End-to-end voice solutions IBM can offer distributed end-to-end voice solutions where some of the speech processing resides on the client device (cell phones etc.) and integrates with additional voice functionality at the server level and large-scale telephony/unified messaging systems. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Enterprise Edition is IBM's flagship speech recognition and text-to-speech product. Representing a dramatic shift in the evolution of embedded speech products, IBM Embedded ViaVoice Enterprise Edition has an unlimited active vocabulary, limited only by the amount of memory on devices, plus a dynamic and swappable 4,000 word phonetic flat list. This feature enables embedded applications to access voice-enabled databases with millions of records. A typical application would be a voice-enabled navigation database of street names for a GPS system in your car. IBM ViaVoice Mobility Suite v.4 at Page 2  The screenshot shows a table titled "IBM Embedded ViaVoice Mobility Suite Features". It has two columns: "Inbox (Mail) Commands" and "Sample Commands". The "Inbox (Mail) Commands" column contains the text "Reads incoming mail; gives you the opportunity to reply and send responses to others." The "Sample Commands" column contains a list of five commands: "Read mail", "Read new mail", "Next message", "Create reply", and "Send reply". <table border="1"> <thead> <tr> <th>Inbox (Mail) Commands</th><th>Sample Commands</th></tr> </thead> <tbody> <tr> <td>Reads incoming mail; gives you the opportunity to reply and send responses to others.</td><td> • "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply" </td></tr> </tbody> </table> IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list.	Inbox (Mail) Commands	Sample Commands	Reads incoming mail; gives you the opportunity to reply and send responses to others.	• "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply"
Inbox (Mail) Commands	Sample Commands				
Reads incoming mail; gives you the opportunity to reply and send responses to others.	• "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply"				

Claims	Prior Art
	• delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message. • read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. IBM ViaVoice Simply Dictation user guide at Page 26 Reading Text Now that you have learned to enter and format text using voice commands, it's time to have ViaVoice read the sentences back to you. Highlight the desired text block and say one of the following voice commands. SAY> BEGIN READING SAY> READ SELECTION ViaVoice reads the highlighted text using the voice selected in the Voices panel in ViaVoice Preferences. On applications with copy and selection capabilities, the Read Selection menu item reads the text using the voice currently selected in the Voices panel in ViaVoice Preferences. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 48 Reading Your Text Provided your application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back text aloud from almost any source, including documents and e-mail messages. 1. Open the application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected system voice. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 49 Changing Voices  You can select the voice used when reading text. 1. Click the ViaVoice Preferences button on the VoiceCenter drawer. 2. Click the Voices icon. 3. From the Voice menu, select a new voice. 4. Press Play to hear a sample of your selection. 5. Click Apply, and close the window. 6. Any text read back from this point uses the voice you just selected.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 61  SpeakPad Menu The SpeakPad menus contain items that are similar in functionality to other word processors. SpeakPad File Edit Format Dictation Window Help About SpeakPad Preferences... Dictation Menu The following items are available on the Dictation Menu. • Playback Selection/Stop Playback - This option allows you to play back your dictated text in your own voice. ViaVoice can play back up to 1,000 words (punctuation is included in the word count). The playback feature is useful in many instances, such as when you are correcting text. For more information, see "Playing Back Dictated Text in SpeakPad" on page 48. • Read Selection/Stop Reading - This option reads any text you have highlighted in a computer-generated voice. You can use this feature to help you proofread your documents. For more information, see "Reading Your Text" on page 48. • Voice Preferences - This option displays the Voices window, which allows you to change the computer-generated voice. For more information, see "Voices Preferences" on page 65. Inside Macintosh Sound_1994 at Page 4–3 This chapter describes the Speech Manager, the part of the Macintosh system software that provides a standardized method for Macintosh applications to generate synthesized speech. Inside Macintosh Sound_1994 at Page 4–3 You need to read this chapter if you want your application to be able to generate speech. For example, you may want your application to incorporate the capability to speak its dialog box messages to the user. A word-processing application might use the Speech Manager to implement a command that speaks a selected section of a document to the user. A multimedia application might use the Speech Manager to provide a narration of a QuickTime movie instead of including sampled-sound data on a movie track. Because sound samples can take up large amounts of room on disk, using text in place of sampled sound is extremely efficient.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–3 This chapter begins with an introduction to the speech generation process and then discusses how you can ▪ check for the availability of the Speech Manager ▪ create and dispose of speech channels ▪ generate speech with different voices ▪ obtain information about and change speech channel settings ▪ start and stop speech production ▪ synchronize speech production with other activities by using callback procedures ▪ embed Speech Manager commands within text to make it more understandable Inside Macintosh Sound_ 1994 at Page 4–4 You can use the Speech Manager to incorporate synthesized speech into your application. This section provides an overview and describes the basic concepts of the Speech Manager, and it outlines the process that the Speech Manager uses to convert text into speech. The Speech Manager converts text into sound data, which it passes to the Sound Manager to play through the current sound output device. The Speech Manager's interaction with the Sound Manager is transparent to your application, so you don't need to be familiar with the Sound Manager to take advantage of the Speech Manager's capabilities. Inside Macintosh Sound_ 1994 at Page 4–4 Figure 4-1 illustrates the speech generation process. Your application can initiate speech generation by passing a string or a buffer of text to the Speech Manager. The Speech Manager is responsible for sending the text to a speech synthesizer, a component that contains executable code that manages all communication between the Speech Manager and the Sound Manager. A synthesizer is usually contained in a resource in a file within the System Folder. A synthesizer is like a speech engine. It uses built-in dictionaries and pronunciation rules to help determine how to pronounce text. You can provide custom pronunciation dictionaries as described in the section “Including Pronunciation Dictionaries” beginning on page 4-36.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4-4 <u>Figure 4-1</u> The speech generation process  <pre> graph LR Text[Text] --> SM[Speech Manager] SM --> SS[Speech synthesizer] SS --> SD[Sound Manager] SD --> AH[Audio hardware] </pre> Inside Macintosh Sound_ 1994 at Page 4-5 <u>Voices</u> Your application can use the system default voice to generate speech or it can specify that the Speech Manager use a particular voice that is available on the current computer system. A voice is a set of characteristics defined in parameters that specify a particular quality of speech. Just as different people's voices have different tonal qualities, so too can different voices have different qualities. A synthesized voice might sound male or female and might sound like an adult or a child. Some voices sound distinctively synthetic, while others sound more like real people. Figure 4-2 shows how the Speech Manager uses speech channels to synthesize speech with different voices. Inside Macintosh Sound_ 1994 at Page 4-6 Most of the time, your application designates the voice that speaks text, and usually that is the default voice. Based on the needs of your users and the way in which you expect them to use voices in your application you can provide access to voices in a number of different ways. You could include access to selecting voices in a dialog box that is available from a menu item such as Voices... Any application that allows users to choose among voices requires additional information about the available voices beyond the information provided by a voice specification record (described in detail on page 4-46), whose data should never be presented to the user. Such additional information might include the name of the voice as well as what script and language it supports.


Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–11 To take advantage of most of the Speech Manager's features, you must allocate a speech channel to pass to Speech Manager functions and dispose of the speech channel when you are finished using it. "Creating, Using, and Disposing of a Speech Channel" beginning on page 4-13 demonstrates how you do this and shows how you can use the <code>SpeakText</code> function to start speech generation from a buffer of text. Some applications permit users to choose a voice from those available to be used for speech generation. The <code>CountVoices</code>, <code>GetIndVoice</code>, and <code>GetVoiceDescription</code> functions support this capability. "Working With Different Voices" beginning on page 4-14 shows how you can use these functions to choose among available voices. Inside Macintosh Sound_ 1994 at Page 4–14 The <code>MyUseSpeechChannel</code> function defined in Listing 4-2 creates a default speech channel using the default system voice. You pass <code>NIL</code> in the first parameter to use the system default voice. You must also pass a global variable to <code>NewSpeechChannel</code> in which is returned a valid speech channel. Once the channel exists, then you can use the <code>SpeakText</code> function to generate speech. To generate synthesized speech, you pass in the channel allocated by <code>NewSpeechChannel</code> in the first parameter, and then you pass a pointer to the text that you want to speak as well as the length of the text that you want the Speech Manager to attempt to speak. That is, you can pass a pointer to a buffer of text that is 500 bytes long, but specify that only the first 10 bytes get spoken. Then <code>MyUseSpeechChannel</code> uses the <code>SpeechBusy</code> function in a <code>WHILE</code> loop to allow the text to be completely spoken before disposing of the channel. Inside Macintosh Sound_ 1994 at Page 4–14 <u>Working With Different Voices</u> When you work with speech channels, you can set a voice for a particular channel. When you set a voice, you may want to filter out certain of its characteristics in order to identify the one you want. For example, in an educational software application for elementary school students, you may want to use only children's voices. In order to choose the voice you want, you get a voice description record that contains information about a voice such as the size of the voice, the name of the voice, the age and gender of the voice, and the synthesizer with which it works. You can get the number of available voices using the <code>CountVoices</code> function. You can cycle through the available voices and identify the one you want to use by using the <code>GetIndVoice</code> function. Then you fill out a voice description record using the <code>GetVoiceDescription</code> function. Listing 4-3 shows how to get identifying information about a voice.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–23 <u>Writing Embedded Speech Commands</u> Embedded speech commands allow you to customize the quality of speech output by fine tuning it. You can make speech much easier to understand than the default way in which text is spoken by a synthesizer. An embedded speech command is a command embedded within a text buffer to be spoken by the Speech Manager that causes the Speech Manager to take a certain action. For example, you could use an embedded speech command to emphasize a particular word in a text string to make it stand out to the user. Inside Macintosh Sound_ 1994 at Page 4–30 <u>Examples of Embedded Speech Commands</u> If you use just a few of the embedded speech commands, you can markedly increase the understandability of text spoken by your application. Your application knows more about the speech being produced than a speech synthesizer does. A synthesizer speaks text according to a predetermined set of rules about language production. Therefore, the voices available on a Macintosh computer with the Speech Manager installed sound very synthetic and sometimes robotic because the pronunciation rules are formalized. You can make the speech produced by the synthesizer sound a lot more human by observing some simple rules of human speech and embedding speech commands in text according to these conventions. The techniques presented in this section could be applied when your application is having a dialog with the user or speaking some error messages or announcements. Inside Macintosh Sound_ 1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ■ generate speech and then pause or stop it ■ obtain information about an individual voice or all voices ■ create and dispose of speech channels ■ obtain the Speech Manager's version and status ■ change the rate or pitch of speech

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–73 <u>Changing Speech Attributes</u> To determine the rate and pitch at which a speech channel is processing text, you can use the <code>GetSpeechRate</code> and <code>GetSpeechPitch</code> functions. The <code>SetSpeechRate</code> and <code>SetSpeechPitch</code> functions allow you to change rate and pitch. The most robust of the Speech Manager's routines are the <code>GetSpeechInfo</code> and <code>SetSpeechInfo</code> functions. These allow you to obtain many types of information about a speech channel and to change many settings of a speech channel. To specify the operation that you wish to perform, you must pass <code>GetSpeechInfo</code> or <code>SetSpeechInfo</code> a selector. A full list of selectors is provided in "Speech Information Selectors" beginning on page 4-39. Inside Macintosh Sound_ 1994 at Page 4–76 <u>SetSpeechPitch</u> You can use the <code>SetSpeechPitch</code> function to set the speech pitch on a designated speech channel. <pre>FUNCTION SetSpeechPitch (chan: SpeechChannel; pitch: Fixed) : OSErr;</pre> <code>chan</code> The speech channel whose pitch you wish to set. <code>pitch</code> The new pitch for the speech channel, expressed as a fixed-point frequency value. A POSITA would have understood that the Apple Prior Art Systems, including the Communication Systems running or in combination with the Speech Recognition Technologies or Speech Synthesis Technologies, are configured to convert text to audio for eventual output and presentation to the user. When implemented with a wireless headset, as explained for 1[pre]-[1c], a POSITA would further have understood the audio output message would have been presented to the user via the wireless headset because the wireless headset would have been a part of, connected to, or combined with the Communication System, for the purpose of presenting audio to a user. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.

Claims	Prior Art
[1e] presenting the incoming audio message for the first user via the wireless headset;	Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “presenting the incoming audio message for the first user via the wireless headset.” Apple Prior Art Systems, including the Communication Systems running or in combination with the Speech Synthesis Technologies and Speech Recognition Technologies, present audio to a user. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. <i>See above claim 1[pre]-[1b]</i> https://support.apple.com/en-us/112502 (iMac Flat Panel Technical Specifications): “Wireless networking – Airport ready. Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b DSSS compliant” “Headphone jack” “Built-in microphone” “Total of five USB 1.1 ports” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink . . . , AOL” https://support.apple.com/en-us/112501 (iMac Technical Specifications): “Support for optional 11-Mbps AirPort Card; IEEE 802.11 DSSS compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply)” “Support for external USB audio devices” “Microsoft Internet Explorer, Microsoft Outlook Express, Netscape Communicator” https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications): “Wireless networking – AirPort ready” “Ready for wireless networking: Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply).” “Support for external third-party USB digital audio peripherals such as microphones, speakers, and MIDI devices” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink”

Claims	Prior Art
	IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message. • read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4-4 You can use the Speech Manager to incorporate synthesized speech into your application. This section provides an overview and describes the basic concepts of the Speech Manager, and it outlines the process that the Speech Manager uses to convert text into speech. The Speech Manager converts text into sound data, which it passes to the Sound Manager to play through the current sound output device. The Speech Manager's interaction with the Sound Manager is transparent to your application, so you don't need to be familiar with the Sound Manager to take advantage of the Speech Manager's capabilities. Inside Macintosh Sound_ 1994 at Page 4-4 Figure 4-1 The speech generation process  <pre> graph LR Text[Text] --> SM[Speech Manager] SM --> SS[Speech synthesizer] SS --> SD[Sound Manager] SD --> AH[Audio hardware] </pre> Kennewick'209 also discloses presenting an audio message converted from text through the speech unit, which is an example of the claimed "wireless headset." The following citations show examples of where Kennewick '209 discloses this limitation. Kennewick '209 at 2:10-20 According to another aspect of the invention, the system comprises a speech unit interface device that receives spoken natural language queries, commands and/or other utterances from a user, and a computer device or system that receives input from the speech unit and processes the input (e.g., retrieves information responsive to the query, takes action consistent with the command and performs other functions as detailed herein). The system may further generate a natural language speech response in response to the user and/or generate a natural language speech message without any prompting from a user such as an alert message. Kennewick '209 at 2:28-46 According to another aspect of the invention, the computer device or system may comprise a stand alone or networked PC, a PDA or other portable computer device, or other computer device or system. For

Claims	Prior Art
	convenience, these and other computer alternatives shall be referred to simply as a computer. One aspect of the invention comprises software that is installed onto the computer, where the software comprises one or more of the following modules: a speech recognition module for capturing the user input; a parser for parsing the input, a text to speech engine module for converting text to speech; a network interface for enabling the computer to interface with one or more networks; a graphical user interface module; and an event manager for managing events. Preferably, the event manager is in communication with a dictionary and phrases module, a user profile module that enables user profiles to be created, modified, and accessed, a personality module that enables various personalities to be created and used, an agent module, an update manager, and one or more databases. Kennewick '209 at 3:8–35 In order to enhance the natural query and response environment, the system may format results in a manner enhancing the understandability to the user. The optimal formatting and presentation depends on the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. Information presented in a rigid, highly formatted, or structured manner seems unnatural to most people. Thus a system simulates some aspects of human “personality”. In some cases, the presentation of the response and the terms used are randomized so they do not appear rigidly formatted or mechanical. The use of other simulated personality characteristics is also desirable. For example, a response that may be upsetting to the user is best presented in a sympathetic manner. Finally, the results of many queries are long text strings, lists, tables or other lengthy sets of data. Natural presentation of this type of information presents particular challenges. Simply reading the long response is generally not preferred. Therefore the system parses the most important sections from the response and, at least initially, only reports these. Determining what parts of a long response are presented depends on the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. At the same time the system gives the user interactive control over what information and how much information is being presented, to stop the response all together, or to take other actions. Kennewick '209 at 5:47–6:6 In addition to the ability to process user questions and commands, the invention may provide local or remote control functions for the system or for other devices. Users may initiate commands locally or remotely. Several techniques for executing remote operations are possible, for example, the use of a telephone or other audio connection. Other remote command techniques may be used. The processing of

Claims	Prior Art
	these commands is performed in a nearly identical manner to a question. One difference being that the result of the command is generally an action rather than a response. Examples of control operations supported by the invention include, finding and playing music, film, games or other entertainment media, control of the appliances or other electronic or electro-mechanical devices, security systems control, audio, text and video conferencing (including Voice over IP conferencing), operating integrated messaging environments and other control operations. By using an optional analog or digital telephone interface or a data network interface, the invention can provide remote control capability of the computer or and auxiliary devices connected to the computer. For example, if the computer is connected via an interface to one or more auxiliary devices, a remote user can use this capability to remotely control or program the operation of the auxiliary device. For example, this enables recording of a television program, checking the status of a home security system, checking voice mail messages, hearing and responding to email messages, and much more. Kennewick '209 at 7:47–49 automatic telephone dialing, reading and sending emails and pages by voice and other communications control functions; Kennewick '209 at 8:17–35 management of network communications and conferencing, including telecommunications, email, instant messaging, Voice over IP communications and conferencing, local and wide area video and audio conferencing, pages and alerts; location, selection, management of play lists and play control of interactive entertainment from local or network sources including, video on demand, digital audio, such as MP3 format material, interactive games, web radio and video broadcasts; organization and calendar management for families, businesses and other groups of users including the management of, meetings, appointments, and events; and interactive educational programs using local and network material, with lesson material level set based on user's profile, and including, interactive multimedia lessons, religious instruction, calculator, dictionary and spelling, language training, foreign language translation and encyclopedias and other reference material. Kennewick '209 at 14:47–15:18 End users may use the data driven agent 156 extension and modification facilities and values stored in user profiles 110 to create special reports, packages of queries, alerts and output formats. A single alert or report can be configured to use multiple data sources, values, and other variables (i.e., time, location, etc.)

Claims	Prior Art
	to condition or otherwise determine when an alert should be sent. For example, an alert can be generated by sampling a stock price every fifteen minutes and sending an alert if the price drops below some value. To create a report, the user first specifies a set of commands or questions. Next, the user creates or selects a format for the report. Finally the user may name the report. A report can have variable parameters. For example, a user may create a company stock report, and execute the report by stating its name and the company name, which gives the user selected information and in a specified format for that company. In another example, a user can create a “morning” report, which presents selected multimedia information from different sources (news, sports, traffic, weather) in the order and formats desired. Alerts and reports can be created using only voice commands and responses, commands and responses through the graphical user interface 114, or a combination of the two. To create a report, alert, or other specialized behavior, the user performs a number of steps including: - specify the command to run a report or alert; - specify the question or questions, including keywords, used for a query; - set the criteria for running the report such as on command or when a particular condition is met; - define preferred information sources; - define preferences for order of result evaluation by source, value, etc.; - specify the presentation medium for a report or alert, such as an email, the text to speech engine, a message to a pager, or a text and graphics display; and - specify the preferred format for the report, such as information to be presented, order of information to be presented, preferred abbreviations or other variable substitutions. Kennewick '209 at 20:42–56 The domain agent 156 may strongly influence how the results are presented. For instance, once the domain agent 156 has created a satisfactory response to a question, the agent 156 may format that response for presentation. Typically, the domain agent 156 formats the response into the markup format used by the text to speech engine 124. The domain agent 156 may also format the result presentation using available format templates and based on the context, the criteria, and the user profile 110. Variable substitutions and transformations may be performed by the agent 156 to produce a response best understood and most natural to the user. Further, the order of presentation of tokens and the exact terminology used to create a more natural response to the user may be varied by the agent 156. The domain agent 156 may also select the presentation personality 108 to be used. Ewer '149 similarly discloses a wireless headset. The following citations show examples of where Ewer '149 discloses this limitation. Ewer '149 at Abstract A system and method for automatic detection and signaling of the connection status of a headset used

Claims	Prior Art
	with telephony or other multimedia application software running on processor-based hosts using digital signaling are disclosed. The signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The signaling module may include a connection status detector and a status signal generating processor. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed. Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the term “connection” utilized herein generally refers to both wired and wireless connections. Jabra FreeSpeak is also a “wireless headset.” The following citations show examples of where Jabra FreeSpeak discloses this limitation.

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak™— cord-free comfort and hands-free convenience in a Bluetooth™ headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators. Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones*. Jabra FreeSpeak Datasheet at Page 2 Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02  • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 2  A POSITA would have understood that the Apple Prior Art Systems, including the Communication Systems running or in combination with the Speech Recognition Technologies or Speech Synthesis Technologies, are configured to convert text to audio and outputting that audio as a presentation to the user. When implemented with a wireless headset, as explained for 1[pre]-[1c], a POSITA would further have understood the audio output message would have been presented to the user via the wireless headset because the wireless headset would have been a part of, connected to, or combined with the Communication System, for the purpose of presenting audio to a user. A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim 1[pre]. Specifically, a POSITA would be motivated to implement a wireless headset as disclosed by Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to the Apple Prior Art Systems such that audio generated through the Apple Prior Art Systems was presented through the wireless headset, and a POSITA would have a reasonable expectation of success in implementing a wireless headset to the Apple Prior Art Systems based on the disclosure of the Jabra FreeSpeak BT200. To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.

Claims	Prior Art
[1f] recording a reply audio message by the first user via the wireless headset;	Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose expressly or inherently disclose “recording a reply audio message by the first user via the wireless headset.” Apple Prior Art Systems, including the Speech Recognition Systems as running on or in combination with the Communication Systems, disclose recording audio. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM Speech assistant (ViaVoice) History at 2 1993--IBM launches the IBM Personal Dictation System (IPDS) for OS/2*, running on an IBM PC using a custom audio adapter card and a conventional microphone. The system takes dictation at about 80 words a minute with 95 percent accuracy, and allows users to "play back" their recording for editing purposes. IPDS supports US English as well as UK English, German, French, Italian and Spanish. IBM Mobility Suite Introduction at 7 • start recording - Starts recording your voice for: • Training the system to pronounce your contact names and program names • Creating a recording as a mail attachment IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message.

Claims	Prior Art
	• read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. IBM ViaVoice Simply Dictation user guide at Page 48 Playing Back Dictated Text in SpeakPad The Playback Selection feature in SpeakPad is useful when you are correcting your dictated text. You can hear your own voice to help determine what you dictated. Playback is limited to 1,000 words of dictated text (punctuation is included in the word count). If your selection is over 1,000 words or contains typed text, you can use the Read Selection feature to hear a computer-generated voice read the text aloud. Kennewick '209 also discloses receiving audio at a “wireless headset.” The following citations show examples of where Kennewick '209 discloses this limitation. Kennewick '209 at 2:10–20 According to another aspect of the invention, the system comprises a speech unit interface device that receives spoken natural language queries, commands and/or other utterances from a user, and a computer device or system that receives input from the speech unit and processes the input (e.g., retrieves information responsive to the query, takes action consistent with the command and performs other functions as detailed herein). The system may further generate a natural language speech response in response to the user and/or generate a natural language speech message without any prompting from a user such as an alert message. Kennewick '209 at 2:28–46 According to another aspect of the invention, the computer device or system may comprise a stand alone or networked PC, a PDA or other portable computer device, or other computer device or system. For convenience, these and other computer alternatives shall be referred to simply as a computer. One aspect of the invention comprises software that is installed onto the computer, where the software comprises one or more of the following modules: a speech recognition module for capturing the user input; a parser for parsing the input, a text to speech engine module for converting text to speech; a network interface for enabling the computer to interface with one or more networks; a graphical user interface module; and an event manager for managing events. Preferably; the event manager is in communication with a dictionary and phrases module, a user profile module that enables user profiles to be created, modified, and accessed, a personality module that enables various personalities to be created and used, an agent module, an update manager, and one or more databases.

Claims	Prior Art
	Kennewick '209 at 3:8–35 In order to enhance the natural query and response environment, the system may format results in a manner enhancing the understandability to the user. The optimal formatting and presentation depends on the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. Information presented in a rigid, highly formatted, or structured manner seems unnatural to most people. Thus a system simulates some aspects of human “personality”. In some cases, the presentation of the response and the terms used are randomized so they do not appear rigidly formatted or mechanical. The use of other simulated personality characteristics is also desirable. For example, a response that may be upsetting to the user is best presented in a sympathetic manner. Finally, the results of many queries are long text strings, lists, tables or other lengthy sets of data. Natural presentation of this type of information presents particular challenges. Simply reading the long response is generally not preferred. Therefore the system parses the most important sections from the response and, at least initially, only reports these. Determining what parts of a long response are presented depends on the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. At the same time the system gives the user interactive control over what information and how much information is being presented, to stop the response all together, or to take other actions. Kennewick '209 at 5:47–6:6 In addition to the ability to process user questions and commands, the invention may provide local or remote control functions for the system or for other devices. Users may initiate commands locally or remotely. Several techniques for executing remote operations are possible, for example, the use of a telephone or other audio connection. Other remote command techniques may be used. The processing of these commands is performed in a nearly identical manner to a question. One difference being that the result of the command is generally an action rather than a response. Examples of control operations supported by the invention include, finding and playing music, film, games or other entertainment media, control of the appliances or other electronic or electro-mechanical devices, security systems control, audio, text and video conferencing (including Voice over IP conferencing), operating integrated messaging environments and other control operations. By using an optional analog or digital telephone interface or a data network interface, the invention can provide remote control capability of the computer or and auxiliary devices connected to the computer. For example, if the computer is connected via an interface to one or more auxiliary devices, a remote user can use this capability to remotely control or program the operation of the auxiliary device. For example, this enables recording of a television program, checking the status of a home security system, checking voice mail messages, hearing and

Claims	Prior Art
	responding to email messages, and much more. Kennewick '209 at 7:47–49 automatic telephone dialing, reading and sending emails and pages by voice and other communications control functions; Kennewick '209 at 8:17–35 management of network communications and conferencing, including telecommunications, email, instant messaging, Voice over IP communications and conferencing, local and wide area video and audio conferencing, pages and alerts; location, selection, management of play lists and play control of interactive entertainment from local or network sources including, video on demand, digital audio, such as MP3 format material, interactive games, web radio and video broadcasts; organization and calendar management for families, businesses and other groups of users including the management of, meetings, appointments, and events; and interactive educational programs using local and network material, with lesson material level set based on user's profile, and including, interactive multimedia lessons, religious instruction, calculator, dictionary and spelling, language training, foreign language translation and encyclopedias and other reference material. Kennewick '209 at 14:47–15:18 End users may use the data driven agent 156 extension and modification facilities and values stored in user profiles 110 to create special reports, packages of queries, alerts and output formats. A single alert or report can be configured to use multiple data sources, values, and other variables (i.e., time, location, etc.) to condition or otherwise determine when an alert should be sent. For example, an alert can be generated by sampling a stock price every fifteen minutes and sending an alert if the price drops below some value. To create a report, the user first specifies a set of commands or questions. Next, the user creates or selects a format for the report. Finally the user may name the report. A report can have variable parameters. For example, a user may create a company stock report, and execute the report by stating its name and the company name, which gives the user selected information and in a specified format for that company. In another example, a user can create a “morning” report, which presents selected multimedia information from different sources (news, sports, traffic, weather) in the order and formats desired. Alerts and reports can be created using only voice commands and responses, commands and responses through the graphical user interface 114, or a combination of the two. To create a report, alert, or other specialized behavior, the user performs a number of steps including:

Claims	Prior Art
	specify the command to run a report or alert; specify the question or questions, including keywords, used for a query; set the criteria for running the report such as on command or when a particular condition is met; define preferred information sources; define preferences for order of result evaluation by source, value, etc.; specify the presentation medium for a report or alert, such as an email, the text to speech engine, a message to a pager, or a text and graphics display; and specify the preferred format for the report, such as information to be presented, order of information to be presented, preferred abbreviations or other variable substitutions. Kennewick '209 at 20:42–56 The domain agent 156 may strongly influence how the results are presented. For instance, once the domain agent 156 has created a satisfactory response to a question, the agent 156 may format that response for presentation. Typically, the domain agent 156 formats the response into the markup format used by the text to speech engine 124. The domain agent 156 may also format the result presentation using available format templates and based on the context, the criteria, and the user profile 110. Variable substitutions and transformations may be performed by the agent 156 to produce a response best understood and most natural to the user. Further, the order of presentation of tokens and the exact terminology used to create a more natural response to the user may be varied by the agent 156. The domain agent 156 may also select the presentation personality 108 to be used. Ewer '149 similarly discloses a wireless headset. The following citations show examples of where Ewer '149 discloses this limitation. Ewer '149 at Abstract A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on processor-based hosts using digital signaling are disclosed. The signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The signaling module may include a connection status detector and a status signal generating processor. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed.

Claims	Prior Art
	Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the term “connection” utilized herein generally refers to both wired and wireless connections. Jabra FreeSpeak is also a “wireless headset.” The following citations show examples of where Jabra FreeSpeak discloses this limitation. Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak®— cord-free comfort and hands-free convenience in a Bluetooth® headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators.

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones.* Jabra FreeSpeak Datasheet at Page 2 Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02  • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide Jabra FreeSpeak Datasheet at Page 2 About Bluetooth Bluetooth is a radio technology developed to deliver short-range wireless mobility. Bluetooth technology eliminates wires between devices, such as mobile phones and headsets, and enables communication between devices that are in close proximity to one another. Bluetooth facilitates fast, secure transmissions of both voice and data. More information is available at www.bluetooth.com. 

Claims	Prior Art
	A POSITA would have understood that the Apple Prior Art Systems, including the Communication Systems running or in combination with the Speech Recognition Technologies or Speech Synthesis Technologies, are configured to receive an audio message of the user. When implemented with a wireless headset, as explained for 1[pre]-[1c], a POSITA would further have understood the audio message of the user would be received via the wireless headset because the wireless headset would have been a part of, connected to, or combined with the Communication System, for the purpose of presenting audio to a user. A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim 1[pre]. Specifically, a POSITA would be motivated to implement a wireless headset as disclosed by Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to the Apple Prior Art Systems such that audio generated through the Apple Prior Art Systems was presented through the wireless headset, and a POSITA would have a reasonable expectation of success in implementing a wireless headset to the Apple Prior Art Systems based on the disclosure of the Jabra FreeSpeak BT200. To the extent that Apple Prior Art Systems, Kennewick '209, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[1g] converting the reply audio message into a reply text message; and	Apple Prior Art Systems expressly or inherently disclose “converting the reply audio message into a reply text message.” Apple Prior Art Systems, including the Speech Recognition Systems as running on or in combination with the Communication Systems, disclose converting a reply message from speech into text through dictation. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM Announces ViaVoice 98 at 1 IBM recently announced the release of ViaVoice 98, the next generation of IBM’s best selling speech recognition software, which features numerous innovations designed to make speech easier and more natural. IBM Announces ViaVoice 98 at 1 ViaVoice 98 Executive Edition is IBM’s most powerful continuous speech software and can be a productivity tool for writers, executives and experienced PC users. ViaVoice Executive offers direct dictation into most popular Windows applications, voice control for the desktop and applications and

Claims	Prior Art
	includes two topics, ViaVoice Computer and ViaVoice Business and Finance, which supplement the basic vocabulary. Topics are special interest vocabularies that add to the existing base vocabulary and increase accuracy. The executive edition is priced at \$149. ViaVoice 98 Office Edition is a continuous speech dictation and voice control software tool designed to help maximize the efficiency of a small office/home office user, faculty member or advanced student. Users can supplement their vocabulary with the included ViaVoice Business and Finance. The price for the Office edition is \$89. ViaVoice 98 Home Edition is IBM's entry-level continuous speech dictation product and is ideal for the family to use to create letters, reports or e-mail at home. ViaVoice 98 Home includes the ViaVoice Cuisine Topic to let users produce original recipes or make a keepsake of family recipes. The cost is \$49.95. A customized, integrated version of IBM ViaVoice is included with Lotus SmartSuite Millennium edition. IBM Speech assistant (ViaVoice) History at 2 1992--IBM introduces its first dictation system, called IBM Speech Server Series (ISSS). IBM Speech assistant (ViaVoice) History at 2 1993--IBM launches the IBM Personal Dictation System (IPDS) for OS/2*, running on an IBM PC using a custom audio adapter card and a conventional microphone. The system takes dictation at about 80 words a minute with 95 percent accuracy, and allows users to "play back" their recording for editing purposes. IPDS supports US English as well as UK English, German, French, Italian and Spanish. IBM ViaVoice 10 at 1 Talk into a microphone and the computer types what you say. That's an easy way to describe voice recognition software like Dragon NaturallySpeaking and IBM ViaVoice. Morphing through various forms since 1994, Dragon NaturallySpeaking is arguably the best voice recognition software available today. IBM's "40 years of commitment to speech research and development" have in part lead to the ViaVoice software. This article compares these two as well as providing general comments on voice recognition technology. (If we want to get picky, "speech recognition" is the process of converting speech to text, and "voice recognition" is the process of identifying a person by their voice alone. However, for this article, I'm using "voice recognition" to describe these software packages since the correct version of this term is rarely used correctly.)


Claims	Prior Art				
	IBM ViaVoice Mobility Suite v.4 at 1 Operate PDA functions using your voice Imagine being able to operate many of the common PDA functions like e-mail, tasks, calendar, phone numbers and addresses, by using voice commands instead of a stylus. IBM Embedded ViaVoice Mobility Suite software is now available to do just that! IBM Embedded ViaVoice Mobility Suite voice-enables many of the most common PIM applications for the handheld PDA, and is now available on the Compaq H3800 Pocket PC. IBM ViaVoice Mobility Suite v.4 at 1 End-to-end voice solutions IBM can offer distributed end-to-end voice solutions where some of the speech processing resides on the client device (cell phones etc.) and integrates with additional voice functionality at the server level and large-scale telephony/unified messaging systems. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Enterprise Edition is IBM's flagship speech recognition and text-to-speech product. Representing a dramatic shift in the evolution of embedded speech products, IBM Embedded ViaVoice Enterprise Edition has an unlimited active vocabulary, limited only by the amount of memory on devices, plus a dynamic and swappable 4,000 word phonetic flat list. This feature enables embedded applications to access voice-enabled databases with millions of records. A typical application would be a voice-enabled navigation database of street names for a GPS system in your car. IBM ViaVoice Mobility Suite v.4 at Page 2 <table border="1"> <thead> <tr> <th data-bbox="638 1068 919 1092">Inbox (Mail) Commands</th><th data-bbox="926 1068 1318 1092">Sample Commands</th></tr> </thead> <tbody> <tr> <td data-bbox="638 1117 919 1187">Reads incoming mail; gives you the opportunity to reply and send responses to others.</td><td data-bbox="926 1117 1318 1252"> • "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply" </td></tr> </tbody> </table> IBM Mobility Suite Introduction at 3 Welcome to IBM Embedded ViaVoice Mobility Suite, a productivity application that controls your Personal Digital Assistant (PDA) by voice. In response to your commands, IBM Embedded ViaVoice	Inbox (Mail) Commands	Sample Commands	Reads incoming mail; gives you the opportunity to reply and send responses to others.	• "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply"
Inbox (Mail) Commands	Sample Commands				
Reads incoming mail; gives you the opportunity to reply and send responses to others.	• "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply"				

Claims	Prior Art
	Mobility Suite will help manage your appointments, contacts, tasks, and mail. IBM Mobility Suite Introduction at 5 • To hear important information, for example, say Read mail or Read my very next appointment. • To hear prompts, messages, or requested information again, say Repeat that. IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message. • read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. IBM ViaVoice Simply Dictation user guide at Page 1 You talk, iType. It's that simple. The computer listens to your words and phrases and then translates them into text or actions that carry out commands. With ViaVoice, you'll give your keyboard and fingertips a well-deserved rest!

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 18 2. When SpeakPad is open, the blinking cursor indicates the place on the window where your dictated text appears. Your voice commands appear in a yellow message screen above the VoiceCenter, as shown in the graphic above.  IBM ViaVoice Simply Dictation user guide at Page 37 Dictating with ViaVoice There are two methods you can use to dictate and process text with ViaVoice: dictating into SpeakPad or using Dictation Stream Redirection (DSR), which allows you to dictate into another application. NOTE Make sure you select the correct speech mode button on the VoiceCenter. You should probably select <i>Dictation</i> and <i>Command</i> mode for most dictation and correction activities. For more information, see page 56. Dictating Into SpeakPad This is the most common method for dictating, since SpeakPad is built upon Apple's TextEdit technology. This integration provides a simple and seamless way to dictate and edit text. Additionally, correction is only available in SpeakPad. Once you dictate into SpeakPad, you can easily edit, format, and transfer your text into other supported applications. 1. Open SpeakPad from the ViaVoice folder. 2. Begin dictating your text.


Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. See above claim 1[pre]-[1b] https://support.apple.com/en-us/112502 (iMac Flat Panel Technical Specifications): “Wireless networking – Airport ready. Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b DSSS compliant” “Headphone jack” “Built-in microphone” “Total of five USB 1.1 ports” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink . . . , AOL” https://support.apple.com/en-us/112501 (iMac Technical Specifications): “Support for optional 11-Mbps AirPort Card; IEEE 802.11 DSSS compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply)” “Support for external USB audio devices” “Microsoft Internet Explorer, Microsoft Outlook Express, Netscape Communicator” https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications): “Wireless networking – AirPort ready” “Ready for wireless networking: Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply).” “Support for external third-party USB digital audio peripherals such as microphones, speakers, and

Claims	Prior Art
	MIDI devices” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink” iListen_15_New_speechrecogniti at Page 39 iLISTEN 1.5 New Speech-Recognition Software for OS X Lacks Polish but Offers Transcription Support iListen_15_New_speechrecogniti at Page 39 It's a happy day when Mac users have two full-featured speech-recognition programs to choose from—programs that let you dictate using natural speech <i>and</i> control your Mac by voice. This technology is critical to people who have limited use of their arms; for everyone else, it offers an exciting new way to interact with a computer. Before late 2001, advanced speech recognition was available only in Windows. That's when IBM released the capable but incomplete ViaVoice for Mac OS X (¶¶); <i>Reviews</i>, April 2002). Now MacSpeech's iListen 1.5.2 joins the Mac OS X speech-recognition ranks. Unfortunately, this program's performance is inconsistent, but its new ability to transcribe audio files may make it worthwhile to some.


Claims	Prior Art
	iListen_15_New_speechrecogniti at Page 39 Most speech-recognition programs require a training period during which you acclimate the program to the unique qualities of your voice. With iListen, this process is easy. To begin, you set up the microphone, which in our case was the VX1 Parrott with Griffin iMic USB adapter (both included). Next, for about ten minutes, iListen leads you through sound-level tests and a sample passage of text. With initial training, iListen was about 78 percent accurate in our dictation tests in OS X and OS X's Classic mode. To improve recognition, you can read 11 other passages. Doing so brought the accuracy up slightly, to about 82 percent. The program can also analyze a set of your documents to learn the vocabulary you use. iListen_15_New_speechrecogniti at Page 39 begin dictating into any application, such as Microsoft Word X or 2001. Using an 800MHz iMac G4 with 512MB of RAM, we saw good performance from the program—there was only a slight delay between speaking the words and seeing them appear on screen. iListen_15_New_speechrecogniti at Page 39  The screenshot shows a Mac OS X desktop environment. A window titled 'Document1' is open, displaying text about 'The light Mile'. The text discusses the experience of watching movies and the complexity of dictation. To the right of the document window is a small iListen control panel with buttons for 'Schedule', 'Transfer & log', and a text box showing 'I heard: leave it hard at the door'. Below the document window, there is a caption: 'Write Anywhere iListen is the only Mac-compatible speech-recognition program with the ability to transcribe Aiff and Wav files imported from a tape recorder.'


Claims	Prior Art
	iListen_15_New_speechrecogniti at Page 39 Macworld's Buying Advice MacSpeech's iListen 1.5.2 is a work in progress. At this point, IBM's ViaVoice—although it also has its flaws—is a better choice for most. The primary benefit of iListen today is the ability to transcribe documents dictated into a tape recorder. □ RATING: ## PROS: Support for audio-file transcription; can learn from dictation corrections in any program. CONS: No speech-controlled correction; frustrating correction process; limited and inconsistent built-in command-and-control abilities; incomplete documentation; some bugs. PRICE: \$99; with headset, \$149; upgrade, \$40 (doesn't include headset) OS COMPATIBILITY: Mac OS 9, Mac OS X COMPANY: MacSpeech, 603/644-7253, www.macspeech.com To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[1h] transmitting the reply text message to the second user.	Apple Prior Art Systems expressly or inherently disclose “transmitting the reply text message to the second user.” Apple Prior Art Systems, including the Speech Recognition Systems as running on or in combination with the Communication Systems, disclose sending text messages like dictated emails to other users. The following citations show examples of where the Apple Prior Art Systems user guide disclose this limitation. See above claim 1[pre]-[1g] https://support.apple.com/en-us/112502 (iMac Flat Panel Technical Specifications): “Wireless networking – Airport ready. Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b DSSS compliant” “Headphone jack” “Built-in microphone”

Claims	Prior Art
	“Total of five USB 1.1 ports” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink . . . , AOL” https://support.apple.com/en-us/112501 (iMac Technical Specifications): “Support for optional 11-Mbps AirPort Card; IEEE 802.11 DSSS compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply)” “Support for external USB audio devices” “Microsoft Internet Explorer, Microsoft Outlook Express, Netscape Communicator” https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications): “Wireless networking – AirPort ready” “Ready for wireless networking: Built-in antennas and card slot for optional 11-Mbps AirPort Card; IEEE 802.11b compliant” “Wireless Internet access requires AirPort Card, AirPort Base Station, and Internet access (fees may apply).” “Support for external third-party USB digital audio peripherals such as microphones, speakers, and MIDI devices” “Mac OS X Mail, Microsoft Internet Explorer, EarthLink” IBM Announces ViaVoice 98 at 1 ViaVoice 98 Executive Edition is IBM’s most powerful continuous speech software and can be a productivity tool for writers, executives and experienced PC users. ViaVoice Executive offers direct dictation into most popular Windows applications, voice control for the desktop and applications and includes two topics, ViaVoice Computer and ViaVoice Business and Finance, which supplement the basic vocabulary. Topics are special interest vocabularies that add to the existing base vocabulary and increase accuracy. The executive edition is priced at \$149. ViaVoice 98 Office Edition is a continuous speech dictation and voice control software tool designed to help maximize the efficiency of a small office/home office user, faculty member or advanced student. Users can supplement their vocabulary with the included ViaVoice Business and Finance. The price for the Office edition is \$89. ViaVoice 98 Home Edition is IBM’s entry-level continuous speech dictation product and is ideal for the family to use to create letters, reports or e-mail at home. ViaVoice 98 Home includes the ViaVoice Cuisine Topic to let users produce original recipes or make a keepsake of family recipes. The cost is

Claims	Prior Art
	\$49.95. A customized, integrated version of IBM ViaVoice is included with Lotus SmartSuite Millennium edition. IBM ViaVoice Mobility Suite v.4 at 1 Operate PDA functions using your voice Imagine being able to operate many of the common PDA functions like e-mail, tasks, calendar, phone numbers and addresses, by using voice commands instead of a stylus. IBM Embedded ViaVoice Mobility Suite software is now available to do just that! IBM Embedded ViaVoice Mobility Suite voice-enables many of the most common PIM applications for the handheld PDA, and is now available on the Compaq H3800 Pocket PC. IBM ViaVoice Mobility Suite v.4 at Page 2  The screenshot shows a table titled "IBM Embedded ViaVoice Mobility Suite Features". It has two columns: "Inbox (Mail) Commands" and "Sample Commands". Under "Inbox (Mail) Commands", it says "Reads incoming mail; gives you the opportunity to reply and send responses to others.". Under "Sample Commands", it lists five items: "Read mail", "Read new mail", "Next message", "Create reply", and "Send reply". IBM Mobility Suite Introduction at 3 Welcome to IBM Embedded ViaVoice Mobility Suite, a productivity application that controls your Personal Digital Assistant (PDA) by voice. In response to your commands, IBM Embedded ViaVoice Mobility Suite will help manage your appointments, contacts, tasks, and mail. IBM Mobility Suite Introduction at 5 • To hear important information, for example, say Read mail or Read my very next appointment. • To hear prompts, messages, or requested information again, say Repeat that. IBM Mobility Suite Introduction at 7 • start recording - Starts recording your voice for: • Training the system to pronounce your contact names and program names • Creating a recording as a mail attachment IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands

Claims	Prior Art
	• Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message. • read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. A POSITA would have understood that Apple's Prior Art Systems, including the Communication Systems, were configured to communicate messages with third-party users on distinct devices. As explained for claim [1b]. A POSITA would have understood that the Communications Systems included in the Apple Prior Art Systems, are wireless communication systems configured to facilitate communication between users over one or more networks, including by sending text and email messages, as was routine and conventional at the time. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
2[pre]. A wireless communication system as recited in claim 1,	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak expressly or inherently disclose "[a] wireless communication system as recited in claim 1." <i>See above at claim 1[pre].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak those elements, as known to them or as disclosed by other prior art identified in the cover pleading.

Claims	Prior Art
[2a] wherein said wireless communication device is configured to perform the operation of determining whether the wireless headset is active, and	Apple Prior Art Systems and Ewer '149 expressly or inherently disclose “wherein said wireless communication device is configured to perform the operation of determining whether the wireless headset is active.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, determine a status of its microphone, including on, off, or sleep, which are examples of the claimed “active” limitation. The following citations show examples of where the Apple Prior Art Systems user guide disclose this limitation. IBM Mobility Suite Introduction at 3 The system default microphone mode is Push-to-Talk. In this mode, you decide how long the microphone is on. Press and hold the Speech button, wait until you hear the tone, then say the command, and then release the button. Another way to use the microphone is Push-to-Activate. Press the Speech button, wait until you hear the tone, then say the command. The system turns off the microphone when you stop speaking. To select the microphone mode and other personal preferences, tap Tools and then Options. IBM ViaVoice Simply Dictation user guide at Page 16  The diagram illustrates the VoiceCenter interface. At the top, there are two microphone icons: one labeled 'Microphone OFF' and another labeled 'Microphone ON'. Below these, the interface is divided into sections. On the left, a 'VoiceCenter - Zoomed Out' view shows a large 'VoiceCenter Drawer' with a grid of buttons. On the right, a 'VoiceCenter - Zoomed In' view shows a smaller set of buttons. A legend on the right side lists the modes: 'Dictation Mode' (speech bubble icon), 'Command Mode' (megaphone icon), and 'Dictation and Command Mode' (combined icon). At the bottom, an 'Audio Level Indicator' shows two bars, one pink and one green. The VoiceCenter contains the Microphone button, conveniently located so you may quickly turn it ON or OFF. It also has standard Mac OS X action buttons to close, zoom, and minimize the VoiceCenter. Click once on the VoiceCenter drawer handle to open and access many ViaVoice components and preferences.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 42 SAY > TURN OFF MICROPHONE When you have finished speaking, always turn your microphone <i>off</i>. If you need to stop dictating for any reason, say the following command. SAY> GO TO SLEEP ViaVoice stops recognizing your speech as dictated text or voice commands until you say the following voice command. SAY> WAKE UP IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 56 VoiceCenter Components Shown below is the VoiceCenter, from which you can access all of the settings and options with a click of your mouse. All VoiceCenter functions, such as ViaVoice Preferences and UserSettings are speech-enabled.  Microphone Button Click to turn ON or OFF. Speech Mode Buttons Switches between the three speech modes: Dictation, Command, and Dictation and Command. The default is Dictation and Command mode, but you can change this setting in VoiceCenter panel in ViaVoice Preferences.

Claims	Prior Art
	Ewer '149 at [0002] More specifically, a system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on personal computers or other processor-based hosts using digital signaling are disclosed. Ewer '149 at [0011] A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on personal computers or other processor-based hosts using digital signaling are disclosed. Ewer '149 at [0012] The digital headset status signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software executed on the host. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The headset may be connected to the signaling module via a quick disconnect connector. The signaling module may include a connection status detector and a status signal generating processor. The detector may actively poll the connector for status or passively monitor the status. The signaling module preferably communicates with the host via USB ports. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed. Ewer '149 at [0013] The method for digitally signaling connection status of a headset to a processor-based host device generally comprises monitoring the headset connection status by a signaling module based on whether the headset is in communication with the signaling module, generating a digital headset status signal by a signal generating processor of the signaling module, and transmitting the digital headset status signal from the signaling module to the host device. The host device is responsive to the digital headset status signal for determining the state of the host and/or the application software. Ewer '149 at [0014] For example, when the host receives a digital headset status signal that indicates that the headset is

Claims	Prior Art
	disconnected, the session will be paused as the trainee is not following the session once the headset is disconnected. When the host receives a digital headset status signal that indicates that the headset is reconnected when the session is paused, the training session will restart to allow the trainee to complete the session from where the training session was paused. Ewer '149 at [0015] For example, when the host receives a digital headset status signal that indicates that the headset is connected when the softphone is idle, the softphone application is placed in a state that allows the softphone application software to receive an incoming call. When the host receives a digital headset status signal that indicates that the headset is disconnected when the softphone is in the available state, the softphone transitions to a state in which the softphone may call a feature in the softphone application or call another application such as an answer phone or speaker phone application upon receiving an incoming call. When the host receives a digital headset status signal that indicates that the headset is disconnected when the softphone is in a state with a call in progress, the softphone transitions to a hold state that places the call in progress on hold. When the host receives a digital headset status signal that indicates that the headset is connected when the softphone is in a hold state, the softphone returns to the call-in-progress state. Ewer '149 at [0025] A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on personal computers or other processor-based hosts using digital signaling are disclosed. Ewer '149 at [0026] In particular, the digital signaling taking place between the adapter 106 and the host 104 for automatic detection of the connection status of the headset 110 is similar regardless of whether the connector 108 is provided. Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be

Claims	Prior Art
	selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the term “connection” utilized herein generally refers to both wired and wireless connections. Ewer '149 at [0028] The host device 104 is generally any suitable processor-based device such as a personal computer (PC), a personal digital assistant (PDA), a digital music player (e.g., MP3 player), a video player (e.g., DVD player), a video game player, and a processor-based telephone. The host device 104 executes application software such as a telephony application software that uses the headset 110, for example, for receiving the user's voice as input and/or for outputting sounds to the user as output. When the user disconnects the headset 110 at the connector 108, the headset adapter 106 transmits a digital flag signal indicating a change in the headset connection status to the host 104 running the application software. Similarly, when the user reconnects the headset 110 at the connector 108, the headset adapter 106 transmits the flag signal to indicate a change in the headset connection status to the application software running on the host 104. Thus, as is evident, the adapter 106 functions at least in part as a headset connection status signaling module to the host 104. Ewer '149 at [0035] Headset connection status of connector 108 is determined by the connector detect circuit 132, i.e., whether the connector 108 is open or closed. Any suitable mechanism such as electronic state and/or mechanical detection mechanism may be employed by the detect circuit 132, the connector 108, and/or the processor 136 to detect the status of the connector 108. In particular, the connector detect circuit 132 or processor 136 either alone or in combination may passively or actively monitor the status of the headset connector 108. As an example of active monitoring, the connector detect circuit 132 may actively monitor the status of an electrical or mechanical switch provided in the connector 108 by polling the connector 108. The switch in the connector 108 may be configured such that the switch is closed when the connector 108 is closed and is open when the connector 108 is open. Thus, when the switch in the connector 108 is opened or closed, the connector detect circuit 132 detects the change as a result of the active polling. Ewer '149 at [0036] Alternatively, the connector detect circuit 132 may passively monitor for changes in the status of the connector 108 such as by detecting a voltage change at the detect circuit 132 as a result of, for example, the switch in the connector 108 being opened or closed. In one preferred embodiment, the system 100 may be configured such that the voltage at the interface between the USB headset adapter 106 and the connector 108 is lower when the headset 110 is connected such that the monitoring of the this voltage

Claims	Prior Art
	level. In addition, depending upon the specific implementation, this voltage change itself may be sufficient to bypass the detection circuit straight into the processor. Alternatively, the status of the connector 108 may be detected by or otherwise passed directly to a pin of the processor 136 in the USB headset adapter 106, as shown by the dashed arrow in FIG. 2. As an example, the voltage change at the interface between the USB headset adapter 106 and the connector 108 may be sufficient to bypass the detection circuit 132 and be passed directly into the pin of the processor 136. Ewer '149 at [0038] the processor 136 of the adapter 106 generates the digital headset connection status signal, the value of which depends upon the connection status of the headset. The digital headset status signal is preferably configured so as to be interpretable by the host and/or various application software products on the host. The value of the signal or flag indicating the headset connection status is changed when the connection status of the headset changes. The flag signal generated by the processor 136 is transmitted to the internal processor 122 of the PC host 104 via the USB ports 120, 130 of the PC host 104 and the USB adapter 106, respectively. Ewer '149 at [0039] application software 124 executed by the internal processor 122 responds in response to changes in the headset connection status. Specifically, the application software 124 is configured to perform certain actions upon occurrence of a corresponding change in the headset connection status flag. In other words, depending on how the application and/or the host PC 104 are configured, the value of the status flag indicating the headset connection status determines which action(s) the application software 124 and/or host PC 104 perform. For example, when the flag signal is transmitted to the internal processor 122 of the PC host 104 indicating that the connector 108 is disconnected, the internal processor 122 may pause the execution of the application software 124 if the application software is running, may default to a screen saver mode regardless of whether the application software is running, and/or may prevent the application software from starting up if the application software is not running. Upon receiving the flag signal indicating that the connector 108 is reconnected, the internal processor 122 may return to actively running or executing the application software 124 if the application software was paused, may return to normal host operation from the screen saver mode if the host PC 104 was in screen saver mode, and/or may allow the application software to be started if the starting up of the application software was prevented. Thus, the processor-based host 104 and/or the application software 124 are responsive to changes in the headset connection status.

Claims	Prior Art
	Ewer '149 at [0040] As noted above, the application software 124 executed by the internal processor 122 is configured to perform certain actions upon occurrence of a corresponding change in the headset connector status. Such actions are typically implemented or otherwise configured in the application software 124 and are performed when the host 104 is informed of the change in the headset connection status, i.e., whether the headset is in communication with the host. Ewer '149 at [0042] The application software 124 may also implement an audio file auto-pause feature. The audio file auto-pause feature automatically pauses the playing of the audio file, e.g., a CD or a .wave file, when the headset is disconnected and automatically resumes the playing of the audio file when the headset is reconnected. Ewer '149 at [0045] The application software 124 may also implement a control module auto-disable feature. The control module auto-disable feature automatically disables in-line controls on a control module when the headset is disconnected and automatically enable the in-line controls when the headset is connected. Typically, the in-line controls on the control module include volume, mute, and hook switch controls. The control module may be integrated with the headset adapter 106 or may be physically separate from the adapter 106. Ewer '149 at [0047] When the application is started, the application and host transition from the idle and available state 150 to an application running state 154. The application and host transition from the application running state 154 to an application paused state 156 when the connection is opened and return to the application running state 154 when the connection is closed. In addition, the application and host may transition from the application running state 154 to the idle state 150 upon exiting the application when the connection is closed or from state 156 to state 152 upon exiting the application when the connection is open. Ewer '149 at [0048] In one preferred exemplary embodiment, the application software is a training application software running a training session on the host where training audio is fed to a trainee through the headset. The training session is in one of a number of possible states and makes state transitions based upon the

Claims	Prior Art
	received headset connection status signals. For example, the training session application is idle and available for start up in state 150 where the headset connection is closed. Upon opening the headset connector, the training session application and host transition to the idle and unavailable state 152 where the training session application is prevented from being started up. Alternatively, the training session application may be configured such that it can be started up even when the headset connection is open Ewer '149 at [0049] When the training session application is started, the training session application and host transition from the idle and available state 150 to the training session running state 154. The training session application and host transition from the training session application running state 154 to the training session paused state 156 when the connection is opened as the trainee is not following the session once the headset is disconnected. The training session application and host return to the training session running state 154 to continue with the training session when the connection is closed or reconnected to allow the trainee to complete the session from where the training session was paused. Ewer '149 at [0053] As is evident, the digital headset connection status signaling system enables the user to interact automatically with the host PC by simply connecting or disconnecting the headset such as at a quick disconnect connector. Such automatic interaction with the host improves the efficiency of the user as well as the effectiveness of various software application products run by the host. A POSITA would have been motivated to combine the teachings of Apple Prior Art Systems and Ewer '149 to meet the claim limitations. Apple Prior Art Systems are configured to utilize a microphone in combination with a processing device (e.g., a computer). Ewer '149 is configured to regulate the interaction between a headset including a microphone and a communication device. A POSITA would have understood that if the microphone of Apple Prior Art Systems were disconnected from the computer, applications on the main unit utilizing the microphone could not function as desired. Apple Prior Art Systems would thus face a problem identified in Ewer '149: the requirement of inconvenient and deliberate “manual intervention by the user in order for the application to be informed of the disconnection and for the application to change its state.” Ewer '149, [0009]. A POSITA would seek to minimize such inconvenience and would thus be motivated to modify Apple Prior Art Systems with the teachings of Ewer '149. A POSITA would have had a reasonable expectation of success in combining the teachings of Apple Prior Art Systems and Ewer '149. Ewer '149's techniques are broadly applicable to “the connection status of a headset

Claims	Prior Art
	used with telephony or other multimedia application software running on personal computers or other processor-based hosts using digital signaling“—analogous to Apple Prior Art Systems’ input devices connected to the computer. Ewer ‘149, [0025]. Further, Apple Prior Art Systems discloses a “sleep” function and an “on/off” function for the microphone which is similar to the monitoring status of Ewer ‘149. Apple Prior Art Systems and Ewer ‘149 are thus easily combined based on similar functions, and it would have been straightforward to modify Apple Prior Art Systems to monitor the connection status. To the extent that Apple Prior Art Systems and Ewer ‘149 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Ewer ‘149, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[2b] wherein said converting of the incoming text message is automatically invoked to convert the incoming text message to the incoming audio message, provided that said determining determines that the wireless headset is active.	Apple Prior Art Systems and Ewer ‘149 expressly or inherently disclose “wherein said converting of the incoming text message is automatically invoked to convert the incoming text message to the incoming audio message, provided that said determining determines that the wireless headset is active.” Ewer ‘149 discloses determining how to present signals based on the active status of the headset. The following citations show examples of where Ewer ‘149 discloses this limitation. Ewer ‘149 at Abstract A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on processor-based hosts using digital signaling are disclosed. The signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The signaling module may include a connection status detector and a status signal generating processor. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed. Ewer ‘149 at [0011] A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on personal computers or other processor-based hosts using digital signaling are disclosed.


Claims	Prior Art
	Ewer '149 at [0012] The digital headset status signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software executed on the host. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The headset may be connected to the signaling module via a quick disconnect connector. The signaling module may include a connection status detector and a status signal generating processor. The detector may actively poll the connector for status or passively monitor the status. The signaling module preferably communicates with the host via USB ports. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed. Ewer '149 at [0013] The method for digitally signaling connection status of a headset to a processor-based host device generally comprises monitoring the headset connection status by a signaling module based on whether the headset is in communication with the signaling module, generating a digital headset status signal by a signal generating processor of the signaling module, and transmitting the digital headset status signal from the signaling module to the host device. The host device is responsive to the digital headset status signal for determining the state of the host and/or the application software. Ewer '149 at [0014] For example, when the host receives a digital headset status signal that indicates that the headset is disconnected, the session will be paused as the trainee is not following the session once the headset is disconnected. When the host receives a digital headset status signal that indicates that the headset is reconnected when the session is paused, the training session will restart to allow the trainee to complete the session from where the training session was paused. Ewer '149 at [0015] For example, when the host receives a digital headset status signal that indicates that the headset is connected when the softphone is idle, the softphone application is placed in a state that allows the softphone application software to receive an incoming call. When the host receives a digital headset status signal that indicates that the headset is disconnected when the softphone is in the available state, the softphone transitions to a state in which the softphone may call a feature in the softphone application or

Claims	Prior Art
	call another application such as an answer phone or speaker phone application upon receiving an incoming call. When the host receives a digital headset status signal that indicates that the headset is disconnected when the softphone is in a state with a call in progress, the softphone transitions to a hold state that places the call in progress on hold. When the host receives a digital headset status signal that indicates that the headset is connected when the softphone is in a hold state, the softphone returns to the call-in-progress state. Ewer '149 at [0025] A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on personal computers or other processor-based hosts using digital signaling are disclosed. Ewer '149 at [0026] In particular, the digital signaling taking place between the adapter 106 and the host 104 for automatic detection of the connection status of the headset 110 is similar regardless of whether the connector 108 is provided. Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the term “connection” utilized herein generally refers to both wired and wireless connections. Ewer '149 at [0028] The host device 104 is generally any suitable processor-based device such as a personal computer (PC), a personal digital assistant (PDA), a digital music player (e.g., MP3 player), a video player (e.g., DVD player), a video game player, and a processor-based telephone. The host device 104 executes application software such as a telephony application software that uses the headset 110, for example, for receiving the user's voice as input and/or for outputting sounds to the user as output. When the user disconnects the headset 110 at the connector 108, the headset adapter 106 transmits a digital flag signal indicating a change in the headset connection status to the host 104 running the application software. Similarly, when

Claims	Prior Art
	the user reconnects the headset 110 at the connector 108, the headset adapter 106 transmits the flag signal to indicate a change in the headset connection status to the application software running on the host 104. Thus, as is evident, the adapter 106 functions at least in part as a headset connection status signaling module to the host 104. Ewer '149 at [0038] the processor 136 of the adapter 106 generates the digital headset connection status signal, the value of which depends upon the connection status of the headset. The digital headset status signal is preferably configured so as to be interpretable by the host and/or various application software products on the host. The value of the signal or flag indicating the headset connection status is changed when the connection status of the headset changes. The flag signal generated by the processor 136 is transmitted to the internal processor 122 of the PC host 104 via the USB ports 120, 130 of the PC host 104 and the USB adapter 106, respectively. Ewer '149 at [0039] application software 124 executed by the internal processor 122 responds in response to changes in the headset connection status. Specifically, the application software 124 is configured to perform certain actions upon occurrence of a corresponding change in the headset connection status flag. In other words, depending on how the application and/or the host PC 104 are configured, the value of the status flag indicating the headset connection status determines which action(s) the application software 124 and/or host PC 104 perform. For example, when the flag signal is transmitted to the internal processor 122 of the PC host 104 indicating that the connector 108 is disconnected, the internal processor 122 may pause the execution of the application software 124 if the application software is running, may default to a screen saver mode regardless of whether the application software is running, and/or may prevent the application software from starting up if the application software is not running. Upon receiving the flag signal indicating that the connector 108 is reconnected, the internal processor 122 may return to actively running or executing the application software 124 if the application software was paused, may return to normal host operation from the screen saver mode if the host PC 104 was in screen saver mode, and/or may allow the application software to be started if the starting up of the application software was prevented. Thus, the processor-based host 104 and/or the application software 124 are responsive to changes in the headset connection status. Ewer '149 at [0042] The application software 124 may also implement an audio file auto-pause feature. The audio file auto-

Claims	Prior Art
	pause feature automatically pauses the playing of the audio file, e.g., a CD or a .wave file, when the headset is disconnected and automatically resumes the playing of the audio file when the headset is reconnected. Ewer '149 at [0045] The application software 124 may also implement a control module auto-disable feature. The control module auto-disable feature automatically disables in-line controls on a control module when the headset is disconnected and automatically enable the in-line controls when the headset is connected. Typically, the in-line controls on the control module include volume, mute, and hook switch controls. The control module may be integrated with the headset adapter 106 or may be physically separate from the adapter 106. Ewer '149 at [0047] When the application is started, the application and host transition from the idle and available state 150 to an application running state 154. The application and host transition from the application running state 154 to an application paused state 156 when the connection is opened and return to the application running state 154 when the connection is closed. In addition, the application and host may transition from the application running state 154 to the idle state 150 upon exiting the application when the connection is closed or from state 156 to state 152 upon exiting the application when the connection is open. Ewer '149 at [0048] In one preferred exemplary embodiment, the application software is a training application software running a training session on the host where training audio is fed to a trainee through the headset. The training session is in one of a number of possible states and makes state transitions based upon the received headset connection status signals. For example, the training session application is idle and available for start up in state 150 where the headset connection is closed. Upon opening the headset connector, the training session application and host transition to the idle and unavailable state 152 where the training session application is prevented from being started up. Alternatively, the training session application may be configured such that it can be started up even when the headset connection is open Ewer '149 at [0049] When the training session application is started, the training session application and host transition from the idle and available state 150 to the training session running state 154. The training session application

Claims	Prior Art
	and host transition from the training session application running state 154 to the training session paused state 156 when the connection is opened as the trainee is not following the session once the headset is disconnected. The training session application and host return to the training session running state 154 to continue with the training session when the connection is closed or reconnected to allow the trainee to complete the session from where the training session was paused. Ewer '149 at [0053] As is evident, the digital headset connection status signaling system enables the user to interact automatically with the host PC by simply connecting or disconnecting the headset such as at a quick disconnect connector. Such automatic interaction with the host improves the efficiency of the user as well as the effectiveness of various software application products run by the host. A POSITA would have been motivated to combine the teachings of Apple Prior Art Systems and Ewer '149 to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim [2a]. Specifically, a POSITA would be motivated to implement a wireless headset as disclosed by Kennewick '209 or Ewer '149 to the Apple Prior Art Systems and to implement a headset detection feature such that audio is presented through the wireless headset only when the wireless headset is on or otherwise “active.” A POSITA would further have a reasonable expectation of success in implementing a wireless headset to the Apple Prior Art Systems based on the disclosure of the Ewer '149 by implementing a signaling module. To the extent that Apple Prior Art Systems and Ewer '149 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Ewer '149, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
3[pre]. A wireless communication system as recited in claim 1, wherein said wireless communication device is configured to perform the operation of:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless communication system as recited in claim 1, wherein said wireless communication device is configured to perform the operation of.” <i>See above at claim 1[pre].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra

Claims	Prior Art
	FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[3a] prompting the first user to provide the reply audio message prior to said recording of the reply audio message.	Apple Prior Art Systems expressly or inherently disclose “prompting the first user to provide the reply audio message prior to said recording of the reply audio message.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, prompt a user to interact with the software, including to provide a reply. The following citations show examples of where the Apple Prior Art Systems user guide disclose this limitation. IBM ViaVoice Mobility Suite v.4 at Page 2  IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working

Claims	Prior Art
	list following the original message. • read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
5[pre]. A wireless communication system as recited in claim 1, wherein said wireless communication device is configured to perform the operation of:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless communication system as recited in claim 1, wherein said wireless communication device is configured to perform the operation of.” <i>See above at claim 1[pre].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[5a] determining whether the incoming text message should be presented to the first user in an audio manner, and	Apple Prior Art Systems and Kennewick '209 expressly or inherently disclose “determining whether the incoming text message should be presented to the first user in an audio manner.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, consider the desirability and optionality of presenting information in an audio and visual manner. The following citations show examples of where Apple Prior Art Systems disclose this limitation. Inside Macintosh Sound_1994 at Page 4–9 Different speech channels can even generate speech simultaneously, subject to processor capabilities and Sound Manager limitations. This capability should be used with restraint, however, because it can be hard for the user to understand any speech when more than one channel is generating speech at a time. In general, your application should generate speech only at the specific request of the user and should allow the user to turn off speech output. At the very least, your application should include an option that allows the user to view text instead of hearing it. Some users might have trouble understanding speech generated by the Speech Manager, and others might have a

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–76 <u>SetSpeechPitch</u> You can use the <code>SetSpeechPitch</code> function to set the speech pitch on a designated speech channel. <pre data-bbox="770 370 1472 418">FUNCTION SetSpeechPitch (chan: SpeechChannel; pitch: Fixed) : OSErr;</pre> <code>chan</code> The speech channel whose pitch you wish to set. <code>pitch</code> The new pitch for the speech channel, expressed as a fixed-point frequency value. Kennewick '209 discloses determining whether to present an audio output of visual output to a user. The following citations show examples of where Kennewick '209 discloses this limitation. Kennewick '209 at 3:8–35 In order to enhance the natural query and response environment, the system may format results in a manner enhancing the understandability to the user. The optimal formatting and presentation depends on the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. Information presented in a rigid, highly formatted, or structured manner seems unnatural to most people. Thus a system simulates some aspects of human “personality”. In some cases, the presentation of the response and the terms used are randomized so they do not appear rigidly formatted or mechanical. The use of other simulated personality characteristics is also desirable. For example, a response that may be upsetting to the user is best presented in a sympathetic manner. Finally, the results of many queries are long text strings, lists, tables or other lengthy sets of data. Natural presentation of this type of information presents particular challenges. Simply reading the long response is generally not preferred. Therefore the system parses the most important sections from the response and, at least initially, only reports these. Determining what parts of a long response are presented depends on the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. At the same time the system gives the user interactive control over what information and how much information is being presented, to stop the response all together, or to take other actions. Kennewick '209 at 11:11–20 The graphical user interface 114 can be used as a substitute or complement to the speech interface. For example, the graphical user interface 114 can be used to view and interact with graphical or tabular

Claims	Prior Art
	information in a manner more easily digested by the user. The graphical user interface can show system state and history in a more concise manner than the speech interface. Users can use the graphical user interface to create or extend agents 106. These operations can include scripting of agents, adding data to the agent or databases 102 used by the agent, adding links to information sources. Kennewick '209 at 20:42–56 The domain agent 156 may strongly influence how the results are presented. For instance, once the domain agent 156 has created a satisfactory response to a question, the agent 156 may format that response for presentation. Typically, the domain agent 156 formats the response into the markup format used by the text to speech engine 124. The domain agent 156 may also format the result presentation using available format templates and based on the context, the criteria, and the user profile 110. Variable substitutions and transformations may be performed by the agent 156 to produce a response best understood and most natural to the user. Further, the order of presentation of tokens and the exact terminology used to create a more natural response to the user may be varied by the agent 156. The domain agent 156 may also select the presentation personality 108 to be used. Kennewick '209 at 20:57–21:3 The domain agent 156, may select the presentation template, determine order of presentation for tokens and determine variable substitutions and transformations using probabilistic or fuzzy set decision methods. The template used to form the presentation can be from the domain agent itself or from the user profile 110. The user profile 110 can completely specify the presentation format or can be used to select and then modify an existing presentation format. Selection and formatting of presentation template can also depend on the presentation personality modules 108. At the same time, the characteristics of the personality used for the response are dynamically determined using probabilities or fuzzy possibilities derived from the context, the criteria, the domain agent itself and the user profile 110. Kennewick '209 at 21:4–21 The domain agent 156 may apply a number of transformations to the tokens before presentation to the user. These variable substitutions and transformations are derived from a number of sources including, domain information carried by the agent, the context, the token values, the criteria, the personality module 108 to be used, and the user profile 110. Examples of variable substitutions and transformations include: - substitution of words for numbers; - substitution of names for acronyms or symbols (i.e., trading symbols); - use of formatting information derived from the information sources (i.e., HTML tags);

Claims	Prior Art
	nature of the response including, text, long text, list, table; possible missing information or errors; units for measurement (i.e., English or metric); and preferred terminology from the user profile 110 or presentation personality 108. Kennewick '209 at 21:38–44 Some information, in formats such as video, pictures and graphics, may be best presented in a displayed format. The domain agents 156 may apply suitable presentation templates in these cases and present the information through the graphical user interface 114. The system agent 150 provides special criteria handlers 152 for presentation and user commands for display presentation and control. Kennewick '209 at 22:28–33 The response may also be formatted to reflect a particular personality or tone to the response in order to make the response more “human.” The response may be relayed to the user and/or third parties as an audio message and/or a visual message displayed on a user interface. A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209 to meet the claim limitations. Apple Prior Art Systems discloses ability to present information in text or through synthesized speech. Kennewick '209 discloses a logic model for determining whether information is best presented in an audio manner or a textual manner. Kennewick '209 at 11:11–20, A POSITA would have understood that a utilizing distinct modes of operation to determine whether to present a received message visually or audibly would automate the process for presenting information and would thus be desirable. A POSITA would thus be motivated to modify Apple Prior Art Systems with the teachings of Kennewick '209 to include distinct modes of operation for how to present received messages. A POSITA would have had a reasonable expectation of success in combining Apple Prior Art Systems with the teachings of Kennewick '209. Kennewick '209's techniques are broadly applicable to computer devices— analogous to the processing devices (e.g., computers) or pocket PCs disclosed in the Apple Prior Art Systems. Kennewick '209 at 1:55–60. Further, Apple Prior Art Systems explains that it provides various settings for synthesized speech. Apple Prior Art Systems user guide at 47. Thus, a POSITA would merely need to incorporate an additional setting to determine a preferred manner of presenting messages to a user similar to the modes of operation disclosed in Kennewick '209. Apple Prior Art Systems and Kennewick '209 are thus easily combined based on familiar components, and it would have been straightforward to implement a setting for text-to-speech synthesis to determine a preferred manner of presentation.

Claims	Prior Art
	To the extent that Apple Prior Art Systems and Kennewick '209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick '209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[5b] wherein said converting of the incoming text message is invoked to convert the incoming text message to the incoming audio message, provide that said determining determines that the incoming text message should be presented to the first user in an audio manner.	Apple Prior Art Systems and Kennewick '209 expressly or inherently disclose “wherein said converting of the incoming text message is invoked to convert the incoming text message to the incoming audio message, provide that said determining determines that the incoming text message should be presented to the first user in an audio manner.” Kennewick '209 discloses presenting audio to a user if the system determines that audio is the best manner to present the information. The following citations show examples of where Kennewick '209 discloses this limitation. Kennewick '209 at 3:8–35 In order to enhance the natural query and response environment, the system may format results in a manner enhancing the understandability to the user. The optimal formatting and presentation depends on the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. Information presented in a rigid, highly formatted, or structured manner seems unnatural to most people. Thus a system simulates some aspects of human “personality”. In some cases, the presentation of the response and the terms used are randomized so they do not appear rigidly formatted or mechanical. The use of other simulated personality characteristics is also desirable. For example, a response that may be upsetting to the user is best presented in a sympathetic manner. Finally, the results of many queries are long text strings, lists, tables or other lengthy sets of data. Natural presentation of this type of information presents particular challenges. Simply reading the long response is generally not preferred. Therefore the system parses the most important sections from the response and, at least initially, only reports these. Determining what parts of a long response are presented depends on the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. At the same time the system gives the user interactive control over what information and how much information is being presented, to stop the response all together, or to take other actions. Kennewick '209 at 14:47–15:18 End users may use the data driven agent 156 extension and modification facilities and values stored in user profiles 110 to create special reports, packages of queries, alerts and output formats. A single alert or

Claims	Prior Art
	report can be configured to use multiple data sources, values, and other variables (i.e., time, location, etc.) to condition or otherwise determine when an alert should be sent. For example, an alert can be generated by sampling a stock price every fifteen minutes and sending an alert if the price drops below some value. To create a report, the user first specifies a set of commands or questions. Next, the user creates or selects a format for the report. Finally the user may name the report. A report can have variable parameters. For example, a user may create a company stock report, and execute the report by stating its name and the company name, which gives the user selected information and in a specified format for that company. In another example, a user can create a “morning” report, which presents selected multimedia information from different sources (news, sports, traffic, weather) in the order and formats desired. Alerts and reports can be created using only voice commands and responses, commands and responses through the graphical user interface 114, or a combination of the two. To create a report, alert, or other specialized behavior, the user performs a number of steps including: - specify the command to run a report or alert; - specify the question or questions, including keywords, used for a query; - set the criteria for running the report such as on command or when a particular condition is met; - define preferred information sources; - define preferences for order of result evaluation by source, value, etc.; - specify the presentation medium for a report or alert, such as an email, the text to speech engine, a message to a pager, or a text and graphics display; and - specify the preferred format for the report, such as information to be presented, order of information to be presented, preferred abbreviations or other variable substitutions. Kennewick '209 at 16:43–51 The system 90, may provide built in training capabilities to help the user learn the best methods to formulate their questions and commands. The interactive training allows the user to audibly or visibly see the machine interpretation of their queries and provides suggestions on how to better structure a query. Using the interactive training a user can quickly become comfortable with the voice query language and at the same time learn how to optimize the amount of information required with each step of a dialog. Kennewick '209 at 20:42–56 The domain agent 156 may strongly influence how the results are presented. For instance, once the domain agent 156 has created a satisfactory response to a question, the agent 156 may format that response for presentation. Typically, the domain agent 156 formats the response into the markup format used by the text to speech engine 124. The domain agent 156 may also format the result presentation using available format templates and based on the context, the criteria, and the user profile 110. Variable

Claims	Prior Art
	substitutions and transformations may be performed by the agent 156 to produce a response best understood and most natural to the user. Further, the order of presentation of tokens and the exact terminology used to create a more natural response to the user may be varied by the agent 156. The domain agent 156 may also select the presentation personality 108 to be used. Kennewick '209 at 20:57–21:3 The domain agent 156, may select the presentation template, determine order of presentation for tokens and determine variable substitutions and transformations using probabilistic or fuzzy set decision methods. The template used to form the presentation can be from the domain agent itself or from the user profile 110. The user profile 110 can completely specify the presentation format or can be used to select and then modify an existing presentation format. Selection and formatting of presentation template can also depend on the presentation personality modules 108. At the same time, the characteristics of the personality used for the response are dynamically determined using probabilities or fuzzy possibilities derived from the context, the criteria, the domain agent itself and the user profile 110. Kennewick '209 at 21:4–21 The domain agent 156 may apply a number of transformations to the tokens before presentation to the user. These variable substitutions and transformations are derived from a number of sources including, domain information carried by the agent, the context, the token values, the criteria, the personality module 108 to be used, and the user profile 110. Examples of variable substitutions and transformations include: - substitution of words for numbers; - substitution of names for acronyms or symbols (i.e., trading symbols); - use of formatting information derived from the information sources (i.e., HTML tags); - nature of the response including, text, long text, list, table; - possible missing information or errors; - units for measurement (i.e., English or metric); and - preferred terminology from the user profile 110 or presentation personality 108. A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209 to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim [5a]. Specifically, a POSITA would be motivated to implement the logic disclosed in Kennewick '209 to determine a best means for presentation, and a POSITA would have a reasonable expectation of success based on the disclosures presented in Kennewick '209, including analyzing user

Claims	Prior Art
	preferences and histories and analyzing context of the date (e.g., how long it would take to present the synthesized audio). To the extent that Apple Prior Art Systems and Kennewick ‘209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick ‘209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
6[pre]. A wireless communication system as recited in claim 5,	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless communication system as recited in claim 5.” <i>See above at claim 1[pre].</i> To the extent that Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[6a] wherein said wireless communication device is configured to perform the operation of determining whether the wireless headset is active, and	Apple Prior Art Systems and Ewer ‘149 expressly or inherently disclose “wherein said wireless communication device is configured to perform the operation of determining whether the wireless headset is active.” <i>See above at claim [2a].</i> To the extent that Apple Prior Art Systems and Ewer ‘149 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Ewer ‘149, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[6b] wherein said converting of the incoming text message is automatically invoked to convert the incoming text message to the incoming audio message, provide that said determining determines that the wireless headset is active.	Apple Prior Art Systems and Ewer ‘149 expressly or inherently disclose “wherein said converting of the incoming text message is automatically invoked to convert the incoming text message to the incoming audio message, provide that said determining determines that the wireless headset is active.” <i>See above at claim [2b].</i> To the extent that Apple Prior Art Systems and Ewer ‘149 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Ewer ‘149, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.

Claims	Prior Art
7[pre]. A wireless communication system as recited in claim 6, wherein said wireless communication device is configured to perform the operation of:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless communication system as recited in claim 6, wherein said wireless communication device is configured to perform the operation of.” <i>See above at claim 1[pre].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[7a] prompting the first user to provide the reply audio message prior to said recording of the reply audio message.	Apple Prior Art Systems expressly or inherently disclose “prompting the first user to provide the reply audio message prior to said recording of the reply audio message.” <i>See above at claim [3a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
8. A wireless communication system as recited in claim 5, wherein said determining is based on a preference setting that is able to be set by the first user.	Apple Prior Art Systems and Kennewick '209 expressly or inherently disclose “[a] wireless communication system as recited in claim 5, wherein said determining is based on a preference setting that is able to be set by the first user.” Kennewick '209 discloses a user profile, which is an example of the claimed “preference setting,” where a user can interact with settings to determine how to present information. The following citations show examples of where Kennewick '209 discloses this limitation.

Claims	Prior Art
	other actions. Kennewick '209 at 4:6–19 The above steps preferably are performed while accounting for the domain of expertise required, the context for the question or command, domain specific information, the history of the user's interaction, user preferences, information sources or commands available, and responses obtained from the sources. At each stage probabilistic or fuzzy set decision and matching methods can be applied to deal with inconsistent, ambiguous, conflicting and incomplete information or responses. In addition, the use of asynchronous queries that may result in rapid and graceful failure of some queries or commands may allow the system to robustly return results quickly and in a manner that seems natural to the user. Kennewick '209 at 4:20–47 Many everyday questions are inherently subjective and result in answers that are a matter of option or consensus as much as fact. Such questions are often ad hoc in their nature, as well. The system, according to another aspect of the invention, may use probabilistic and fuzzy set decision and matching methods to first identify the subjective nature of the question and to evaluate a range of possible answers, selecting the one answer or few answers that best represent the type of result desired by the user. The context and expected results from a particular question may be highly dependent on the individual asking the question. Therefore, the system may create, store and use extensive personal profile information for each user. Information in the profile may be added and updated automatically as the user uses the system or may be manually added or updated by the user or by others. Domain specific agents may collect, store and use specific profile information, as may be required for optimal operations. Users can create commands for regularly used reports, automatically generated alerts, and other queries and for the formatting and presentation of results. The system may use profile data in interpreting questions, formulating queries, interpreting results of queries and presenting answers to the user. Examples of information in a user profile includes, history of questions asked, session histories, formatting and presentation preferences, special word spelling, terms of interest, special data sources of interest, age, sex, education, location or address, place of business, type of business, investments, hobbies, sports interests, news interests and other profile data. Kennewick '209 at 10:65–11:10 The agents 106 return results of questions as responses to users. The response may be created using the results of information queries, the system personality 108 and the user preferences or other data in the user profile 110. The agents 106 generally present these results using the speech unit 128. The agents 106

Claims	Prior Art
	create a response string, which is sent to the text to speech engine 124. The text to speech engine 124 creates the required utterances, which may be encoded and compressed by the speech coder 122. Once coded, the utterances are transmitted from the main unit 98 by the transceiver 126 to the transceiver 130 on the speech unit 128. The utterance is then decoded and decompressed by the speech coder 138 and output by the speaker 136. Kennewick '209 at 14:15–46 Commands and questions are interpreted, queries formulated, responses created and results presented based on the users personal or user profile 110 values. Personal profiles may include information specific to the individual, their interests, their special use of terminology, the history of their interactions with the system, and domains of interest. The personal profile data may be used by the agents 106, the speech recognition engine 120, the text to speech engine 124, and the parser 118. Preferences can include, special (modified) commands, past behavior or history, questions, information sources, formats, reports, and alerts. User profile data can be manually entered by the user and/or can be learned by the system 90 based on user behavior. User profile values may include: - spelling preferences; - date of birth for user, family and friends; - income level; - gender; - occupation; - location information such as, home address, neighborhood, and business address; - car type; - telecommunications and other service providers and services; - financial and investment information; - synonyms (i.e., a nick name for someone); - special spelling; - keywords; - transformation or substitution variables; - domains of interest; and - other values. Kennewick '209 at 14:47–15:18 End users may use the data driven agent 156 extension and modification facilities and values stored in user profiles 110 to create special reports, packages of queries, alerts and output formats. A single alert or report can be configured to use multiple data sources, values, and other variables (i.e., time, location, etc.)

Claims	Prior Art
	to condition or otherwise determine when an alert should be sent. For example, an alert can be generated by sampling a stock price every fifteen minutes and sending an alert if the price drops below some value. To create a report, the user first specifies a set of commands or questions. Next, the user creates or selects a format for the report. Finally the user may name the report. A report can have variable parameters. For example, a user may create a company stock report, and execute the report by stating its name and the company name, which gives the user selected information and in a specified format for that company. In another example, a user can create a “morning” report, which presents selected multimedia information from different sources (news, sports, traffic, weather) in the order and formats desired. Alerts and reports can be created using only voice commands and responses, commands and responses through the graphical user interface 114, or a combination of the two. To create a report, alert, or other specialized behavior, the user performs a number of steps including: - specify the command to run a report or alert; - specify the question or questions, including keywords, used for a query; - set the criteria for running the report such as on command or when a particular condition is met; - define preferred information sources; - define preferences for order of result evaluation by source, value, etc.; - specify the presentation medium for a report or alert, such as an email, the text to speech engine, a message to a pager, or a text and graphics display; and - specify the preferred format for the report, such as information to be presented, order of information to be presented, preferred abbreviations or other variable substitutions. Kennewick '209 at 18:8–31 Once the context for the question or command has been determined, the parser 118 can invoke the correct agent 156, 150. To formulate a question or command in the regular grammar used by agents, the parser 118 may determine required and optional values for the criteria or parameters. These criteria may have been explicitly supplied by the user or may need to be inferred. The parser 118 makes use of the criteria handlers 152 supplied by the system agent 150. The criteria handlers 152 provide context sensitive procedures for extracting the criteria or parameters from the user's question or command. Some criteria are determined by executing algorithms in the agent, while others may be determined by applying probabilistic or fuzzy reasoning to tables of possible values. Prior probabilities or fuzzy possibilities and associated values are received from a number of sources including the history of the dialog, the user profile 110, and the agent. Based on user responses, the prior probabilities or fuzzy possibilities are updated as the system learns the desired behavior. For a weather context, examples of criteria include, location, date and time. Other criteria can include command criteria (i.e., yes/no, on/off, pause, stop), and spelling. Special criteria handlers are available from the system agent for processing lists, tables, barge-in

Claims	Prior Art
	commands, long strings of text and system commands. Kennewick '209 at 20:42–56 The domain agent 156 may strongly influence how the results are presented. For instance, once the domain agent 156 has created a satisfactory response to a question, the agent 156 may format that response for presentation. Typically, the domain agent 156 formats the response into the markup format used by the text to speech engine 124. The domain agent 156 may also format the result presentation using available format templates and based on the context, the criteria, and the user profile 110. Variable substitutions and transformations may be performed by the agent 156 to produce a response best understood and most natural to the user. Further, the order of presentation of tokens and the exact terminology used to create a more natural response to the user may be varied by the agent 156. The domain agent 156 may also select the presentation personality 108 to be used. Kennewick '209 at 20:57–21:3 The domain agent 156, may select the presentation template, determine order of presentation for tokens and determine variable substitutions and transformations using probabilistic or fuzzy set decision methods. The template used to form the presentation can be from the domain agent itself or from the user profile 110. The user profile 110 can completely specify the presentation format or can be used to select and then modify an existing presentation format. Selection and formatting of presentation template can also depend on the presentation personality modules 108. At the same time, the characteristics of the personality used for the response are dynamically determined using probabilities or fuzzy possibilities derived from the context, the criteria, the domain agent itself and the user profile 110. Kennewick '209 at 21:4–21 The domain agent 156 may apply a number of transformations to the tokens before presentation to the user. These variable substitutions and transformations are derived from a number of sources including, domain information carried by the agent, the context, the token values, the criteria, the personality module 108 to be used, and the user profile 110. Examples of variable substitutions and transformations include: - substitution of words for numbers; - substitution of names for acronyms or symbols (i.e., trading symbols); - use of formatting information derived from the information sources (i.e., HTML tags); - nature of the response including, text, long text, list, table; - possible missing information or errors;

Claims	Prior Art
	units for measurement (i.e., English or metric); and preferred terminology from the user profile 110 or presentation personality 108. A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209 to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim [5a]. Specifically, a POSITA would be motivated to implement the logic disclosed in Kennewick '209 to determine a best means for presentation, and a POSITA would have a reasonable expectation of success based on the disclosures presented in Kennewick '209, including analyzing user preferences and histories and analyzing context of the date (e.g., how long it would take to present the synthesized audio). To the extent that Apple Prior Art Systems and Kennewick '209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick '209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
9. A wireless communication system as recited in claim 5, wherein said determining is based on a user indication.	Apple Prior Art Systems and Kennewick '209 expressly or inherently disclose “[a] wireless communication system as recited in claim 5, wherein said determining is based on a user indication.” Kennewick '209 discloses allowing a user to indicate how to present information, which is an example of the claimed “user indication.” The following citations show examples of where Kennewick '209 discloses this limitation.

Claims

Prior Art

Kennewick '209 at FIG. 1




FIG. 1

Kennewick '209 at 2:28–46

According to another aspect of the invention, the computer device or system may comprise a stand alone or networked PC, a PDA or other portable computer device, or other computer device or system. For convenience, these and other computer alternatives shall be referred to simply as a computer. One aspect of the invention comprises software that is installed onto the computer, where the software comprises one or more of the following modules: a speech recognition module for capturing the user input; a parser for parsing the input, a text to speech engine module for converting text to speech; a network interface for enabling the computer to interface with one or more networks; a graphical user interface module; and an event manager for managing events. Preferably, the event manager is in communication with a dictionary and phrases module, a user profile module that enables user profiles to be created, modified, and accessed, a personality module that enables various personalities to be created and used, an agent module, an update manager, and one or more databases.

Kennewick '209 at 3:8–35

In order to enhance the natural query and response environment, the system may format results in a manner enhancing the understandability to the user. The optimal formatting and presentation depends on

Claims	Prior Art
	the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. Information presented in a rigid, highly formatted, or structured manner seems unnatural to most people. Thus a system simulates some aspects of human “personality”. In some cases, the presentation of the response and the terms used are randomized so they do not appear rigidly formatted or mechanical. The use of other simulated personality characteristics is also desirable. For example, a response that may be upsetting to the user is best presented in a sympathetic manner. Finally, the results of many queries are long text strings, lists, tables or other lengthy sets of data. Natural presentation of this type of information presents particular challenges. Simply reading the long response is generally not preferred. Therefore the system parses the most important sections from the response and, at least initially, only reports these. Determining what parts of a long response are presented depends on the context of the questions, the contents of the response being presented, the history of the interaction with the user, the user's preferences and interests and the nature of the domain. At the same time the system gives the user interactive control over what information and how much information is being presented, to stop the response all together, or to take other actions. Kennewick '209 at 4:6–19 The above steps preferably are performed while accounting for the domain of expertise required, the context for the question or command, domain specific information, the history of the user's interaction, user preferences, information sources or commands available, and responses obtained from the sources. At each stage probabilistic or fuzzy set decision and matching methods can be applied to deal with inconsistent, ambiguous, conflicting and incomplete information or responses. In addition, the use of asynchronous queries that may result in rapid and graceful failure of some queries or commands may allow the system to robustly return results quickly and in a manner that seems natural to the user. Kennewick '209 at 4:20–47 Many everyday questions are inherently subjective and result in answers that are a matter of option or consensus as much as fact. Such questions are often ad hoc in their nature, as well. The system, according to another aspect of the invention, may use probabilistic and fuzzy set decision and matching methods to first identify the subjective nature of the question and to evaluate a range of possible answers, selecting the one answer or few answers that best represent the type of result desired by the user. The context and expected results from a particular question may be highly dependent on the individual asking the question. Therefore, the system may create, store and use extensive personal profile information for each user. Information in the profile may be added and updated automatically as the user uses the system or may be


Claims	Prior Art
	manually added or updated by the user or by others. Domain specific agents may collect, store and use specific profile information, as may be required for optimal operations. Users can create commands for regularly used reports, automatically generated alerts, and other queries and for the formatting and presentation of results. The system may use profile data in interpreting questions, formulating queries, interpreting results of queries and presenting answers to the user. Examples of information in a user profile includes, history of questions asked, session histories, formatting and presentation preferences, special word spelling, terms of interest, special data sources of interest, age, sex, education, location or address, place of business, type of business, investments, hobbies, sports interests, news interests and other profile data. Kennewick '209 at 10:65–11:10 The agents 106 return results of questions as responses to users. The response may be created using the results of information queries, the system personality 108 and the user preferences or other data in the user profile 110. The agents 106 generally present these results using the speech unit 128. The agents 106 create a response string, which is sent to the text to speech engine 124. The text to speech engine 124 creates the required utterances, which may be encoded and compressed by the speech coder 122. Once coded, the utterances are transmitted from the main unit 98 by the transceiver 126 to the transceiver 130 on the speech unit 128. The utterance is then decoded and decompressed by the speech coder 138 and output by the speaker 136. Kennewick '209 at 14:15–46 Commands and questions are interpreted, queries formulated, responses created and results presented based on the users personal or user profile 110 values. Personal profiles may include information specific to the individual, their interests, their special use of terminology, the history of their interactions with the system, and domains of interest. The personal profile data may be used by the agents 106, the speech recognition engine 120, the text to speech engine 124, and the parser 118. Preferences can include, special (modified) commands, past behavior or history, questions, information sources, formats, reports, and alerts. User profile data can be manually entered by the user and/or can be learned by the system 90 based on user behavior. User profile values may include: - spelling preferences; - date of birth for user, family and friends; - income level; - gender; - occupation;


Claims	Prior Art
	location information such as, home address, neighborhood, and business address; car type; telecommunications and other service providers and services; financial and investment information; synonyms (i.e., a nick name for someone); special spelling; keywords; transformation or substitution variables; domains of interest; and other values. Kennewick '209 at 14:47–15:18 End users may use the data driven agent 156 extension and modification facilities and values stored in user profiles 110 to create special reports, packages of queries, alerts and output formats. A single alert or report can be configured to use multiple data sources, values, and other variables (i.e., time, location, etc.) to condition or otherwise determine when an alert should be sent. For example, an alert can be generated by sampling a stock price every fifteen minutes and sending an alert if the price drops below some value. To create a report, the user first specifies a set of commands or questions. Next, the user creates or selects a format for the report. Finally the user may name the report. A report can have variable parameters. For example, a user may create a company stock report, and execute the report by stating its name and the company name, which gives the user selected information and in a specified format for that company. In another example, a user can create a “morning” report, which presents selected multimedia information from different sources (news, sports, traffic, weather) in the order and formats desired. Alerts and reports can be created using only voice commands and responses, commands and responses through the graphical user interface 114, or a combination of the two. To create a report, alert, or other specialized behavior, the user performs a number of steps including: specify the command to run a report or alert; specify the question or questions, including keywords, used for a query; set the criteria for running the report such as on command or when a particular condition is met; define preferred information sources; define preferences for order of result evaluation by source, value, etc.; specify the presentation medium for a report or alert, such as an email, the text to speech engine, a message to a pager, or a text and graphics display; and specify the preferred format for the report, such as information to be presented, order of information to be presented, preferred abbreviations or other variable substitutions.

Claims	Prior Art
	Kennewick '209 at 20:42–56 The domain agent 156 may strongly influence how the results are presented. For instance, once the domain agent 156 has created a satisfactory response to a question, the agent 156 may format that response for presentation. Typically, the domain agent 156 formats the response into the markup format used by the text to speech engine 124. The domain agent 156 may also format the result presentation using available format templates and based on the context, the criteria, and the user profile 110. Variable substitutions and transformations may be performed by the agent 156 to produce a response best understood and most natural to the user. Further, the order of presentation of tokens and the exact terminology used to create a more natural response to the user may be varied by the agent 156. The domain agent 156 may also select the presentation personality 108 to be used. Kennewick '209 at 20:57–21:3 The domain agent 156, may select the presentation template, determine order of presentation for tokens and determine variable substitutions and transformations using probabilistic or fuzzy set decision methods. The template used to form the presentation can be from the domain agent itself or from the user profile 110. The user profile 110 can completely specify the presentation format or can be used to select and then modify an existing presentation format. Selection and formatting of presentation template can also depend on the presentation personality modules 108. At the same time, the characteristics of the personality used for the response are dynamically determined using probabilities or fuzzy possibilities derived from the context, the criteria, the domain agent itself and the user profile 110. Kennewick '209 at 21:4–21 The domain agent 156 may apply a number of transformations to the tokens before presentation to the user. These variable substitutions and transformations are derived from a number of sources including, domain information carried by the agent, the context, the token values, the criteria, the personality module 108 to be used, and the user profile 110. Examples of variable substitutions and transformations include: substitution of words for numbers; substitution of names for acronyms or symbols (i.e., trading symbols); use of formatting information derived from the information sources (i.e., HTML tags); nature of the response including, text, long text, list, table; possible missing information or errors; units for measurement (i.e., English or metric); and preferred terminology from the user profile 110 or presentation personality 108.

Claims	Prior Art
	A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209 to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim [5a]. Specifically, a POSITA would be motivated to implement the logic disclosed in Kennewick '209 to determine a best means for presentation, and a POSITA would have a reasonable expectation of success based on the disclosures presented in Kennewick '209, including analyzing user preferences and histories and analyzing context of the date (e.g., how long it would take to present the synthesized audio). To the extent that Apple Prior Art Systems and Kennewick '209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick '209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
10. A wireless communication system as recited in claim 5, wherein said determining comprises determining whether the wireless headset is active.	Apple Prior Art Systems and Ewer '149 expressly or inherently disclose “[a] wireless communication system as recited in claim 5, wherein said determining comprises determining whether the wireless headset is active.” <i>See</i> above at claim [2a]. A POSITA would have been motivated to combine the teachings of Apple Prior Art Systems and Ewer '149 to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim [2a]. To the extent that Apple Prior Art Systems and Ewer '149 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Ewer '149, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
11. A wireless communication system as recited in claim 1, wherein said converting of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless communication system as recited in claim 1, wherein said converting of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, include settings (e.g., preferences) to set a voice for converting text to audio such as voice and frequency adjustments, which is an example of the claimed “user-selected speech characteristic.” The following citations show examples of where the Apple Prior Art Systems disclose this limitation.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 26 Reading Text Now that you have learned to enter and format text using voice commands, it's time to have ViaVoice read the sentences back to you. Highlight the desired text block and say one of the following voice commands. SAY> BEGIN READING SAY> READ SELECTION ViaVoice reads the highlighted text using the voice selected in the Voices panel in ViaVoice Preferences. On applications with copy and selection capabilities, the Read Selection menu item reads the text using the voice currently selected in the Voices panel in ViaVoice Preferences. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 48 Reading Your Text Provided your application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back text aloud from almost any source, including documents and e-mail messages. 1. Open the application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected system voice. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 49 Changing Voices  You can select the voice used when reading text. 1. Click the ViaVoice Preferences button on the VoiceCenter drawer. 2. Click the Voices icon. 3. From the Voice menu, select a new voice. 4. Press Play to hear a sample of your selection. 5. Click Apply, and close the window. 6. Any text read back from this point uses the voice you just selected. IBM ViaVoice Simply Dictation user guide at Page 61  SpeakPad Menu The SpeakPad menus contain items that are similar in functionality to other word processors. SpeakPad File Edit Format Dictation Window Help About SpeakPad Preferences... Dictation Menu The following items are available on the Dictation Menu. • Playback Selection/Stop Playback - This option allows you to play back your dictated text in your own voice. ViaVoice can play back up to 1,000 words (punctuation is included in the word count). The playback feature is useful in many instances, such as when you are correcting text. For more information, see "Playing Back Dictated Text in SpeakPad" on page 48. • Read Selection/Stop Reading - This option reads any text you have highlighted in a computer-generated voice. You can use this feature to help you proofread your documents. For more information, see "Reading Your Text" on page 48. • Voice Preferences - This option displays the Voices window, which allows you to change the computer-generated voice. For more information, see "Voices Preferences" on page 65. IBM ViaVoice Simply Dictation user guide at Page 65 Voices Preferences This page allows you to select the voice that reads your text and hear a sample played. From the Speech, Pitch, and Voice Quality tabs, adjust the sliders until you reach the desired settings. From the Options tab, select the hot key for reading text.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4-3 This chapter begins with an introduction to the speech generation process and then discusses how you can ▪ check for the availability of the Speech Manager ▪ create and dispose of speech channels ▪ generate speech with different voices ▪ obtain information about and change speech channel settings ▪ start and stop speech production ▪ synchronize speech production with other activities by using callback procedures ▪ embed Speech Manager commands within text to make it more understandable Inside Macintosh Sound_ 1994 at Page 4-5 Voices Your application can use the system default voice to generate speech or it can specify that the Speech Manager use a particular voice that is available on the current computer system. A voice is a set of characteristics defined in parameters that specify a particular quality of speech. Just as different people's voices have different tonal qualities, so too can different voices have different qualities. A synthesized voice might sound male or female and might sound like an adult or a child. Some voices sound distinctively synthetic, while others sound more like real people. Figure 4-2 shows how the Speech Manager uses speech channels to synthesize speech with different voices. Inside Macintosh Sound_ 1994 at Page 4-5 Figure 4-2 The Speech Manager and multiple voices  <pre> graph TD SM[Speech Manager] --> A[Speech synthesizer A] SM --> B[Speech synthesizer B] SM --> C[Speech synthesizer C] A --> V1[Voice 1] A --> V2[Voice 2] B --> V3[Voice 3] C --> V4[Voice 4] </pre> The diagram illustrates the architecture of the Speech Manager and its interaction with multiple voices. At the top, a box labeled "Speech Manager" is connected by three downward arrows to three separate speech synthesizers, labeled A, B, and C. These synthesizers are grouped within a shaded rectangular area labeled "Speech synthesizers" on the left. Below this group, four boxes represent individual voices: "Voice 1", "Voice 2", "Voice 3", and "Voice 4". Arrows indicate the flow of data from the synthesizers to the voices: synthesizer A connects to Voice 1 and Voice 2; synthesizer B connects to Voice 3; and synthesizer C connects to Voice 4.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–6 Most of the time, your application designates the voice that speaks text, and usually that is the default voice. Based on the needs of your users and the way in which you expect them to use voices in your application you can provide access to voices in a number of different ways. You could include access to selecting voices in a dialog box that is available from a menu item such as Voices... Any application that allows users to choose among voices requires additional information about the available voices beyond the information provided by a voice specification record (described in detail on page 4-46), whose data should never be presented to the user. Such additional information might include the name of the voice as well as what script and language it supports. Inside Macintosh Sound_ 1994 at Page 4–6 Applications can use the <code>GetVoiceDescription</code> function (described in detail on page 4-66) with a voice specification record to obtain such information in a voice description record (described in detail on page 4-47). You might provide access to voices through a control panel. For information about implementing control panels, see <i>Inside Macintosh: More Macintosh Toolbox</i>. Or, you could implement a voices menu in your application's main menu bar, if you think that users will want to change the voice often and you have the room available. It's not a good idea to implement a hierarchical Voices menu since hierarchical menus are harder to use. For more information about choosing a user interface for your application, see <i>Macintosh Human Interface Guidelines</i>.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–6 Speech Attributes Any given person has only one voice, but can alter the characteristics of his or her speech in a number of different ways. For example, a person can speak slowly or quickly and with a low or a high pitch. Similarly, the Speech Manager provides routines that allow you to modify these and other speech attributes, regardless of which voice is in use. A speech attribute is a setting defined for a class of voices or for all voices that affects the quality of speech produced by the Speech Manager. The Speech Manager provides routines to directly alter two speech attributes—speech rate and speech pitch. These routines are described in the section “Changing Speech Attributes” beginning on page 4-73. You can change two other speech attributes—pitch modulation and speech volume—by using the mechanism of speech information selectors, which is described in the section “Speech Information Selectors” beginning on page 4-39. The speech rate of a speech channel is the approximate number of words of text that the synthesizer should say in one minute. Slower speech rates make the speech easier to understand, but can be annoyingly tedious to listen to. Some applications, such as aids for the visually impaired, require very fast speech rates. Speech rates are expressed as fixed-point values. Each speech synthesizer determines its own range of speech rates. The speech pitch of a speech channel represents the middle pitch of the voice, roughly corresponding to the key in which a song is played. It is a fixed-point value in the range of 0.000 through 127.000, where 60.000 corresponds to middle C on a conventional piano. Each 1.000-unit change in a value corresponds to a musical half step. This is the same scale used in specifying MIDI note values, as described in the chapter “Sound Manager” Inside Macintosh Sound_ 1994 at Page 4–7 You can determine the current speech pitch on a speech channel by calling the <code>GetSpeechPitch</code> function, described on page 4-75. You can change the current pitch by calling the <code>SetSpeechPitch</code> function, described on page 4-76. You can also determine the current speech rate and change it by using the <code>GetSpeechRate</code> function, described on page 4-73 and the <code>SetSpeechRate</code> function, described on page 4-74. Changes in speech pitch and speech rate are effective immediately (as soon as the synthesizer can respond), even if they occur in the middle of a word. Inside Macintosh Sound_ 1994 at Page 4–9 You can obtain the pitch modulation by using the <code>GetSpeechInfo</code> function with the <code>soPitchMod</code> speech information selector, and you can change the pitch modulation by using the <code>SetSpeechInfo</code> function with the same selector. Speech information selectors are described in “Speech Information Selectors” beginning on page 4-39.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–11 To take advantage of most of the Speech Manager's features, you must allocate a speech channel to pass to Speech Manager functions and dispose of the speech channel when you are finished using it. "Creating, Using, and Disposing of a Speech Channel" beginning on page 4-13 demonstrates how you do this and shows how you can use the <code>SpeakText</code> function to start speech generation from a buffer of text. Some applications permit users to choose a voice from those available to be used for speech generation. The <code>CountVoices</code>, <code>GetIndVoice</code>, and <code>GetVoiceDescription</code> functions support this capability. "Working With Different Voices" beginning on page 4-14 shows how you can use these functions to choose among available voices. Inside Macintosh Sound_1994 at Page 4–14 <u>Working With Different Voices</u> When you work with speech channels, you can set a voice for a particular channel. When you set a voice, you may want to filter out certain of its characteristics in order to identify the one you want. For example, in an educational software application for elementary school students, you may want to use only children's voices. In order to choose the voice you want, you get a voice description record that contains information about a voice such as the size of the voice, the name of the voice, the age and gender of the voice, and the synthesizer with which it works. You can get the number of available voices using the <code>CountVoices</code> function. You can cycle through the available voices and identify the one you want to use by using the <code>GetIndVoice</code> function. Then you fill out a voice description record using the <code>GetVoiceDescription</code> function. Listing 4-3 shows how to get identifying information about a voice. Inside Macintosh Sound_1994 at Page 4–16 <u>Adjusting Speech Attributes</u> Speech attributes are settings defined for a class of voices or for all voices that affect the quality of speech produced by the Speech Manager. In general, an application should not try to second-guess the developers of a voice or synthesizer by arbitrarily setting a speech attribute. However, there are some cases in which you would want to adjust the rate of speech (how many words per minute are spoken) or the speech pitch (the listener's subjective interpretation of speech's average frequency). Listing 4-4 shows how to adjust the speech pitch and speech rate of a particular channel.


Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–23 <u>Writing Embedded Speech Commands</u> Embedded speech commands allow you to customize the quality of speech output by fine tuning it. You can make speech much easier to understand than the default way in which text is spoken by a synthesizer. An embedded speech command is a command embedded within a text buffer to be spoken by the Speech Manager that causes the Speech Manager to take a certain action. For example, you could use an embedded speech command to emphasize a particular word in a text string to make it stand out to the user. Inside Macintosh Sound_ 1994 at Page 4–30 <u>Examples of Embedded Speech Commands</u> If you use just a few of the embedded speech commands, you can markedly increase the understandability of text spoken by your application. Your application knows more about the speech being produced than a speech synthesizer does. A synthesizer speaks text according to a predetermined set of rules about language production. Therefore, the voices available on a Macintosh computer with the Speech Manager installed sound very synthetic and sometimes robotic because the pronunciation rules are formalized. You can make the speech produced by the synthesizer sound a lot more human by observing some simple rules of human speech and embedding speech commands in text according to these conventions. The techniques presented in this section could be applied when your application is having a dialog with the user or speaking some error messages or announcements.


Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–47 <u>Voice Description Records</u> By calling the <code>GetVoiceDescription</code> function, you can obtain information about a voice in a voice description record. The <code>VoiceDescription</code> data type defines a voice description record. <pre data-bbox="766 370 1528 821"> TYPE VoiceDescription = RECORD length: LongInt; {size of record} voice: VoiceSpec; {voice synthesizer and ID info} version: LongInt; {version number of voice} name: Str63; {name of voice} comment: Str255; {text information about voice} gender: Integer; {neuter, male, or female} age: Integer; {approximate age in years} script: Integer; {script code of text voice can process} language: Integer; {language code of voice output} region: Integer; {region code of voice output} reserved1: LongInt; {always 0--reserved for future use} reserved2: LongInt; {always 0--reserved for future use} reserved3: LongInt; {always 0--reserved for future use} reserved4: LongInt; {always 0--reserved for future use} END; </pre> Inside Macintosh Sound_ 1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ■ generate speech and then pause or stop it ■ obtain information about an individual voice or all voices ■ create and dispose of speech channels ■ obtain the Speech Manager's version and status ■ change the rate or pitch of speech

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–73 <u>Changing Speech Attributes</u> To determine the rate and pitch at which a speech channel is processing text, you can use the <code>GetSpeechRate</code> and <code>GetSpeechPitch</code> functions. The <code>SetSpeechRate</code> and <code>SetSpeechPitch</code> functions allow you to change rate and pitch. The most robust of the Speech Manager's routines are the <code>GetSpeechInfo</code> and <code>SetSpeechInfo</code> functions. These allow you to obtain many types of information about a speech channel and to change many settings of a speech channel. To specify the operation that you wish to perform, you must pass <code>GetSpeechInfo</code> or <code>SetSpeechInfo</code> a selector. A full list of selectors is provided in "Speech Information Selectors" beginning on page 4-39. Inside Macintosh Sound_ 1994 at Page 4–76 <u>SetSpeechPitch</u> You can use the <code>SetSpeechPitch</code> function to set the speech pitch on a designated speech channel. <pre>FUNCTION SetSpeechPitch (chan: SpeechChannel; pitch: Fixed) : OSErr;</pre> <code>chan</code> The speech channel whose pitch you wish to set. <code>pitch</code> The new pitch for the speech channel, expressed as a fixed-point frequency value. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
12. A wireless communication system as recited in claim 11, wherein the at least one user-selected speech characteristic is set based on at least one user preference that is able to be set by the first user.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless communication system as recited in claim 11, wherein the at least one user-selected speech characteristic is set based on at least one user preference that is able to be set by the first user.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, include settings (e.g., preferences) to set or modify a voice for converting text to audio, which is an example of the claimed “user–selected speech characteristic.” The following citations show examples of where the Apple Prior Art Systems disclose this limitation.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 26 Reading Text Now that you have learned to enter and format text using voice commands, it's time to have ViaVoice read the sentences back to you. Highlight the desired text block and say one of the following voice commands. SAY> BEGIN READING SAY> READ SELECTION ViaVoice reads the highlighted text using the voice selected in the Voices panel in ViaVoice Preferences. On applications with copy and selection capabilities, the Read Selection menu item reads the text using the voice currently selected in the Voices panel in ViaVoice Preferences. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 48 Reading Your Text Provided your application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back text aloud from almost any source, including documents and e-mail messages. 1. Open the application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected system voice. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 49 Changing Voices  You can select the voice used when reading text. 1. Click the ViaVoice Preferences button on the VoiceCenter drawer. 2. Click the Voices icon. 3. From the Voice menu, select a new voice. 4. Press Play to hear a sample of your selection. 5. Click Apply, and close the window. 6. Any text read back from this point uses the voice you just selected. IBM ViaVoice Simply Dictation user guide at Page 57 VoiceCenter Drawer The VoiceCenter drawer menu allows easy access to most of the ViaVoice options, functions, and information.  • What Can I Say - Displays a list of available voice commands. • Correction - Displays the Correction window. • Preferences - Use to customize functions of ViaVoice. • Macros - Displays the Macro Editor window so you can create and edit dictation macros. • Manage My Vocabulary - Displays the Manage My Vocabulary window so you can re-train the pronunciation of words. • UserSettings - Use to add Voice Models or change other User information. • Information Central - Displays Information Central, the ViaVoice online help file. • SetupAssistant - Use to set up your audio, create a user or analyze documents. • About ViaVoice - Displays the ViaVoice About box.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 61  SpeakPad Menu The SpeakPad menus contain items that are similar in functionality to other word processors. SpeakPad File Edit Format Dictation Window Help About SpeakPad Preferences... Dictation Menu The following items are available on the Dictation Menu. • Playback Selection/Stop Playback - This option allows you to play back your dictated text in your own voice. ViaVoice can play back up to 1,000 words (punctuation is included in the word count). The playback feature is useful in many instances, such as when you are correcting text. For more information, see "Playing Back Dictated Text in SpeakPad" on page 48. • Read Selection/Stop Reading - This option reads any text you have highlighted in a computer-generated voice. You can use this feature to help you proofread your documents. For more information, see "Reading Your Text" on page 48. • Voice Preferences - This option displays the Voices window, which allows you to change the computer-generated voice. For more information, see "Voices Preferences" on page 65. IBM ViaVoice Simply Dictation user guide at Page 65 Voices Preferences This page allows you to select the voice that reads your text and hear a sample played. From the Speech, Pitch, and Voice Quality tabs, adjust the sliders until you reach the desired settings. From the Options tab, select the hot key for reading text. Inside Macintosh Sound_1994 at Page 4–3 This chapter begins with an introduction to the speech generation process and then discusses how you can ▪ check for the availability of the Speech Manager ▪ create and dispose of speech channels ▪ generate speech with different voices ▪ obtain information about and change speech channel settings ▪ start and stop speech production ▪ synchronize speech production with other activities by using callback procedures ▪ embed Speech Manager commands within text to make it more understandable

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4-5 <u>Voices</u> Your application can use the system default voice to generate speech or it can specify that the Speech Manager use a particular voice that is available on the current computer system. A voice is a set of characteristics defined in parameters that specify a particular quality of speech. Just as different people's voices have different tonal qualities, so too can different voices have different qualities. A synthesized voice might sound male or female and might sound like an adult or a child. Some voices sound distinctively synthetic, while others sound more like real people. Figure 4-2 shows how the Speech Manager uses speech channels to synthesize speech with different voices. Inside Macintosh Sound_ 1994 at Page 4-5 Figure 4-2 The Speech Manager and multiple voices  Inside Macintosh Sound_ 1994 at Page 4-6 Applications can use the <code>GetVoiceDescription</code> function (described in detail on page 4-66) with a voice specification record to obtain such information in a voice description record (described in detail on page 4-47). You might provide access to voices through a control panel. For information about implementing control panels, see <i>Inside Macintosh: More Macintosh Toolbox</i>. Or, you could implement a voices menu in your application's main menu bar, if you think that users will want to change the voice often and you have the room available. It's not a good idea to implement a hierarchical Voices menu since hierarchical menus are harder to use. For more information about choosing a user interface for your application, see <i>Macintosh Human Interface Guidelines</i>.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–6 Speech Attributes Any given person has only one voice, but can alter the characteristics of his or her speech in a number of different ways. For example, a person can speak slowly or quickly and with a low or a high pitch. Similarly, the Speech Manager provides routines that allow you to modify these and other speech attributes, regardless of which voice is in use. A speech attribute is a setting defined for a class of voices or for all voices that affects the quality of speech produced by the Speech Manager. The Speech Manager provides routines to directly alter two speech attributes—speech rate and speech pitch. These routines are described in the section “Changing Speech Attributes” beginning on page 4-73. You can change two other speech attributes—pitch modulation and speech volume—by using the mechanism of speech information selectors, which is described in the section “Speech Information Selectors” beginning on page 4-39. The speech rate of a speech channel is the approximate number of words of text that the synthesizer should say in one minute. Slower speech rates make the speech easier to understand, but can be annoyingly tedious to listen to. Some applications, such as aids for the visually impaired, require very fast speech rates. Speech rates are expressed as fixed-point values. Each speech synthesizer determines its own range of speech rates. The speech pitch of a speech channel represents the middle pitch of the voice, roughly corresponding to the key in which a song is played. It is a fixed-point value in the range of 0.000 through 127.000, where 60.000 corresponds to middle C on a conventional piano. Each 1.000-unit change in a value corresponds to a musical half step. This is the same scale used in specifying MIDI note values, as described in the chapter “Sound Manager” Inside Macintosh Sound_ 1994 at Page 4–6 <i>More of my voice...</i> Most of the time, your application designates the voice that speaks text, and usually that is the default voice. Based on the needs of your users and the way in which you expect them to use voices in your application you can provide access to voices in a number of different ways. You could include access to selecting voices in a dialog box that is available from a menu item such as Voices... Any application that allows users to choose among voices requires additional information about the available voices beyond the information provided by a voice specification record (described in detail on page 4-46), whose data should never be presented to the user. Such additional information might include the name of the voice as well as what script and language it supports. Inside Macintosh Sound_ 1994 at Page 4–7 You can determine the current speech pitch on a speech channel by calling the <code>GetSpeechPitch</code> function, described on page 4-75. You can change the current pitch by calling the <code>SetSpeechPitch</code> function, described on page 4-76. You can also determine the current speech rate and change it by using the <code>GetSpeechRate</code> function, described on page 4-73 and the <code>SetSpeechRate</code> function, described on page 4-74. Changes in speech pitch and speech rate are effective immediately (as soon as the synthesizer can respond), even if they occur in the middle of a word.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–9 You can obtain the pitch modulation by using the <code>GetSpeechInfo</code> function with the <code>soPitchMod</code> speech information selector, and you can change the pitch modulation by using the <code>SetSpeechInfo</code> function with the same selector. Speech information selectors are described in “Speech Information Selectors” beginning on page 4-39. Inside Macintosh Sound_1994 at Page 4–11 To take advantage of most of the Speech Manager’s features, you must allocate a speech channel to pass to Speech Manager functions and dispose of the speech channel when you are finished using it. “Creating, Using, and Disposing of a Speech Channel” beginning on page 4-13 demonstrates how you do this and shows how you can use the <code>SpeakText</code> function to start speech generation from a buffer of text. Some applications permit users to choose a voice from those available to be used for speech generation. The <code>CountVoices</code>, <code>GetIndVoice</code>, and <code>GetVoiceDescription</code> functions support this capability. “Working With Different Voices” beginning on page 4-14 shows how you can use these functions to choose among available voices. Inside Macintosh Sound_1994 at Page 4–14 <u>Working With Different Voices</u> When you work with speech channels, you can set a voice for a particular channel. When you set a voice, you may want to filter out certain of its characteristics in order to identify the one you want. For example, in an educational software application for elementary school students, you may want to use only children’s voices. In order to choose the voice you want, you get a voice description record that contains information about a voice such as the size of the voice, the name of the voice, the age and gender of the voice, and the synthesizer with which it works. You can get the number of available voices using the <code>CountVoices</code> function. You can cycle through the available voices and identify the one you want to use by using the <code>GetIndVoice</code> function. Then you fill out a voice description record using the <code>GetVoiceDescription</code> function. Listing 4-3 shows how to get identifying information about a voice. Inside Macintosh Sound_1994 at Page 4–16 <u>Adjusting Speech Attributes</u> Speech attributes are settings defined for a class of voices or for all voices that affect the quality of speech produced by the Speech Manager. In general, an application should not try to second-guess the developers of a voice or synthesizer by arbitrarily setting a speech attribute. However, there are some cases in which you would want to adjust the rate of speech (how many words per minute are spoken) or the speech pitch (the listener’s subjective interpretation of speech’s average frequency). Listing 4-4 shows how to adjust the speech pitch and speech rate of a particular channel.


Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–23 <u>Writing Embedded Speech Commands</u> Embedded speech commands allow you to customize the quality of speech output by fine tuning it. You can make speech much easier to understand than the default way in which text is spoken by a synthesizer. An embedded speech command is a command embedded within a text buffer to be spoken by the Speech Manager that causes the Speech Manager to take a certain action. For example, you could use an embedded speech command to emphasize a particular word in a text string to make it stand out to the user. Inside Macintosh Sound_ 1994 at Page 4–30 <u>Examples of Embedded Speech Commands</u> If you use just a few of the embedded speech commands, you can markedly increase the understandability of text spoken by your application. Your application knows more about the speech being produced than a speech synthesizer does. A synthesizer speaks text according to a predetermined set of rules about language production. Therefore, the voices available on a Macintosh computer with the Speech Manager installed sound very synthetic and sometimes robotic because the pronunciation rules are formalized. You can make the speech produced by the synthesizer sound a lot more human by observing some simple rules of human speech and embedding speech commands in text according to these conventions. The techniques presented in this section could be applied when your application is having a dialog with the user or speaking some error messages or announcements.


Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–47 <u>Voice Description Records</u> By calling the <code>GetVoiceDescription</code> function, you can obtain information about a voice in a voice description record. The <code>VoiceDescription</code> data type defines a voice description record. <pre data-bbox="766 370 1528 821"> TYPE VoiceDescription = RECORD length: LongInt; {size of record} voice: VoiceSpec; {voice synthesizer and ID info} version: LongInt; {version number of voice} name: Str63; {name of voice} comment: Str255; {text information about voice} gender: Integer; {neuter, male, or female} age: Integer; {approximate age in years} script: Integer; {script code of text voice can process} language: Integer; {language code of voice output} region: Integer; {region code of voice output} reserved1: LongInt; {always 0--reserved for future use} reserved2: LongInt; {always 0--reserved for future use} reserved3: LongInt; {always 0--reserved for future use} reserved4: LongInt; {always 0--reserved for future use} END; </pre> Inside Macintosh Sound_1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ■ generate speech and then pause or stop it ■ obtain information about an individual voice or all voices ■ create and dispose of speech channels ■ obtain the Speech Manager's version and status ■ change the rate or pitch of speech

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–73 <u>Changing Speech Attributes</u> To determine the rate and pitch at which a speech channel is processing text, you can use the <code>GetSpeechRate</code> and <code>GetSpeechPitch</code> functions. The <code>SetSpeechRate</code> and <code>SetSpeechPitch</code> functions allow you to change rate and pitch. The most robust of the Speech Manager's routines are the <code>GetSpeechInfo</code> and <code>SetSpeechInfo</code> functions. These allow you to obtain many types of information about a speech channel and to change many settings of a speech channel. To specify the operation that you wish to perform, you must pass <code>GetSpeechInfo</code> or <code>SetSpeechInfo</code> a selector. A full list of selectors is provided in "Speech Information Selectors" beginning on page 4-39. Inside Macintosh Sound_ 1994 at Page 4–76 <u>SetSpeechPitch</u> You can use the <code>SetSpeechPitch</code> function to set the speech pitch on a designated speech channel. <pre>FUNCTION SetSpeechPitch (chan: SpeechChannel; pitch: Fixed) : OSErr;</pre> <code>chan</code> The speech channel whose pitch you wish to set. <code>pitch</code> The new pitch for the speech channel, expressed as a fixed-point frequency value. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
13. A wireless communication system as recited in claim 11, wherein the at least one user-selected speech characteristic is set based on at least one user setting provided at said wireless communication device.	Apple Prior Art Systems expressly or inherently disclose "[a] wireless communication system as recited in claim 11, wherein the at least one user-selected speech characteristic is set based on at least one user setting provided at said wireless communication device." Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, include settings (e.g., preferences) that can be set at the computer or PDA (e.g., the "wireless communication device") to set a voice for converting text to audio, which is an example of the claimed "user-selected speech characteristic." The following citations show examples of where the Apple Prior Art Systems disclose this limitation.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 26 Reading Text Now that you have learned to enter and format text using voice commands, it's time to have ViaVoice read the sentences back to you. Highlight the desired text block and say one of the following voice commands. SAY> BEGIN READING SAY> READ SELECTION ViaVoice reads the highlighted text using the voice selected in the Voices panel in ViaVoice Preferences. On applications with copy and selection capabilities, the Read Selection menu item reads the text using the voice currently selected in the Voices panel in ViaVoice Preferences. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 48 Reading Your Text Provided your application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back text aloud from almost any source, including documents and e-mail messages. 1. Open the application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected system voice. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 49 Changing Voices  You can select the voice used when reading text. 1. Click the ViaVoice Preferences button on the VoiceCenter drawer. 2. Click the Voices icon. 3. From the Voice menu, select a new voice. 4. Press Play to hear a sample of your selection. 5. Click Apply, and close the window. 6. Any text read back from this point uses the voice you just selected. IBM ViaVoice Simply Dictation user guide at Page 57 VoiceCenter Drawer The VoiceCenter drawer menu allows easy access to most of the ViaVoice options, functions, and information.  • What Can I Say - Displays a list of available voice commands. • Correction - Displays the Correction window. • Preferences - Use to customize functions of ViaVoice. • Macros - Displays the Macro Editor window so you can create and edit dictation macros. • Manage My Vocabulary - Displays the Manage My Vocabulary window so you can re-train the pronunciation of words. • UserSettings - Use to add Voice Models or change other User information. • Information Central - Displays Information Central, the ViaVoice online help file. • SetupAssistant - Use to set up your audio, create a user or analyze documents. • About ViaVoice - Displays the ViaVoice About box.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 61  SpeakPad Menu The SpeakPad menus contain items that are similar in functionality to other word processors. SpeakPad File Edit Format Dictation Window Help About SpeakPad Preferences... Dictation Menu The following items are available on the Dictation Menu. • Playback Selection/Stop Playback - This option allows you to play back your dictated text in your own voice. ViaVoice can play back up to 1,000 words (punctuation is included in the word count). The playback feature is useful in many instances, such as when you are correcting text. For more information, see "Playing Back Dictated Text in SpeakPad" on page 48. • Read Selection/Stop Reading - This option reads any text you have highlighted in a computer-generated voice. You can use this feature to help you proofread your documents. For more information, see "Reading Your Text" on page 48. • Voice Preferences - This option displays the Voices window, which allows you to change the computer-generated voice. For more information, see "Voices Preferences" on page 65. IBM ViaVoice Simply Dictation user guide at Page 65 Voices Preferences This page allows you to select the voice that reads your text and hear a sample played. From the Speech, Pitch, and Voice Quality tabs, adjust the sliders until you reach the desired settings. From the Options tab, select the hot key for reading text. Inside Macintosh Sound_1994 at Page 4–3 This chapter begins with an introduction to the speech generation process and then discusses how you can ▪ check for the availability of the Speech Manager ▪ create and dispose of speech channels ▪ generate speech with different voices ▪ obtain information about and change speech channel settings ▪ start and stop speech production ▪ synchronize speech production with other activities by using callback procedures ▪ embed Speech Manager commands within text to make it more understandable

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4-5 <u>Voices</u> Your application can use the system default voice to generate speech or it can specify that the Speech Manager use a particular voice that is available on the current computer system. A voice is a set of characteristics defined in parameters that specify a particular quality of speech. Just as different people's voices have different tonal qualities, so too can different voices have different qualities. A synthesized voice might sound male or female and might sound like an adult or a child. Some voices sound distinctively synthetic, while others sound more like real people. Figure 4-2 shows how the Speech Manager uses speech channels to synthesize speech with different voices. Inside Macintosh Sound_1994 at Page 4-5 Figure 4-2 The Speech Manager and multiple voices  Inside Macintosh Sound_1994 at Page 4-6 Most of the time, your application designates the voice that speaks text, and usually that is the default voice. Based on the needs of your users and the way in which you expect them to use voices in your application you can provide access to voices in a number of different ways. You could include access to selecting voices in a dialog box that is available from a menu item such as Voices... Any application that allows users to choose among voices requires additional information about the available voices beyond the information provided by a voice specification record (described in detail on page 4-46), whose data should never be presented to the user. Such additional information might include the name of the voice as well as what script and language it supports.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–6 Applications can use the <code>GetVoiceDescription</code> function (described in detail on page 4-66) with a voice specification record to obtain such information in a voice description record (described in detail on page 4-47). You might provide access to voices through a control panel. For information about implementing control panels, see <i>Inside Macintosh: More Macintosh Toolbox</i>. Or, you could implement a voices menu in your application's main menu bar, if you think that users will want to change the voice often and you have the room available. It's not a good idea to implement a hierarchical Voices menu since hierarchical menus are harder to use. For more information about choosing a user interface for your application, see <i>Macintosh Human Interface Guidelines</i>. Inside Macintosh Sound_1994 at Page 4–6 <u>Speech Attributes</u> Any given person has only one voice, but can alter the characteristics of his or her speech in a number of different ways. For example, a person can speak slowly or quickly and with a low or a high pitch. Similarly, the Speech Manager provides routines that allow you to modify these and other speech attributes, regardless of which voice is in use. A speech attribute is a setting defined for a class of voices or for all voices that affects the quality of speech produced by the Speech Manager. The Speech Manager provides routines to directly alter two speech attributes—speech rate and speech pitch. These routines are described in the section “Changing Speech Attributes” beginning on page 4-73. You can change two other speech attributes—pitch modulation and speech volume—by using the mechanism of speech information selectors, which is described in the section “Speech Information Selectors” beginning on page 4-39. The speech rate of a speech channel is the approximate number of words of text that the synthesizer should say in one minute. Slower speech rates make the speech easier to understand, but can be annoyingly tedious to listen to. Some applications, such as aids for the visually impaired, require very fast speech rates. Speech rates are expressed as fixed-point values. Each speech synthesizer determines its own range of speech rates. The speech pitch of a speech channel represents the middle pitch of the voice, roughly corresponding to the key in which a song is played. It is a fixed-point value in the range of 0.000 through 127.000, where 60.000 corresponds to middle C on a conventional piano. Each 1.000-unit change in a value corresponds to a musical half step. This is the same scale used in specifying MIDI note values, as described in the chapter “Sound Manager” Inside Macintosh Sound_1994 at Page 4–7 You can determine the current speech pitch on a speech channel by calling the <code>GetSpeechPitch</code> function, described on page 4-75. You can change the current pitch by calling the <code>SetSpeechPitch</code> function, described on page 4-76. You can also determine the current speech rate and change it by using the <code>GetSpeechRate</code> function, described on page 4-73 and the <code>SetSpeechRate</code> function, described on page 4-74. Changes in speech pitch and speech rate are effective immediately (as soon as the synthesizer can respond), even if they occur in the middle of a word.


Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–9 You can obtain the pitch modulation by using the <code>GetSpeechInfo</code> function with the <code>soPitchMod</code> speech information selector, and you can change the pitch modulation by using the <code>SetSpeechInfo</code> function with the same selector. Speech information selectors are described in “Speech Information Selectors” beginning on page 4-39. Inside Macintosh Sound_1994 at Page 4–11 To take advantage of most of the Speech Manager’s features, you must allocate a speech channel to pass to Speech Manager functions and dispose of the speech channel when you are finished using it. “Creating, Using, and Disposing of a Speech Channel” beginning on page 4-13 demonstrates how you do this and shows how you can use the <code>SpeakText</code> function to start speech generation from a buffer of text. Some applications permit users to choose a voice from those available to be used for speech generation. The <code>CountVoices</code>, <code>GetIndVoice</code>, and <code>GetVoiceDescription</code> functions support this capability. “Working With Different Voices” beginning on page 4-14 shows how you can use these functions to choose among available voices. Inside Macintosh Sound_1994 at Page 4–14 <u>Working With Different Voices</u> When you work with speech channels, you can set a voice for a particular channel. When you set a voice, you may want to filter out certain of its characteristics in order to identify the one you want. For example, in an educational software application for elementary school students, you may want to use only children’s voices. In order to choose the voice you want, you get a voice description record that contains information about a voice such as the size of the voice, the name of the voice, the age and gender of the voice, and the synthesizer with which it works. You can get the number of available voices using the <code>CountVoices</code> function. You can cycle through the available voices and identify the one you want to use by using the <code>GetIndVoice</code> function. Then you fill out a voice description record using the <code>GetVoiceDescription</code> function. Listing 4-3 shows how to get identifying information about a voice. Inside Macintosh Sound_1994 at Page 4–16 <u>Adjusting Speech Attributes</u> Speech attributes are settings defined for a class of voices or for all voices that affect the quality of speech produced by the Speech Manager. In general, an application should not try to second-guess the developers of a voice or synthesizer by arbitrarily setting a speech attribute. However, there are some cases in which you would want to adjust the rate of speech (how many words per minute are spoken) or the speech pitch (the listener’s subjective interpretation of speech’s average frequency). Listing 4-4 shows how to adjust the speech pitch and speech rate of a particular channel.


Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–23 <u>Writing Embedded Speech Commands</u> Embedded speech commands allow you to customize the quality of speech output by fine tuning it. You can make speech much easier to understand than the default way in which text is spoken by a synthesizer. An embedded speech command is a command embedded within a text buffer to be spoken by the Speech Manager that causes the Speech Manager to take a certain action. For example, you could use an embedded speech command to emphasize a particular word in a text string to make it stand out to the user. Inside Macintosh Sound_ 1994 at Page 4–30 <u>Examples of Embedded Speech Commands</u> If you use just a few of the embedded speech commands, you can markedly increase the understandability of text spoken by your application. Your application knows more about the speech being produced than a speech synthesizer does. A synthesizer speaks text according to a predetermined set of rules about language production. Therefore, the voices available on a Macintosh computer with the Speech Manager installed sound very synthetic and sometimes robotic because the pronunciation rules are formalized. You can make the speech produced by the synthesizer sound a lot more human by observing some simple rules of human speech and embedding speech commands in text according to these conventions. The techniques presented in this section could be applied when your application is having a dialog with the user or speaking some error messages or announcements.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–47 <u>Voice Description Records</u> By calling the <code>GetVoiceDescription</code> function, you can obtain information about a voice in a voice description record. The <code>VoiceDescription</code> data type defines a voice description record. <pre data-bbox="766 370 1528 821"> TYPE VoiceDescription = RECORD length: LongInt; {size of record} voice: VoiceSpec; {voice synthesizer and ID info} version: LongInt; {version number of voice} name: Str63; {name of voice} comment: Str255; {text information about voice} gender: Integer; {neuter, male, or female} age: Integer; {approximate age in years} script: Integer; {script code of text voice can process} language: Integer; {language code of voice output} region: Integer; {region code of voice output} reserved1: LongInt; {always 0--reserved for future use} reserved2: LongInt; {always 0--reserved for future use} reserved3: LongInt; {always 0--reserved for future use} reserved4: LongInt; {always 0--reserved for future use} END; </pre> Inside Macintosh Sound_1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ■ generate speech and then pause or stop it ■ obtain information about an individual voice or all voices ■ create and dispose of speech channels ■ obtain the Speech Manager's version and status ■ change the rate or pitch of speech

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–73 <u>Changing Speech Attributes</u> To determine the rate and pitch at which a speech channel is processing text, you can use the <code>GetSpeechRate</code> and <code>GetSpeechPitch</code> functions. The <code>SetSpeechRate</code> and <code>SetSpeechPitch</code> functions allow you to change rate and pitch. The most robust of the Speech Manager's routines are the <code>GetSpeechInfo</code> and <code>SetSpeechInfo</code> functions. These allow you to obtain many types of information about a speech channel and to change many settings of a speech channel. To specify the operation that you wish to perform, you must pass <code>GetSpeechInfo</code> or <code>SetSpeechInfo</code> a selector. A full list of selectors is provided in "Speech Information Selectors" beginning on page 4-39. Inside Macintosh Sound_ 1994 at Page 4–76 <u>SetSpeechPitch</u> You can use the <code>SetSpeechPitch</code> function to set the speech pitch on a designated speech channel. <pre>FUNCTION SetSpeechPitch (chan: SpeechChannel; pitch: Fixed) : OSErr;</pre> <code>chan</code> The speech channel whose pitch you wish to set. <code>pitch</code> The new pitch for the speech channel, expressed as a fixed-point frequency value. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
14. A wireless communication system as recited in claim 11, wherein the at least one user-selected speech characteristic pertains to at least a voice for use in producing the incoming audio message.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless communication system as recited in claim 11, wherein the at least one user-selected speech characteristic pertains to at least a voice for use in producing the incoming audio message.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, includes settings (e.g., preferences) to set a voice for converting text to audio, which is an example of the claimed “user–selected speech characteristic.” The following citations show examples of where the Apple Prior Art Systems disclose this limitation.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 26 Reading Text Now that you have learned to enter and format text using voice commands, it's time to have ViaVoice read the sentences back to you. Highlight the desired text block and say one of the following voice commands. SAY> BEGIN READING SAY> READ SELECTION ViaVoice reads the highlighted text using the voice selected in the Voices panel in ViaVoice Preferences. On applications with copy and selection capabilities, the Read Selection menu item reads the text using the voice currently selected in the Voices panel in ViaVoice Preferences. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 48 Reading Your Text Provided your application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back text aloud from almost any source, including documents and e-mail messages. 1. Open the application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected system voice. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 49 Changing Voices  You can select the voice used when reading text. 1. Click the ViaVoice Preferences button on the VoiceCenter drawer. 2. Click the Voices icon. 3. From the Voice menu, select a new voice. 4. Press Play to hear a sample of your selection. 5. Click Apply, and close the window. 6. Any text read back from this point uses the voice you just selected. IBM ViaVoice Simply Dictation user guide at Page 61  SpeakPad Menu The SpeakPad menus contain items that are similar in functionality to other word processors. SpeakPad File Edit Format Dictation Window Help About SpeakPad Preferences... Dictation Menu The following items are available on the Dictation Menu. • Playback Selection/Stop Playback - This option allows you to play back your dictated text in your own voice. ViaVoice can play back up to 1,000 words (punctuation is included in the word count). The playback feature is useful in many instances, such as when you are correcting text. For more information, see "Playing Back Dictated Text in SpeakPad" on page 48. • Read Selection/Stop Reading - This option reads any text you have highlighted in a computer-generated voice. You can use this feature to help you proofread your documents. For more information, see "Reading Your Text" on page 48. • Voice Preferences - This option displays the Voices window, which allows you to change the computer-generated voice. For more information, see "Voices Preferences" on page 65. IBM ViaVoice Simply Dictation user guide at Page 65 Voices Preferences This page allows you to select the voice that reads your text and hear a sample played. From the Speech, Pitch, and Voice Quality tabs, adjust the sliders until you reach the desired settings. From the Options tab, select the hot key for reading text.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4-3 This chapter begins with an introduction to the speech generation process and then discusses how you can ▪ check for the availability of the Speech Manager ▪ create and dispose of speech channels ▪ generate speech with different voices ▪ obtain information about and change speech channel settings ▪ start and stop speech production ▪ synchronize speech production with other activities by using callback procedures ▪ embed Speech Manager commands within text to make it more understandable Inside Macintosh Sound_ 1994 at Page 4-5 Voices Your application can use the system default voice to generate speech or it can specify that the Speech Manager use a particular voice that is available on the current computer system. A voice is a set of characteristics defined in parameters that specify a particular quality of speech. Just as different people's voices have different tonal qualities, so too can different voices have different qualities. A synthesized voice might sound male or female and might sound like an adult or a child. Some voices sound distinctively synthetic, while others sound more like real people. Figure 4-2 shows how the Speech Manager uses speech channels to synthesize speech with different voices. Inside Macintosh Sound_ 1994 at Page 4-5 Figure 4-2 The Speech Manager and multiple voices  <pre> graph TD SM[Speech Manager] --> A[Speech synthesizer A] SM --> B[Speech synthesizer B] SM --> C[Speech synthesizer C] A --> V1[Voice 1] A --> V2[Voice 2] B --> V3[Voice 3] C --> V4[Voice 4] </pre> The diagram illustrates the architecture of the Speech Manager and its interaction with multiple voices. At the top, a box labeled "Speech Manager" is connected by three downward arrows to three separate speech synthesizers, labeled A, B, and C. These synthesizers are grouped within a shaded rectangular area labeled "Speech synthesizers" on the left. Below this group, four boxes represent individual voices: "Voice 1", "Voice 2", "Voice 3", and "Voice 4". Arrows indicate the flow of data: from synthesizer A to Voice 1 and Voice 2; from synthesizer B to Voice 3; and from synthesizer C to Voice 4.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–6 Most of the time, your application designates the voice that speaks text, and usually that is the default voice. Based on the needs of your users and the way in which you expect them to use voices in your application you can provide access to voices in a number of different ways. You could include access to selecting voices in a dialog box that is available from a menu item such as Voices... Any application that allows users to choose among voices requires additional information about the available voices beyond the information provided by a voice specification record (described in detail on page 4-46), whose data should never be presented to the user. Such additional information might include the name of the voice as well as what script and language it supports. Inside Macintosh Sound_ 1994 at Page 4–6 Applications can use the <code>GetVoiceDescription</code> function (described in detail on page 4-66) with a voice specification record to obtain such information in a voice description record (described in detail on page 4-47). You might provide access to voices through a control panel. For information about implementing control panels, see <i>Inside Macintosh: More Macintosh Toolbox</i>. Or, you could implement a voices menu in your application's main menu bar, if you think that users will want to change the voice often and you have the room available. It's not a good idea to implement a hierarchical Voices menu since hierarchical menus are harder to use. For more information about choosing a user interface for your application, see <i>Macintosh Human Interface Guidelines</i>.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–6 Speech Attributes Any given person has only one voice, but can alter the characteristics of his or her speech in a number of different ways. For example, a person can speak slowly or quickly and with a low or a high pitch. Similarly, the Speech Manager provides routines that allow you to modify these and other speech attributes, regardless of which voice is in use. A speech attribute is a setting defined for a class of voices or for all voices that affects the quality of speech produced by the Speech Manager. The Speech Manager provides routines to directly alter two speech attributes—speech rate and speech pitch. These routines are described in the section “Changing Speech Attributes” beginning on page 4-73. You can change two other speech attributes—pitch modulation and speech volume—by using the mechanism of speech information selectors, which is described in the section “Speech Information Selectors” beginning on page 4-39. The speech rate of a speech channel is the approximate number of words of text that the synthesizer should say in one minute. Slower speech rates make the speech easier to understand, but can be annoyingly tedious to listen to. Some applications, such as aids for the visually impaired, require very fast speech rates. Speech rates are expressed as fixed-point values. Each speech synthesizer determines its own range of speech rates. The speech pitch of a speech channel represents the middle pitch of the voice, roughly corresponding to the key in which a song is played. It is a fixed-point value in the range of 0.000 through 127.000, where 60.000 corresponds to middle C on a conventional piano. Each 1.000-unit change in a value corresponds to a musical half step. This is the same scale used in specifying MIDI note values, as described in the chapter “Sound Manager” Inside Macintosh Sound_ 1994 at Page 4–11 To take advantage of most of the Speech Manager’s features, you must allocate a speech channel to pass to Speech Manager functions and dispose of the speech channel when you are finished using it. “Creating, Using, and Disposing of a Speech Channel” beginning on page 4-13 demonstrates how you do this and shows how you can use the SpeakText function to start speech generation from a buffer of text. Some applications permit users to choose a voice from those available to be used for speech generation. The CountVoices, GetIndVoice, and GetVoiceDescription functions support this capability. “Working With Different Voices” beginning on page 4-14 shows how you can use these functions to choose among available voices.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–14 <u>Working With Different Voices</u> When you work with speech channels, you can set a voice for a particular channel. When you set a voice, you may want to filter out certain of its characteristics in order to identify the one you want. For example, in an educational software application for elementary school students, you may want to use only children's voices. In order to choose the voice you want, you get a voice description record that contains information about a voice such as the size of the voice, the name of the voice, the age and gender of the voice, and the synthesizer with which it works. You can get the number of available voices using the <code>CountVoices</code> function. You can cycle through the available voices and identify the one you want to use by using the <code>GetIndVoice</code> function. Then you fill out a voice description record using the <code>GetVoiceDescription</code> function. Listing 4-3 shows how to get identifying information about a voice. Inside Macintosh Sound_1994 at Page 4–16 <u>Adjusting Speech Attributes</u> Speech attributes are settings defined for a class of voices or for all voices that affect the quality of speech produced by the Speech Manager. In general, an application should not try to second-guess the developers of a voice or synthesizer by arbitrarily setting a speech attribute. However, there are some cases in which you would want to adjust the rate of speech (how many words per minute are spoken) or the speech pitch (the listener's subjective interpretation of speech's average frequency). Listing 4-4 shows how to adjust the speech pitch and speech rate of a particular channel. Inside Macintosh Sound_1994 at Page 4–30 <u>Examples of Embedded Speech Commands</u> If you use just a few of the embedded speech commands, you can markedly increase the understandability of text spoken by your application. Your application knows more about the speech being produced than a speech synthesizer does. A synthesizer speaks text according to a predetermined set of rules about language production. Therefore, the voices available on a Macintosh computer with the Speech Manager installed sound very synthetic and sometimes robotic because the pronunciation rules are formalized. You can make the speech produced by the synthesizer sound a lot more human by observing some simple rules of human speech and embedding speech commands in text according to these conventions. The techniques presented in this section could be applied when your application is having a dialog with the user or speaking some error messages or announcements.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–47 <u>Voice Description Records</u> By calling the <code>GetVoiceDescription</code> function, you can obtain information about a voice in a voice description record. The <code>VoiceDescription</code> data type defines a voice description record. <pre> TYPE VoiceDescription = RECORD length: LongInt; {size of record} voice: VoiceSpec; {voice synthesizer and ID info} version: LongInt; {version number of voice} name: Str63; {name of voice} comment: Str255; {text information about voice} gender: Integer; {neuter, male, or female} age: Integer; {approximate age in years} script: Integer; {script code of text voice can process} language: Integer; {language code of voice output} region: Integer; {region code of voice output} reserved1: LongInt; {always 0--reserved for future use} reserved2: LongInt; {always 0--reserved for future use} reserved3: LongInt; {always 0--reserved for future use} reserved4: LongInt; {always 0--reserved for future use} END; </pre> Inside Macintosh Sound_1994 at Page 4–73 <u>Changing Speech Attributes</u> To determine the rate and pitch at which a speech channel is processing text, you can use the <code>GetSpeechRate</code> and <code>GetSpeechPitch</code> functions. The <code>SetSpeechRate</code> and <code>SetSpeechPitch</code> functions allow you to change rate and pitch. The most robust of the Speech Manager's routines are the <code>GetSpeechInfo</code> and <code>SetSpeechInfo</code> functions. These allow you to obtain many types of information about a speech channel and to change many settings of a speech channel. To specify the operation that you wish to perform, you must pass <code>GetSpeechInfo</code> or <code>SetSpeechInfo</code> a selector. A full list of selectors is provided in "Speech Information Selectors" beginning on page 4-39. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
15. A wireless communication system as recited in claim 11,	Apple Prior Art Systems expressly or inherently disclose "[a] wireless communication system as recited in claim 11, wherein said presenting of the incoming audio message comprises playing the incoming audio

Claims	Prior Art
wherein said presenting of the incoming audio message comprises playing the incoming audio message, with the playing of the incoming audio message being dependent on at least one user-selected speech characteristic.	message, with the playing of the incoming audio message being dependent on at least one user-selected speech characteristic.” <i>See above at claim 11.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
16. A wireless communication system as recited in claim 1, wherein the incoming text message is a SMS message.	Apple Prior Art Systems and Kennewick ‘209 expressly or inherently disclose “[a] wireless communication system as recited in claim 1, wherein the incoming text message is a SMS message.” iListen_15_New_speechrecogniti at Page 39 begin dictating into any application, such as Microsoft Word X or 2001. Using an 800MHz iMac G4 with 512MB of RAM, we saw good performance from the program—there was only a slight delay between speaking the words and seeing them appear on screen. Kennewick ’209 discloses using telephone connections to send and receive texting and integrated messaging, which at least suggests “SMS.” The following citations show examples of where Kennewick ‘209 discloses this limitation. Kennewick ’209 at 5:47–6:6 In addition to the ability to process user questions and commands, the invention may provide local or remote control functions for the system or for other devices. Users may initiate commands locally or remotely. Several techniques for executing remote operations are possible, for example, the use of a telephone or other audio connection. Other remote command techniques may be used. The processing of these commands is performed in a nearly identical manner to a question. One difference being that the result of the command is generally an action rather than a response. Examples of control operations supported by the invention include, finding and playing music, film, games or other entertainment media, control of the appliances or other electronic or electro-mechanical devices, security systems control, audio, text and video conferencing (including Voice over IP conferencing), operating integrated messaging environments and other control operations. By using an optional analog or digital telephone interface or a data network interface, the invention can provide remote control capability of the computer or and auxiliary devices connected to the computer. For example, if the computer is connected via an

Claims	Prior Art
	interface to one or more auxiliary devices, a remote user can use this capability to remotely control or program the operation of the auxiliary device. For example, this enables recording of a television program, checking the status of a home security system, checking voice mail messages, hearing and responding to email messages, and much more. A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick ‘209 to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim 1[pre]. Specifically, a POSITA would be motivated to implement Apple Prior Art Systems on an SMS enabled device such as the device with integrated messaging in Kennewick ‘209, and a POSITA would have a reasonable expectation of success since the Speech Synthesis Technologies and Speech Recognition Technologies are compatible with devices such as PDAs. To the extent that Apple Prior Art Systems and Kennewick ‘209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick ‘209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
17. A wireless communication system as recited in claim 1, wherein the incoming text message is a SMS message, and wherein the reply text message is a SMS message.	Apple Prior Art Systems and Kennewick ‘209 expressly or inherently disclose “[a] wireless communication system as recited in claim 1, wherein the incoming text message is a SMS message, and wherein the reply text message is a SMS message.” <i>See above at claim 16.</i> To the extent that Apple Prior Art Systems and Kennewick ‘209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick ‘209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
18[pre]. A wireless communication system as recited in claim 1, wherein said converting of the incoming text message to the incoming audio message comprises:	To the extent the preamble is limiting, Apple Prior Art Systems expressly or inherently disclose “[a] wireless communication system as recited in claim 1, wherein said converting of the incoming text message to the incoming audio message comprises.” <i>See above at claim [1d].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.

Claims	Prior Art
[18a] translating at least a portion of the incoming text message from a condensed format to an uncondensed text format; and	Apple Prior Art Systems expressly or inherently disclose “translating at least a portion of the incoming text message from a condensed format to an uncondensed text format.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, include vocabulary settings where users are able to set their own spellings and dictations for abbreviations and acronyms, which is an example of “translating ... from a condensed format to an uncondensed format.” The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM Announces ViaVoice 98 at 1 ViaVoice 98 Executive Edition is IBM’s most powerful continuous speech software and can be a productivity tool for writers, executives and experienced PC users. ViaVoice Executive offers direct dictation into most popular Windows applications, voice control for the desktop and applications and includes two topics, ViaVoice Computer and ViaVoice Business and Finance, which supplement the basic vocabulary. Topics are special interest vocabularies that add to the existing base vocabulary and increase accuracy. The executive edition is priced at \$149. ViaVoice 98 Office Edition is a continuous speech dictation and voice control software tool designed to help maximize the efficiency of a small office/home office user, faculty member or advanced student. Users can supplement their vocabulary with the included ViaVoice Business and Finance. The price for the Office edition is \$89. ViaVoice 98 Home Edition is IBM’s entry-level continuous speech dictation product and is ideal for the family to use to create letters, reports or e-mail at home. ViaVoice 98 Home includes the ViaVoice Cuisine Topic to let users produce original recipes or make a keepsake of family recipes. The cost is \$49.95. A customized, integrated version of IBM ViaVoice is included with Lotus SmartSuite Millennium edition. IBM Announces ViaVoice 98 at 2 Large active vocabulary - A 128,000 word active vocabulary includes 64,000 active base vocabulary and 64,000 more user-added words using Vocabulary Expander and Vocabulary Manager. Users can include multiple topics (each topic can contain up to 64,000 words) and a backup vocabulary of 260,000 words. Industry specific vocabularies - ViaVoice 98 supports industry-specific vocabularies. The first vocabularies will be ViaVoice Medical and ViaVoice Legal in U.S. English. These are specialized vocabularies with 25,000 or more words for use by medical and legal professionals. They cost \$149. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Enterprise Edition is IBM’s flagship speech recognition and text-to-speech product. Representing a dramatic shift in the evolution of embedded speech products, IBM Embedded

Claims	Prior Art
	ViaVoice Enterprise Edition has an unlimited active vocabulary, limited only by the amount of memory on devices, plus a dynamic and swappable 4,000 word phonetic flat list. This feature enables embedded applications to access voice-enabled databases with millions of records. A typical application would be a voice-enabled navigation database of street names for a GPS system in your car. IBM ViaVoice Simply Dictation user guide at Page 38 Correcting Text in SpeakPad  Correcting words as you dictate in SpeakPad updates your personal vocabulary, which helps improve recognition accuracy by immediately teaching ViaVoice new words and phrases. This way, if you say the word again while you are dictating the same document, ViaVoice is more likely to recognize it correctly. IBM ViaVoice Simply Dictation user guide at Page 49 Using Manage My Vocabulary The Manage My Vocabulary window allows you to manage your personal vocabulary by deleting or training how you say specific words. When you start Manage My Vocabulary, it shows you all of the words you have added to your personal vocabulary. Words shown in this window were added during dictation, correction, or when creating voice commands. 1. SAY> MANAGE MY VOCABULARY or, click the Manage My Vocabulary button from the VoiceCenter drawer. 2. Select the word to be re-trained. 3. Click the Train button. The Train Word window appears. 4. Click the Record button and speak the word. 5. Click OK when you have trained the word. If you want to check the quality of your recording, click the Play button.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 57 VoiceCenter Drawer The VoiceCenter drawer menu allows easy access to most of the ViaVoice options, functions, and information.  • What Can I Say - Displays a list of available voice commands. • Correction - Displays the Correction window. • Preferences - Use to customize functions of ViaVoice. • Macros - Displays the Macro Editor window so you can create and edit dictation macros. • Manage My Vocabulary - Displays the Manage My Vocabulary window so you can re-train the pronunciation of words. • UserSettings - Use to add Voice Models or change other User information. • Information Central - Displays Information Central, the ViaVoice online help file. • SetupAssistant - Use to set up your audio, create a user or analyze documents. • About ViaVoice - Displays the ViaVoice About box. IBM ViaVoice Simply Dictation user guide at Page 69 Recording a Pronunciation 1. If the word does not sound like it is written, type the word or phrase the way that you pronounce it in the <i>Sounds-like</i> spelling field. Remember to place a hyphen between the words and letters. For example, if you pronounce the word NCAA as NC double A, enter N-C-double-A. If you pronounce NCAA as NCAA, enter N-C-A-A. IBM ViaVoice Simply Dictation user guide at Page 71 Capitalize This - Capitalizes the selected text. Lowercase This - Lowercases the selected text. Uppercase This - Uppercases the selected text. Numeric - Changes numeric text to digits. Spell Out - Changes digits or abbreviations to words.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–36 <u>Including Pronunciation Dictionaries</u> No matter how sophisticated a speech synthesis system is, there will always be words that it does not automatically pronounce correctly. A clear instance of words that are often mispronounced is the class of proper nouns (names of people, place names, and so on). The Speech Manager supports pronunciation dictionaries which allow applications to override the default pronunciations of words. A pronunciation dictionary is a list of words along with their associated pronunciations stored in a resource of resource type 'dict'. The application is free to store dictionaries in either the resource fork or the data fork of a file. The application is responsible for loading the individual dictionaries into RAM and then passing a handle to the dictionary data to the Speech Manager. The initial release of the Speech Manager, however, does not include any routines that can add entries to dictionaries or manipulate them in other ways. The Speech Manager does include a routine, the <code>UseDictionary</code> function, that you can use to install one or more pronunciation dictionaries in a speech channel. Inside Macintosh Sound_ 1994 at Page 4–38 <u>Listing 4-10</u> A sample pronunciation dictionary resource <pre data-bbox="768 776 1520 1105"> resource 'dict' (1, "TestDict") { smRoman, langEnglish, verUS, ThisSecond, { pron, {tx, "ROOSEVELT", ph, "_1EHf_d1IY_1AAr"}, pron, {tx, "CHELSEA", ph, "_C1EHls2IY"}, pron, {tx, "AMHERST", ph, "_2UXmAKrst"}, pron, {tx, "REDSOX", ph, "_r1EHd_1AAks"}, pron, {tx, "HALLOWEEN", ph, "_h1AA12OW_w1IYn"}, pron, {tx, "FELIX", ph, "_f1IY12IHks_D2UX_k1AEt"}, pron, {tx, "WEDNESDAY", ph, "_m1IHd_w1IYk"}, }, }; </pre> Note that you are not restricted to using pronunciations similar to those of the words listed. Typically, however, pronunciation dictionaries contain entries for words that the Speech Manager pronounces unsatisfactorily.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ■ generate speech and then pause or stop it ■ obtain information about an individual voice or all voices ■ create and dispose of speech channels ■ obtain the Speech Manager's version and status ■ change the rate or pitch of speech Inside Macintosh Sound_1994 at Page 4–55 ■ convert textual into phonetic data ■ install a pronunciation dictionary into a speech channel With the exception of the <code>SpeechManagerVersion</code>, <code>SpeechBusy</code>, and <code>SpeechBusySystemWide</code> functions, all Speech Manager routines return a result code to indicate whether an error has occurred. The section “Application-Defined Routines” beginning on page 4-82 describes the syntax and operation of application-defined callback procedures. Inside Macintosh Sound_1994 at Page 4–80 <u>Installing a Pronunciation Dictionary</u> Pronunciation dictionaries allow your application to override the default Speech Manager pronunciations of individual words, such as names with quirky spellings. The <code>UseDictionary</code> function allows your application to load a pronunciation dictionary into a speech channel.


Claims	Prior Art
	iListen_15_New_speechrecogniti at Page 39 Most speech-recognition programs require a training period during which you acclimate the program to the unique qualities of your voice. With iListen, this process is easy. To begin, you set up the microphone, which in our case was the VX1 Parrott with Griffin iMic USB adapter (both included). Next, for about ten minutes, iListen leads you through sound-level tests and a sample passage of text. With initial training, iListen was about 78 percent accurate in our dictation tests in OS X and OS X's Classic mode. To improve recognition, you can read 11 other passages. Doing so brought the accuracy up slightly, to about 82 percent. The program can also analyze a set of your documents to learn the vocabulary you use. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[18b] converting at least the portion of the incoming text message in the uncondensed text format to at least a portion of the incoming audio message.	Apple Prior Art Systems expressly or inherently disclose “converting at least the portion of the incoming text message in the uncondensed text format to at least a portion of the incoming audio message.” <i>See</i> above at claim [18a]. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
19. A wireless communication system as recited in claim 18, wherein said converting of the incoming text message to the incoming audio message depends	Apple Prior Art Systems expressly or inherently disclose “[a] wireless communication system as recited in claim 18, wherein said converting of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic.”

Claims	Prior Art
on at least one user-selected speech characteristic.	Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, include settings (e.g., preferences) to set a voice for converting text to audio, which is an example of the claimed “user–selected speech characteristic.” The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM Announces ViaVoice 98 at 1 ViaVoice 98 Executive Edition is IBM’s most powerful continuous speech software and can be a productivity tool for writers, executives and experienced PC users. ViaVoice Executive offers direct dictation into most popular Windows applications, voice control for the desktop and applications and includes two topics, ViaVoice Computer and ViaVoice Business and Finance, which supplement the basic vocabulary. Topics are special interest vocabularies that add to the existing base vocabulary and increase accuracy. The executive edition is priced at \$149. ViaVoice 98 Office Edition is a continuous speech dictation and voice control software tool designed to help maximize the efficiency of a small office/home office user, faculty member or advanced student. Users can supplement their vocabulary with the included ViaVoice Business and Finance. The price for the Office edition is \$89. ViaVoice 98 Home Edition is IBM’s entry-level continuous speech dictation product and is ideal for the family to use to create letters, reports or e-mail at home. ViaVoice 98 Home includes the ViaVoice Cuisine Topic to let users produce original recipes or make a keepsake of family recipes. The cost is \$49.95. A customized, integrated version of IBM ViaVoice is included with Lotus SmartSuite Millennium edition. IBM Announces ViaVoice 98 at 2 Large active vocabulary - A 128,000 word active vocabulary includes 64,000 active base vocabulary and 64,000 more user-added words using Vocabulary Expander and Vocabulary Manager. Users can include multiple topics (each topic can contain up to 64,000 words) and a backup vocabulary of 260,000 words. Industry specific vocabularies - ViaVoice 98 supports industry-specific vocabularies. The first vocabularies will be ViaVoice Medical and ViaVoice Legal in U.S. English. These are specialized vocabularies with 25,000 or more words for use by medical and legal professionals. They cost \$149. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Enterprise Edition is IBM’s flagship speech recognition and text-to-speech product. Representing a dramatic shift in the evolution of embedded speech products, IBM Embedded ViaVoice Enterprise Edition has an unlimited active vocabulary, limited only by the amount of memory on devices, plus a dynamic and swappable 4,000 word phonetic flat list. This feature enables embedded applications to access voice-enabled databases with millions of records. A typical application would be a

Claims	Prior Art
	voice-enabled navigation database of street names for a GPS system in your car. IBM ViaVoice Simply Dictation user guide at Page 26 Reading Text Now that you have learned to enter and format text using voice commands, it's time to have ViaVoice read the sentences back to you. Highlight the desired text block and say one of the following voice commands. SAY> BEGIN READING SAY> READ SELECTION ViaVoice reads the highlighted text using the voice selected in the Voices panel in ViaVoice Preferences. On applications with copy and selection capabilities, the Read Selection menu item reads the text using the voice currently selected in the Voices panel in ViaVoice Preferences. IBM ViaVoice Simply Dictation user guide at Page 38 Correcting Text in SpeakPad Correcting words as you dictate in SpeakPad updates your personal vocabulary, which helps improve recognition accuracy by immediately teaching ViaVoice new words and phrases. This way, if you say the word again while you are dictating the same document, ViaVoice is more likely to recognize it correctly. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 48 Reading Your Text Provided your application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back text aloud from almost any source, including documents and e-mail messages. 1. Open the application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected system voice. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 49 Changing Voices  You can select the voice used when reading text. 1. Click the ViaVoice Preferences button on the VoiceCenter drawer. 2. Click the Voices icon. 3. From the Voice menu, select a new voice. 4. Press Play to hear a sample of your selection. 5. Click Apply, and close the window. 6. Any text read back from this point uses the voice you just selected. IBM ViaVoice Simply Dictation user guide at Page 49 Using Manage My Vocabulary The Manage My Vocabulary window allows you to manage your personal vocabulary by deleting or training how you say specific words. When you start Manage My Vocabulary, it shows you all of the words you have added to your personal vocabulary. Words shown in this window were added during dictation, correction, or when creating voice commands. 1. SAY> MANAGE MY VOCABULARY or, click the Manage My Vocabulary button from the VoiceCenter drawer. 2. Select the word to be re-trained. 3. Click the Train button. The Train Word window appears. 4. Click the Record button and speak the word. 5. Click OK when you have trained the word. If you want to check the quality of your recording, click the Play button.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 57 VoiceCenter Drawer The VoiceCenter drawer menu allows easy access to most of the ViaVoice options, functions, and information.  The diagram shows a central 3x3 grid of icons. Arrows point from text labels to specific icons: 'Correction' points to the top-center icon (checkmark), 'Preferences' points to the top-right icon (gears), 'UserSettings' points to the middle-right icon (person), 'About ViaVoice' points to the bottom-right icon (speech bubble), 'SetupAssistant' points to the bottom-center icon (wrench), 'Information Central' points to the bottom-left icon (question mark), 'Manage My Vocabulary' points to the middle-left icon (book), 'Macros' points to the top-left icon (document), and 'What Can I Say' points to the top-left icon (question mark). • What Can I Say - Displays a list of available voice commands. • Correction - Displays the Correction window. • Preferences - Use to customize functions of ViaVoice. • Macros - Displays the Macro Editor window so you can create and edit dictation macros. • Manage My Vocabulary - Displays the Manage My Vocabulary window so you can re-train the pronunciation of words. • UserSettings - Use to add Voice Models or change other User information. • Information Central - Displays Information Central, the ViaVoice online help file. • SetupAssistant - Use to set up your audio, create a user or analyze documents. • About ViaVoice - Displays the ViaVoice About box.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 61  SpeakPad Menu The SpeakPad menus contain items that are similar in functionality to other word processors. SpeakPad File Edit Format Dictation Window Help About SpeakPad Preferences... Dictation Menu The following items are available on the Dictation Menu. • Playback Selection/Stop Playback - This option allows you to play back your dictated text in your own voice. ViaVoice can play back up to 1,000 words (punctuation is included in the word count). The playback feature is useful in many instances, such as when you are correcting text. For more information, see "Playing Back Dictated Text in SpeakPad" on page 48. • Read Selection/Stop Reading - This option reads any text you have highlighted in a computer-generated voice. You can use this feature to help you proofread your documents. For more information, see "Reading Your Text" on page 48. • Voice Preferences - This option displays the Voices window, which allows you to change the computer-generated voice. For more information, see "Voices Preferences" on page 65. IBM ViaVoice Simply Dictation user guide at Page 65 Voices Preferences This page allows you to select the voice that reads your text and hear a sample played. From the Speech, Pitch, and Voice Quality tabs, adjust the sliders until you reach the desired settings. From the Options tab, select the hot key for reading text. IBM ViaVoice Simply Dictation user guide at Page 69 Recording a Pronunciation 1. If the word does not sound like it is written, type the word or phrase the way that you pronounce it in the <i>Sounds-like</i> spelling field. Remember to place a hyphen between the words and letters. For example, if you pronounce the word NCAA as NC double A, enter N-C-double-A. If you pronounce NCAA as NCAA, enter N-C-A-A.


Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 71 Capitalize This - Capitalizes the selected text. Lowercase This - Lowercases the selected text. Uppercase This - Uppercases the selected text. Numeric - Changes numeric text to digits. Spell Out - Changes digits or abbreviations to words. Inside Macintosh Sound_ 1994 at Page 4–3 This chapter begins with an introduction to the speech generation process and then discusses how you can ▪ check for the availability of the Speech Manager ▪ create and dispose of speech channels ▪ generate speech with different voices ▪ obtain information about and change speech channel settings ▪ start and stop speech production ▪ synchronize speech production with other activities by using callback procedures ▪ embed Speech Manager commands within text to make it more understandable Inside Macintosh Sound_ 1994 at Page 4–5 <u>Voices</u> Your application can use the system default voice to generate speech or it can specify that the Speech Manager use a particular voice that is available on the current computer system. A voice is a set of characteristics defined in parameters that specify a particular quality of speech. Just as different people's voices have different tonal qualities, so too can different voices have different qualities. A synthesized voice might sound male or female and might sound like an adult or a child. Some voices sound distinctively synthetic, while others sound more like real people. Figure 4-2 shows how the Speech Manager uses speech channels to synthesize speech with different voices.

Claims

Prior Art

Inside Macintosh Sound_1994 at Page 4-5

Figure 4-2 The Speech Manager and multiple voices



Inside Macintosh Sound_1994 at Page 4-6

Most of the time,

your application designates the voice that speaks text, and usually that is the default voice. Based on the needs of your users and the way in which you expect them to use voices in your application you can provide access to voices in a number of different ways. You could include access to selecting voices in a dialog box that is available from a menu item such as Voices... Any application that allows users to choose among voices requires additional information about the available voices beyond the information provided by a voice specification record (described in detail on page 4-46), whose data should never be presented to the user. Such additional information might include the name of the voice as well as what script and language it supports.

Inside Macintosh Sound_1994 at Page 4-6

Applications can use the `GetVoiceDescription` function (described in detail on page 4-66) with a voice specification record to obtain such information in a voice description record (described in detail on page 4-47). You might provide access to voices through a control panel. For information about implementing control panels, see *Inside Macintosh: More Macintosh Toolbox*. Or, you could implement a voices menu in your application's main menu bar, if you think that users will want to change the voice often and you have the room available. It's not a good idea to implement a hierarchical Voices menu since hierarchical menus are harder to use. For more information about choosing a user interface for your application, see *Macintosh Human Interface Guidelines*.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–6 Speech Attributes Any given person has only one voice, but can alter the characteristics of his or her speech in a number of different ways. For example, a person can speak slowly or quickly and with a low or a high pitch. Similarly, the Speech Manager provides routines that allow you to modify these and other speech attributes, regardless of which voice is in use. A speech attribute is a setting defined for a class of voices or for all voices that affects the quality of speech produced by the Speech Manager. The Speech Manager provides routines to directly alter two speech attributes—speech rate and speech pitch. These routines are described in the section “Changing Speech Attributes” beginning on page 4-73. You can change two other speech attributes—pitch modulation and speech volume—by using the mechanism of speech information selectors, which is described in the section “Speech Information Selectors” beginning on page 4-39. The speech rate of a speech channel is the approximate number of words of text that the synthesizer should say in one minute. Slower speech rates make the speech easier to understand, but can be annoyingly tedious to listen to. Some applications, such as aids for the visually impaired, require very fast speech rates. Speech rates are expressed as fixed-point values. Each speech synthesizer determines its own range of speech rates. The speech pitch of a speech channel represents the middle pitch of the voice, roughly corresponding to the key in which a song is played. It is a fixed-point value in the range of 0.000 through 127.000, where 60.000 corresponds to middle C on a conventional piano. Each 1.000-unit change in a value corresponds to a musical half step. This is the same scale used in specifying MIDI note values, as described in the chapter “Sound Manager” Inside Macintosh Sound_ 1994 at Page 4–7 You can determine the current speech pitch on a speech channel by calling the GetSpeechPitch function, described on page 4-75. You can change the current pitch by calling the SetSpeechPitch function, described on page 4-76. You can also determine the current speech rate and change it by using the GetSpeechRate function, described on page 4-73 and the SetSpeechRate function, described on page 4-74. Changes in speech pitch and speech rate are effective immediately (as soon as the synthesizer can respond), even if they occur in the middle of a word. Inside Macintosh Sound_ 1994 at Page 4–9 You can obtain the pitch modulation by using the GetSpeechInfo function with the soPitchMod speech information selector, and you can change the pitch modulation by using the SetSpeechInfo function with the same selector. Speech information selectors are described in “Speech Information Selectors” beginning on page 4-39.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–11 To take advantage of most of the Speech Manager's features, you must allocate a speech channel to pass to Speech Manager functions and dispose of the speech channel when you are finished using it. "Creating, Using, and Disposing of a Speech Channel" beginning on page 4-13 demonstrates how you do this and shows how you can use the <code>SpeakText</code> function to start speech generation from a buffer of text. Some applications permit users to choose a voice from those available to be used for speech generation. The <code>CountVoices</code>, <code>GetIndVoice</code>, and <code>GetVoiceDescription</code> functions support this capability. "Working With Different Voices" beginning on page 4-14 shows how you can use these functions to choose among available voices. Inside Macintosh Sound_1994 at Page 4–14 <u>Working With Different Voices</u> When you work with speech channels, you can set a voice for a particular channel. When you set a voice, you may want to filter out certain of its characteristics in order to identify the one you want. For example, in an educational software application for elementary school students, you may want to use only children's voices. In order to choose the voice you want, you get a voice description record that contains information about a voice such as the size of the voice, the name of the voice, the age and gender of the voice, and the synthesizer with which it works. You can get the number of available voices using the <code>CountVoices</code> function. You can cycle through the available voices and identify the one you want to use by using the <code>GetIndVoice</code> function. Then you fill out a voice description record using the <code>GetVoiceDescription</code> function. Listing 4-3 shows how to get identifying information about a voice. Inside Macintosh Sound_1994 at Page 4–16 <u>Adjusting Speech Attributes</u> Speech attributes are settings defined for a class of voices or for all voices that affect the quality of speech produced by the Speech Manager. In general, an application should not try to second-guess the developers of a voice or synthesizer by arbitrarily setting a speech attribute. However, there are some cases in which you would want to adjust the rate of speech (how many words per minute are spoken) or the speech pitch (the listener's subjective interpretation of speech's average frequency). Listing 4-4 shows how to adjust the speech pitch and speech rate of a particular channel.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–23 <u>Writing Embedded Speech Commands</u> Embedded speech commands allow you to customize the quality of speech output by fine tuning it. You can make speech much easier to understand than the default way in which text is spoken by a synthesizer. An embedded speech command is a command embedded within a text buffer to be spoken by the Speech Manager that causes the Speech Manager to take a certain action. For example, you could use an embedded speech command to emphasize a particular word in a text string to make it stand out to the user. Inside Macintosh Sound_ 1994 at Page 4–30 <u>Examples of Embedded Speech Commands</u> If you use just a few of the embedded speech commands, you can markedly increase the understandability of text spoken by your application. Your application knows more about the speech being produced than a speech synthesizer does. A synthesizer speaks text according to a predetermined set of rules about language production. Therefore, the voices available on a Macintosh computer with the Speech Manager installed sound very synthetic and sometimes robotic because the pronunciation rules are formalized. You can make the speech produced by the synthesizer sound a lot more human by observing some simple rules of human speech and embedding speech commands in text according to these conventions. The techniques presented in this section could be applied when your application is having a dialog with the user or speaking some error messages or announcements. Inside Macintosh Sound_ 1994 at Page 4–36 <u>Including Pronunciation Dictionaries</u> No matter how sophisticated a speech synthesis system is, there will always be words that it does not automatically pronounce correctly. A clear instance of words that are often mispronounced is the class of proper nouns (names of people, place names, and so on). The Speech Manager supports pronunciation dictionaries which allow applications to override the default pronunciations of words. A pronunciation dictionary is a list of words along with their associated pronunciations stored in a resource of resource type 'dict'. The application is free to store dictionaries in either the resource fork or the data fork of a file. The application is responsible for loading the individual dictionaries into RAM and then passing a handle to the dictionary data to the Speech Manager. The initial release of the Speech Manager, however, does not include any routines that can add entries to dictionaries or manipulate them in other ways. The Speech Manager does include a routine, the <code>UseDictionary</code> function, that you can use to install one or more pronunciation dictionaries in a speech channel.

Inside Macintosh Sound_1994 at Page 4-38

Listing 4-10 A sample pronunciation dictionary resource

```
resource 'dict' (1, "TestDict") {
    smRoman, langEnglish, verUS, ThisSecond,
    {
        pron, {tx, "ROOSEVELT",    ph, "_1EHf_d1IY_1AAr"},
        pron, {tx, "CHELSEA",      ph, "_C1EHls2IY"},
        pron, {tx, "AMHERST",      ph, "_2UXmAKrst"},
        pron, {tx, "REDSOX",       ph, "_r1EHd_s1AAks"},
        pron, {tx, "HALLOWEEN",    ph, "_h1AA12OW_w1IYn"},
        pron, {tx, "FELIX",        ph, "_f1IYl2IHks_D2UX_k1AEt"},
        pron, {tx, "WEDNESDAY",    ph, "_m1IHd_w1IYk"},
    },
};
```

Note that you are not restricted to using pronunciations similar to those of the words listed. Typically, however, pronunciation dictionaries contain entries for words that the Speech Manager pronounces unsatisfactorily.

Inside Macintosh Sound_1994 at Page 4-47

Voice Description Records

By calling the GetVoiceDescription function, you can obtain information about a voice in a voice description record. The VoiceDescription data type defines a voice description record.

```
TYPE VoiceDescription =
RECORD
    length:      LongInt;      {size of record}
    voice:       VoiceSpec;    {voice synthesizer and ID info}
    version:     LongInt;      {version number of voice}
    name:        Str63;        {name of voice}
    comment:     Str255;       {text information about voice}
    gender:      Integer;      {neuter, male, or female}
    age:         Integer;      {approximate age in years}
    script:      Integer;      {script code of text voice can process}
    language:    Integer;      {language code of voice output}
    region:      Integer;      {region code of voice output}
    reserved1:   LongInt;      {always 0--reserved for future use}
    reserved2:   LongInt;      {always 0--reserved for future use}
    reserved3:   LongInt;      {always 0--reserved for future use}
    reserved4:   LongInt;      {always 0--reserved for future use}
END;
```

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ■ generate speech and then pause or stop it ■ obtain information about an individual voice or all voices ■ create and dispose of speech channels ■ obtain the Speech Manager's version and status ■ change the rate or pitch of speech Inside Macintosh Sound_1994 at Page 4–55 ■ convert textual into phonetic data ■ install a pronunciation dictionary into a speech channel With the exception of the <code>SpeechManagerVersion</code>, <code>SpeechBusy</code>, and <code>SpeechBusySystemWide</code> functions, all Speech Manager routines return a result code to indicate whether an error has occurred. The section “Application-Defined Routines” beginning on page 4-82 describes the syntax and operation of application-defined callback procedures. Inside Macintosh Sound_1994 at Page 4–73 <u>Changing Speech Attributes</u> To determine the rate and pitch at which a speech channel is processing text, you can use the <code>GetSpeechRate</code> and <code>GetSpeechPitch</code> functions. The <code>SetSpeechRate</code> and <code>SetSpeechPitch</code> functions allow you to change rate and pitch. The most robust of the Speech Manager's routines are the <code>GetSpeechInfo</code> and <code>SetSpeechInfo</code> functions. These allow you to obtain many types of information about a speech channel and to change many settings of a speech channel. To specify the operation that you wish to perform, you must pass <code>GetSpeechInfo</code> or <code>SetSpeechInfo</code> a selector. A full list of selectors is provided in “Speech Information Selectors” beginning on page 4-39.


Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–76 SetSpeechPitch You can use the <code>SetSpeechPitch</code> function to set the speech pitch on a designated speech channel. <pre data-bbox="770 370 1472 418">FUNCTION SetSpeechPitch (chan: SpeechChannel; pitch: Fixed) : OSErr;</pre> chan The speech channel whose pitch you wish to set. pitch The new pitch for the speech channel, expressed as a fixed-point frequency value. Inside Macintosh Sound_1994 at Page 4–80 Installing a Pronunciation Dictionary Pronunciation dictionaries allow your application to override the default Speech Manager pronunciations of individual words, such as names with quirky spellings. The <code>UseDictionary</code> function allows your application to load a pronunciation dictionary into a speech channel. iListen_15_New_speechrecogniti at Page 39 Most speech-recognition programs require a training period during which you acclimate the program to the unique qualities of your voice. With iListen, this process is easy. To begin, you set up the microphone, which in our case was the VXI Parrott with Griffin iMic USB adapter (both included). Next, for about ten minutes, iListen leads you through sound-level tests and a sample passage of text. With initial training, iListen was about 78 percent accurate in our dictation tests in OS X and OS X's Classic mode. To improve recognition, you can read 11 other passages. Doing so brought the accuracy up slightly, to about 82 percent. The program can also analyze a set of your documents to learn the vocabulary you use. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior

Claims	Prior Art
	Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
20. A wireless communication system as recited in claim 19, wherein the at least one user-selected speech characteristic is set based on at least one user setting provided at said wireless communication device.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless communication system as recited in claim 19, wherein the at least one user-selected speech characteristic is set based on at least one user setting provided at said wireless communication device.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, includes settings (e.g., preferences) that can be set at the computer or PDA (e.g., the “wireless communication device”) to set a voice for converting text to audio, which is an example of the claimed “user–selected speech characteristic.” The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM Announces ViaVoice 98 at 1 ViaVoice 98 Executive Edition is IBM’s most powerful continuous speech software and can be a productivity tool for writers, executives and experienced PC users. ViaVoice Executive offers direct dictation into most popular Windows applications, voice control for the desktop and applications and includes two topics, ViaVoice Computer and ViaVoice Business and Finance, which supplement the basic vocabulary. Topics are special interest vocabularies that add to the existing base vocabulary and increase accuracy. The executive edition is priced at \$149. ViaVoice 98 Office Edition is a continuous speech dictation and voice control software tool designed to help maximize the efficiency of a small office/home office user, faculty member or advanced student. Users can supplement their vocabulary with the included ViaVoice Business and Finance. The price for the Office edition is \$89. ViaVoice 98 Home Edition is IBM’s entry-level continuous speech dictation product and is ideal for the family to use to create letters, reports or e-mail at home. ViaVoice 98 Home includes the ViaVoice Cuisine Topic to let users produce original recipes or make a keepsake of family recipes. The cost is \$49.95. A customized, integrated version of IBM ViaVoice is included with Lotus SmartSuite Millennium edition. IBM Announces ViaVoice 98 at 2 Large active vocabulary - A 128,000 word active vocabulary includes 64,000 active base vocabulary and 64,000 more user-added words using Vocabulary Expander and Vocabulary Manager. Users can include multiple topics (each topic can contain up to 64,000 words) and a backup vocabulary of 260,000 words. Industry specific vocabularies - ViaVoice 98 supports industry-specific vocabularies. The first vocabularies will be ViaVoice Medical and ViaVoice Legal in U.S. English. These are specialized vocabularies with 25,000 or more words for use by medical and legal professionals. They cost \$149.

Claims	Prior Art
	IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Enterprise Edition is IBM's flagship speech recognition and text-to-speech product. Representing a dramatic shift in the evolution of embedded speech products, IBM Embedded ViaVoice Enterprise Edition has an unlimited active vocabulary, limited only by the amount of memory on devices, plus a dynamic and swappable 4,000 word phonetic flat list. This feature enables embedded applications to access voice-enabled databases with millions of records. A typical application would be a voice-enabled navigation database of street names for a GPS system in your car. IBM ViaVoice Simply Dictation user guide at Page 26 Reading Text Now that you have learned to enter and format text using voice commands, it's time to have ViaVoice read the sentences back to you. Highlight the desired text block and say one of the following voice commands. SAY> BEGIN READING SAY> READ SELECTION ViaVoice reads the highlighted text using the voice selected in the Voices panel in ViaVoice Preferences. On applications with copy and selection capabilities, the Read Selection menu item reads the text using the voice currently selected in the Voices panel in ViaVoice Preferences. IBM ViaVoice Simply Dictation user guide at Page 38 Correcting Text in SpeakPad  Correcting words as you dictate in SpeakPad updates your personal vocabulary, which helps improve recognition accuracy by immediately teaching ViaVoice new words and phrases. This way, if you say the word again while you are dictating the same document, ViaVoice is more likely to recognize it correctly. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 48 Reading Your Text Provided your application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back text aloud from almost any source, including documents and e-mail messages. 1. Open the application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected system voice. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 49 Changing Voices  You can select the voice used when reading text. 1. Click the ViaVoice Preferences button on the VoiceCenter drawer. 2. Click the Voices icon. 3. From the Voice menu, select a new voice. 4. Press Play to hear a sample of your selection. 5. Click Apply, and close the window. 6. Any text read back from this point uses the voice you just selected. IBM ViaVoice Simply Dictation user guide at Page 49 Using Manage My Vocabulary The Manage My Vocabulary window allows you to manage your personal vocabulary by deleting or training how you say specific words. When you start Manage My Vocabulary, it shows you all of the words you have added to your personal vocabulary. Words shown in this window were added during dictation, correction, or when creating voice commands. 1. SAY> MANAGE MY VOCABULARY or, click the Manage My Vocabulary button from the VoiceCenter drawer. 2. Select the word to be re-trained. 3. Click the Train button. The Train Word window appears. 4. Click the Record button and speak the word. 5. Click OK when you have trained the word. If you want to check the quality of your recording, click the Play button.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 57 VoiceCenter Drawer The VoiceCenter drawer menu allows easy access to most of the ViaVoice options, functions, and information.  The diagram shows a central 3x3 grid of icons. Arrows point from text labels to specific icons: 'Correction' points to the top-middle icon (checkmark), 'Preferences' points to the top-right icon (gears), 'UserSettings' points to the middle-right icon (person), 'About ViaVoice' points to the bottom-right icon (document with 'w'), 'SetupAssistant' points to the bottom-middle icon (megaphone), 'Information Central' points to the bottom-left icon (question mark), 'Manage My Vocabulary' points to the middle-left icon (document with 'v'), 'Macros' points to the top-left icon (document with 'm'), and 'What Can I Say' points to the top-left icon (question mark). The entire grid is enclosed in a blue square border. • What Can I Say - Displays a list of available voice commands. • Correction - Displays the Correction window. • Preferences - Use to customize functions of ViaVoice. • Macros - Displays the Macro Editor window so you can create and edit dictation macros. • Manage My Vocabulary - Displays the Manage My Vocabulary window so you can re-train the pronunciation of words. • UserSettings - Use to add Voice Models or change other User information. • Information Central - Displays Information Central, the ViaVoice online help file. • SetupAssistant - Use to set up your audio, create a user or analyze documents. • About ViaVoice - Displays the ViaVoice About box.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 61  SpeakPad Menu The SpeakPad menus contain items that are similar in functionality to other word processors. SpeakPad File Edit Format Dictation Window Help About SpeakPad Preferences... Dictation Menu The following items are available on the Dictation Menu. • Playback Selection/Stop Playback - This option allows you to play back your dictated text in your own voice. ViaVoice can play back up to 1,000 words (punctuation is included in the word count). The playback feature is useful in many instances, such as when you are correcting text. For more information, see "Playing Back Dictated Text in SpeakPad" on page 48. • Read Selection/Stop Reading - This option reads any text you have highlighted in a computer-generated voice. You can use this feature to help you proofread your documents. For more information, see "Reading Your Text" on page 48. • Voice Preferences - This option displays the Voices window, which allows you to change the computer-generated voice. For more information, see "Voices Preferences" on page 65. IBM ViaVoice Simply Dictation user guide at Page 65 Voices Preferences This page allows you to select the voice that reads your text and hear a sample played. From the Speech, Pitch, and Voice Quality tabs, adjust the sliders until you reach the desired settings. From the Options tab, select the hot key for reading text. IBM ViaVoice Simply Dictation user guide at Page 69 Recording a Pronunciation 1. If the word does not sound like it is written, type the word or phrase the way that you pronounce it in the <i>Sounds-like</i> spelling field. Remember to place a hyphen between the words and letters. For example, if you pronounce the word NCAA as NC double A, enter N-C-double-A. If you pronounce NCAA as NCAA, enter N-C-A-A.


Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 71 Capitalize This - Capitalizes the selected text. Lowercase This - Lowercases the selected text. Uppercase This - Uppercases the selected text. Numeric - Changes numeric text to digits. Spell Out - Changes digits or abbreviations to words. Inside Macintosh Sound_ 1994 at Page 4–3 This chapter begins with an introduction to the speech generation process and then discusses how you can ▪ check for the availability of the Speech Manager ▪ create and dispose of speech channels ▪ generate speech with different voices ▪ obtain information about and change speech channel settings ▪ start and stop speech production ▪ synchronize speech production with other activities by using callback procedures ▪ embed Speech Manager commands within text to make it more understandable Inside Macintosh Sound_ 1994 at Page 4–5 <u>Voices</u> Your application can use the system default voice to generate speech or it can specify that the Speech Manager use a particular voice that is available on the current computer system. A voice is a set of characteristics defined in parameters that specify a particular quality of speech. Just as different people's voices have different tonal qualities, so too can different voices have different qualities. A synthesized voice might sound male or female and might sound like an adult or a child. Some voices sound distinctively synthetic, while others sound more like real people. Figure 4-2 shows how the Speech Manager uses speech channels to synthesize speech with different voices.

Claims

Prior Art

Inside Macintosh Sound_1994 at Page 4-5

Figure 4-2 The Speech Manager and multiple voices



Inside Macintosh Sound_1994 at Page 4-6

Most of the time,

your application designates the voice that speaks text, and usually that is the default voice. Based on the needs of your users and the way in which you expect them to use voices in your application you can provide access to voices in a number of different ways. You could include access to selecting voices in a dialog box that is available from a menu item such as Voices... Any application that allows users to choose among voices requires additional information about the available voices beyond the information provided by a voice specification record (described in detail on page 4-46), whose data should never be presented to the user. Such additional information might include the name of the voice as well as what script and language it supports.

Inside Macintosh Sound_1994 at Page 4-6

Applications can use the `GetVoiceDescription` function (described in detail on page 4-66) with a voice specification record to obtain such information in a voice description record (described in detail on page 4-47). You might provide access to voices through a control panel. For information about implementing control panels, see *Inside Macintosh: More Macintosh Toolbox*. Or, you could implement a voices menu in your application's main menu bar, if you think that users will want to change the voice often and you have the room available. It's not a good idea to implement a hierarchical Voices menu since hierarchical menus are harder to use. For more information about choosing a user interface for your application, see *Macintosh Human Interface Guidelines*.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–6 <u>Speech Attributes</u> Any given person has only one voice, but can alter the characteristics of his or her speech in a number of different ways. For example, a person can speak slowly or quickly and with a low or a high pitch. Similarly, the Speech Manager provides routines that allow you to modify these and other speech attributes, regardless of which voice is in use. A speech attribute is a setting defined for a class of voices or for all voices that affects the quality of speech produced by the Speech Manager. The Speech Manager provides routines to directly alter two speech attributes—speech rate and speech pitch. These routines are described in the section “Changing Speech Attributes” beginning on page 4-73. You can change two other speech attributes—pitch modulation and speech volume—by using the mechanism of speech information selectors, which is described in the section “Speech Information Selectors” beginning on page 4-39. The speech rate of a speech channel is the approximate number of words of text that the synthesizer should say in one minute. Slower speech rates make the speech easier to understand, but can be annoyingly tedious to listen to. Some applications, such as aids for the visually impaired, require very fast speech rates. Speech rates are expressed as fixed-point values. Each speech synthesizer determines its own range of speech rates. The speech pitch of a speech channel represents the middle pitch of the voice, roughly corresponding to the key in which a song is played. It is a fixed-point value in the range of 0.000 through 127.000, where 60.000 corresponds to middle C on a conventional piano. Each 1.000-unit change in a value corresponds to a musical half step. This is the same scale used in specifying MIDI note values, as described in the chapter “Sound Manager” Inside Macintosh Sound_ 1994 at Page 4–7 You can determine the current speech pitch on a speech channel by calling the <code>GetSpeechPitch</code> function, described on page 4-75. You can change the current pitch by calling the <code>SetSpeechPitch</code> function, described on page 4-76. You can also determine the current speech rate and change it by using the <code>GetSpeechRate</code> function, described on page 4-73 and the <code>SetSpeechRate</code> function, described on page 4-74. Changes in speech pitch and speech rate are effective immediately (as soon as the synthesizer can respond), even if they occur in the middle of a word. Inside Macintosh Sound_ 1994 at Page 4–9 You can obtain the pitch modulation by using the <code>GetSpeechInfo</code> function with the <code>soPitchMod</code> speech information selector, and you can change the pitch modulation by using the <code>SetSpeechInfo</code> function with the same selector. Speech information selectors are described in “Speech Information Selectors” beginning on page 4-39.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–11 To take advantage of most of the Speech Manager's features, you must allocate a speech channel to pass to Speech Manager functions and dispose of the speech channel when you are finished using it. "Creating, Using, and Disposing of a Speech Channel" beginning on page 4-13 demonstrates how you do this and shows how you can use the <code>SpeakText</code> function to start speech generation from a buffer of text. Some applications permit users to choose a voice from those available to be used for speech generation. The <code>CountVoices</code>, <code>GetIndVoice</code>, and <code>GetVoiceDescription</code> functions support this capability. "Working With Different Voices" beginning on page 4-14 shows how you can use these functions to choose among available voices. Inside Macintosh Sound_ 1994 at Page 4–14 <u>Working With Different Voices</u> When you work with speech channels, you can set a voice for a particular channel. When you set a voice, you may want to filter out certain of its characteristics in order to identify the one you want. For example, in an educational software application for elementary school students, you may want to use only children's voices. In order to choose the voice you want, you get a voice description record that contains information about a voice such as the size of the voice, the name of the voice, the age and gender of the voice, and the synthesizer with which it works. You can get the number of available voices using the <code>CountVoices</code> function. You can cycle through the available voices and identify the one you want to use by using the <code>GetIndVoice</code> function. Then you fill out a voice description record using the <code>GetVoiceDescription</code> function. Listing 4-3 shows how to get identifying information about a voice. Inside Macintosh Sound_ 1994 at Page 4–16 <u>Adjusting Speech Attributes</u> Speech attributes are settings defined for a class of voices or for all voices that affect the quality of speech produced by the Speech Manager. In general, an application should not try to second-guess the developers of a voice or synthesizer by arbitrarily setting a speech attribute. However, there are some cases in which you would want to adjust the rate of speech (how many words per minute are spoken) or the speech pitch (the listener's subjective interpretation of speech's average frequency). Listing 4-4 shows how to adjust the speech pitch and speech rate of a particular channel.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–23 <u>Writing Embedded Speech Commands</u> Embedded speech commands allow you to customize the quality of speech output by fine tuning it. You can make speech much easier to understand than the default way in which text is spoken by a synthesizer. An embedded speech command is a command embedded within a text buffer to be spoken by the Speech Manager that causes the Speech Manager to take a certain action. For example, you could use an embedded speech command to emphasize a particular word in a text string to make it stand out to the user. Inside Macintosh Sound_ 1994 at Page 4–30 <u>Examples of Embedded Speech Commands</u> If you use just a few of the embedded speech commands, you can markedly increase the understandability of text spoken by your application. Your application knows more about the speech being produced than a speech synthesizer does. A synthesizer speaks text according to a predetermined set of rules about language production. Therefore, the voices available on a Macintosh computer with the Speech Manager installed sound very synthetic and sometimes robotic because the pronunciation rules are formalized. You can make the speech produced by the synthesizer sound a lot more human by observing some simple rules of human speech and embedding speech commands in text according to these conventions. The techniques presented in this section could be applied when your application is having a dialog with the user or speaking some error messages or announcements. Inside Macintosh Sound_ 1994 at Page 4–36 <u>Including Pronunciation Dictionaries</u> No matter how sophisticated a speech synthesis system is, there will always be words that it does not automatically pronounce correctly. A clear instance of words that are often mispronounced is the class of proper nouns (names of people, place names, and so on). The Speech Manager supports pronunciation dictionaries which allow applications to override the default pronunciations of words. A pronunciation dictionary is a list of words along with their associated pronunciations stored in a resource of resource type 'dict'. The application is free to store dictionaries in either the resource fork or the data fork of a file. The application is responsible for loading the individual dictionaries into RAM and then passing a handle to the dictionary data to the Speech Manager. The initial release of the Speech Manager, however, does not include any routines that can add entries to dictionaries or manipulate them in other ways. The Speech Manager does include a routine, the <code>UseDictionary</code> function, that you can use to install one or more pronunciation dictionaries in a speech channel.

Inside Macintosh Sound_1994 at Page 4-38

Listing 4-10 A sample pronunciation dictionary resource

```
resource 'dict' (1, "TestDict") {
    smRoman, langEnglish, verUS, ThisSecond,
    {
        pron, {tx, "ROOSEVELT",    ph, "_1EHf_d1IY_1AAr"},
        pron, {tx, "CHELSEA",      ph, "_C1EHls2IY"},
        pron, {tx, "AMHERST",      ph, "_2UXmAKrst"},
        pron, {tx, "REDSOX",       ph, "_r1EHd_s1AAks"},
        pron, {tx, "HALLOWEEN",    ph, "_h1AA12OW_w1IYn"},
        pron, {tx, "FELIX",        ph, "_f1IYl2IHks_D2UX_k1AEt"},
        pron, {tx, "WEDNESDAY",    ph, "_m1IHd_w1IYk"},
    },
};
```

Note that you are not restricted to using pronunciations similar to those of the words listed. Typically, however, pronunciation dictionaries contain entries for words that the Speech Manager pronounces unsatisfactorily.

Inside Macintosh Sound_1994 at Page 4-47

Voice Description Records

By calling the GetVoiceDescription function, you can obtain information about a voice in a voice description record. The VoiceDescription data type defines a voice description record.

```
TYPE VoiceDescription =
RECORD
    length:      LongInt;      {size of record}
    voice:       VoiceSpec;    {voice synthesizer and ID info}
    version:     LongInt;      {version number of voice}
    name:        Str63;        {name of voice}
    comment:     Str255;       {text information about voice}
    gender:      Integer;      {neuter, male, or female}
    age:         Integer;      {approximate age in years}
    script:      Integer;      {script code of text voice can process}
    language:    Integer;      {language code of voice output}
    region:      Integer;      {region code of voice output}
    reserved1:   LongInt;      {always 0--reserved for future use}
    reserved2:   LongInt;      {always 0--reserved for future use}
    reserved3:   LongInt;      {always 0--reserved for future use}
    reserved4:   LongInt;      {always 0--reserved for future use}
END;
```

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ■ generate speech and then pause or stop it ■ obtain information about an individual voice or all voices ■ create and dispose of speech channels ■ obtain the Speech Manager's version and status ■ change the rate or pitch of speech Inside Macintosh Sound_1994 at Page 4–55 ■ convert textual into phonetic data ■ install a pronunciation dictionary into a speech channel With the exception of the <code>SpeechManagerVersion</code>, <code>SpeechBusy</code>, and <code>SpeechBusySystemWide</code> functions, all Speech Manager routines return a result code to indicate whether an error has occurred. The section “Application-Defined Routines” beginning on page 4-82 describes the syntax and operation of application-defined callback procedures. Inside Macintosh Sound_1994 at Page 4–73 <u>Changing Speech Attributes</u> To determine the rate and pitch at which a speech channel is processing text, you can use the <code>GetSpeechRate</code> and <code>GetSpeechPitch</code> functions. The <code>SetSpeechRate</code> and <code>SetSpeechPitch</code> functions allow you to change rate and pitch. The most robust of the Speech Manager's routines are the <code>GetSpeechInfo</code> and <code>SetSpeechInfo</code> functions. These allow you to obtain many types of information about a speech channel and to change many settings of a speech channel. To specify the operation that you wish to perform, you must pass <code>GetSpeechInfo</code> or <code>SetSpeechInfo</code> a selector. A full list of selectors is provided in “Speech Information Selectors” beginning on page 4-39.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–76 <u>SetSpeechPitch</u> You can use the <code>SetSpeechPitch</code> function to set the speech pitch on a designated speech channel. <pre>FUNCTION SetSpeechPitch (chan: SpeechChannel; pitch: Fixed) : OSErr;</pre> <code>chan</code> The speech channel whose pitch you wish to set. <code>pitch</code> The new pitch for the speech channel, expressed as a fixed-point frequency value. Inside Macintosh Sound_1994 at Page 4–80 <u>Installing a Pronunciation Dictionary</u> Pronunciation dictionaries allow your application to override the default Speech Manager pronunciations of individual words, such as names with quirky spellings. The <code>UseDictionary</code> function allows your application to load a pronunciation dictionary into a speech channel. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
21. A wireless communication system as recited in claim 1, wherein at least a portion of the incoming text message is in a chat language, and wherein said converting of the incoming text message to the incoming audio message comprises expanding the portion of the incoming text message that is in the chat language to a plurality of words.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless communication system as recited in claim 1, wherein at least a portion of the incoming text message is in a chat language, and wherein said converting of the incoming text message to the incoming audio message comprises expanding the portion of the incoming text message that is in the chat language to a plurality of words.” <i>See above at claim [18a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.


Claims	Prior Art
22. A wireless communication system as recited in claim 1, wherein said wireless communication device is a mobile telephone.	Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and Jabra FreeSpeak expressly or inherently disclose “[a] wireless communication system as recited in claim 1, wherein said wireless communication device is a mobile telephone.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, are compatible with PDAs, which at least suggests compatibility with a mobile phone. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM ViaVoice Mobility Suite v.4 at 1 Operate PDA functions using your voice Imagine being able to operate many of the common PDA functions like e-mail, tasks, calendar, phone numbers and addresses, by using voice commands instead of a stylus. IBM Embedded ViaVoice Mobility Suite software is now available to do just that! IBM Embedded ViaVoice Mobility Suite voice-enables many of the most common PIM applications for the handheld PDA, and is now available on the Compaq H3800 Pocket PC. IBM ViaVoice Mobility Suite v.4 at 1 End-to-end voice solutions IBM can offer distributed end-to-end voice solutions where some of the speech processing resides on the client device (cell phones etc.) and integrates with additional voice functionality at the server level and large-scale telephony/unified messaging systems. IBM Mobility Suite Introduction at 3 Welcome to IBM Embedded ViaVoice Mobility Suite, a productivity application that controls your Personal Digital Assistant (PDA) by voice. In response to your commands, IBM Embedded ViaVoice Mobility Suite will help manage your appointments, contacts, tasks, and mail. IBM Mobility Suite Introduction at 5 • To hear important information, for example, say Read mail or Read my very next appointment. • To hear prompts, messages, or requested information again, say Repeat that. To the extent Apple Prior Art Systems do not explicitly disclose a “cell phone,” Kennewick ’209 discloses that the main unit may be a wireless telephone, which is an example of the claimed “mobile telephone.” The following citations show examples of where Kennewick ’209 discloses this limitation. Kennewick ’209 at 5:47–6:6 In addition to the ability to process user questions and commands, the invention may provide local or

Claims	Prior Art
	remote control functions for the system or for other devices. Users may initiate commands locally or remotely. Several techniques for executing remote operations are possible, for example, the use of a telephone or other audio connection. Other remote command techniques may be used. The processing of these commands is performed in a nearly identical manner to a question. One difference being that the result of the command is generally an action rather than a response. Examples of control operations supported by the invention include, finding and playing music, film, games or other entertainment media, control of the appliances or other electronic or electro-mechanical devices, security systems control, audio, text and video conferencing (including Voice over IP conferencing), operating integrated messaging environments and other control operations. By using an optional analog or digital telephone interface or a data network interface, the invention can provide remote control capability of the computer or and auxiliary devices connected to the computer. For example, if the computer is connected via an interface to one or more auxiliary devices, a remote user can use this capability to remotely control or program the operation of the auxiliary device. For example, this enables recording of a television program, checking the status of a home security system, checking voice mail messages, hearing and responding to email messages, and much more. Kennewick '209 at 6:24–37 The system may be used in conjunction with a wide range of platform environments. These environments include, as a peripheral to a PC, a portable or wearable computer or other computer, embedded in a PC or other computer, on a personal digital assistant (e.g., a PDA with wireless networking capability, in a specialized device such as a bar code scanner or payment terminal, on a wireless telephone, or other platforms). If the invention is used with a mobile or portable device that has position location capability, the location data can be used by the invention as part of the context for user questions. A user may use the system on one or more devices. In this case history and profile information for the user may be synchronized between the multiple devices on a periodic basis or other basis. Ewer '149 similarly discloses a host device with telephony applications, which at least suggests a “cell phone.” Ewer '149 at [0028] The host device 104 is generally any suitable processor-based device such as a personal computer (PC), a personal digital assistant (PDA), a digital music player (e.g., MP3 player), a video player (e.g., DVD player), a video game player, and a processor-based telephone. The host device 104 executes application software such as a telephony application software that uses the headset 110, for example, for receiving the user's voice as input and/or for outputting sounds to the user as output. When the user disconnects the

Claims	Prior Art
	headset 110 at the connector 108, the headset adapter 106 transmits a digital flag signal indicating a change in the headset connection status to the host 104 running the application software. Similarly, when the user reconnects the headset 110 at the connector 108, the headset adapter 106 transmits the flag signal to indicate a change in the headset connection status to the application software running on the host 104. Thus, as is evident, the adapter 106 functions at least in part as a headset connection status signaling module to the host 104. Jabra FreeSpeak discloses communications between a headset and communication device such as a phone. The following citations show examples of where Jabra FreeSpeak discloses this limitation. Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak™— cord-free comfort and hands-free convenience in a Bluetooth® headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators. Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones.*

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 2 Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02  • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide Jabra FreeSpeak Datasheet at Page 2 About Bluetooth Bluetooth is a radio technology developed to deliver short-range wireless mobility. Bluetooth technology eliminates wires between devices, such as mobile phones and headsets, and enables communication between devices that are in close proximity to one another. Bluetooth facilitates fast, secure transmissions of both voice and data. More information is available at www.bluetooth.com.  A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim 1[pre]. Specifically, a POSITA would be motivated to implement a wireless headset as disclosed by Kennewick '209, Ewer '149, and/or Jabra FreeSpeak with the Apple Prior Art Systems. A POSITA would have been further motivated to provide the connectivity on a mobile device because the Speech Synthesis Technologies and Speech Recognition Technologies are already compatible with devices such as PDAs. A POSITA would further have a reasonable expectation of success in implementing a wireless headset with the Apple Prior Art Systems based on the disclosure of Jabra FreeSpeak BT200, which can provide wireless connectivity to devices that do not have Bluetooth capabilities through its adapter. To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in

Claims	Prior Art
	the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
23[pre]. A wireless headset that at least supports messaging for a first user, said wireless headset operatively connects to a mobile electronic device having wireless communication capability, and wherein said wireless headset is configured to at least:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless headset that at least supports messaging for a first user, said wireless headset operatively connects to a mobile electronic device having wireless communication capability, and wherein said wireless headset is configured to at least.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, have wireless communication capability and are compatible with and use wireless headsets as explained for claim 1[pre]-[1b]. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. iListen_15_New_speechrecogniti at Page 39 Most speech-recognition programs require a training period during which you acclimate the program to the unique qualities of your voice. With iListen, this process is easy. To begin, you set up the microphone, which in our case was the VX1 Parrott with Griffin iMic USB adapter (both included). Next, for about ten minutes, iListen leads you through sound-level tests and a sample passage of text. With initial training, iListen was about 78 percent accurate in our dictation tests in OS X and OS X's Classic mode. To improve recognition, you can read 11 other passages. Doing so brought the accuracy up slightly, to about 82 percent. The program can also analyze a set of your documents to learn the vocabulary you use.

Claims	Prior Art
	iListen_15_New_speechrecogniti at Page 39  The screenshot shows a window titled 'Document1' containing text about 'The Night Mile'. To the right is a control panel for 'iListen' with buttons for 'Listen', 'Write', and 'Transcribe'. Below these buttons is a text area labeled 'I heard:' containing the text 'leave it hard at the door'. iListen_15_New_speechrecogniti at Page 39 The Freedom of Transcription Despite its many drawbacks, iListen 1.5.2 has one outstanding new feature: it's the only Mac speech-recognition software able to transcribe AIF and WAV audio files. This feature alone may appeal to doctors, novelists, and others who prefer to do their writing far from the keyboard. Simply dictate into your digital tape recorder, download the file, push a button, and watch while iListen types what you said.

Claims	Prior Art
	iListen_15_New_speechrecogniti at Page 39 Macworld's Buying Advice MacSpeech's iListen 1.5.2 is a work in progress. At this point, IBM's ViaVoice—although it also has its flaws—is a better choice for most. The primary benefit of iListen today is the ability to transcribe documents dictated into a tape recorder. □ RATING: ♂♂ PROS: Support for audio-file transcription; can learn from dictation corrections in any program. CONS: No speech-controlled correction; frustrating correction process; limited and inconsistent built-in command-and-control abilities; incomplete documentation; some bugs. PRICE: \$99; with headset, \$149; upgrade, \$40 (doesn't include headset) OS COMPATIBILITY: Mac OS 9, Mac OS X COMPANY: MacSpeech, 603/644-7253, www.macspeech.com Kennewick '209 also discloses that the speech unit (wireless headset) supports messaging and is operatively connected to the main unit. The following citations show examples of where Kennewick '209 discloses this limitation. <i>See</i> claim [1pre]-[1b].

Claims

Prior Art

Kennewick '209 at FIG. 1




FIG. 1

Kennewick '209 at 5:47–6:6

In addition to the ability to process user questions and commands, the invention may provide local or remote control functions for the system or for other devices. Users may initiate commands locally or remotely. Several techniques for executing remote operations are possible, for example, the use of a telephone or other audio connection. Other remote command techniques may be used. The processing of these commands is performed in a nearly identical manner to a question. One difference being that the result of the command is generally an action rather than a response. Examples of control operations supported by the invention include, finding and playing music, film, games or other entertainment media, control of the appliances or other electronic or electro-mechanical devices, security systems control, audio, text and video conferencing (including Voice over IP conferencing), operating integrated messaging environments and other control operations. By using an optional analog or digital telephone interface or a data network interface, the invention can provide remote control capability of the computer or and auxiliary devices connected to the computer. For example, if the computer is connected via an interface to one or more auxiliary devices, a remote user can use this capability to remotely control or program the operation of the auxiliary device. For example, this enables recording of a television program, checking the status of a home security system, checking voice mail messages, hearing and

Claims	Prior Art
	responding to email messages, and much more. Kennewick '209 at 7:47–49 automatic telephone dialing, reading and sending emails and pages by voice and other communications control functions; Kennewick '209 at 9:66–10:7 The system 90 may include a main unit 98 and a speech unit 128. Alternatively, the system 98 may only comprise of the main unit 98, the speech unit 128 being a completely separate system. The event manager 100 may mediate interactions between other components of the main unit 98. The event manager 100 provides a multi-threaded environment allowing the system 98 to operate on multiple commands or questions from multiple user sessions without conflict and in an efficient manner, maintaining real-time response capabilities. Kennewick '209 at 10:21–32 The main user interface for the system 90, according to an embodiment of the invention, is through one or more speech units 128. The speech unit 128 includes one or more microphones, for example array microphone 134, to receive the utterances of the user. The speech received at the microphone 134 may be processed by filter 132 and passed to the speech coder 138 for encoding and compression. In one embodiment, a transceiver module 130 transmits the coded speech to the main unit 98. Coded speech received from the main unit 98 is detected by the transceiver 130, then decoded and decompressed by the speech coder 138 and annunciated by the speaker 136. Kennewick '209 at 10:33–48 The one or more speech units 128 and the main unit 98 may communicate over a communication link. The communication link can include a wired or wireless link. According to one embodiment, the communication link comprises an RF link. The transceiver 130 on the speech unit communicates coded speech data bi-directionally over the communication link with the transceiver 126 on the main unit 98. According to another embodiment, RF link may use any standard local area wireless data protocols including the IEEE 802.11, Bluetooth or other standards. Alternatively, an infrared data link conforming to any suitable standard such as IrDA or other infrared standards can be used. In an alternative embodiment, wires connect the speech unit 128 and the main unit 98, eliminating the need for one speech


Claims	Prior Art
	coder 138. Other wired or wireless analog or digital transmission techniques can be used. Kennewick '209 at 10:65–11:10 The agents 106 return results of questions as responses to users. The response may be created using the results of information queries, the system personality 108 and the user preferences or other data in the user profile 110. The agents 106 generally present these results using the speech unit 128. The agents 106 create a response string, which is sent to the text to speech engine 124. The text to speech engine 124 creates the required utterances, which may be encoded and compressed by the speech coder 122. Once coded, the utterances are transmitted from the main unit 98 by the transceiver 126 to the transceiver 130 on the speech unit 128. The utterance is then decoded and decompressed by the speech coder 138 and output by the speaker 136. Kennewick '209 at 16:17–23 In an alternative embodiment, the array microphone can be replaced by a set of physically distributed microphones or a headset worn by the user. The distributed microphones can be placed in different parts of a room or in different rooms. The distributed microphones can create a three-dimensional array to improve signal to noise ratio. The headset may use a wireless or wired connection. Ewer '149 similarly discloses its headset being wirelessly connected to a host device (e.g., “mobile electronic device”). <i>See</i> claim [1pre]-[1b]. The following citations show examples of where Ewer '149 discloses this limitation. Ewer '149 at Abstract A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on processor-based hosts using digital signaling are disclosed. The signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The signaling module may include a connection status detector and a status signal generating processor. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed.

Claims	Prior Art
	Ewer '149 at [0011] A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on personal computers or other processor-based hosts using digital signaling are disclosed. Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the term “connection” utilized herein generally refers to both wired and wireless connections. Ewer '149 at [0028] The host device 104 is generally any suitable processor-based device such as a personal computer (PC), a personal digital assistant (PDA), a digital music player (e.g., MP3 player), a video player (e.g., DVD player), a video game player, and a processor-based telephone. The host device 104 executes application software such as a telephony application software that uses the headset 110, for example, for receiving the user's voice as input and/or for outputting sounds to the user as output. When the user disconnects the headset 110 at the connector 108, the headset adapter 106 transmits a digital flag signal indicating a change in the headset connection status to the host 104 running the application software. Similarly, when the user reconnects the headset 110 at the connector 108, the headset adapter 106 transmits the flag signal to indicate a change in the headset connection status to the application software running on the host 104. Thus, as is evident, the adapter 106 functions at least in part as a headset connection status signaling module to the host 104. Jabra FreeSpeak discloses communications between a headset and communication device such as a phone. <i>See</i> claim [1pre]-[1b]. The following citations show examples of where Jabra FreeSpeak discloses this limitation.


Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak™— cord-free comfort and hands-free convenience in a Bluetooth™ headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators. Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones*.

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 2  Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02 • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide Jabra FreeSpeak Datasheet at Page 2 About Bluetooth Bluetooth is a radio technology developed to deliver short-range wireless mobility. Bluetooth technology eliminates wires between devices, such as mobile phones and headsets, and enables communication between devices that are in close proximity to one another. Bluetooth facilitates fast, secure transmissions of both voice and data. More information is available at www.bluetooth.com.  A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim 1[pre]-[1b]. Specifically, a POSITA would be motivated to implement a wireless headset as disclosed by Kennewick '209, Ewer '149, and/or Jabra FreeSpeak with the Apple Prior Art Systems, and a POSITA would have had a reasonable expectation of success in implementing a wireless headset with the Apple Prior Art Systems based on the disclosure of the Jabra FreeSpeak BT200. To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.


Claims	Prior Art
[23a] wirelessly receive an incoming audio message from the mobile electronic device, the incoming audio message being produced by converting an incoming text message from another wireless electronic device, the incoming text message being received at the mobile electronic device via a wireless network;	Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “wirelessly receive an incoming audio message from the mobile electronic device, the incoming audio message being produced by converting an incoming text message from another wireless electronic device, the incoming text message being received at the mobile electronic device via a wireless network.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, works with devices compatible with messaging functionality such as email. Apple Prior Art Systems can further “read” (e.g., convert the text to audio) the emails to a user. <i>See</i> claim [1pre]-[1d]. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Mobility Suite gives you voice access to your PIM applications. It allows you to use simple, natural voice commands such as "Go to Contacts" and "Read John Smith's home phone number". When you "Go to Mail", you can have all of your mail read to you, or you can read only selected messages. In the "Calendar" function, you can check your calendar for your next appointment, or add new ones. You can easily keep track of your busy schedule, all by using voice commands. In addition, this IBM software gives you the added convenience of being able to listen to rather than read information and messages. IBM ViaVoice Mobility Suite v.4 at 1 End-to-end voice solutions IBM can offer distributed end-to-end voice solutions where some of the speech processing resides on the client device (cell phones etc.) and integrates with additional voice functionality at the server level and large-scale telephony/unified messaging systems. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Enterprise Edition is IBM's flagship speech recognition and text-to-speech product. Representing a dramatic shift in the evolution of embedded speech products, IBM Embedded ViaVoice Enterprise Edition has an unlimited active vocabulary, limited only by the amount of memory on devices, plus a dynamic and swappable 4,000 word phonetic flat list. This feature enables embedded applications to access voice-enabled databases with millions of records. A typical application would be a voice-enabled navigation database of street names for a GPS system in your car.

Claims	Prior Art
	IBM ViaVoice Mobility Suite v.4 at Page 2  The screenshot shows a table titled "IBM Embedded ViaVoice Mobility Suite Features". It has two columns: "Inbox (Mail) Commands" and "Sample Commands". Under "Inbox (Mail) Commands", it says "Reads incoming mail; gives you the opportunity to reply and send responses to others." Under "Sample Commands", it lists: "Read mail", "Read new mail", "Next message", "Create reply", and "Send reply". IBM Mobility Suite Introduction at 3 Welcome to IBM Embedded ViaVoice Mobility Suite, a productivity application that controls your Personal Digital Assistant (PDA) by voice. In response to your commands, IBM Embedded ViaVoice Mobility Suite will help manage your appointments, contacts, tasks, and mail. IBM Mobility Suite Introduction at 5 • To hear important information, for example, say Read mail or Read my very next appointment. • To hear prompts, messages, or requested information again, say Repeat that. IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message.

Claims	Prior Art
	• read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. IBM ViaVoice Simply Dictation user guide at Page 26 Reading Text Now that you have learned to enter and format text using voice commands, it's time to have ViaVoice read the sentences back to you. Highlight the desired text block and say one of the following voice commands. SAY> BEGIN READING SAY> READ SELECTION ViaVoice reads the highlighted text using the voice selected in the Voices panel in ViaVoice Preferences. On applications with copy and selection capabilities, the Read Selection menu item reads the text using the voice currently selected in the Voices panel in ViaVoice Preferences. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 48 Reading Your Text Provided your application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back text aloud from almost any source, including documents and e-mail messages. 1. Open the application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected system voice. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 49 Changing Voices You can select the voice used when reading text. 1. Click the ViaVoice Preferences button on the VoiceCenter drawer. 2. Click the Voices icon. 3. From the Voice menu, select a new voice. 4. Press Play to hear a sample of your selection. 5. Click Apply, and close the window. 6. Any text read back from this point uses the voice you just selected. IBM ViaVoice Simply Dictation user guide at Page 61  The screenshot shows two sections from a user guide. The top section, titled 'SpeakPad Menu', explains that these menus contain items similar to other word processors and shows a menu bar with 'SpeakPad', 'File', 'Edit', 'Format', 'Dictation', 'Window', and 'Help'. The 'SpeakPad' menu is open, showing 'About SpeakPad' and 'Preferences...'. The bottom section, titled 'Dictation Menu', lists three items: 'Playback Selection/Stop Playback' (allows playing back dictated text in your own voice), 'Read Selection/Stop Reading' (reads highlighted text in a computer-generated voice), and 'Voice Preferences' (displays the Voices window to change the computer-generated voice). Inside Macintosh Sound_ 1994 at Page 4–3 This chapter describes the Speech Manager, the part of the Macintosh system software that provides a standardized method for Macintosh applications to generate synthesized speech.


Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–3 You need to read this chapter if you want your application to be able to generate speech. For example, you may want your application to incorporate the capability to speak its dialog box messages to the user. A word-processing application might use the Speech Manager to implement a command that speaks a selected section of a document to the user. A multimedia application might use the Speech Manager to provide a narration of a QuickTime movie instead of including sampled-sound data on a movie track. Because sound samples can take up large amounts of room on disk, using text in place of sampled sound is extremely efficient. Inside Macintosh Sound_ 1994 at Page 4–3 This chapter begins with an introduction to the speech generation process and then discusses how you can ■ check for the availability of the Speech Manager ■ create and dispose of speech channels ■ generate speech with different voices ■ obtain information about and change speech channel settings ■ start and stop speech production ■ synchronize speech production with other activities by using callback procedures ■ embed Speech Manager commands within text to make it more understandable Inside Macintosh Sound_ 1994 at Page 4–4 You can use the Speech Manager to incorporate synthesized speech into your application. This section provides an overview and describes the basic concepts of the Speech Manager, and it outlines the process that the Speech Manager uses to convert text into speech. The Speech Manager converts text into sound data, which it passes to the Sound Manager to play through the current sound output device. The Speech Manager's interaction with the Sound Manager is transparent to your application, so you don't need to be familiar with the Sound Manager to take advantage of the Speech Manager's capabilities. Inside Macintosh Sound_ 1994 at Page 4–4 Figure 4-1 illustrates the speech generation process. Your application can initiate speech generation by passing a string or a buffer of text to the Speech Manager. The Speech Manager is responsible for sending the text to a speech synthesizer, a component that contains executable code that manages all communication between the Speech Manager and the Sound Manager. A synthesizer is usually contained in a resource in a file within the System Folder. A synthesizer is like a speech engine. It uses built-in dictionaries and pronunciation rules to help determine how to pronounce text. You can provide custom pronunciation dictionaries as described in the section “Including Pronunciation Dictionaries” beginning on page 4-36.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4-4 <u>Figure 4-1</u> The speech generation process  <pre> graph LR Text[Text] --> SM[Speech Manager] SM --> SS[Speech synthesizer] SS --> SMgr[Sound Manager] SMgr --> AH[Audio hardware] </pre> Inside Macintosh Sound_ 1994 at Page 4-5 <u>Voices</u> Your application can use the system default voice to generate speech or it can specify that the Speech Manager use a particular voice that is available on the current computer system. A voice is a set of characteristics defined in parameters that specify a particular quality of speech. Just as different people's voices have different tonal qualities, so too can different voices have different qualities. A synthesized voice might sound male or female and might sound like an adult or a child. Some voices sound distinctively synthetic, while others sound more like real people. Figure 4-2 shows how the Speech Manager uses speech channels to synthesize speech with different voices. Inside Macintosh Sound_ 1994 at Page 4-6 Most of the time, your application designates the voice that speaks text, and usually that is the default voice. Based on the needs of your users and the way in which you expect them to use voices in your application you can provide access to voices in a number of different ways. You could include access to selecting voices in a dialog box that is available from a menu item such as Voices... Any application that allows users to choose among voices requires additional information about the available voices beyond the information provided by a voice specification record (described in detail on page 4-46), whose data should never be presented to the user. Such additional information might include the name of the voice as well as what script and language it supports.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–11 To take advantage of most of the Speech Manager's features, you must allocate a speech channel to pass to Speech Manager functions and dispose of the speech channel when you are finished using it. "Creating, Using, and Disposing of a Speech Channel" beginning on page 4-13 demonstrates how you do this and shows how you can use the <code>SpeakText</code> function to start speech generation from a buffer of text. Some applications permit users to choose a voice from those available to be used for speech generation. The <code>CountVoices</code>, <code>GetIndVoice</code>, and <code>GetVoiceDescription</code> functions support this capability. "Working With Different Voices" beginning on page 4-14 shows how you can use these functions to choose among available voices. Inside Macintosh Sound_1994 at Page 4–14 The <code>MyUseSpeechChannel</code> function defined in Listing 4-2 creates a default speech channel using the default system voice. You pass <code>NIL</code> in the first parameter to use the system default voice. You must also pass a global variable to <code>NewSpeechChannel</code> in which is returned a valid speech channel. Once the channel exists, then you can use the <code>SpeakText</code> function to generate speech. To generate synthesized speech, you pass in the channel allocated by <code>NewSpeechChannel</code> in the first parameter, and then you pass a pointer to the text that you want to speak as well as the length of the text that you want the Speech Manager to attempt to speak. That is, you can pass a pointer to a buffer of text that is 500 bytes long, but specify that only the first 10 bytes get spoken. Then <code>MyUseSpeechChannel</code> uses the <code>SpeechBusy</code> function in a <code>WHILE</code> loop to allow the text to be completely spoken before disposing of the channel. Inside Macintosh Sound_1994 at Page 4–14 <u>Working With Different Voices</u> When you work with speech channels, you can set a voice for a particular channel. When you set a voice, you may want to filter out certain of its characteristics in order to identify the one you want. For example, in an educational software application for elementary school students, you may want to use only children's voices. In order to choose the voice you want, you get a voice description record that contains information about a voice such as the size of the voice, the name of the voice, the age and gender of the voice, and the synthesizer with which it works. You can get the number of available voices using the <code>CountVoices</code> function. You can cycle through the available voices and identify the one you want to use by using the <code>GetIndVoice</code> function. Then you fill out a voice description record using the <code>GetVoiceDescription</code> function. Listing 4-3 shows how to get identifying information about a voice.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–23 <u>Writing Embedded Speech Commands</u> Embedded speech commands allow you to customize the quality of speech output by fine tuning it. You can make speech much easier to understand than the default way in which text is spoken by a synthesizer. An embedded speech command is a command embedded within a text buffer to be spoken by the Speech Manager that causes the Speech Manager to take a certain action. For example, you could use an embedded speech command to emphasize a particular word in a text string to make it stand out to the user. Inside Macintosh Sound_ 1994 at Page 4–30 <u>Examples of Embedded Speech Commands</u> If you use just a few of the embedded speech commands, you can markedly increase the understandability of text spoken by your application. Your application knows more about the speech being produced than a speech synthesizer does. A synthesizer speaks text according to a predetermined set of rules about language production. Therefore, the voices available on a Macintosh computer with the Speech Manager installed sound very synthetic and sometimes robotic because the pronunciation rules are formalized. You can make the speech produced by the synthesizer sound a lot more human by observing some simple rules of human speech and embedding speech commands in text according to these conventions. The techniques presented in this section could be applied when your application is having a dialog with the user or speaking some error messages or announcements. Inside Macintosh Sound_ 1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ■ generate speech and then pause or stop it ■ obtain information about an individual voice or all voices ■ create and dispose of speech channels ■ obtain the Speech Manager's version and status ■ change the rate or pitch of speech

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–73 <u>Changing Speech Attributes</u> To determine the rate and pitch at which a speech channel is processing text, you can use the <code>GetSpeechRate</code> and <code>GetSpeechPitch</code> functions. The <code>SetSpeechRate</code> and <code>SetSpeechPitch</code> functions allow you to change rate and pitch. The most robust of the Speech Manager's routines are the <code>GetSpeechInfo</code> and <code>SetSpeechInfo</code> functions. These allow you to obtain many types of information about a speech channel and to change many settings of a speech channel. To specify the operation that you wish to perform, you must pass <code>GetSpeechInfo</code> or <code>SetSpeechInfo</code> a selector. A full list of selectors is provided in "Speech Information Selectors" beginning on page 4-39. Inside Macintosh Sound_ 1994 at Page 4–76 <u>SetSpeechPitch</u> You can use the <code>SetSpeechPitch</code> function to set the speech pitch on a designated speech channel. <pre>FUNCTION SetSpeechPitch (chan: SpeechChannel; pitch: Fixed) : OSErr;</pre> <code>chan</code> The speech channel whose pitch you wish to set. <code>pitch</code> The new pitch for the speech channel, expressed as a fixed-point frequency value.

Claims	Prior Art
	iListen_15_New_speechrecogniti at Page 39 Most speech-recognition programs require a training period during which you acclimate the program to the unique qualities of your voice. With iListen, this process is easy. To begin, you set up the microphone, which in our case was the VX1 Parrott with Griffin iMic USB adapter (both included). Next, for about ten minutes, iListen leads you through sound-level tests and a sample passage of text. With initial training, iListen was about 78 percent accurate in our dictation tests in OS X and OS X's Classic mode. To improve recognition, you can read 11 other passages. Doing so brought the accuracy up slightly, to about 82 percent. The program can also analyze a set of your documents to learn the vocabulary you use. iListen_15_New_speechrecogniti at Page 39  Write Anywhere iListen is the only Mac-compatible speech-recognition program with the ability to transcribe AIFF and WAV files imported from a tape recorder.

Claims	Prior Art
	iListen_15_New_speechrecogniti at Page 39 The Freedom of Transcription Despite its many drawbacks, iListen 1.5.2 has one outstanding new feature: it's the only Mac speech-recognition software able to transcribe AIFF and WAV audio files. This feature alone may appeal to doctors, novelists, and others who prefer to do their writing far from the keyboard. Simply dictate into your digital tape recorder, download the file, push a button, and watch while iListen types what you said. iListen_15_New_speechrecogniti at Page 39 Macworld's Buying Advice MacSpeech's iListen 1.5.2 is a work in progress. At this point, IBM's ViaVoice—although it also has its flaws—is a better choice for most. The primary benefit of iListen today is the ability to transcribe documents dictated into a tape recorder. □ RATING: ## PROS: Support for audio-file transcription; can learn from dictation corrections in any program. CONS: No speech-controlled correction; frustrating correction process; limited and inconsistent built-in command-and-control abilities; incomplete documentation; some bugs. PRICE: \$99; with headset, \$149; upgrade, \$40 (doesn't include headset) OS COMPATIBILITY: Mac OS 9, Mac OS X COMPANY: MacSpeech, 603/644-7253, www.macspeech.com Kennewick '209 also discloses receiving a text message and presenting the text message to the user as audio through a speech unit (e.g., “wireless headset”). The following citations show examples of where Kennewick '209 discloses this limitation.

Claims

Prior Art

Kennewick '209 at FIG. 1




FIG. 1

Kennewick '209 at 5:47–6:6

In addition to the ability to process user questions and commands, the invention may provide local or remote control functions for the system or for other devices. Users may initiate commands locally or remotely. Several techniques for executing remote operations are possible, for example, the use of a telephone or other audio connection. Other remote command techniques may be used. The processing of these commands is performed in a nearly identical manner to a question. One difference being that the result of the command is generally an action rather than a response. Examples of control operations supported by the invention include, finding and playing music, film, games or other entertainment media, control of the appliances or other electronic or electro-mechanical devices, security systems control, audio, text and video conferencing (including Voice over IP conferencing), operating integrated messaging environments and other control operations. By using an optional analog or digital telephone interface or a data network interface, the invention can provide remote control capability of the computer or and auxiliary devices connected to the computer. For example, if the computer is connected via an interface to one or more auxiliary devices, a remote user can use this capability to remotely control or program the operation of the auxiliary device. For example, this enables recording of a television program, checking the status of a home security system, checking voice mail messages, hearing and

Claims	Prior Art
	responding to email messages, and much more. Kennewick '209 at 7:47–49 automatic telephone dialing, reading and sending emails and pages by voice and other communications control functions; Kennewick '209 at 9:66–10:7 The system 90 may include a main unit 98 and a speech unit 128. Alternatively, the system 98 may only comprise of the main unit 98, the speech unit 128 being a completely separate system. The event manager 100 may mediate interactions between other components of the main unit 98. The event manager 100 provides a multi-threaded environment allowing the system 98 to operate on multiple commands or questions from multiple user sessions without conflict and in an efficient manner, maintaining real-time response capabilities. Kennewick '209 at 10:21–32 The main user interface for the system 90, according to an embodiment of the invention, is through one or more speech units 128. The speech unit 128 includes one or more microphones, for example array microphone 134, to receive the utterances of the user. The speech received at the microphone 134 may be processed by filter 132 and passed to the speech coder 138 for encoding and compression. In one embodiment, a transceiver module 130 transmits the coded speech to the main unit 98. Coded speech received from the main unit 98 is detected by the transceiver 130, then decoded and decompressed by the speech coder 138 and annunciated by the speaker 136. Kennewick '209 at 10:33–48 The one or more speech units 128 and the main unit 98 may communicate over a communication link. The communication link can include a wired or wireless link. According to one embodiment, the communication link comprises an RF link. The transceiver 130 on the speech unit communicates coded speech data bi-directionally over the communication link with the transceiver 126 on the main unit 98. According to another embodiment, RF link may use any standard local area wireless data protocols including the IEEE 802.11, Bluetooth or other standards. Alternatively, an infrared data link conforming to any suitable standard such as IrDA or other infrared standards can be used. In an alternative embodiment, wires connect the speech unit 128 and the main unit 98, eliminating the need for one speech

Claims	Prior Art
	coder 138. Other wired or wireless analog or digital transmission techniques can be used. Kennewick '209 at 10:65–11:10 The agents 106 return results of questions as responses to users. The response may be created using the results of information queries, the system personality 108 and the user preferences or other data in the user profile 110. The agents 106 generally present these results using the speech unit 128. The agents 106 create a response string, which is sent to the text to speech engine 124. The text to speech engine 124 creates the required utterances, which may be encoded and compressed by the speech coder 122. Once coded, the utterances are transmitted from the main unit 98 by the transceiver 126 to the transceiver 130 on the speech unit 128. The utterance is then decoded and decompressed by the speech coder 138 and output by the speaker 136. Ewer '149 similarly discloses communications between a headset and communication device such as a phone. The following citations show examples of where Ewer '149 discloses this limitation. Ewer '149 at Abstract A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on processor-based hosts using digital signaling are disclosed. The signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The signaling module may include a connection status detector and a status signal generating processor. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed. Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the term “connection” utilized herein generally refers to both wired and wireless connections.


Claims	Prior Art
	Jabra FreeSpeak similarly discloses communications between a headset and communication device such as a phone. The following citations show examples of where Jabra FreeSpeak discloses this limitation. Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak™— cord-free comfort and hands-free convenience in a Bluetooth® headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators. Jabra FreeSpeak Datasheet at Page 2 Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02  • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide 

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 2 About Bluetooth Bluetooth is a radio technology developed to deliver short-range wireless mobility. Bluetooth technology eliminates wires between devices, such as mobile phones and headsets, and enables communication between devices that are in close proximity to one another. Bluetooth facilitates fast, secure transmissions of both voice and data. More information is available at www.bluetooth.com.  Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones*. A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim 1[pre]. Specifically, a POSITA would be motivated to implement a wireless headset as disclosed by Kennewick '209, Ewer '149, and/or Jabra FreeSpeak with the Apple Prior Art Systems, and a POSITA would have a reasonable expectation of success in implementing a wireless headset to the Apple Prior Art Systems based on the disclosure of the Jabra FreeSpeak. To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra

Claims	Prior Art
	FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[23b] output the incoming audio message for the first user;	Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “output the incoming audio message for the first user.” <i>See above at claim [1e].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[23c] receive a reply audio message from the first user in response to the incoming audio message; and	Apple Prior Art Systems expressly or inherently disclose “receive a reply audio message from the first user in response to the incoming audio message.” <i>See above at claim [1f].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[23d] wirelessly transmit the reply audio message to the mobile electronic device.	Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “wirelessly transmit the reply audio message to the mobile electronic device.” <i>See above at claim [1b], [1g], [1h].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
24. A wireless headset according to claim 23, wherein the reply audio message thereafter being converted to a reply text message	Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset according to claim 23, wherein the reply audio message thereafter being converted to a reply text message and transmitted to the another wireless electronic device via the wireless network.”

Claims	Prior Art
and transmitted to the another wireless electronic device via the wireless network.	Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, work with devices compatible with messaging functionality such as email. Apple Prior Art Systems can further send messages to other devices. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM Announces ViaVoice 98 at 1 IBM recently announced the release of ViaVoice 98, the next generation of IBM's best selling speech recognition software, which features numerous innovations designed to make speech easier and more natural. IBM Announces ViaVoice 98 at 1 ViaVoice 98 Executive Edition is IBM's most powerful continuous speech software and can be a productivity tool for writers, executives and experienced PC users. ViaVoice Executive offers direct dictation into most popular Windows applications, voice control for the desktop and applications and includes two topics, ViaVoice Computer and ViaVoice Business and Finance, which supplement the basic vocabulary. Topics are special interest vocabularies that add to the existing base vocabulary and increase accuracy. The executive edition is priced at \$149. ViaVoice 98 Office Edition is a continuous speech dictation and voice control software tool designed to help maximize the efficiency of a small office/home office user, faculty member or advanced student. Users can supplement their vocabulary with the included ViaVoice Business and Finance. The price for the Office edition is \$89. ViaVoice 98 Home Edition is IBM's entry-level continuous speech dictation product and is ideal for the family to use to create letters, reports or e-mail at home. ViaVoice 98 Home includes the ViaVoice Cuisine Topic to let users produce original recipes or make a keepsake of family recipes. The cost is \$49.95. A customized, integrated version of IBM ViaVoice is included with Lotus SmartSuite Millennium edition. IBM Speech assistant (ViaVoice) History at 2 1992--IBM introduces its first dictation system, called IBM Speech Server Series (ISSS). IBM Speech assistant (ViaVoice) History at 2 1993--IBM launches the IBM Personal Dictation System (IPDS) for OS/2*, running on an IBM PC using a custom audio adapter card and a conventional microphone. The system takes dictation at about 80 words a minute with 95 percent accuracy, and allows users to "play back" their recording for editing purposes.


Claims	Prior Art
	IPDS supports US English as well as UK English, German, French, Italian and Spanish. IBM ViaVoice 10 at 1 Talk into a microphone and the computer types what you say. That's an easy way to describe voice recognition software like Dragon NaturallySpeaking and IBM ViaVoice. Morphing through various forms since 1994, Dragon NaturallySpeaking is arguably the best voice recognition software available today. IBM's "40 years of commitment to speech research and development" have in part lead to the ViaVoice software. This article compares these two as well as providing general comments on voice recognition technology. (If we want to get picky, "speech recognition" is the process of converting speech to text, and "voice recognition" is the process of identifying a person by their voice alone. However, for this article, I'm using "voice recognition" to describe these software packages since the correct version of this term is rarely used correctly.) IBM ViaVoice Mobility Suite v.4 at 1 Operate PDA functions using your voice Imagine being able to operate many of the common PDA functions like e-mail, tasks, calendar, phone numbers and addresses, by using voice commands instead of a stylus. IBM Embedded ViaVoice Mobility Suite software is now available to do just that! IBM Embedded ViaVoice Mobility Suite voice-enables many of the most common PIM applications for the handheld PDA, and is now available on the Compaq H3800 Pocket PC. IBM ViaVoice Mobility Suite v.4 at 1 End-to-end voice solutions IBM can offer distributed end-to-end voice solutions where some of the speech processing resides on the client device (cell phones etc.) and integrates with additional voice functionality at the server level and large-scale telephony/unified messaging systems. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Enterprise Edition is IBM's flagship speech recognition and text-to-speech product. Representing a dramatic shift in the evolution of embedded speech products, IBM Embedded ViaVoice Enterprise Edition has an unlimited active vocabulary, limited only by the amount of memory on devices, plus a dynamic and swappable 4,000 word phonetic flat list. This feature enables embedded applications to access voice-enabled databases with millions of records. A typical application would be a voice-enabled navigation database of street names for a GPS system in your car.

Claims	Prior Art
	IBM ViaVoice Mobility Suite v.4 at Page 2  The screenshot shows a table titled "IBM Embedded ViaVoice Mobility Suite Features". It has two columns: "Inbox (Mail) Commands" and "Sample Commands". Under "Inbox (Mail) Commands", it says "Reads incoming mail; gives you the opportunity to reply and send responses to others." Under "Sample Commands", it lists: "Read mail", "Read new mail", "Next message", "Create reply", and "Send reply". IBM Mobility Suite Introduction at 3 Welcome to IBM Embedded ViaVoice Mobility Suite, a productivity application that controls your Personal Digital Assistant (PDA) by voice. In response to your commands, IBM Embedded ViaVoice Mobility Suite will help manage your appointments, contacts, tasks, and mail. IBM Mobility Suite Introduction at 5 • To hear important information, for example, say Read mail or Read my very next appointment. • To hear prompts, messages, or requested information again, say Repeat that. IBM Mobility Suite Introduction at 7 • start recording - Starts recording your voice for: • Training the system to pronounce your contact names and program names • Creating a recording as a mail attachment IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the

Claims	Prior Art
	Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message. • read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. IBM ViaVoice Simply Dictation user guide at Page 1 You talk, iType. It's that simple. The computer listens to your words and phrases and then translates them into text or actions that carry out commands. With ViaVoice, you'll give your keyboard and fingertips a well-deserved rest! IBM ViaVoice Simply Dictation user guide at Page 18 2. When SpeakPad is open, the blinking cursor indicates the place on the window where your dictated text appears. Your voice commands appear in a yellow message screen above the VoiceCenter, as shown in the graphic above. 

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 37 Dictating with ViaVoice There are two methods you can use to dictate and process text with ViaVoice: dictating into SpeakPad or using Dictation Stream Redirection (DSR), which allows you to dictate into another application. NOTE Make sure you select the correct speech mode button on the VoiceCenter. You should probably select <i>Dictation</i> and <i>Command</i> mode for most dictation and correction activities. For more information, see page 56. Dictating Into SpeakPad This is the most common method for dictating, since SpeakPad is built upon Apple's TextEdit technology. This integration provides a simple and seamless way to dictate and edit text. Additionally, correction is only available in SpeakPad. Once you dictate into SpeakPad, you can easily edit, format, and transfer your text into other supported applications. 1. Open SpeakPad from the ViaVoice folder. 2. Begin dictating your text. IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. iListen_15_New_speechrecogniti at Page 39 iLISTEN 1.5 New Speech-Recognition Software for OS X Lacks Polish but Offers Transcription Support

Claims	Prior Art
	iListen_15_New_speechrecogniti at Page 39 It's a happy day when Mac users have two full-featured speech-recognition programs to choose from—programs that let you dictate using natural speech <i>and</i> control your Mac by voice. This technology is critical to people who have limited use of their arms; for everyone else, it offers an exciting new way to interact with a computer. Before late 2001, advanced speech recognition was available only in Windows. That's when IBM released the capable but incomplete ViaVoice for Mac OS X (¶¶); <i>Reviews</i>, April 2002). Now MacSpeech's iListen 1.5.2 joins the Mac OS X speech-recognition ranks. Unfortunately, this program's performance is inconsistent, but its new ability to transcribe audio files may make it worthwhile to some. iListen_15_New_speechrecogniti at Page 39 Most speech-recognition programs require a training period during which you acclimate the program to the unique qualities of your voice. With iListen, this process is easy. To begin, you set up the microphone, which in our case was the VX1 Parrott with Griffin iMic USB adapter (both included). Next, for about ten minutes, iListen leads you through sound-level tests and a sample passage of text. With initial training, iListen was about 78 percent accurate in our dictation tests in OS X and OS X's Classic mode. To improve recognition, you can read 11 other passages. Doing so brought the accuracy up slightly, to about 82 percent. The program can also analyze a set of your documents to learn the vocabulary you use.

Claims	Prior Art
	iListen_15_New_speechrecogniti at Page 39 begin dictating into any application, such as Microsoft Word X or 2001. Using an 800MHz iMac G4 with 512MB of RAM, we saw good performance from the program—there was only a slight delay between speaking the words and seeing them appear on screen. iListen_15_New_speechrecogniti at Page 39  Write Anywhere iListen is the only Mac-compatible speech-recognition program with the ability to transcribe Aiff and WAV files imported from a tape recorder. iListen_15_New_speechrecogniti at Page 39 Macworld's Buying Advice MacSpeech's iListen 1.5.2 is a work in progress. At this point, IBM's ViaVoice—although it also has its flaws—is a better choice for most. The primary benefit of iListen today is the ability to transcribe documents dictated into a tape recorder. □ RATING: ■■ PROS: Support for audio-file transcription; can learn from dictation corrections in any program. CONS: No speech-controlled corrections; frustrating correction process; limited and inconsistent built-in command-and-control abilities; incomplete documentation; some bugs. PRICE: \$99, with headset, \$149; upgrade, \$40 (doesn't include headset) OS COMPATIBILITY: Mac OS 9, Mac OS X COMPANY: MacSpeech, 603/644-7253, www.macspeech.com

Claims	Prior Art
	To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
26. A wireless headset as recited in claim 24, wherein the incoming text message is a SMS message, and wherein the reply text message is a SMS message.	Apple Prior Art Systems and Kennewick ‘209 expressly or inherently disclose “[a] wireless headset as recited in claim 24, wherein the incoming text message is a SMS message, and wherein the reply text message is a SMS message.” <i>See above at claim 16.</i> To the extent that Apple Prior Art Systems and Kennewick ‘209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick ‘209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
27[pre]. A wireless headset as recited in claim 23, wherein said wireless headset is configured to:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless headset as recited in claim 23, wherein said wireless headset is configured to.” <i>See above at claim [1a].</i> To the extent that Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[27a] prompt the first user to provide the reply audio message prior to said receiving of the reply audio message.	Apple Prior Art Systems expressly or inherently disclose “prompt the first user to provide the reply audio message prior to said receiving of the reply audio message.” <i>See above at claim [3a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.

Claims	Prior Art
28. A wireless headset as recited in claim 23, wherein the conversion of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset as recited in claim 23, wherein the conversion of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic.” <i>See above at claim 11.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
29. A wireless headset as recited in claim 28, wherein the at least one user-selected speech characteristic is set based on at least one user preference able to be set by or for the first user.	Apple Prior Art. Systems expressly or inherently disclose “[a] wireless headset as recited in claim 28, wherein the at least one user-selected speech characteristic is set based on at least one user preference able to be set by or for the first user.” <i>See above at claim 12.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
30. A wireless headset as recited in claim 28, wherein the at least one user-selected speech characteristic is set based on at least one user setting able to be provided at the mobile electronic device.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset as recited in claim 28, wherein the at least one user-selected speech characteristic is set based on at least one user setting able to be provided at the mobile electronic device.” <i>See above at claim 13.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
31. A wireless headset as recited in claim 28, wherein the at least one user-selected speech characteristic pertains to at least a voice for use in producing the incoming audio message.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset as recited in claim 28, wherein the at least one user-selected speech characteristic pertains to at least a voice for use in producing the incoming audio message.” <i>See above at claim 14.</i>

Claims	Prior Art
	To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
32[pre]. A wireless headset as recited in claim 23,	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak expressly or inherently disclose “[a] wireless headset as recited in claim 23.” <i>See above at claim [1a].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[32a] wherein said wireless headset includes at least one speaker,	Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “wherein said wireless headset includes at least one speaker.” Kennewick '209 discloses a speaker in the speech unit. The following citations show examples of where Kennewick '209 discloses this limitation.

Claims

Prior Art

Kennewick '209 at FIG. 1




FIG. 1

Kennewick '209 at 10:21–32

The main user interface for the system 90, according to an embodiment of the invention, is through one or more speech units 128. The speech unit 128 includes one or more microphones, for example array microphone 134, to receive the utterances of the user. The speech received at the microphone 134 may be processed by filter 132 and passed to the speech coder 138 for encoding and compression. In one embodiment, a transceiver module 130 transmits the coded speech to the main unit 98. Coded speech received from the main unit 98 is detected by the transceiver 130, then decoded and decompressed by the speech coder 138 and annunciated by the speaker 136.

Kennewick '209 at 10:65–11:10

The agents 106 return results of questions as responses to users. The response may be created using the results of information queries, the system personality 108 and the user preferences or other data in the user profile 110. The agents 106 generally present these results using the speech unit 128. The agents 106 create a response string, which is sent to the text to speech engine 124. The text to speech engine 124 creates the required utterances, which may be encoded and compressed by the speech coder 122. Once coded, the utterances are transmitted from the main unit 98 by the transceiver 126 to the transceiver 130

Claims	Prior Art
	on the speech unit 128. The utterance is then decoded and decompressed by the speech coder 138 and output by the speaker 136. Kennewick '209 at 15:25–31 According to an embodiment of the invention, the microphone array, filters and speech coder 138 are physically separated from the main unit 98 into a speech unit 128, and connected using a wireless link. Since bandwidth on a wireless connection is at a premium, the speech coder dynamically adapts the digitization rate and compression of the captured speech. Ewer '149 similarly discloses a system with speakers. The following citations show examples of where Ewer '149 discloses this limitation. Ewer '149 at [0029] As shown, the system 100 may additionally include any number of input interfaces 112, typically external interfaces such as keyboard and mouse as well as any number of external applications 102 such as a display, DVD player, CD-ROM drive, CD-RW drive, microphone, and speakers. Jabra FreeSpeak similarly discloses communications between a headset and communication device such as a phone. The following citations show examples of where Jabra FreeSpeak discloses this limitation. Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak™— cord-free comfort and hands-free convenience in a Bluetooth® headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators.

Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 2  Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02 • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide Jabra FreeSpeak Datasheet at Page 2 About Bluetooth Bluetooth is a radio technology developed to deliver short-range wireless mobility. Bluetooth technology eliminates wires between devices, such as mobile phones and headsets, and enables communication between devices that are in close proximity to one another. Bluetooth facilitates fast, secure transmissions of both voice and data. More information is available at www.bluetooth.com.  Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones.*

Claims	Prior Art
	A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim 1[pre]. Specifically, a POSITA would be motivated to implement a wireless headset with a speaker as disclosed by Kennewick '209, Ewer '149, and/or Jabra FreeSpeak to the Apple Prior Art Systems, and a POSITA would have a reasonable expectation of success in implementing a wireless headset with a speaker to the Apple Prior Art Systems based on Apple Prior Art Systems' compatibility with wired headsets that include speakers for presenting audio. <i>See</i> https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications). To the extent that Apple Prior Art Systems, Kennewick '209, and Ewer '149 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, and Ewer '149, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[32b] wherein the output of the incoming audio message comprises directing the incoming audio message to the at least one speaker in said wireless headset, and	Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “wherein the output of the incoming audio message comprises directing the incoming audio message to the at least one speaker in said wireless headset.” <i>See</i> above at claim [1b]. To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[32c] wherein the output of the incoming audio message is dependent on at least one user-selected speech characteristic.	Apple Prior Art Systems expressly or inherently disclose “wherein the output of the incoming audio message is dependent on at least one user-selected speech characteristic.” <i>See</i> above at claim 11. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.


Claims	Prior Art
33. A wireless headset as recited in claim 23, wherein the incoming text message is a SMS message.	Apple Prior Art Systems and Kennewick ‘209 expressly or inherently disclose “[a] wireless headset as recited in claim 23, wherein the incoming text message is a SMS message.” <i>See above at claim 16.</i> To the extent that Apple Prior Art Systems and Kennewick ‘209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick ‘209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
34[pre]. A wireless headset as recited in claim 23, wherein the conversion of the incoming text message to the incoming audio message comprises:	To the extent the preamble is limiting, Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset as recited in claim 23, wherein the conversion of the incoming text message to the incoming audio message comprises.” <i>See above at claim [1d].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[34a] translation of at least a portion of the incoming text message from a condensed format to an uncondensed text format; and	Apple Prior Art Systems expressly or inherently disclose “translation of at least a portion of the incoming text message from a condensed format to an uncondensed text format.” <i>See above at claim [18a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[34b] conversion of at least the portion of the incoming text message in the uncondensed text format to at least a portion of the incoming audio message.	Apple Prior Art Systems expressly or inherently disclose “conversion of at least the portion of the incoming text message in the uncondensed text format to at least a portion of the incoming audio message.” <i>See above at claim [18a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.


Claims	Prior Art
35. A wireless headset as recited in claim 34, wherein the conversion of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset as recited in claim 34, wherein the conversion of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic.” <i>See above at claim 19.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
36. A wireless headset as recited in claim 35, wherein the at least one user-selected speech characteristic is set based on at least one user setting able to be provided at the mobile electronic device.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset as recited in claim 35, wherein the at least one user-selected speech characteristic is set based on at least one user setting able to be provided at the mobile electronic device.” <i>See above at claim 20.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
37. A wireless headset as recited in claim 23, wherein at least a portion of the incoming text message is in a chat language, and wherein the conversion of the incoming text message to the incoming audio message comprises expanding the portion of the incoming text message that is in the chat language to a plurality of words.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset as recited in claim 23, wherein at least a portion of the incoming text message is in a chat language, and wherein the conversion of the incoming text message to the incoming audio message comprises expanding the portion of the incoming text message that is in the chat language to a plurality of words.” <i>See above at claim [18a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
39. A wireless headset as recited in claim 23, wherein the mobile electronic device is a mobile telephone.	Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless headset as recited in claim 23, wherein the mobile electronic device is a mobile telephone.” <i>See above at claim 22.</i>

Claims	Prior Art
	To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
40[pre]. A wireless headset as recited in claim 23, wherein said wireless headset is configured to:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless headset as recited in claim 23, wherein said wireless headset is configured to.” <i>See above at claim [1a].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[40a] prompt the first user to provide the reply audio message prior to said receiving of the reply audio message,	Apple Prior Art Systems “prompt the first user to provide the reply audio message prior to said receiving of the reply audio message.” <i>See above at claim [3a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[40b] wherein the reply audio message thereafter being converted to a reply text message and transmitted to the another wireless electronic device via the wireless network,	Apple Prior Art Systems expressly or inherently disclose “wherein the reply audio message thereafter being converted to a reply text message and transmitted to the another wireless electronic device via the wireless network.” <i>See above at claim 24.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.

Claims	Prior Art
[40c] wherein the conversion of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic,	Apple Prior Art Systems expressly or inherently disclose “wherein the conversion of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic.” <i>See above at claim 11.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[40d] wherein the at least one user-selected speech characteristic is set based on at least one user setting able to be provided at the mobile electronic device, and	Apple Prior Art Systems expressly or inherently disclose “wherein the at least one user-selected speech characteristic is set based on at least one user setting able to be provided at the mobile electronic device.” <i>See above at claim 9.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[40e] wherein the incoming text message is a SMS message, and wherein the reply text message is a SMS message.	Apple Prior Art Systems and Kennewick ‘209 expressly or inherently disclose “wherein the incoming text message is a SMS message, and wherein the reply text message is a SMS message.” <i>See above at claim 16.</i> To the extent that Apple Prior Art Systems and Kennewick ‘209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick ‘209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
41. A wireless headset as recited in claim 40, wherein at least a portion of the incoming text message is in a chat language, and wherein the conversion of the incoming text message to the incoming audio message comprises expanding the portion of the incoming text	Apple Prior Art Systems and Kennewick ‘209 expressly or inherently disclose “[a] wireless headset as recited in claim 40, wherein at least a portion of the incoming text message is in a chat language, and wherein the conversion of the incoming text message to the incoming audio message comprises expanding the portion of the incoming text message that is in the chat language to a plurality of words.” Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, include vocabulary settings where users are able to set their own spellings and dictations for abbreviations and acronyms, which is an example of

Claims	Prior Art
message that is in the chat language to a plurality of words.	“translating ... from a condensed format to an uncondensed format.” The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM Announces ViaVoice 98 at 1 ViaVoice 98 Executive Edition is IBM’s most powerful continuous speech software and can be a productivity tool for writers, executives and experienced PC users. ViaVoice Executive offers direct dictation into most popular Windows applications, voice control for the desktop and applications and includes two topics, ViaVoice Computer and ViaVoice Business and Finance, which supplement the basic vocabulary. Topics are special interest vocabularies that add to the existing base vocabulary and increase accuracy. The executive edition is priced at \$149. ViaVoice 98 Office Edition is a continuous speech dictation and voice control software tool designed to help maximize the efficiency of a small office/home office user, faculty member or advanced student. Users can supplement their vocabulary with the included ViaVoice Business and Finance. The price for the Office edition is \$89. ViaVoice 98 Home Edition is IBM’s entry-level continuous speech dictation product and is ideal for the family to use to create letters, reports or e-mail at home. ViaVoice 98 Home includes the ViaVoice Cuisine Topic to let users produce original recipes or make a keepsake of family recipes. The cost is \$49.95. A customized, integrated version of IBM ViaVoice is included with Lotus SmartSuite Millennium edition. IBM Announces ViaVoice 98 at 2 Large active vocabulary - A 128,000 word active vocabulary includes 64,000 active base vocabulary and 64,000 more user-added words using Vocabulary Expander and Vocabulary Manager. Users can include multiple topics (each topic can contain up to 64,000 words) and a backup vocabulary of 260,000 words. Industry specific vocabularies - ViaVoice 98 supports industry-specific vocabularies. The first vocabularies will be ViaVoice Medical and ViaVoice Legal in U.S. English. These are specialized vocabularies with 25,000 or more words for use by medical and legal professionals. They cost \$149. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Enterprise Edition is IBM’s flagship speech recognition and text-to-speech product. Representing a dramatic shift in the evolution of embedded speech products, IBM Embedded ViaVoice Enterprise Edition has an unlimited active vocabulary, limited only by the amount of memory on devices, plus a dynamic and swappable 4,000 word phonetic flat list. This feature enables embedded applications to access voice-enabled databases with millions of records. A typical application would be a voice-enabled navigation database of street names for a GPS system in your car.

Claims	Prior Art				
	IBM ViaVoice Mobility Suite v.4 at Page 2  <table border="1"> <thead> <tr> <th>Inbox (Mail) Commands</th><th>Sample Commands</th></tr> </thead> <tbody> <tr> <td>Reads incoming mail; gives you the opportunity to reply and send responses to others.</td><td> • "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply" </td></tr> </tbody> </table> IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message. • read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. IBM ViaVoice Simply Dictation user guide at Page 38 Correcting Text in SpeakPad  Correcting words as you dictate in SpeakPad updates your personal vocabulary, which helps improve recognition accuracy by immediately teaching ViaVoice new words and phrases. This way, if you say the word again while you are dictating the same document, ViaVoice is more likely to recognize it correctly.	Inbox (Mail) Commands	Sample Commands	Reads incoming mail; gives you the opportunity to reply and send responses to others.	• "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply"
Inbox (Mail) Commands	Sample Commands				
Reads incoming mail; gives you the opportunity to reply and send responses to others.	• "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply"				

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 49 Using Manage My Vocabulary The Manage My Vocabulary window allows you to manage your personal vocabulary by deleting or training how you say specific words. When you start Manage My Vocabulary, it shows you all of the words you have added to your personal vocabulary. Words shown in this window were added during dictation, correction, or when creating voice commands. 1. SAY> MANAGE MY VOCABULARY or, click the Manage My Vocabulary button from the VoiceCenter drawer. 2. Select the word to be re-trained. 3. Click the Train button. The Train Word window appears. 4. Click the Record button and speak the word. 5. Click OK when you have trained the word. If you want to check the quality of your recording, click the Play button. IBM ViaVoice Simply Dictation user guide at Page 57 VoiceCenter Drawer The VoiceCenter drawer menu allows easy access to most of the ViaVoice options, functions, and information.  • What Can I Say - Displays a list of available voice commands. • Correction - Displays the Correction window. • Preferences - Use to customize functions of ViaVoice. • Macros - Displays the Macro Editor window so you can create and edit dictation macros. • Manage My Vocabulary - Displays the Manage My Vocabulary window so you can re-train the pronunciation of words. • UserSettings - Use to add Voice Models or change other User information. • Information Central - Displays Information Central, the ViaVoice online help file. • SetupAssistant - Use to set up your audio, create a user or analyze documents. • About ViaVoice - Displays the ViaVoice About box.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 69 Recording a Pronunciation 1. If the word does not sound like it is written, type the word or phrase the way that you pronounce it in the <i>Sounds-like</i> spelling field. Remember to place a hyphen between the words and letters. For example, if you pronounce the word NCAA as NC double A, enter N-C-double-A. If you pronounce NCAA as NCAA, enter N-C-A-A. IBM ViaVoice Simply Dictation user guide at Page 71 Capitalize This - Capitalizes the selected text. Lowercase This - Lowercases the selected text. Uppercase This - Uppercases the selected text. Numeric - Changes numeric text to digits. Spell Out - Changes digits or abbreviations to words. Inside Macintosh Sound_1994 at Page 4–36 <u>Including Pronunciation Dictionaries</u> No matter how sophisticated a speech synthesis system is, there will always be words that it does not automatically pronounce correctly. A clear instance of words that are often mispronounced is the class of proper nouns (names of people, place names, and so on). The Speech Manager supports pronunciation dictionaries which allow applications to override the default pronunciations of words. A pronunciation dictionary is a list of words along with their associated pronunciations stored in a resource of resource type 'dict'. The application is free to store dictionaries in either the resource fork or the data fork of a file. The application is responsible for loading the individual dictionaries into RAM and then passing a handle to the dictionary data to the Speech Manager. The initial release of the Speech Manager, however, does not include any routines that can add entries to dictionaries or manipulate them in other ways. The Speech Manager does include a routine, the <code>UseDictionary</code> function, that you can use to install one or more pronunciation dictionaries in a speech channel.


Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–38 <u>Listing 4-10</u> A sample pronunciation dictionary resource <pre data-bbox="768 285 1520 613"> resource 'dict' (1, "TestDict") { smRoman, langEnglish, verUS, ThisSecond, { pron, {tx, "ROOSEVELT", ph, "_1EHf_d1IY_1AAr"}, pron, {tx, "CHELSEA", ph, "_C1EHls2IY"}, pron, {tx, "AMHERST", ph, "_2UXmAKrst"}, pron, {tx, "REDSOX", ph, "_r1EHd_s1AAks"}, pron, {tx, "HALLOWEEN", ph, "_h1AA12OW_w1IYn"}, pron, {tx, "FELIX", ph, "_f1IY12IHks_D2UX_k1AEt"}, pron, {tx, "WEDNESDAY", ph, "_m1IHd_w1IYk"}, }, }; </pre> Note that you are not restricted to using pronunciations similar to those of the words listed. Typically, however, pronunciation dictionaries contain entries for words that the Speech Manager pronounces unsatisfactorily. Inside Macintosh Sound_1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ▪ generate speech and then pause or stop it ▪ obtain information about an individual voice or all voices ▪ create and dispose of speech channels ▪ obtain the Speech Manager's version and status ▪ change the rate or pitch of speech Inside Macintosh Sound_1994 at Page 4–55 ▪ convert textual into phonetic data ▪ install a pronunciation dictionary into a speech channel With the exception of the <code>SpeechManagerVersion</code>, <code>SpeechBusy</code>, and <code>SpeechBusySystemWide</code> functions, all Speech Manager routines return a result code to indicate whether an error has occurred. The section “Application-Defined Routines” beginning on page 4-82 describes the syntax and operation of application-defined callback procedures.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–80 <u>Installing a Pronunciation Dictionary</u> Pronunciation dictionaries allow your application to override the default Speech Manager pronunciations of individual words, such as names with quirky spellings. The <code>UseDictionary</code> function allows your application to load a pronunciation dictionary into a speech channel. iListen_15_New_speechrecogniti at Page 39 Most speech-recognition programs require a training period during which you acclimate the program to the unique qualities of your voice. With iListen, this process is easy. To begin, you set up the microphone, which in our case was the VX1 Parrott with Griffin iMic USB adapter (both included). Next, for about ten minutes, iListen leads you through sound-level tests and a sample passage of text. With initial training, iListen was about 78 percent accurate in our dictation tests in OS X and OS X's Classic mode. To improve recognition, you can read 11 other passages. Doing so brought the accuracy up slightly, to about 82 percent. The program can also analyze a set of your documents to learn the vocabulary you use. Kennewick also '209 discloses a dictionary and phrases module that allows users to incorporate special spellings, translations, and other indications of interpreting inputs, including substituting abbreviations. The following citations show examples of where Kennewick '209 discloses this limitation.

Claims

Prior Art

Kennewick '209 at FIG. 1



Kennewick '209 at 10:33–48

The one or more speech units 128 and the main unit 98 may communicate over a communication link. The communication link can include a wired or wireless link. According to one embodiment, the communication link comprises an RF link. The transceiver 130 on the speech unit communicates coded speech data bi-directionally over the communication link with the transceiver 126 on the main unit 98. According to another embodiment, RF link may use any standard local area wireless data protocols including the IEEE 802.11, Bluetooth or other standards. Alternatively, an infrared data link conforming to any suitable standard such as IrDA or other infrared standards can be used. In an alternative embodiment, wires connect the speech unit 128 and the main unit 98, eliminating the need for one speech coder 138. Other wired or wireless analog or digital transmission techniques can be used.

Kennewick '209 at 10:65–11:10

The agents 106 return results of questions as responses to users. The response may be created using the results of information queries, the system personality 108 and the user preferences or other data in the user profile 110. The agents 106 generally present these results using the speech unit 128. The agents 106 create a response string, which is sent to the text to speech engine 124. The text to speech engine 124

Claims	Prior Art
	creates the required utterances, which may be encoded and compressed by the speech coder 122. Once coded, the utterances are transmitted from the main unit 98 by the transceiver 126 to the transceiver 130 on the speech unit 128. The utterance is then decoded and decompressed by the speech coder 138 and output by the speaker 136. Kennewick '209 at 15:25–31 According to an embodiment of the invention, the microphone array, filters and speech coder 138 are physically separated from the main unit 98 into a speech unit 128, and connected using a wireless link. Since bandwidth on a wireless connection is at a premium, the speech coder dynamically adapts the digitization rate and compression of the captured speech. Kennewick '209 at 16:17–23 In an alternative embodiment, the array microphone can be replaced by a set of physically distributed microphones or a headset worn by the user. The distributed microphones can be placed in different parts of a room or in different rooms. The distributed microphones can create a three-dimensional array to improve signal to noise ratio. The headset may use a wireless or wired connection. Ewer '149 similarly discloses communications between a headset and communication device such as a phone. The following citations show examples of where Ewer '149 discloses this limitation. Ewer '149 at Abstract A system and method for automatic detection and signaling of the connection status of a headset used with telephony or other multimedia application software running on processor-based hosts using digital signaling are disclosed. The signaling system includes a signaling module for communicating with the host and a headset selectively in communication with the signaling module for use with the application software. The signaling module monitors the headset connection status, generates a status signal, and transmits the status signal to the host for determining the state of host and/or the application software. The signaling module may include a connection status detector and a status signal generating processor. A signal indicating that the headset is disconnected may pause or terminate an executing application or prevent the application from being executed. A signal indicating that the headset is connected may resume a paused application or allow the application to be executed. Ewer '149 at [0027] In addition, although wired connections are typically and preferably employed, such as between the USB

Claims	Prior Art
	headset adapter 106 and the host 104 and between the adapter 106 and the headset 110, wireless connections may alternatively be employed. For example, the headset 110 may be a cordless headset in wireless communication with the adapter 106, e.g., using RF technology. The headset 110 can be selectively powered on or off and thus be selectively in communication with the adapter 106. Thus, the term “connection” utilized herein generally refers to both wired and wireless connections. Jabra FreeSpeak similarly discloses communications between a headset and communication device such as a phone. The following citations show examples of where Jabra FreeSpeak discloses this limitation. Jabra FreeSpeak Datasheet at Page 1 Introducing FreeSpeak™— cord-free comfort and hands-free convenience in a Bluetooth® headset. JABRA FreeSpeak is the ultimate hands-free and wire-free solution that connects your mobile phone to your headset via Bluetooth radio technology instead of a wire. FreeSpeak works up to 30 feet away from your mobile phone – talk freely even if your phone is inside a pocket or across the room. All calls are transmitted wirelessly to your earpiece. FreeSpeak gives you easy, intuitive access to volume controls, and features visible and audible status indicators. Jabra FreeSpeak Datasheet at Page 2 Headset and Adapter for Non-Bluetooth Microsoft-Powered Pocket PC Phones and Smartphones* 100-91130000-02 • JABRA FreeSpeak headset • Multi-adapter with modified 2.5mm, 4-conductor audio plug (charger/holder/Bluetooth adapter with belt clip) • AC power adapter (115 VAC, 60Hz) • 3 one-size-fits-all MiniGels • 1 stylus • Illustrated User Guide in 3 languages: English, Spanish, French • Quick Start Guide 


Claims	Prior Art
	Jabra FreeSpeak Datasheet at Page 2 About Bluetooth Bluetooth is a radio technology developed to deliver short-range wireless mobility. Bluetooth technology eliminates wires between devices, such as mobile phones and headsets, and enables communication between devices that are in close proximity to one another. Bluetooth facilitates fast, secure transmissions of both voice and data. More information is available at www.bluetooth.com.  Jabra FreeSpeak Datasheet at Page 1 Bluetooth adapter that also functions as a charger and holder. The multi-adapter plugs directly into your non-Bluetooth phone to enable Bluetooth wireless connectivity. JABRA provides versions of the multi-adapter for use with Nokia phones, phones with a 2.5mm jack (connection varies with phone model), and many Pocket PC phones and Smartphones.* A POSITA would have been motivated to combine Apple Prior Art Systems with the teachings of Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak to meet the claim limitations with reasonable expectation of success for the reasons described above with respect to claim 1[pre]-[1h]. For example, a POSITA would be motivated to implement a wireless headset as disclosed by Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak to the Apple Prior Art Systems, and a POSITA would have a reasonable expectation of success in implementing a wireless headset to the Apple Prior Art Systems based on Apple Prior Art Systems’ compatibility with wired headsets that include speakers for presenting audio. <i>See</i> https://support.apple.com/en-us/112450 (Power Mac G4 Technical Specifications). Moreover, a POSITA would have recognized that it was known to convert or expand messages into a plurality of words and would

Claims	Prior Art
	have found it obvious to do so in the Apple Prior Art Systems. Such a conversion or expansion would have been nothing more than the application of known methods to prior art systems to yield the predictable and desired result of a system that provided accurate and understandable audio messages to the user. To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
42[pre]. A wireless headset that at least supports messaging for a first user, said wireless headset operatively connects to a mobile electronic device having wireless communication capability, and wherein said wireless headset is configured to at least:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless headset that at least supports messaging for a first user, said wireless headset operatively connects to a mobile electronic device having wireless communication capability, and wherein said wireless headset is configured to at least.” <i>See above at claim 23[pre].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[42a] play an incoming audio message for the first user via the wireless headset, the incoming audio message being produced by converting an incoming text message from a second user, the incoming text message being received at the mobile electronic device via a wireless network; and	Apple Prior Art Systems expressly or inherently disclose “play an incoming audio message for the first user via the wireless headset, the incoming audio message being produced by converting an incoming text message from a second user, the incoming text message being received at the mobile electronic device via a wireless network.” <i>See above claim [23a], [23b].</i> Apple Prior Art Systems, including the Communication Systems running or combined with the Speech Recognition Technologies and Speech Synthesis Technologies, work with devices compatible with messaging functionality such as email. Apple Prior Art Systems can further “read” (e.g., convert the text to audio) the emails to a user. The following citations show examples of where the Apple Prior Art Systems disclose this limitation. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Mobility Suite gives you voice access to your PIM applications. It allows you to use simple, natural voice commands such as "Go to Contacts" and "Read John Smith's home phone number". When you "Go to Mail", you can have all of your mail read to you, or you can read only selected messages. In the "Calendar" function, you can check your calendar for your next appointment, or


Claims	Prior Art				
	add new ones. You can easily keep track of your busy schedule, all by using voice commands. In addition, this IBM software gives you the added convenience of being able to listen to rather than read information and messages. IBM ViaVoice Mobility Suite v.4 at 1 End-to-end voice solutions IBM can offer distributed end-to-end voice solutions where some of the speech processing resides on the client device (cell phones etc.) and integrates with additional voice functionality at the server level and large-scale telephony/unified messaging systems. IBM ViaVoice Mobility Suite v.4 at 1 IBM Embedded ViaVoice Enterprise Edition is IBM's flagship speech recognition and text-to-speech product. Representing a dramatic shift in the evolution of embedded speech products, IBM Embedded ViaVoice Enterprise Edition has an unlimited active vocabulary, limited only by the amount of memory on devices, plus a dynamic and swappable 4,000 word phonetic flat list. This feature enables embedded applications to access voice-enabled databases with millions of records. A typical application would be a voice-enabled navigation database of street names for a GPS system in your car. IBM ViaVoice Mobility Suite v.4 at Page 2 <table border="1" data-bbox="636 902 1318 1162"> <thead> <tr> <th data-bbox="636 954 919 987">Inbox (Mail) Commands</th><th data-bbox="919 954 1318 987">Sample Commands</th></tr> </thead> <tbody> <tr> <td data-bbox="636 1003 919 1076">Reads incoming mail; gives you the opportunity to reply and send responses to others.</td><td data-bbox="919 1003 1318 1141"> • "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply" </td></tr> </tbody> </table> IBM Mobility Suite Introduction at 3 Welcome to IBM Embedded ViaVoice Mobility Suite, a productivity application that controls your Personal Digital Assistant (PDA) by voice. In response to your commands, IBM Embedded ViaVoice Mobility Suite will help manage your appointments, contacts, tasks, and mail.	Inbox (Mail) Commands	Sample Commands	Reads incoming mail; gives you the opportunity to reply and send responses to others.	• "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply"
Inbox (Mail) Commands	Sample Commands				
Reads incoming mail; gives you the opportunity to reply and send responses to others.	• "Read mail" • "Read new mail" • "Next message" • "Create reply" • "Send reply"				

Claims	Prior Art
	IBM Mobility Suite Introduction at 5 • To hear important information, for example, say Read mail or Read my very next appointment. • To hear prompts, messages, or requested information again, say Repeat that. IBM Mobility Suite Introduction at 10 10.0 Mail Commands Tips for Mail Commands • Say go to inbox or go to mail to enable mail commands. • The default working list is the list of new (unread) messages in the inbox. The first message in this list is the default message. • read [new] (messages mail) - Reads a list of messages from the inbox, numbering them as message 1, 2, ... n. • [read] (next previous) message - Reads the details of the message before or after the current message. • [read] message <1-99> - Reads details for a message number in the working list. • delete message [<1-99>] - Deletes the current message or a specified message from the working list. This command requires confirmation. • create reply [to all] [with history] - Prompts you to record an attachment (by pressing and holding the Record button) as a reply to the current message. The to all option addresses the reply to the sender and all recipients of the original message. The with history option includes the body from the original message in the reply. After recording, the reply becomes the new working message and is inserted into the working list following the original message. • read reply - Reads the most recent reply. A reply is no longer available after it is sent. • send reply - Sends a reply message and moves it from the inbox to the outbox. IBM ViaVoice Simply Dictation user guide at Page 26 Reading Text Now that you have learned to enter and format text using voice commands, it's time to have ViaVoice read the sentences back to you. Highlight the desired text block and say one of the following voice commands. SAY> BEGIN READING SAY> READ SELECTION ViaVoice reads the highlighted text using the voice selected in the Voices panel in ViaVoice Preferences. On applications with copy and selection capabilities, the Read Selection menu item reads the text using the voice currently selected in the Voices panel in ViaVoice Preferences.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 47 Reading E-mail Provided your e-mail application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back the text in your e-mail. 1. Open your e-mail application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected voice in the Voices panel in ViaVoice Preferences. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 48 Reading Your Text Provided your application supports the <i>Select All</i> and <i>Copy to Clipboard</i> functions, you can have ViaVoice read back text aloud from almost any source, including documents and e-mail messages. 1. Open the application. 2. SAY> READ ALL or, highlight the text you want to read and SAY> READ SELECTION. The text selected is read using the currently selected system voice. 3. Turn on the microphone and SAY> STOP READING to stop playback at any time. IBM ViaVoice Simply Dictation user guide at Page 49 Changing Voices  You can select the voice used when reading text. 1. Click the ViaVoice Preferences button on the VoiceCenter drawer. 2. Click the Voices icon. 3. From the Voice menu, select a new voice. 4. Press Play to hear a sample of your selection. 5. Click Apply, and close the window. 6. Any text read back from this point uses the voice you just selected.

Claims	Prior Art
	IBM ViaVoice Simply Dictation user guide at Page 61  Inside Macintosh Sound_1994 at Page 4–3 This chapter describes the Speech Manager, the part of the Macintosh system software that provides a standardized method for Macintosh applications to generate synthesized speech. Inside Macintosh Sound_1994 at Page 4–3 You need to read this chapter if you want your application to be able to generate speech. For example, you may want your application to incorporate the capability to speak its dialog box messages to the user. A word-processing application might use the Speech Manager to implement a command that speaks a selected section of a document to the user. A multimedia application might use the Speech Manager to provide a narration of a QuickTime movie instead of including sampled-sound data on a movie track. Because sound samples can take up large amounts of room on disk, using text in place of sampled sound is extremely efficient.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–3 This chapter begins with an introduction to the speech generation process and then discusses how you can ▪ check for the availability of the Speech Manager ▪ create and dispose of speech channels ▪ generate speech with different voices ▪ obtain information about and change speech channel settings ▪ start and stop speech production ▪ synchronize speech production with other activities by using callback procedures ▪ embed Speech Manager commands within text to make it more understandable Inside Macintosh Sound_ 1994 at Page 4–4 You can use the Speech Manager to incorporate synthesized speech into your application. This section provides an overview and describes the basic concepts of the Speech Manager, and it outlines the process that the Speech Manager uses to convert text into speech. The Speech Manager converts text into sound data, which it passes to the Sound Manager to play through the current sound output device. The Speech Manager's interaction with the Sound Manager is transparent to your application, so you don't need to be familiar with the Sound Manager to take advantage of the Speech Manager's capabilities. Inside Macintosh Sound_ 1994 at Page 4–4 Figure 4-1 illustrates the speech generation process. Your application can initiate speech generation by passing a string or a buffer of text to the Speech Manager. The Speech Manager is responsible for sending the text to a speech synthesizer, a component that contains executable code that manages all communication between the Speech Manager and the Sound Manager. A synthesizer is usually contained in a resource in a file within the System Folder. A synthesizer is like a speech engine. It uses built-in dictionaries and pronunciation rules to help determine how to pronounce text. You can provide custom pronunciation dictionaries as described in the section “Including Pronunciation Dictionaries” beginning on page 4-36.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4-4 <u>Figure 4-1</u> The speech generation process  <pre> graph LR Text[Text] --> SM[Speech Manager] SM --> SS[Speech synthesizer] SS --> SD[Sound Manager] SD --> AH[Audio hardware] </pre> Inside Macintosh Sound_1994 at Page 4-5 <u>Voices</u> Your application can use the system default voice to generate speech or it can specify that the Speech Manager use a particular voice that is available on the current computer system. A voice is a set of characteristics defined in parameters that specify a particular quality of speech. Just as different people's voices have different tonal qualities, so too can different voices have different qualities. A synthesized voice might sound male or female and might sound like an adult or a child. Some voices sound distinctively synthetic, while others sound more like real people. Figure 4-2 shows how the Speech Manager uses speech channels to synthesize speech with different voices. Inside Macintosh Sound_1994 at Page 4-6 Most of the time, your application designates the voice that speaks text, and usually that is the default voice. Based on the needs of your users and the way in which you expect them to use voices in your application you can provide access to voices in a number of different ways. You could include access to selecting voices in a dialog box that is available from a menu item such as Voices... Any application that allows users to choose among voices requires additional information about the available voices beyond the information provided by a voice specification record (described in detail on page 4-46), whose data should never be presented to the user. Such additional information might include the name of the voice as well as what script and language it supports.

Claims	Prior Art
	Inside Macintosh Sound_1994 at Page 4–11 To take advantage of most of the Speech Manager's features, you must allocate a speech channel to pass to Speech Manager functions and dispose of the speech channel when you are finished using it. "Creating, Using, and Disposing of a Speech Channel" beginning on page 4-13 demonstrates how you do this and shows how you can use the <code>SpeakText</code> function to start speech generation from a buffer of text. Some applications permit users to choose a voice from those available to be used for speech generation. The <code>CountVoices</code>, <code>GetIndVoice</code>, and <code>GetVoiceDescription</code> functions support this capability. "Working With Different Voices" beginning on page 4-14 shows how you can use these functions to choose among available voices. Inside Macintosh Sound_1994 at Page 4–14 The <code>MyUseSpeechChannel</code> function defined in Listing 4-2 creates a default speech channel using the default system voice. You pass <code>NIL</code> in the first parameter to use the system default voice. You must also pass a global variable to <code>NewSpeechChannel</code> in which is returned a valid speech channel. Once the channel exists, then you can use the <code>SpeakText</code> function to generate speech. To generate synthesized speech, you pass in the channel allocated by <code>NewSpeechChannel</code> in the first parameter, and then you pass a pointer to the text that you want to speak as well as the length of the text that you want the Speech Manager to attempt to speak. That is, you can pass a pointer to a buffer of text that is 500 bytes long, but specify that only the first 10 bytes get spoken. Then <code>MyUseSpeechChannel</code> uses the <code>SpeechBusy</code> function in a <code>WHILE</code> loop to allow the text to be completely spoken before disposing of the channel. Inside Macintosh Sound_1994 at Page 4–14 <u>Working With Different Voices</u> When you work with speech channels, you can set a voice for a particular channel. When you set a voice, you may want to filter out certain of its characteristics in order to identify the one you want. For example, in an educational software application for elementary school students, you may want to use only children's voices. In order to choose the voice you want, you get a voice description record that contains information about a voice such as the size of the voice, the name of the voice, the age and gender of the voice, and the synthesizer with which it works. You can get the number of available voices using the <code>CountVoices</code> function. You can cycle through the available voices and identify the one you want to use by using the <code>GetIndVoice</code> function. Then you fill out a voice description record using the <code>GetVoiceDescription</code> function. Listing 4-3 shows how to get identifying information about a voice.

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–23 <u>Writing Embedded Speech Commands</u> Embedded speech commands allow you to customize the quality of speech output by fine tuning it. You can make speech much easier to understand than the default way in which text is spoken by a synthesizer. An embedded speech command is a command embedded within a text buffer to be spoken by the Speech Manager that causes the Speech Manager to take a certain action. For example, you could use an embedded speech command to emphasize a particular word in a text string to make it stand out to the user. Inside Macintosh Sound_ 1994 at Page 4–30 <u>Examples of Embedded Speech Commands</u> If you use just a few of the embedded speech commands, you can markedly increase the understandability of text spoken by your application. Your application knows more about the speech being produced than a speech synthesizer does. A synthesizer speaks text according to a predetermined set of rules about language production. Therefore, the voices available on a Macintosh computer with the Speech Manager installed sound very synthetic and sometimes robotic because the pronunciation rules are formalized. You can make the speech produced by the synthesizer sound a lot more human by observing some simple rules of human speech and embedding speech commands in text according to these conventions. The techniques presented in this section could be applied when your application is having a dialog with the user or speaking some error messages or announcements. Inside Macintosh Sound_ 1994 at Page 4–54 <u>Speech Manager Routines</u> This section describes the routines provided by the Speech Manager. You can use these routines to ■ generate speech and then pause or stop it ■ obtain information about an individual voice or all voices ■ create and dispose of speech channels ■ obtain the Speech Manager's version and status ■ change the rate or pitch of speech

Claims	Prior Art
	Inside Macintosh Sound_ 1994 at Page 4–73 <u>Changing Speech Attributes</u> To determine the rate and pitch at which a speech channel is processing text, you can use the <code>GetSpeechRate</code> and <code>GetSpeechPitch</code> functions. The <code>SetSpeechRate</code> and <code>SetSpeechPitch</code> functions allow you to change rate and pitch. The most robust of the Speech Manager's routines are the <code>GetSpeechInfo</code> and <code>SetSpeechInfo</code> functions. These allow you to obtain many types of information about a speech channel and to change many settings of a speech channel. To specify the operation that you wish to perform, you must pass <code>GetSpeechInfo</code> or <code>SetSpeechInfo</code> a selector. A full list of selectors is provided in "Speech Information Selectors" beginning on page 4-39. Inside Macintosh Sound_ 1994 at Page 4–76 <u>SetSpeechPitch</u> You can use the <code>SetSpeechPitch</code> function to set the speech pitch on a designated speech channel. <pre>FUNCTION SetSpeechPitch (chan: SpeechChannel; pitch: Fixed) : OSErr;</pre> <code>chan</code> The speech channel whose pitch you wish to set. <code>pitch</code> The new pitch for the speech channel, expressed as a fixed-point frequency value. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[42b] receive a reply audio message from the first user via the wireless headset in response to the incoming audio message, the reply audio message thereafter being converted to a reply text message and transmitted to the second user via the wireless network.	Apple Prior Art Systems expressly or inherently disclose “receive a reply audio message from the first user via the wireless headset in response to the incoming audio message, the reply audio message thereafter being converted to a reply text message and transmitted to the second user via the wireless network.” See above at claims [23c], [23d], 24. To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.

Claims	Prior Art
43[pre]. A wireless headset as recited in claim 42,	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless headset as recited in claim 42.” <i>See above at claim [1a].</i> To the extent that Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[43a] wherein the conversion of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic, and	Apple Prior Art Systems expressly or inherently disclose “wherein the conversion of the incoming text message to the incoming audio message depends on at least one user-selected speech characteristic.” <i>See above at claim 11.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[43b] wherein the at least one user-selected speech characteristic is set based on at least one user setting able to be provided at the mobile electronic device.	Apple Prior Art Systems and Kennewick ‘209 expressly or inherently disclose “wherein the at least one user-selected speech characteristic is set based on at least one user setting able to be provided at the mobile electronic device.” <i>See above at claim 9.</i> To the extent that Apple Prior Art Systems and Kennewick ‘209 do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems and Kennewick ‘209, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
44[pre]. A wireless headset as recited in claim 43, wherein said wireless headset is configured to:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless headset as recited in claim 43, wherein said wireless headset is configured to.” <i>See above at claim [1a].</i> To the extent that Apple Prior Art Systems, Kennewick ‘209, Ewer ‘149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art

Claims	Prior Art
	to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[44a] prompt the first user to provide the reply audio message prior to the reply audio message being received at said wireless headset,	Apple Prior Art Systems expressly or inherently disclose “prompt the first user to provide the reply audio message prior to the reply audio message being received at said wireless headset.” <i>See above at claim [3a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
45. A wireless headset as recited in claim 43, wherein at least a portion of the incoming text message is in a chat language, and wherein the conversion of the incoming text message to the incoming audio message comprises expanding the portion of the incoming text message that is in the chat language to a plurality of words.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset as recited in claim 43, wherein at least a portion of the incoming text message is in a chat language, and wherein the conversion of the incoming text message to the incoming audio message comprises expanding the portion of the incoming text message that is in the chat language to a plurality of words.” <i>See above at claim [18a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
46[pre]. A wireless headset as recited in claim 43, wherein the conversion of the incoming text message to the incoming audio message comprises:	To the extent the preamble is limiting, Apple Prior Art Systems, Kennewick '209, Ewer '149, and/or Jabra FreeSpeak expressly or inherently disclose “[a] wireless headset as recited in claim 43, wherein the conversion of the incoming text message to the incoming audio message comprises.” <i>See above at claim [1a].</i> To the extent that Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, Kennewick '209, Ewer '149, and Jabra FreeSpeak, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.

Claims	Prior Art
[46a] translation at least a portion of the incoming text message from a condensed format to an uncondensed text format; and	Apple Prior Art Systems expressly or inherently disclose “translation at least a portion of the incoming text message from a condensed format to an uncondensed text format.” <i>See above at claim [18a].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
[46b] conversion of at least the portion of the incoming text message in the uncondensed text format to at least a portion of the incoming audio message.	Apple Prior Art Systems expressly or inherently disclose “conversion of at least the portion of the incoming text message in the uncondensed text format to at least a portion of the incoming audio message.” <i>See above at claim [1d].</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.
47. A wireless headset as recited in claim 43, wherein the reply audio message is converted to the reply text message and transmitted to the another wireless electronic device by the mobile electronic device.	Apple Prior Art Systems expressly or inherently disclose “[a] wireless headset as recited in claim 43, wherein the reply audio message is converted to the reply text message and transmitted to the another wireless electronic device by the mobile electronic device.” <i>See above at claim 24.</i> To the extent that Apple Prior Art Systems do not disclose any elements of this claim limitation, it would have been obvious to a person having ordinary skill in the art to combine, with the teachings of Apple Prior Art Systems, those elements, as known to them or as disclosed by other prior art identified in the cover pleading.