

Digital Voice: The Next New Mode?

Interest in digital voice systems is on the rise. Do they have a place in Amateur Radio? Come on a brief tour of the technology and see for yourself.

Why Digital Speech?

These days, it seems communications systems are going digital everywhere they can. Why are we doing it? What's wrong with well-established analog techniques?

Well, nothing much is wrong with them; in fact, they will always represent the most straightforward ways for the transmission and perfect reproduction of speech signals. But propagation paths for radio signals may be far from perfect and that's where digital voice comes in.

Digital modes offer certain advantages over their analog counterparts. Foremost among those is that digital detectors have a very clear-cut decision to make. In principle, it's easier to decide whether a received signal represents a binary zero or one than to decide exactly what analog voltage it represents. With appropriate restrictions, that's also true in practice. A second big advantage of digital modes is that errors in transmission may be made relatively easy to detect and correct. Coding schemes have been devised that produce very robust performance, even through poor propagation media. Finally, digital signals lend themselves to some advanced processing techniques that would be incredibly complex in analog. Those techniques generally achieve performance levels not otherwise possible.

In many cases, the advantages mentioned above have made it very worthwhile to employ digital transmission and processing of analog signals. Commercially, digital high-definition TV (DTV) and cellular phones have begun to show that. The resounding surge in DSP-based transceivers is certainly evidence of what's possible with signal processing; but here, I'd like to discuss how analog signals—specifically speech signals—may be transmitted and received in digi-

tal format. A look back at the history of digital speech modes reveals a lot about both how and why.

A Brief History of Digital Voice Modes

The public switched telephone network (PSTN), the communications medium to which the most people have access, went digital a long time ago. Engineers realized that to obtain the best performance over a large area, many repeaters and switches are required. Analog amplifiers, repeaters and switches introduce noise; that makes it difficult to maintain acceptable signal-to-noise ratios

(SNRs) over long distances. As against that, digital signals received at a repeater or amplifier may be cleanly detected and a new, noise-free copy of those signals may be retransmitted. A digital transmission format was therefore chosen for the PSTN around WW2 time.

The first task for those working on the problem was to decide on a way to convert analog speech signals to digital. The device doing that job is aptly known as an *analog-to-digital converter (ADC)*. The job itself is called *sampling*. Samples are taken at regularly spaced intervals and the result is a string of numbers that represent the analog voltage at those discrete

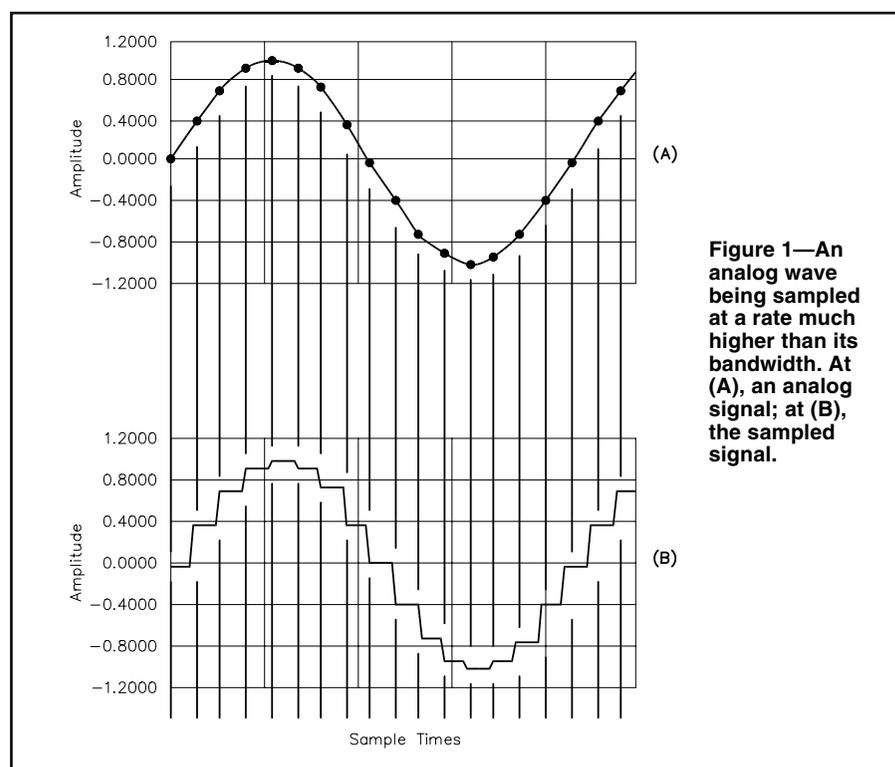


Figure 1—An analog wave being sampled at a rate much higher than its bandwidth. At (A), an analog signal; at (B), the sampled signal.

times. Each voltage sample is converted to a binary number proportional to the voltage. To get an accurate representation, many samples per second must be taken so that the voltage doesn't change much between samples. See Figure 1. The number of voltages that can be represented, therefore, is determined by the number of binary digits or *bits* available. For example, if eight bits are available, then 2^8 or 256 voltage levels are possible.

One of the first things discovered about such a scheme is that since only 256 levels are possible, the binary number chosen at any particular sample time may not correspond exactly to the actual analog voltage; it's only the closest of those available. For a large signal over time, errors are just as likely to be positive as negative; they are also just as likely to be small as large, within certain limits. Errors therefore show up as *quantization noise* in the sampled signal, which limits the total range of signal amplitudes. That range is called the *dynamic range*.

Telephone engineers recognized that if they used more bits for the smaller signals, and fewer for the large, they could achieve an increase in dynamic range. The system now in use on the PSTN in North America and Japan, called *μ-law coding*, does exactly that and extends dynamic range by quite a lot.¹ The chief penalty is that the maximum SNR is reduced slightly—not a bad trade-off. Other countries use *A-law coding*, which is slightly different.

The sampling rate must be at least twice the bandwidth of the signal being sampled.² The phone company decided that about 3 kHz of bandwidth was good enough for speech and so chose a sampling rate of 8000 samples per second. With eight bits per sample, the transmission rate is $8 \times 8000 = 64,000$ bits/second (bps). The system provides what is generally known as *toll-quality* speech and it preserves most of the important characteristics of a person's voice.

The US space program also had need for voice communications and NASA, too, recognized the value of digital transmission modes. During the 1960s, designers found that certain digital coding schemes gave them the ability to determine the transit time between transmitter and receiver, hence the distance between the two, while using a continuously transmitted digital signal. They also knew that square-wave limiting (clipping) of human voice signals increases the talk power of those signals. Clipped speech signals resemble digital waveforms, so they reasoned that they could

¹Notes appear on page 32.

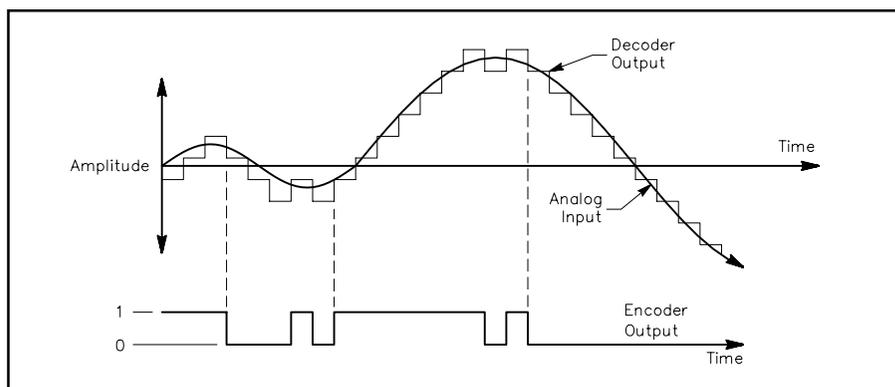


Figure 2—A representation of delta modulation (DM).

How Do I Sound?

That seems like an innocent question and it's easy to slip into non-technical terms, like "scratchy," "warm" and so forth. If you are serious about giving a meaningful response, though, some forethought is required. For scientific voice-quality evaluation, a uniform system that gauges subjective responses is necessary.

A wide variety of factors influences perceived voice quality, including amplitude and frequency distortion, echoes and noise. Anything detracting from the naturalness of speech increases the effort a listener must exert to understand what is being said. For signals that are significantly impaired, the annoyance experienced by a listener may be rated on a linear scale called *mean opinion score (MOS)*.

The MOS scale is shown below:

MOS	Quality	Impairment
5	Excellent	Imperceptible
4	Good	Perceptible, but not annoying
3	Fair	Slightly annoying
2	Poor	Annoying
1	Bad	Very annoying
0	Unusable	Total

Non-integer scores like 3.5 are possible. An MOS of 3.0 is generally referred to as "toll quality," meaning "good enough to pay for." Digital voice users may tolerate MOS levels less than three if they get additional benefits, such as simultaneous voice and data services.

While evaluation of voice systems may be made based on test-bench measurements, they must ultimately relate to the perception of the listener. A large body of voice-system evaluations exists based on MOS. Comparisons among systems are therefore readily made. MOS relates well to the readability figures commonly used in Amateur Radio signal reports.

Comparison is always part of subjective analysis. In fact, comparison is absolutely necessary to remove all bias in voice-quality evaluation. Most often, a listener is presented with two audio samples in succession; he or she is not informed beforehand which sample is the one being evaluated. Several repetitions using many different listeners may be averaged to mitigate the effects of individual listening talents. For digital voice systems, MOS may be correlated with the bit error rate (BER) on the communications link. Performance in hostile environments—those containing high levels of environmental and man-made noise—may thereby be quantified.—*Doug Smith, KF6DX*

use them as such. Combined with range-determining codes to produce a single digital bit stream, they found that gave them both voice communications and the distance information they sought.

Such a system was used by NASA for the Apollo program.³ It's very clever; but if you think it crude by today's standards, remember that at the time, LED displays had not yet been perfected and digital numerical readouts aboard the spacecraft were provided by Nixie tubes! Since then other, more-sophisticated schemes have been developed and some even existed

before the space age.

Just after WW2, researchers discovered a waveform-coding system known as *delta modulation (DM)*. In it, when an analog input wave's voltage is increasing, a binary one is transmitted; when the analog voltage is decreasing, a zero is transmitted. See Figure 2. A fixed amount of voltage change is associated with each bit so that the analog waveform may be reconstructed at the receiver through integration. It's a very simple system and it works reasonably well, but it has an inherent problem: It can't represent analog

waves having slopes exceeding the maximum voltage change per bit. In the 1970s, others found that limitation could be overcome by incorporating a greater slope when several ones or zeros occurred in a row.⁴ Their system, called continuously variable slope-delta modulation (CVSD), produces toll quality at bit rates significantly lower than those on the PSTN and it's more immune to errors in the bit stream. Its maximum SNR, though, is generally not as good as what you get over the telephone.

Other schemes, such as adaptive differential pulse code modulation (ADPCM), have achieved some measure of success.⁵ Over the last 30 years, a lot of experimentation has gone into finding better ways to characterize voice signals than those of the *waveform coders* described above. Driving that research is the need to minimize the number of bits transmitted and thus, the occupied bandwidth of digital voice signals, as well as the complexity of modems used to do it.

Intense investigation about the nature of human speech production and hearing began in earnest in the 1930s.⁶ Many things discovered then remain relevant to this day.

On the Nature of Human Speech and Hearing

Investigators of human speech have found that it may be modeled as a source of excitation (wind from the lungs) followed by a filter (the voice tract).⁷ They've also discovered that certain properties of a person's voice may be characterized and extracted from voice signals that lend themselves to efficient digital coding.⁸ Those characteristics relate to the basic nature of human speech sounds and physical factors in their production.

Some voice coders make use of a source-filter model to achieve good speech reproduction at low bit rates. Instead of transmitting information about the wave shape of speech, they transmit spectral information about the source and the frequency response of the vocal-tract filter. That approach is a winner, largely because the spectrum of speech changes relatively slowly. That is, the frequency content of speech may be considered constant over short time frames of, say, 20 ms or so. Even over time frames longer than that, the source spectrum may remain reasonably constant. Those sorts of speech characteristics allow *parametric* speech coders a large measure of efficiency.

Human hearing has evolved so that it's good at distinguishing human speech sounds. Auditory research has revealed some interesting things about the ear-brain combination that are relevant to speech coders and decoders (codecs).

Such research is conducted subjectively; that is, what someone hears (or doesn't hear) can only be determined by asking questions of the observer and attempting to infer something from his or her answers. For that reason, we define physical and perceptual parameters of sounds differently and separately.⁹

Intensity is the physical measure of sound amplitude. *Loudness* is the corresponding perceptual magnitude; it is arbitrarily defined with respect to a fixed-frequency tone at a certain intensity. We have no guarantee that two listeners will say that any particular sound has the same loudness; however, controlled experiments have shown that observers agree closely on whether one sound is twice as loud as another. So the perception of loudness may be scaled in an orderly way from soft to loud.

Frequency is, of course, the physical measure of cycles per second of a sound. The corresponding perceptual measure is known as *pitch*. This term is not to be confused with the base frequency of a person's voice. Pitch is to frequency as loudness is to intensity.

Having separate perceptual measures for sound characteristics might seem useless at first, but research has shown that loudness is not independent of frequency.¹⁰ By now, it's fairly well-known that human hearing is most sensitive to frequencies in the range of 2-3 kHz. For instance, a 2-kHz tone sounds louder than a 500-Hz tone of the same intensity. Also, pitch is not independent of intensity. You may demonstrate that to yourself by turning up the intensity on a pair of headphones and comparing the pitch of what you hear when they're on your head to what you hear as you move them away. Don't turn the intensity up too much, though, because researchers have also found that permanent hearing loss may occur at intensity levels far below those causing significant discomfort.¹¹

Human hearing seems to have certain thresholds that come into play during recognition of speech, music and other sounds. One important threshold of hearing is the ability to tell whether one sound is louder than another. In the presence of multi-frequency or *polyphonic* sounds, that threshold is influenced by how close in frequency the sounds are. For example, a quiet sound that is close in frequency to a louder sound might not be audible at all. Such *masking* is important in speech coding because it implies that the number of discrete intensities and frequencies to be represented may be reduced.

Another threshold of hearing is the ability to tell whether one sound is higher or lower in frequency than another. Although it's influenced by intensity, ex-

periments generally find that threshold increases as the frequencies of sounds increase. In other words, it's harder to discern subtle differences in frequency among higher-frequency sounds. The significance of that in speech coders is that the number of discrete frequencies that have to be represented may be reduced.¹²

Much of the energy in human speech above 3 kHz is produced by sounds like "p" and "f," which are inherently noisy. It's therefore no surprise that our hearing has not developed good frequency discernment up there: Not much useful information is contained in those frequencies. There may be physical reasons for that as well, but it's interesting that our ability to understand speech closely matches our ability to communicate verbally.¹³ For example: The fastest talker can go about 300 wpm, which is about the limit of most listeners' comprehension.

Technical Goals of Digital Voice Systems

All the above directly relates to our desire and ability to reduce the data rate of digital speech signals. Lower data rates are good because they may be transmitted in smaller bandwidths and recovered with higher SNRs using narrower receiver bandwidths. A definite trade-off exists, though, between data rate and speech quality. To illustrate what's possible, consider the following example that draws on several key concepts in speech coding.

Let's say we want to build a speech coder—for a single language only—that uses a bit rate approaching the minimum possible bit rate. We may not know what that minimum is, but we want to see if we can find it. Let's also say that cost and complexity aren't big concerns. Occupied bandwidth is our chief concern; other goals are secondary.

We decide to employ a speech-recognition engine at the transmitter that identifies individual words from the talker. That's already being done with much success, so it's not a big technical leap of faith. We assume that a vocabulary of about 65,000 words is enough to support all the sentences the speaker is likely to construct. Each word may then be represented by a 16-bit code, since $65,000 \approx 2^{16}$. The speech-recognition engine looks up a 16-bit code for each word and puts them together into a serial bit stream. Ignoring the requirements for synchronization, pauses between words, error detection and correction, a person talking at 150 wpm generates data at a rate of $(150 \text{ wpm})(16 \text{ bits/word})(1/60 \text{ minutes/second}) = 40 \text{ bps}$. Many languages have a heck of a lot more words than just 65,000 and some people might talk faster, but you get the idea.

Now that signal can be coded into an

analog format that occupies very little bandwidth. The inverse process is employed at the receiver, terminating in a speech synthesizer that drives an audio power amplifier and loudspeaker. See Figure 3.

What are the drawbacks of this scheme? Well first of all, it's rather elaborate and expensive. Secondly, the software has to be different for each language supported. You'd have to know which language was being used ahead of time to correctly decode messages. Finally, the listener at the receiver can't tell who is speaking unless he or she reveals it; none of the speaker's emotions or inflection is transmitted. The listener can't tell if the person has a stuffy nose or whether there are any other voices or sounds in the background. Speech from the decoder sounds robotic and it's difficult to listen to; comprehension has been sacrificed to some extent because of the lack of important speech properties. The conclusion is that we have reduced the bit rate too much and traded off too many important speech characteristics. The bit rate must obviously be increased to improve things. That brings us to some definitions about what is acceptable for digital speech in Amateur Radio. The following restrictions ultimately determine the lower bit-rate limit.

For Apollo astronauts or military personnel, it's not always very important to be able to tell who is speaking, so long as the information is communicated. Amateur Radio is a different story, because how something is said and how it sounds is sometimes as important as what's being said. We may deduce, then, that digital voice for hams must be of high quality so that it's difficult to tell the speech was coded.

Amateurs often work with signals near the SNR limit of detection. In that regard, digital voice systems need to perform at least as well as existing analog formats to become popular. Digital coding opens some interesting possibilities for redundant transmission, such as sending the data many times and comparing data sets to achieve a large measure of forward error correction. Data transmission rates may also be artificially slowed to aid reception, then sped back up at the receiver after all the data have been received. How that kind of thing will affect phone contests and distance records is open to speculation.

I suspect that many hams would like to try digital voice without having to buy a new transceiver. That means digital voice systems may initially take the form of external boxes that interface to existing transceivers at the audio level. Such boxes are already being developed.¹⁴ Aside from speech-quality goals, certain other benefits may come to digital voice users. The ability to embed certain iden-

tifiers in a digital voice transmission provides significant benefits. Transmissions may be automatically identified as to their source, destination, protocol, and other parameters. As that kind of thing is made possible, cellular and trunking systems come within reach.

Is Digital Voice Legal on the Amateur Bands? If So, What Frequencies and Emissions May Be Used?

Part 97 of the FCC rules states that phone signals—whether analog or digital—must remain in the phone subbands.¹⁵ That's mainly a concern for the eight HF bands where phone is used. In the VHF bands above 10 meters, phone is legal for US-licensed amateurs at all allocated frequencies, with the exception of 50-50.1, 144-144.1 and 219-220 MHz. The rules also say that no transmission "... shall occupy more bandwidth than necessary for the information rate and emission type being transmitted, in accordance with good amateur practice."¹⁶ That's purposefully vague: The Amateur Radio Service is free to experiment with almost any mode you can think of, as long as it's not wasteful of bandwidth. You can take it to mean that a digital voice transmission should not occupy more than the equivalent SSB transmission on congested bands or the equivalent AM or FM transmission on sparsely occupied bands, such as 10 meters. While the symbol rate (baud rate) of digital data transmissions is limited on many US ham bands, the baud rate of digital phone transmissions is unlimited!¹⁷

What is the emission designator for digital voice? Well, the first symbol of an emission designator tells what modulation format is being used. For an SSB transmitter, that is letter "J." For an FM

or PM transmitter, the letter is "F" or "G." The second symbol tells about the nature of the modulating (baseband) signal. The most likely situation in amateur operation is the application of a modulated audio signal to the input of a transmitter. The symbol for that is numeral "2." The third symbol tells about the type of information being transmitted. That would be letter "E" for phone. So the most likely emission designators for digital voice would be J2E or F2E.

It may be weird to hear digital signals on the phone bands and courtesy dictates that operators explain—using analog phone—what's going on until general understanding is reached on the use of digital phone. The same kind of situation occurs during HF slow-scan television operation (SSTV, designator J3F) and it's been handled admirably by practitioners. Note that digital video is also perfectly legal on the HF phone bands (designator J2F), although it hasn't seen much use.

What is the State of the Art Now? Where Does Amateur Radio Come In?

International bodies have drafted several standards for audio codecs and modems; many are seeing use on the Internet and elsewhere.¹⁸ Work continues in commercial and academic sectors, as well as in Amateur Radio. Those efforts are making it easier for more amateurs to get involved—and involved we are.

The ARRL is making a significant commitment to digital voice and several other developing technologies. Those technologies relate to one another well; they reflect global trends toward more effective use of our radio communications spectrum. They also represent excellent opportunities for Amateur Radio to make significant contributions to the advancement of the communications art. The possibilities are very exciting, since they may constitute the next big changes in our service.

The FCC is very interested in amateur work in this field. They recognize that the Amateur Radio Service is an ideal place for experimentation with and testing of those concepts. Since we're a large and organized force of dedicated communicators, we belong at the forefront of their development. That notion is alive and well.

Considerable work is already being done by amateurs. A couple of years ago, Charles Brain, G4GUO, and Andy Talbot, G4JNT, started working with it. They produced a system satisfying the technical goals outlined above that was described in a paper summarizing their accomplishments (see Note 17). Tucson Amateur Packet Radio (TAPR) is producing a kit of this digital voice codec that's now available.¹⁹ It helps you to get started in digital voice with a minimal invest-

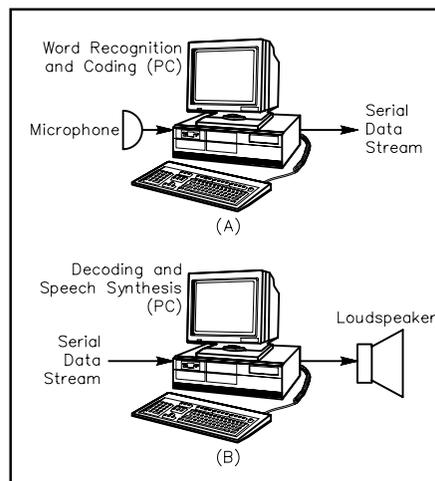


Figure 3—A digital speech system occupying very little bandwidth in transmission—but you have to know what language is in use.

ment in time and hardware.

The system employs a digital speech coding scheme known as *advanced multi-band excitation coding (AMBE)*.²⁰ Data rates up to 9600 bps are supported and the rate may be changed for experimentation. Coupled to a suitable modem and transceiver, it supports digital voice operation in both half-duplex and full-duplex modes. While AMBE is a complex algorithm, the significant details of its operation are in the public domain.

AMBE codecs provide high recovered speech quality and they've won spots in some very prominent systems, including Iridium and APCO 25. APCO 25 is a project to provide reliable digital voice communications to the public-service community.

Where Do We Go From Here?

Even with a digital voice codec in hand, you're going to need a modem that supports 2400-9600 bps: Many TNCs can do it. Those rates are relatively easy to achieve using audio frequency-shift keying (AFSK) and audio phase-shift keying (APSK) when 15 kHz or more of bandwidth is available, such as at VHF and above. Because of dispersive propagation on HF, though, those rates are difficult to sustain and some innovative techniques must be employed. Therefore,

high-speed HF modem design is one area that invites further work.

Some of us are working toward a single DSP system for digital voice that incorporates both the codec and the modem in software or firmware. The work is being undertaken on DSP development platforms that have data-conversion hardware (ADCs and DACs) included. Others have suggested that fast PCs, equipped with sound cards, might be capable of digital voice operation meeting the goals outlined above. That is another area ripe for experimentation.

Digital repeaters or "digipeaters" may be desirable on VHF and above to extend the range of digital voice communications. It might even be possible to build digipeaters that simultaneously handle more than one QSO.

Summary

I guess there's no going back now that we've identified and proven the benefits of digital communications technology. There may be other, as-yet-unidentified fruits to harvest in the quest for practical digital voice systems.

For more information about digital voice, point your browser to www.arrrl.org/tis/info/digivoice.html and take a look at some of the information and links

provided there. Reports of the TTF, TWG and DVC are available at www.arrrl.org/announce/reports-01/tt.html. League comments on so-called "software-defined radios" may be found at www.arrrl.org/fcc/arrrldocs/et-0047.pdf.

Doug Smith, KF6DX, a member of the engineering staff of Ten-Tec Corporation, serves as chair of the ARRL Digital Voice Committee. He is also editor of QEX/Communications Quarterly and author of the DSP chapter of the ARRL Handbook for Radio Amateurs. He can be reached c/o ARRL Headquarters, 225 Main St, Newington, CT 06111; kf6dx@arrrl.org.

Notes

- ¹J. C. Bellamy, *Digital Telephony*, Wiley and Sons, New York, 1982.
- ²A. V. Oppenheim and R. W. Schaefer, *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ, 1975.
- ³E. J. Nossen, "The RCA VHF ranging system for Apollo," *RCA Engineering*, Vol 19, Dec 1973/Jan 1974.
- ⁴J. A. Greefkes and K. Riemens, "Code Modulation with Digitally Controlled Companding for Speech Transmission," *Philips Technical Review*, 1970.
- ⁵G. 721: *Adaptive Differential Pulse Code Modulation*, ITU, Geneva, Switzerland, 1984.
- ⁶H. Fletcher and W. A. Munson, "Relation Between Loudness and Masking," *Journal of the Acoustical Society of America*, Vol 9, 1937; also see S. S. Stevens and H. W. Davis, *Hearing*, Wiley and Sons, New York, 1938.
- ⁷L. R. Rabiner and R. W. Schaefer, *Digital Processing of Speech Signals*, Prentice-Hall, 1978.
- ⁸N. S. Jayant and P. Noll, *Digital Coding of Waveforms: Principles and Applications to Speech and Video*, Prentice-Hall, 1984.
- ⁹J. L. Hall, "Auditory Psychophysics for Coding Applications," *The Digital Signal Processing Handbook*, V. K. Madisetti and D. B. Williams, Eds., CRC Press, Boca Raton, FL, 1998.
- ¹⁰B. Moore, *An Introduction to the Psychology of Hearing*, Academic Press, London, 1989.
- ¹¹C. M. Harris, Ed., *Handbook of Acoustic Measurements and Noise Control*, McGraw-Hill, New York, 1991.
- ¹²D. Smith, KF6DX, "PTC: Perceptual Transform Coding for Bandwidth Reduction of Speech in the Analog Domain," *QEX/Communications Quarterly*, Part 1, May/June 2000; Part 2, Mar/Apr 2001. The article may be found on *ARRLWeb*, www.arrrl.org/tis/info/digivoice.html.
- ¹³R. C. Stauffer, Ed., *Charles Darwin's Natural Selection*, Cambridge University Press, 1987.
- ¹⁴The G4GUO digital voice codec is a prime example. See Note 17 for more information.
- ¹⁵47 CFR 97.305.
- ¹⁶47 CFR 97.307(a).
- ¹⁷P. Rinaldo, W4RI, "Is Digital Voice Permissible under Part 97?" sidebar to C. Brain, G4GUO, and A. Talbot, G4JNT, "Practical HF Digital Voice," *QEX/Communications Quarterly*, May/June 2000. The article may be found on *ARRLWeb*, www.arrrl.org/tis/info/digivoice.html. "The Help Desk," elsewhere in this issue, contains an updated list of HF band plans.
- ¹⁸See, for example, G. 723.1, ITU.
- ¹⁹For details, visit the TAPR Web site, www.tapr.org.
- ²⁰Information and audio samples are available at www.dvsinc.com.

A Continuing Legacy of Innovation

Around the turn of the last century, experimenters began working with electromagnetic waves. This gave birth to Amateur Radio and wireless communications by a mode known as "spark."

It didn't take long for amateurs to find better and more efficient modes of communicating via wireless. Spark soon gave way to CW, then to AM voice. As time progressed, technology advanced and SSB brought spectral efficiency beyond the capabilities of AM. While amateurs have utilized RTTY techniques for many years, the explosion in interest did not occur until the computer became a popular tool in amateur stations, spawning a variety of digital modes. Now, at the turn of another century, it is time for us once again to lead the challenge for new modes in the Amateur Radio Service.

Early in 2000, the ARRL Board of Directors unanimously approved a recommendation from its Technology Task Force to create a Digital Voice Working Group. The TTF's Technology Working Group had performed a survey of radio amateurs throughout the world, seeking input on new technologies for the Amateur Service. The survey revealed that digital voice was one of the top recommendations. Subsequently, ARRL President Jim Haynie, W5JBP, appointed a Digital Voice Working Group with the objective of paving the road for digital voice to become a reality in the Amateur Service.

For a new mode to be widely accepted, participation from a wide geographical area must be sought. The working group involves radio amateurs knowledgeable in relevant techniques from the United States and Europe, where significant digital-voice work in the Amateur Service has already been performed.

Under the guidance of this working group, many amateurs should soon be enjoying yet another new mode of communication. Yet to come will be two additional working groups with similar objective assignments: high-speed digital networks and multimedia, and software-defined radio.

Moving from spark to CW and from AM to SSB were important events. The next generation of changes should be equally outstanding. For those who say nothing new comes from our Service anymore, and that the technology train left the amateur station years ago, I say "Listen up!" The interesting thing about that train is that it always comes back to the station looking for new passengers, and the Amateur Service has a long, continuing tradition of loading the train to capacity each time!—Joel Harrison, W5ZN, ARRL First Vice President, Chair, Technology Task Force