Bluetooth: sufficient fidelity even for average listeners?

APPLICATIONS REQUIRE VIDEO SUPPORT AND ADDITIONAL AUDIO CHANNELS TO COPE WITH STEREO, DRIVING THE NEED FOR LOW-LATENCY CODECS. USING A BLUETOOTH CHIP SET SIMPLIFIES DEVELOPMENT AND ENABLES THE MANUFACTURER TO LEVERAGE THE BRAND RECOGNITION THAT BLUETOOTH NOW ENJOYS AMONG CONSUMERS.

D elivering seamless-quality audio in real time using wireless technology is one of the great challenges facing audio engineers. Bandwidth constraints, coding delays, and the introduction of bit errors can all hamper wireless-audio transfers, thereby causing significant audio-quality degradation.

The mobile-audio-device market has grown rapidly over recent years, producing consumers well-used to scrutinizing audio and less willing to compromise on quality. "Average Joe has grown a set of golden ears," you might say. Reinforcing this claim, in April 2007, Apple announced that, through the iTunes Store, it would be offering tracks at higher-quality, 256-kbps AAC (advanced audio coding), a move signifying the mass market's increasing appreciation of quality audio.

Along with this market trend, real-time wireless audio is experiencing escalating demand from emerging consumer markets, a demand that manufacturers have so far struggled to satisfy. And the convergence of audio with video across the spectrum of consumer devices means that wireless audio has an additional issue to overcome. Any delay on delivery results in lip-synch issues and an unsatisfactory user experience. Wireless headsets for mobile TV, video playback, and gaming, along with wireless speakers for stereo and 5.1-channel surround connected to a video source, require real-time-audio delivery.

Uncompressed CD-quality stereo audio uses a 1.411-Mbps bandwidth. For most wireless applications, this full bandwidth is impractical. Issues of design, efficiency, power optimization, and error resilience put pressure on available data rates. Also, in many standards and protocols, such bandwidth is simply not available. Bluetooth, for example, stipulates a maximum available bandwidth for A2DP (advanced audio-distribution profile) of 768 kbps. So, for high-quality stereo audio, it is necessary to use some form of audio coding to reduce the required data rate.

WIRELESS AUDIO TODAY

The proliferation of wireless technologies, such as Bluetooth and Wi-Fi, has given the consumer the ability to wirelessly receive digital audio wherever he may be and however desired—in the home or on the move, by streaming audio over Wi-Fi from a Mac or PC, or by connecting a transmitter dongle to a mobile-audio device and listening with wireless headphones. However, with every technical advance, there is often a bottleneck in which one aspect of the technology advances beyond the capabilities of another.

Personal wireless audio has experienced such an issue. Bandwidth limitations are an obvious problem for wireless-system applications as manufacturers strive for ultralow-power consumption in mobile devices. For live streaming, audio-coding delays are again prohibitive constraints. Such delays have implications for video applications requiring lip-synching—for example, when using a wireless-stereo headset with a video-playback-supported iPod or a mobile-TV receiver.

In bidding to be free of the wire, designers can take many approaches to handling wireless-stereo audio. For personal audio streaming, the predominant radio frequency is the license-free, 2.4-GHz spectrum because it can provide sufficient bandwidth, range, and power consumption. Bluetooth and other proprietary RF technologies operate in this frequency.

The Bluetooth SIG (Special Interest Group) ratified the A2DP to manage the transfer of stereo audio, and the consumer market has subsequently experienced the arrival of A2DP-enabled products on both the audio-source and the headset sides. Motorola, for example, has come to market with products such
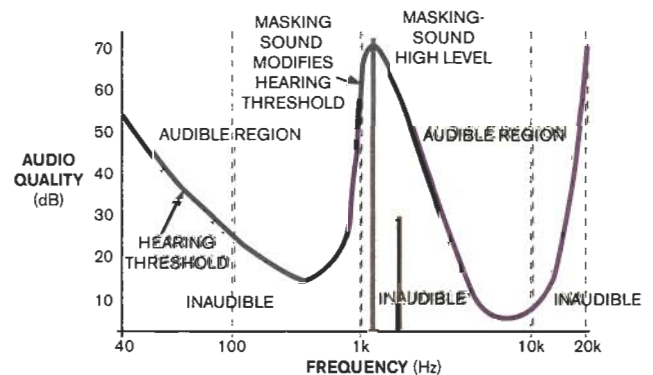
Figure Psychoacoustic techniques enable discarding of substantial portions of the audio information without adversely affecting quality.

as the S9 Bluetooth stereo-active headphones. However, an A2DP-supportive transmitter dongle will be necessary to ensure connectivity to most audio players because audio sources do not yet widely use A2DP. The reluctance of consumer-audio companies to integrate the A2DP profile has predominantly been due to issues of audio quality and coding delay.

The industry regards 16-bit audio as the entry-level quality requirement for audio systems now on the market, along with a minimum sample rate of 44.1 kHz to match that of the venerable audio CD. Consider the dynamic-range capabilities of various sample sizes: 96.32, 120.4, and 144.5 dB for 16-, 20-, and 24-bit digital audio, respectively. To achieve CD-quality dynamic range in bandwidth-limited applications, such as Bluetooth-stereo headsets, necessitates the use of at least 16-bit audio as the raw input; a compression technology that can reproduce virtually all of the original dynamic range subsequently transforms this audio (**Table 1**). The challenge is to find an algorithm that can deliver this quality level with low corresponding latency and maintain efficient processing power to prevent excessive battery drain.

The main difficulty for live audio is the coding-plus-decoding delay of compression technology. Although, in most wired systems, the lengthy video-decoding delay masks the audio-coding delay, wireless-system applications have no such luxury. Developing the ability to lip-synch audio to decoded video after audio encoding, packetization, passage over a wireless link, and decoding is indeed a significant challenge.

When dealing with wireless speakers for high-definition home theater, a delay greater than 10 msec can negatively impact the desirable seamless, full-surround-sound experience for discerning viewers. For gaming applications, 10 msec would again be the target because gamers' reaction times allow no room for delay. It is one thing to hear audio while viewing video, but it is another thing to expect to hear audio and instantly react to it.

Wireless-stereo headsets interacting only with audio sources can accept delays because audio-only applications have no need for lip-synch. However, when using a wireless headset with a video source, the difference is clear. Depending on screen size and distance from the device, the delay target for the industry is currently 40 msec or less. In most applications, the radio has its own inherent delay characteristics and complies with a standard. If you assume that radio delays and the packing and unpacking delays associated with the RF protocol are fixed, you have only the audio-compression delays to work with.

Bluetooth is robust and ensures accurate signal delivery, but this focus on resilience results in fundamental delay issues. Bluetooth uses a series of fixed-size transmission and reception slots, which therefore have response-time limitations. The Bluetooth protocol can retransmit packets to correct errors in the transmitted stream. If you could minimize the retransmissions by using a more robust algorithm, you could improve the system response. In addition, avoiding frame-based algorithms, which require filling an
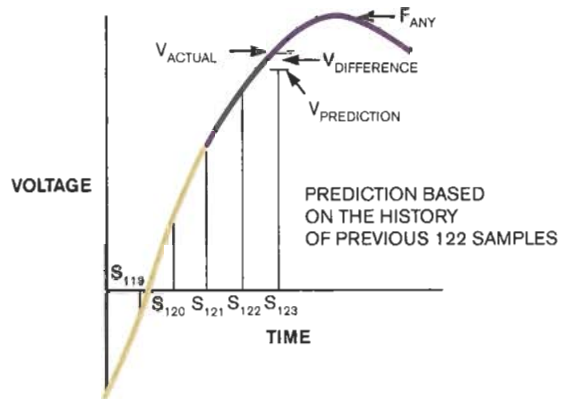

Figure 2 Variable-step-size quantization and differential coding are at the core of the ADPCM compression approach.

entire frame of audio samples before decoding, further minimizes delay.

THE NEED FOR COMPRESSION

Bluetooth A2DP has a maximum available bandwidth of 768 kbps. So, audio compression is necessary to deliver two-channel digital-stereo sound. Myriad compression technologies are currently available, each targeting and offering benefits in specific applications. However, most of them derive from two fundamental audio-compression processes: perceptual techniques based on psychoacoustic models of hearing (**Figure 1**) and predictive techniques, which as their name implies, employ a system of predictive coding. They are therefore known as ADPCM (adaptive-differential-pulse-code-modulation) codecs (**Figure 2**).

Generally, the higher the compression ratio, the more audio content you lose. With perceptual codecs, such as MP3, AAC, and their derivatives, analysis of the frequency spectrum results in the removal of any content the technology deems imperceptible to the human ear. This technique requires buffering of an audio sample of approximately 512 bytes to perform the analysis. Buffering is often the fundamental source of coding delay. The complexity of the audio can also affect the delay of the encoding process. The psychoacoustic procedure, with its ability to produce high compression ratios and retain reasonably high audio quality, is processor-intensive and therefore not a good approach for power-efficient, battery-powered devices.

ADPCM codecs operate in a different manner, due to their unique characteristics. PCM is the digital representation of an analog signal, wherein regular sampling of the signal magnitude at uniform levels results in quantization to a series of symbols in a digital code. CDs are examples of the implementation of PCM audio. ADPCM involves audio-value encoding as the difference between the current and the previous val-

TABLE 1 KEY AUDIO METRICS

Sampling frequency (kHz)	Resolution (bits)	Per-channel bit rate	4-to-1 compression ratio (bits/sec)
44.1	16	705,600	176,400
44.1	20	882,000	220,500
44.1	24	1,058,400	264,600

ues, and the quantization step size varies to allow a bandwidth reduction for a given SNR (signal-to-noise ratio).

The quantization process is by nature lossy, and, depending on the accuracy of the linear predictor and inverse quantization you use, it can produce small errors in the reproduced audio. However, removal of audio content does not occur, and PCM is therefore a popular technique in applications in which issues with tandem coding or transcoding would otherwise exist. ADPCM-based algorithms range from International G.711, G.722, and G.726 codecs for low-bit-rate voice to professional broadcast standards, such as apt-X, for high-quality, multichannel audio (**Table 2**). Their shared feature is their low delay, which enables real-time two-way communication. As the ADPCM technique does not buffer a frame of audio and analyze the full-audio spectrum with each encoding step, the processing delay is also a fraction of that you find in the alternative perceptual-coding approach.

AUDIO-CODING OPTIONS

Some years ago, the Bluetooth SIG selected the SBC (smart-bit-rate-control) compression algorithm, which Philips

TABLE 2 AUDIO-COMPRESSION OPTIONS			
Algorithm	Compression ratio	Bit rate (kbps)	Encode-to-decode delay (msec)
MP3	12-to-1	128	326 at 128 kbps
AAC	Variable	128 or 256	172
AAC low density	Variable	128	60
SBC	Variable	366 (bit pool: 55)	2.6
apt-X	4-to-1	384	1.87

developed, as a mandatory codec to ensure interoperability for Bluetooth products. The SIG chose this codec for a number of reasons. It was freely available to the Bluetooth SIG, it has low complexity in processing overhead, and it has better encoding-and-decoding latency than alternative compression algorithms, such as MP3 and AAC. With the arrival of Bluetooth-stereo headsets, however, widespread concerns arose regarding SBC's ability to deliver full-bandwidth, high-quality audio. Additionally, providers of Bluetooth A2DP devices claim that, using SBC, their devices could only occasionally achieve the industry target of 40-msec delay for lip-synching.

The wireless link is typically not robust enough to achieve low latency, and high processing and power consumption also do not make SBC viable in certain situations. For these reasons and because of substantial A2DP demand, several fabless-semiconductor companies have brought to market proprietary technologies that offer full-bandwidth uncompressed audio operating over a 2.4-GHz RF spectrum. These approaches aim to match CD-audio-quality requirements and have real-time-transfer ability, but they have drawbacks in power con-

sumption—because the radio must transmit full uncompressed audio—and Bluetooth-standard compatibility.

Many mobile devices integrate Bluetooth chips, offering monophonic-headset interaction for voice, and several handsets offer the A2DP for stereo streaming. An additional proprietary approach not only requires an additional chip, but also introduces compatibility issues due to the fact that there is no agreed-upon compliance standard for audio transfer between consumer devices.

The ideal scenario for many consumer-electronic and mobile-device companies would be to use Bluetooth and provide full-bandwidth stereo-audio-quality streaming in real time. Only a few companies currently provide products to fulfill this need. Given the issues regarding using psychoacoustic algorithms, you should discount MP3 as a viable technology for wireless transfers. Therefore, you must look at ADPCM-based alternatives. US-based Open Interface North America, for example, in 2003 launched Soundabout eSBC (Enhanced SBC). Based on the same principles as SBC, eSBC allows a 510-kbps data rate and, hence, some quality benefit. However, this higher data rate comes at the expense of power consumption, which can have a significant impact on battery life, and the algorithm offers no latency improvement.

Last year, UK-based Audio Processing Technology partnered with a leading Bluetooth-chip provider to provide an SBC al-

ternative. The company's apt-X audio algorithm also uses ADPCM principles but incorporates additional techniques for accurate linear prediction and inverse quantization to retain optimal audio quality. Matching uncompressed CD quality, apt-X offers a dynamic range greater than 92 dB and runs at 384 kbps. The overall framework of apt-X ensures a robust connection, which enables optimization of the overall Bluetooth system latency. The technology can also synchronize within 3 msec on startup or in response to a dropout, and the algorithmic coding delay is less than 2 msec to ensure real-time connections. With such offerings now available, it is likely that Bluetooth-stereo headsets will mature in performance and quality throughout 2008 and that the market for these devices will experience **strong growth.EDN**

AUTHOR'S BIOGRAPHY

Stephen Wray joined Audio Processing Technology as licensing-business manager in 2005. Responsible for the implementation and execution of licensing-business strategies, he cultivates liaisons with all major licensees, attends trade shows and industry symposiums, and represents APT licensing at industry-networking events. He is a graduate of the University of Edinburgh with an honors degree in electrical and electronic engineering. Wray is a member of the IET (Institute of Engineering and Technology), the AES (Audio Engineering Society), and the CIM (Chartered Institute of Marketing).