BLUETOOTH and wireless networking

a primer for audio engineers

Wireless networking, using radio frequency links instead of wires, is becoming increasingly popular as a means of connecting digital devices and computers. It provides greater flexibility and mobility to the user and does away with those tangled strands of wires. And wireless networks can be integrated with wired systems, thereby acting as extensions of existing networks. Audio information is increasingly transferred over networks instead of dedicated digital interfaces or analog cabling, making it important that audio engineers have an appreciation of the issues involved.

Bluetooth is a technology for wireless personal area networks (WPANs) that has been widely publicized recently, but it is only one of a number of options for wireless data communication. This article puts Bluetooth into context and describes wireless data communications in relation to audio systems. This field is changing fast like most aspects of computing, making it hard for standardization to keep pace with technology development. A number of wireless networking standards have been or are being standardized by the IEEE, but alternative solutions also exist that overlap with or differ from IEEE standards. As a rule it seems that IEEE has attempted where possible to adopt commercial technology in response to calls for suitable solutions.

WIRELESS PERSONAL AREA NETWORKS

WPANs connect desktop digital devices, usually over short distances within a so-called personal operating space of 10 m enveloping the person. WPANs are intended mainly for indoor and fixed outdoor applications and operate at data rates of up to about 1 Mbit/s. They are a useful next step from infrared communications. Bluetooth is the prime example of a WPAN. Home RF is potentially similar, but it sits somewhere between a PAN and a LAN (local area network). Fig. 1 shows an example of the interrelationship between PANs, LANs, and WANs. WPANs are broadly covered in IEEE standards work by the 802.15 groups. The primary concerns when designing successful WPAN technology are power consumption, simplicity, and product size. ➥

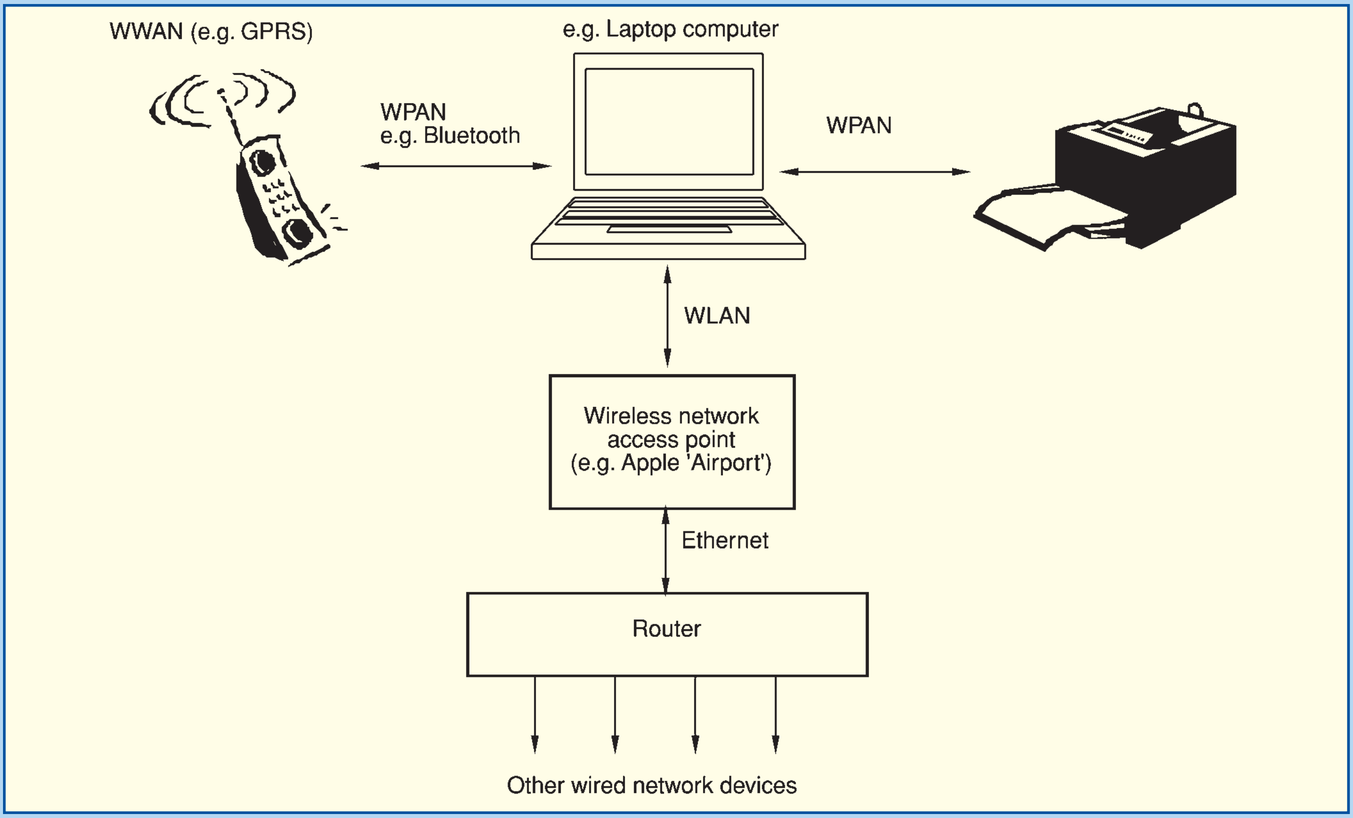

Fig. 1. Example of WPAN, WLAN, and WWAN interacting.

WIRELESS LOCAL AREA NETWORKS

WLANs are used over larger distances, up to 500 m, in areas where people may move around with digital devices but usually remain stationary while using the devices. They typically operate around hot-spot access points, for example in offices, homes, and schools. Data rates are typically higher than WPANs, between 2 and about 50 Mbit/s. WLANs are covered in IEEE standards activity by the 802.11 groups. Typically above this data rate are WMANs (wireless metropolitan area networks) that interconnect LANs over large regions.

WIRELESS WIDE AREA NETWORKS

WWANs have more to do with global telecommunications systems, for example mobile phone operators using GPRS (General Packet Radio Service) equipment, and are mainly designed for applications in which people move around while using devices. They are not featured in this article as they have little to do with potential audio applications in consumer and professional environments.

General radio frequency issues

Wireless networks typically use low-power, spread-spectrum, and frequency-hopping techniques to ensure robust communications in the face of interference and fading. They typically use coding techniques to make the data signal look like noise or very short-term impulsive interference to a narrow-band receiver in the radio frequency (RF) domain, as shown in Fig. 2. Most current systems operate in a license-free region of the RF spectrum between 2.4 and 5 GHz, relying principally on the low-power and interference-avoiding mechanisms of systems to operate satisfactorily in this free-for-all region.

The 2.4-GHz band is becoming increasingly crowded owing to Bluetooth devices, wireless networks, and telephones, so there can be significant RF interference in this band. The 5-GHz band is somewhat quieter in the United States, but signals have a much shorter range at this frequency and are more easily obstructed by walls, doors, and other objects. Communication therefore tends to be over line-of-sight paths or short distances. Some chipsets are appearing that will operate in both bands for IEEE 802.11 standards. The issue of interference is covered in greater detail later in this article.

EXAMPLES OF WPAN TECHNOLOGY

Bluetooth operates in the 2.4-GHz band at data rates up to 1 Mbit/s over distances of up to 10 m or 100 m depending on the implementation. A relatively simple binary FM modulation method is used because Bluetooth has to be implementable in basic devices. Communication between devices can be either point-to-point or point-to-multipoint. In the latter case one device acts as master, synchronizing, controlling, and sharing the channel with as many as seven slaves in a so-called piconet of up to eight devices. Frequency hopping between 23 or 79 RF channels is employed to increase robustness, and data packets are each transmitted on a different hop frequency. The IEEE 802.15.1 WPAN standard is based on Bluetooth.

The primary modes for data communication are ACL (asynchronous connectionless) links and SCO (synchronous connection-oriented) links. The latter are limited to a data rate of 64 kbit/s each, and up to three concurrent SCO links are allowed per Bluetooth channel. They are set up between a master and a single Bluetooth device at a time for time-critical purposes such as voice audio. SCO links are established as real-time links between two points for the duration of the connection. Packets are never retransmitted, owing to their real-time nature. In the time remaining within each transmission slot the master can set up an asynchronous communication with one or more slaves. ACL communications operate rather like conventional packet-switched networks, with the packet header being used to indicate the destination on the piconet.

It is possible to operate either one asynchronous and three synchronous voice channels (bidirectional) or one channel that handles both asynchronous data and synchronous voice

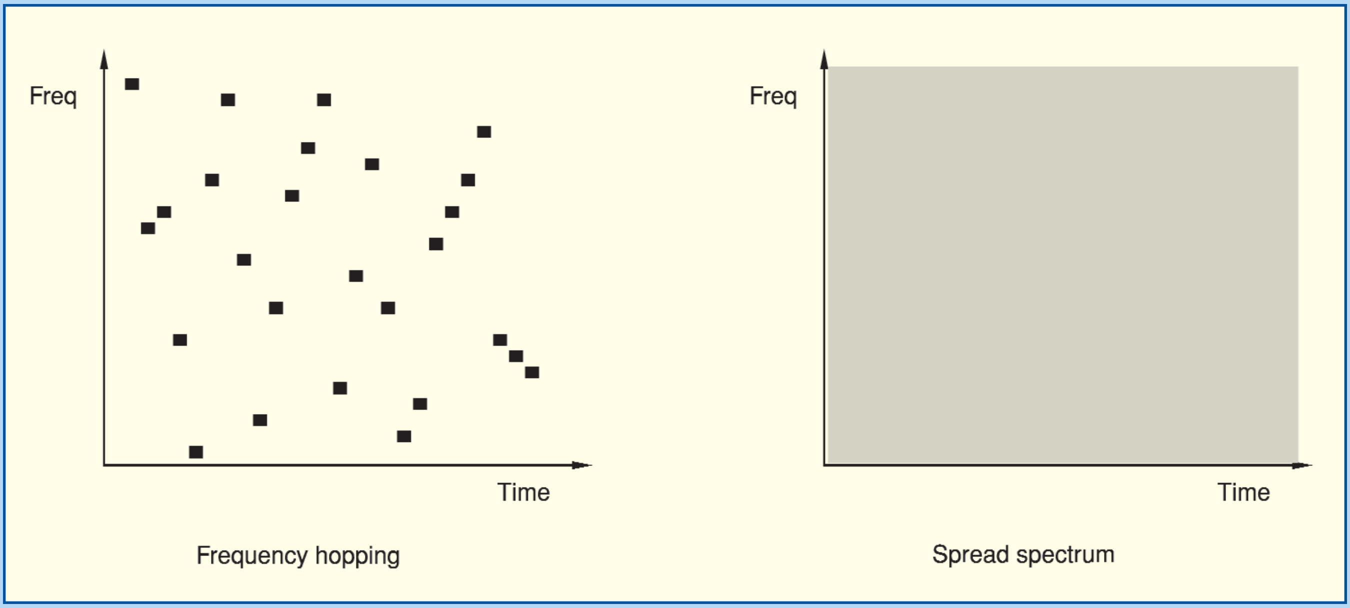

Fig. 2. Wireless networks use low-power, spread-spectrum, and frequency-hopping techniques for robust communications.

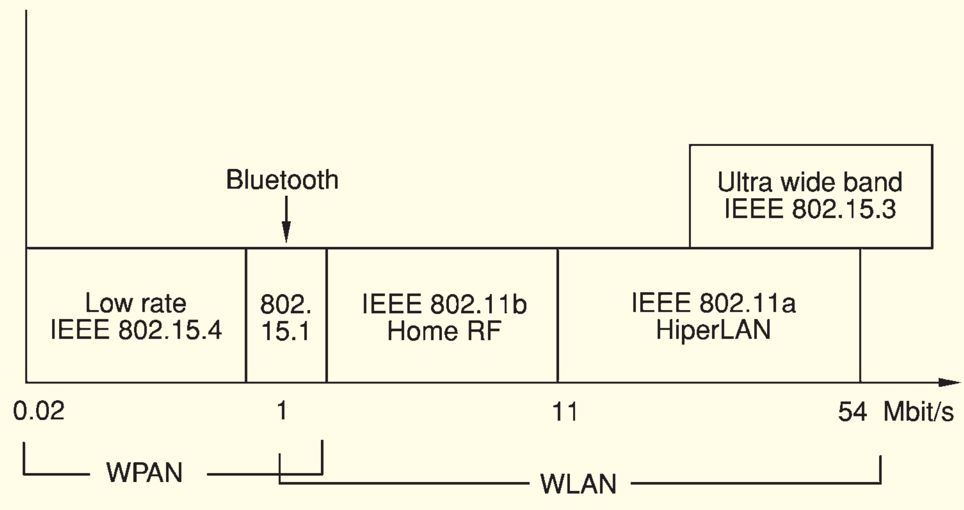

Fig. 3. Data rates of common wireless technologies.

(this is essentially made up of single SCO packets with two parts). The maximum asynchronous data rate is 723.2 kbit/s in one direction, with a return rate of 57.6 kbit/s, or 433.3 kbit/s symmetrically.

In addition to the Bluetooth Core Specification that describes the principles of communication, the Bluetooth SIG (special interest group) publishes a number of profile documents that describe specific approaches to data transfer for particular applications. The documents of primary interest to audio engineers are the Generic Audio/Video Distribution Profile (GAVDP) and the Advanced Audio Distribution Profile (A2DP). These are discussed later in this article.

The long-term future of Bluetooth is not certain, but the same can be said of virtually any data communications technology today. For example, Apple and a number of mobile-telephone manufacturers have significant investment in Bluetooth, and this may help to ensure its future.

Many commentators appear enthusiastic about the market potential for the technology. It is a relatively low-rate system but is intended for low-cost, low-power implementation in highly portable devices. Some commentators cite the relative advantages of broadband technology such as UWB (ultrawide band) that may ultimately become more useful in high-rate multimedia environments. Some concern has also been expressed about the interoperability of Bluetooth devices, as peer-to-peer networking interoperability is apparently not mandatory for device certification. Success may therefore come down to a question of which Bluetooth devices actually work together and at what level of sophistication.

High-rate and ultrawide band systems

The IEEE 802.15.3 group is studying high-rate applications above 20 Mbit/s. These could ultimately be up to 100 times faster than Bluetooth. UWB is also being developed for applications requiring very high data rates. UWB uses very wide band transmissions at extremely low power, requiring receivers and transmitters to be tightly synchronized. The U.S. Federal Communications Commission (FCC) has approved UWB within certain low-power limits, but other countries have not yet done the same. IEEE 802.15 Study Group 3a is studying alternative physical layer options for high-rate WPANs.

Low-rate, low-power consumption systems

IEEE 802.15.4 is devising a lower rate standard for simple devices, such as joysticks, that do not need high data rates but need small size and long battery life. Data rates of 20, 40, and 250 kbit/s are targeted for use within RF bands at 2.4 GHz, 915, and 868 MHz. This clearly lies below the typical Bluetooth application range of data rates, whereas high rate and UWB lie above the Bluetooth range, as shown in Fig. 3.

EXAMPLES OF WLAN TECHNOLOGY

Broadband WLAN technology is developing so fast that the regulatory and commercial situation changes almost weekly. However, the following is an attempt to summarize the current situation as clearly as possible.

IEEE 802.11

IEEE 802.11a and IEEE 802.11b are wireless LAN standards and the basis of wireless Ethernet or WiFi. WiFi is a mark for compatible products awarded by the Wireless Ethernet Compatibility Alliance (WECA). Apple uses this technology in its AirPort products, as do many of the consumer wireless network systems currently marketed in computer stores.

The two standards define different physical layers, where 802.11b transmits at 2.4 GHz and 802.11a transmits at 5 GHz. The standards do not interoperate directly with each other, although bridges between them are possible and chipsets are appearing that operate in either mode. 802.11 deals with the physical and MAC (medium access control) layers of networking protocol, but the remaining layers are virtually identical to Ethernet, which is IEEE 802.3. In network terminology layers are different levels in the ➥

USEFUL WEB SITES

Bluetooth specifications:
www.bluetooth.com/dev/specifications.asp

HiperLAN:
www.hiperlan2.com

HomeRF:
www.homerf.org

IEEE 802.11:
grouper.ieee.org/groups/802/11/

IEEE 802.15:
grouper.ieee.org/groups/802/15/

Ultra wide band:
www.uwb.org

Wireless networking links:
www.wireless-nets.com/links.htm

somewhat arcane hierarchy of data communications, with the application at the top and the physical medium at the bottom. As with cell phones, users can roam between wireless access points (APs), using the beacon of the AP as a judge of signal strength. Collision avoidance and carrier detection are used for medium access, rather like Ethernet, and authentication or full encryption can be used for security purposes.

Both FHSS (frequency hopping spread spectrum) and DSSS (direct sequence spread spectrum) approaches are specified for 1- and 2-Mbit/s rates, but only DSSS is allowed for 11-Mbit/s communications. 5.5 Mbit/s is also allowed and even higher bit-rate standards were approved in 1999. In fact it appears that most developers actually went for DSSS even at the lower rates, for future compatibility with the higher rates. Generally, the higher the radio frequency (and therefore the greater the potential for high data rates), the shorter the range that may be covered. For example the 802.11a standard can operate at 54 Mbit/s but only within a 50-meter range. Alternatively, 802.11b operates at a maximum speed of 11 Mbit/s but can cover 300 meters outdoors (100 meters indoors). As with most packet-switched RF networks, contention, interference, and overheads can reduce the realistic data rate closer to half the maximum.

Contrary to what may have been anticipated, 802.11b is already in wide use but 802.11a products are only just beginning to appear. In the UK an internal agreement is being formulated to allow 802.11a to operate in a limited fashion in the region from 5.15 to 5.25 GHz. This allows four access points as opposed to the eight possible if the system occupies the full band from 5.15 to 5.35 GHz. Full approval from the European Telecommunications Standards Institute (ETSI) is not yet finalized because of military and government use of the 5-GHz band. But other countries are striking individual agreements rather like the UK. Dynamic frequency selection (DFS) and transmission power control (TPC) have been required for conformity with European requirements, and these are now virtually completed for 802.11a (this project was originally coded 802.11h). However, as discussed in the next section, full ETSI approval of wireless networking in the 5-GHz band does appear to be feasible, and this band is apparently being cleared in many countries for broadband data communications.

The range of 802.11 activities is wide and fast moving. See http://grouper.ieee.org/groups/802/11/ for the latest information.

HiperLAN

HiperLAN operates in the 5-GHz band up to 54 Mbit/s and requires approximately 330 MHz of bandwidth. These rates and frequencies are the same as 802.11a, and at the physical layer the two are almost identical. HiperLAN has been developed by ETSI and the Broadband Radio Access Network (BRAN), whose founding members are Tenovis (Bosch), Dell, Ericsson, Nokia, Telia, and Texas Instruments. Typically, the operating ranges are 30 m indoors and 150 m outdoors. In Europe specific bands from 5.15 to 5.35 GHz and 5.470 to 5.725 GHz seem to have been approved already for dedicated use by HiperLAN. In the U.S. the 5.15–5.35-GHz and 5.725–5.825-GHz bands are unlicensed and usable.

The main difference between HiperLAN and 802.11a is at the media access (MAC) level. Whereas 802.11a is an extension of wireless Ethernet, primarily based on contention mechanisms and packet switching, HiperLAN is claimed to offer connection-oriented communication and is compatible with other circuit-switched network standards such as ATM. As a result it offers quality-of-service (QoS) guarantees and therefore will be useful for real-time streaming applications.

HomeRF

HomeRF is yet another competing technology in the 2.4-GHz band. It offers a data rate of 10 Mbit/s and combines cordless telephone links (up to eight concurrently), wireless networking, and data streaming for home entertainment products.

Up to eight simultaneous real-time streaming sessions are possible, either two-way, multicast, or receive only. Priority is given to real-time voice communications over streaming applications as the default condition. Asynchronous packets take last place in the queue. The IEEE describes HomeRF as a sort of trimmed-down 802.11.

RF INTERFERENCE ISSUES

Interference is a particular problem in the crowded 2.4-GHz band, which is license free in most parts of the world. Industrial, scientific, and medical applications (ISM) use it, and microwave ovens are the primary culprits when it comes to interference.

In order to limit the effects of interference, frequency hopping spread spectrum (FHSS) techniques are used by Bluetooth and HomeRF. Hopping reduces channel efficiency in favor of robustness and is simpler to implement than the DSSS techniques described below. Such techniques usually employ at least 75 frequencies with a maximum dwell time per frequency of 400 ms. This ensures that communication does not spend too long on each frequency and the signal looks like random impulsive interference to a narrow-band receiver. Both transmitter and receiver follow the same pseudorandom pattern of hops, and adaptive hopping can avoid frequencies that are known to be blocked. Typical FHSS products occupy about 1 MHz of bandwidth. They achieve higher power within each frequency band at any one time and may therefore result in a better instantaneous RF signal-to-noise ratio than other spread-spectrum techniques.

The Bluetooth hopping rate is quite fast (1,600 times per second) compared with some other systems. Per second, in any one band, it is said that one is more likely to encounter an 802.11b transmission than a Bluetooth signal. There is some concern over the colocation of Bluetooth and 802.11b devices, as Bluetooth transmitters have been shown to reduce the performance of the other network considerably if placed close to a receiver. Bluetooth has limited error handling and retry mechanisms beyond the frequency-hopping physical layer. So it is not particularly sophisticated for robust wireless LANs, for which it was not designed, intended instead as a simple and cheap short-distance cable-replacement technology.

The direct sequence spread spectrum (DSSS) approach, as used by 802.11b for example, uses more bandwidth than strictly required by the data rate. Data are coded onto chips (simply a pattern of data bits) that have redundant portions spread over the frequencies in the band. The data are exclusively ordered within an 11-bit Barker code (a pseudorandom sequence), the bits of which make up the chips, and there are usually an integral number of chips per bit. These are modulated onto the carrier using differential phase-shift keying (DPSK), so a narrowband receiver perceives this signal as low-level noise. Even if some parts are lost during transmission owing to interference, error correction can be used to recover the data. The so-called spreading ratio is related to the degree of redundancy employed (number of chips per bit); the spreading ratio of most wireless LANs is around eight. This ratio makes the use of the band reasonably efficient at the expense of higher robustness. Sometimes the band is split between more than one network; a maximum of three networks is possible in typical current implementations. Typical current DSSS products occupy 20 to 22 MHz of bandwidth no matter what the data rate. The power within any one band is relatively low, making the RF signal potentially less robust than with frequency-hopping techniques, but performance in practice depends on the coding scheme and spreading ratio employed.

BLUETOOTH AUDIO

The basic audio functionality in the Bluetooth Core Specification is really only suitable for telecommunications applications. The bandwidth is typically limited to 4 kHz for telephony, and the sampling rate is correspondingly 8 kHz. Relevant standards are ITU-T P.71, G.711, 712, and 714. The 64-kbit/s voice channels use µlaw or A-law logarithmic PCM coding or CVSDM (continuous variable slope delta modulation). CVSDM is said to be preferable with respect to quality and robustness to errors. Audio error correction depends on which packet type carries the audio transmission: HV3 (High-quality Voice 3) packets have no forward error correction (FEC) and contain 3.75 ms of audio at 64 kbit/s, whereas HV1 and 2 packets have some error correction and contain shorter durations of audio, 1.25 and 2.5 ms respectively.

Bluetooth audio streaming using ACL packets

New draft profiles have been created for audio streaming using asynchronous (ACL) packets that can occupy the full remaining bit rate of Bluetooth (721 kbit/s, after the voice quality streams are taken into account). These use real-time protocol (RTP) for streaming. RTP was originally developed for managing streaming connections on asynchronous packet-switched networks such as parts of the Internet.

Either point-to-point or point-to-multipoint (such as one transmitter to multiple Bluetooth loudspeakers) are allowed. QoS is not guaranteed with ACL connections although RTP does provide for some real-time requirements provided that buffering is used at the receiver. Better QoS provision for streaming has been requested by the A/V working group of Bluetooth SIG in the next revision of the Bluetooth data-link specification. Although Bluetooth supports isochronous communication (the type required for clock-dependent processes) for streaming applications through the use of higher-level L2CAP (logical link control and adaptation protocol) connections, this is always only on a best-effort basis. The only truly synchronous reserved slots at the baseband level are for SCO packets, basically speech audio.

AVCTP is the Audio/Video Control Transport Protocol that can be used for conveying messages intended for controlling Bluetooth A/V devices. The AVDTP (Audio/Video Distribution Transport Protocol) uses L2CAP packets to transfer audio with connections established between a transmitter and a receiver. This protocol provides a mechanism for reporting QoS, optimizing the use of bandwidth, minimizing delays, and attaching time-stamp data to packets for synchronization ➥

purposes. The AVDTP document provides some advice in an appendix relating to the synchronization of devices either to each other or to a separate network clock. It also briefly mentions an approach to measuring timing jitter. Broadly, the approach relies on the timing mechanisms inherent in RTP and RTCP (real-time control protocol).

The A2DP (Advanced Audio Distribution Profile) describes a configuration of layers in the Bluetooth stack and higher application layers that can be used for the conveyance of audio. It was authored by the audio/video working group consisting of members from Sony, Toshiba, Nokia, Philips, Ericsson, and Matsushita.

Considering the bit rate, there is really no way that uncompressed high-quality audio can be carried, requiring that some form of low bit-rate coding be employed. Specified in A2DP is a mandatory audio codec, which is a low-complexity, subband codec (SBC), whereas other codec types (MPEG and ATRAC in the current version) are listed as optional. The SBC codec was developed for Bluetooth but based upon an earlier Philips system described by de Bont et al.[1] NonA2DP codec types can be accommodated, although the rubric is rather confusing in relation to this eventuality. This is supposed to be achieved either by upgrading the codec to optional status within A2DP (this requires the manufacturer to submit clear definitions of certain required characteristics) or by transcoding the audio data to SBC if the receiver does not support the decoding of the data type; this way interoperability is maintained as much as possible.

The profile document does not define anything in relation to nonA2DP codecs except that the vendor is supposed to use a Bluetooth Assigned Number to identify itself, and its parameters must be signalled within the standard packet headers. Audio may be encrypted for content protection or not, but this is application dependent. The mandatory subband codec should use at least one of 44.1- or 48-kHz sampling frequencies, and other lower rates can be specified. The encoder should be able to handle at least mono and one stereo mode (such as dual channel, stereo, joint stereo) and the receiver should be able to decode all of these. Similar requirements exist for MPEG 1, Layer I, II or III, for MPEG 2 and 4-AAC, and for ATRAC codecs. MPEG codecs are also allowed variable bit-rate (VBR) encoding in the profile.

Quality and Robustness of Bluetooth Audio Streaming

For adequate audio quality the A2DP profile requires that the audio data rate be sufficiently lower than the available link data rate to allow for the retransmissions that will avoid packet loss. The margin allowed for this obviously depends on the robustness expected of the application in the face of interference or long distances. The profile limits the SBC bit rate to a maximum of 320 kbit/s for mono and 512 kbit/s for stereo. The maximum available bit rate is 721 kbit/s. The data overhead when one transmitter is communicating multiple ACL streams to different receivers can lower the overall bandwidth available for audio, hence the number of channels is tightly limited even with data compression. An alternative is to use broadcast mode in which the full data stream for all audio channels is broadcast to all receivers, requiring them to separate the channels themselves (they therefore need to have a means of channel identification). The maximum number of audio channels in this mode is seven, that is the maximum number of connections to a master device.

There are some problems with point-to-multipoint connection for audio using ACL and RTP because different retransmission rates will apply on each connection and possibly affect inter-channel synchronization. The result of this is timing differences between the audio channels or phase distortion. This can be ameliorated using adequate buffering and resynchronization. Broadcasting gets around this problem, but packet loss can be encountered and not recovered because there is no dynamic retransmission method for broadcast mode. However, there is a fixed retransmission option for broadcast mode, which appears to act rather like a form of permanent redundancy, whereby packets are always retransmitted at the expense of a reduction in available bit rate on the channel. Floros et al.[2] found broadcast mode with no retransmissions only just acceptable for audio streaming owing to packet loss on the wireless link. One retransmission reduced audio data loss to 2 to 3% compared with 7 to 8%, but reduced the effective bit rate from 551 to 325 kbit/s. Because the losses were in compressed audio data, the resulting uncompressed audio was badly affected. They claim that retrieval mechanisms in the Bluetooth lower levels and application layer could be adapted to minimize the effects of packet loss on audio quality, but have some reservations about the use of the approach for synchronous compressed multichannel audio. They concluded that one could easily transmit stereo audio at 256 kbit/s per channel plus control information, within the bandwidth available, using two separate ACL links. The best data rate was obtained with DH5 packets, getting close to the upper limit of 721 kbit/s, but all DH packets have no forward error correction and so are more prone to data loss on the link. DM packets have forward error correction (FEC) and a correspondingly lower overall data rate.

Whether Bluetooth stands the test of time remains to be seen, but wireless technology for audio transmission will undoubtedly continue to expand. See the box on p. 982 for useful web sites with more information.

FURTHER READING

Specification of the Bluetooth System, version 1.1 (2001).

Bluetooth Audio/Video Working Group (2002), Advanced Audio Distribution Profile 0.95b.

Bluetooth Audio/Video Working Group (2002), Generic Audio/Video Distribution Profile 0.95b.

[1] F. de Bont, M. Groenewegen, and W. Oomen, "A High Quality Audio-Coding System at 128 kb/s," presented at the 98th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 43, p. 388 (1995 May), preprint 3937.

[2] A. Floros, M. Koutroubas, N. Alexander-Tatlas, and J. Mourjopoulos, "A Study of Wireless Compressed Digital Audio Transmission," presented at the 112th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 50, p. 498 (2002 June), preprint 5516.