ABSTRACT

The recent development of the V.34 modem standard permits full-duplex transmission at rates up to 33.6 kb/s in the ordinary general switched telephone network (GSTN). This article briefly describes the technologies that are used to make these dramatically increased bit rates possible. This new high-speed modem enables various new multimedia modem applications.

The V.34 High-Speed Modem Standard

G. David Forney, Jr., Les Brown, M. Vedat Eyuboglu, John L. Moran III, Motorola, Inc.

International Telecommunications Union — Telecommunications Standardization Sector (ITU-T) Recommendation V.34 [1] is a new standard for full-duplex data transmission over the general switched telephone network (GSTN) at bit rates up to 28.8 kb/s (recently extended to 33.6 kb/s). The standard was completed in the summer of 1994, and V.34-compliant modems were introduced by many vendors shortly thereafter. V.34 modems have quickly become low-priced consumer products, and are rapidly displacing 14.4 kb/s V.32bis [2] modems (1990) in popular applications such as remote access to corporate networks, on-line services, and the Internet.

The dramatically increased bit rates of V.34 modems, combined with recent advances in digital voice coding which provide near toll quality at rates of the order of 8 kb/s or less, and similar advances in very-low-bit-rate video coding which provide acceptable video quality for certain applications at rates below 20 kb/s, allow the simultaneous transmission of voice, data, and video over ordinary voice-grade GSTN lines. These advances have recently led to the development of new multimedia modem standards such as H.324 and V.70, described elsewhere in this issue.

This article provides a brief overview of the technology embodied in the V.34 standard. For more comprehensive treatments the reader may consult recent digital communications textbooks and *IEEE Transactions* papers, although often the best references are the actual standards contributions.

The V.34 standard incorporates proposals from many contributors, and represents the state of the art of digital transmission over bandlimited channels. Indeed, the technology in V.34 is superior to that proposed by any single company, and in that sense V.34 may be regarded as a triumph of the standards process.

FUNDAMENTAL PHILOSOPHY: ADAPTIVITY

One of the key factors contributing to the development of V.34 is the general upgrading of the GSTN throughout the world. On most connections today, the transmission medium is almost entirely digital, with a fairly short analog local loop connecting the subscriber at each end to this medium via a digital central office. The principal impairments encountered on most intracontinental connections are the quantization noise in the 8-bit μ-law or A-law pulse code modulation (PCM) conversion and the bandlimiting to about 3700 Hz by linear anti-aliasing and interpolation filters. The signal-to-noise ratio (SNR) on such a good digital channel is typically of the order of 34–38 dB.

On the other hand, the GSTN also incorporates many other types of carrier facilities, which can vary enormously in transmission characteristics. For example, on intercontinental connections the use of adaptive differential pulse code modulation (ADPCM) at 32 or 40 kb/s for improved voice transmission efficiency has become commonplace. ITU-T standards for ADPCM support about the same bandwidth as PCM but provide a reduced SNR: about 21 dB at 32 kb/s (G.721), or about 28 dB at 40 kb/s (G.726). Proprietary 32 kb/s ADPCM encoders/decoders (codecs) that support a reduced bandwidth of less than 3200 Hz at an SNR of about 28 dB are also in common use.

More generally, older analog carrier equipment typically has lower usable bandwidth and SNR than newer digital facilities. For example, the North American GSTN still incorporates some 40-year-old *N*-carrier systems that typically support less than 3000 Hz of bandwidth at an SNR of 24–28 dB, and are also susceptible to other impairments such as frequency offset, phase jitter, hits, and dropouts.

Consequently, the fundamental design philosophy of V.34 involves not only the latest modulation technology but also a much higher level of optional capabilities, intelligence, and adaptivity than in V.32bis or previous modem standards, in order to make the best use of these various types of connections. Both of these advances depend on the availability of inexpensive, programmable digital signal processing.

A V.34 transmitter and receiver jointly agree during initial startup on the bandwidth and bit rate to be used, as well as whether to use various modulation options that are specified in the V.34 standard, and the parameters of these options.

The result of this design philosophy is that a V.34 modem is a "best-effort" modem: it will transmit at as high a bit rate as possible, given the characteristics of the actual connection and the capabilities implemented in a particular modem. The percentage of lines over which a given bit rate such as 28.8 kb/s is actually achieved is then a statistical question, which we address at the end of this article.

0163-6804/96/$05.00 © 1996 IEEE

Cerner Corp. - Ex. 1023
Page 28

Symbol rate			Low carrier			High carrier		
S	a	b	Frequency (Hz)	c	d	Frequency (Hz)	c	d
2400	1	1	1600	2	3	1800	3	4
2743	8	7	1646	3	5	1829	2	3
2800	7	6	1680	3	5	1867	2	3
3000	5	4	1800	3	5	2000	2	3
3200	4	3	1829	4	7	1920	3	5
3429	10	7	1959	4	7	1959	4	7

■ **Table 1.** *Symbol rates and carrier frequencies in V. 34.*

Adaptive Bandwidth

The single most significant factor contributing to the increased bit rates of V.34 is the use of the maximum possible bandwidth permitted by the channel.

Like all previous high-speed modem standards, V.34 uses quadrature amplitude modulation (QAM), in which the two components of a two-dimensional symbol are amplitude-modulated on in-phase and quadrature sinusoidal carriers at a common carrier frequency. The nominal Nyquist bandwidth is therefore equal to the symbol rate in Hertz, and the center of the band is at the carrier frequency.

In earlier modem standards such as V.32bis, the nominal bandwidth (symbol rate) was fixed at 2400 Hz with a fixed carrier frequency of 1800 Hz, resulting in a nominal transmission band of 600–3000 Hz. In V.34, however, the bandwidth and carrier frequency are both adaptive, with a maximum bandwidth of about 3429 Hz.

For example, at a rate of 8 b/s/Hz (8 bits per QAM symbol), a symbol rate of 2400 translates to a bit rate of 19,200 b/s, whereas a symbol rate of 3429 Hz translates to 27,429 b/s. Thus, at high bit rates there is a very large payoff from using the greatest possible bandwidth.

There are six symbol rates specified in V.34: 2400, 2743, 2800, 3000, 3200, and 3429 Hz. Of these, the three rates 2400, 3000, and 3200 are required in all V.34-compliant modems, while the remaining three are optional (although most modem manufacturers appear to have implemented 3429 Hz). These symbol rates are all of the form (a/b) x 2400 Hz, where a and b are small integers (see Table 1).

For each possible symbol rate S, one or two carrier frequencies of the form $(c/d)S$ are specified, in Table 1.

The symbol rate and carrier frequency are chosen during initial training. The transmitter sends a "line probing" sequence that generates a set of tones across the maximum possible band so that the signal-to-noise ratio can be measured as a function of frequency. The symbol rate and carrier frequency are then selected according to the results of this probing and the available symbol rates.

Adaptive Bit Rates

For consistency with earlier modem standards and to avoid excessive rate granularity, V.34 supports bit rates that are integer multiples of 2.4 kb/s, initially up to 28.8 kb/s (12 x 2400) and more recently up to 33.6 kb/s (14 x 2400).

Thus, the bit rate is usually not equal to an integer number of bits per symbol. (For example, 33.6 kb/s at a symbol rate of 3429 requires the transmission of 8.4 b/symbol.) A mapping technique involving a large "superframe" is used to accommodate all possible combinations of bit rate and symbol rate.

The bit rate is selected during training, according to the receiver's estimate of the maximum bit rate that can be supported at a reasonably low bit error probability such as 10^{-5}–10^{-6}. During data transmission, there are mechanisms for falling forward or back in bit rate according to the observed apparent error rate. The bit rates in the two directions of transmission may be different when both modems have the ability to support asymmetric bit rates.

Trellis Coding

One of the principal innovations in the V.32 9.6 kb/s modem standard (1984) was the use of trellis-coded modulation (TCM). At that time TCM was new, and a simple eight-state two-dimensional (2-D) trellis code was selected. This code, due to Wei [3], achieved an effective coding gain of about 3.6 dB.

Almost all of the trellis codes proposed for V.34 were four-dimensional (4-D) codes. 4-D codes have a smaller constellation expansion, which helps against certain non-Gaussian impairments, such as unequalized intersymbol interference (ISI). Also, certain 4-D codes have very good coding gains for their complexity. In particular, the 16-state 4-D trellis code of Wei [4] has an effective coding gain of about 4.2 dB, or 0.6 dB better than the V.32 code, with about the same decoding complexity. This code was used in all hardware prototypes for which performance results were submitted during standardization, and was eventually adopted for V.34.

Two additional, more powerful trellis codes were included in V.34 as further options: a 32-state 4-D code due to Williams [5], with an effective coding gain of about 4.5 dB, and a 64-state 4-D code, due again to Wei [6], with an effective coding gain of about 4.7 dB. The 32-state 4-D code involves a novel 4-D lattice partition, and is probably the nicest new code to come out of V.34 development. The 64-state 4-D code is a variant of Wei's original 64-state 4-D code [4], which was redesigned by Wei after Rossin *et al.* [7] discovered a flaw in the original code.

A V.34 transmitter is required to support all three encoders (since encoding is simple); however, a V.34 receiver may support whichever code(s) it likes.

In view of the relatively small returns in coding gain versus decoding complexity of more complex codes, it may be questioned why all three codes were included. The answer has to do partly with the relatively trivial cost in hardware of including multiple encoders, partly with the increased immunity of the more complex codes to impairments other than Gaussian noise, and partly with the dynamics of the standards process.

Shaping via Shell Mapping

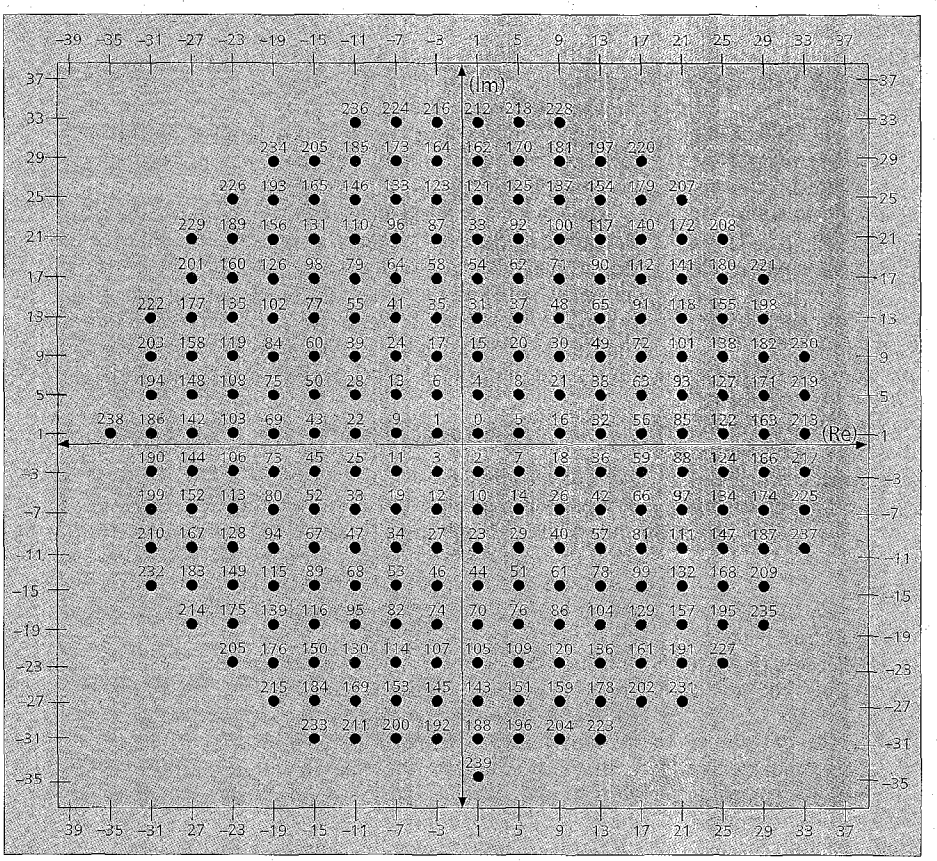
Another advance since V.32 has been the recognition that there is a modest but not insignificant gain in forming signal constellations in high-dimensional spaces to minimize average signal power, quite independent of coding gain. This so-called shaping gain can never be greater than a factor of $\pi e/6$ (1.53 dB); however, it is not difficult to achieve shaping gains on the order of 1 dB.

Regulatory restrictions limit the average power of the signal transmitted on a telephone line. By minimizing average signal power, shaping maximizes the noise margin subject to these restrictions, at the cost of a larger QAM constellation and an increased peak-to-average ratio (PAR), which can lead to greater susceptibility to nonlinear impairments. The optimum shape in theory is a sphere in a high number of dimensions, but spherical constellations are difficult to implement and yield excessive constellation expansion and PAR.

After consideration of several shaping methods, a technique called shell mapping was eventually included in V.34. Shell mapping, which has been developed by many authors [8–12], is an algorithmic method of achieving near-spherical constellation shaping in a high number of dimensions with bounded QAM constellation expansion. V.34 specifies shell mapping in 16 dimensions with QAM constellation expansion limited to about 25 percent, which yields a shaping gain of about 0.8 dB. It also includes a shaping option with essentially no constellation expansion, which still achieves a shaping gain of about 0.2 dB.

The mapping algorithm also supports any integer number of bits per 16 dimensions (8 QAM symbols) up to a certain maximum, which in combination with appropriate framing and switching supports all the combinations of symbol rates and bit rates specified in V.34. V.34 uses nested signal constellations which can be generated easily as a subset of a single QAM superconstellation of 960 points, consisting of the 240 points shown in Fig. 1 and their rotations by 90°, 180°, and 270°. The recent increase of the maximum V.34 bit rate to 33.6 kb/s was accomplished simply by increasing the maximum constellation size to 1664 points.

Finally, V.34 also includes an option called "nonlinear encoding," which is designed to combat PCM quantization noise and nonlinear distortion [13]. These impairments typically cause larger perturbations in the higher-energy signal points, usually radially. If nonlinear encoding is enabled, a memoryless nonlinearity increases the distance between outer signal points at the cost of a slight decrease in distance between inner points.



■ **Figure 1.** *V.34 quarter-superconstellation with 240 signal points. The full superconstellation is obtained by rotating these points by 0°, 90°, 180°, and 270°.*

Equalization and Precoding

All previous high-speed modems, such as V.32bis, use adaptive linear equalizers in the receiver to combat ISI. In these modems, the transmission band is confined to a "sweet spot" of 2400 Hz or less in which it is known a priori that channel attenuation will not be too severe.

In contrast, in V.34 every effort is made to make use of all available bandwidth, including frequencies near the band edges where there can be attenuation of as much as 10–20 dB. In such a situation it is well known that linear equalizers (which essentially invert the channel frequency response) cause significant "noise enhancement."

It is also well known that a decision-feedback equalizer (DFE) is well suited to such channels. However, it is not possible to combine coding with a DFE straightforwardly because decision feedback requires immediate decisions, whereas coding inherently involves decoding delay. The solution to this problem involves putting the feedback part of the DFE into the transmitter. Techniques of this type, which were devised many years ago for uncoded transmission by Tomlinson and Harashima [14–16], are called "equalization via precoding" or simply "precoding."

In the course of the V.34 development, a series of alternative forms of precoding were developed, each superior to its predecessor [17–22]. This development was an outstanding example of different contributors building on each others' work to develop a series of improvements that none would likely have arrived at individually.

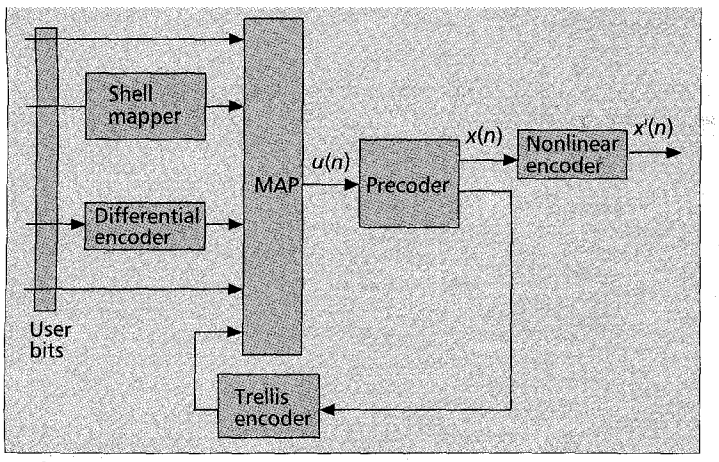
The V.34 standard provides for a simple three-tap precoding filter (representing the "feedback filter" in a DFE) whose coefficients are determined during initial training by the receiver and sent to the transmitter. The "feedforward filter" in the DFE is realized as an adaptive linear equalizer in the receiver and continues to adapt during data transmission.

To combat nonlinear impairments, V.34 also provides for optional transmitter pre-emphasis, which tends to reduce the signal level when the signal reaches the nonlinearity, and thereby to reduce the distortion seen by the receiver in situations where the nonlinearity is after the channel filter. V.34 also allows the receiver to request the transmitter to reduce the transmit power below its maximum allowed level to reduce the effect of nonlinearities.

V.34 Transmitter

We now give an overview of the connections between these various elements in a V.34 transmitter, as illustrated in Fig. 2.

Given a particular bit rate and symbol rate (e.g., 28.8 kb/s at 3200 Hz), the transmitter must send a certain number of bits per QAM symbol (e.g., 9 in this case). If shaping is selected, the QAM constellation expansion is 75 percent over the constellation required to sup-



■ **Figure 2.** *V.34 transmitter.*

port uncoded transmission, 50 percent due to trellis coding and 25 percent to shaping. For 9 b/symbol, for example, an 896-point signal constellation is required rather than the 512-point constellation required with no coding.

To implement shell mapping for this case, the signal constellation is partitioned into 14 equal-sized concentric rings of 64 points each. The encoder collects incoming user bits in frames of 72 bits over 8 symbol periods. Of these, 28 user bits are used by the shell mapper, which selects one of the 14 rings for each of the 8 QAM symbols in the frame; the rings are chosen so that the 16-dimensional signal point lies in a quasi-spherical 16-D region. The remaining 44 user bits, together with 4 coded bits generated by the trellis encoder, are then used 6 at a time to select one of the 64 transmitted signal points $u(n)$ in each selected ring. Some of the user bits are differentially encoded to ensure that the system is rotationally invariant.

The precoder and the trellis encoder are connected in a novel feedback arrangement, shown in Fig. 2. The current output of the trellis encoder determines the entire set of valid 4-D signal point pairs that may be selected by the precoder; the precoder selects one such 4-D pair, and the encoder then deduces the encoder state transition corresponding to that pair, thus determining its next state.

This arrangement minimizes the "dither" $d(n) = x(n) - u(n)$ that needs to be added to the signal point $u(n)$ by the precoder to generate the transmitted signal point $x(n)$. The dither sequence is selected so that after transmission through a channel with known linear distortion, the output sequence is equal to a valid trellis-coded sequence plus noise, and therefore can be decoded by a standard trellis decoder. The precoder minimizes the average transmitted power subject to this constraint, and ensures that the input data may be recovered from the decoded data in the receiver by a simple feedbackfree operation, so error propagation is limited.

After precoding, the sequence $x(n)$ may be optionally further modified by a nonlinear encoder to counter the effects of nonlinear distortion. The resulting sequence $x'(n)$ is typically filtered by a pulse-shaping filter, which, depending on the selections made by the receiver, may simply be a standard square-root-of-Nyquist filter with no spectral shaping, or a filter that provides pre-emphasis according to one of five spectral shapes defined in V.34. The output of the pulse-shaping filter is then modulated at the selected carrier frequency and power level for transmission over the telephone line.

Start-up and Operating Procedures

Recommendation V.34 specifies both duplex and half-duplex operating procedures. Duplex operation is for typical simultaneous two-way data applications, while half-duplex operation (one direction at a time) is primarily for facsimile. Duplex operating procedures include startup, retrain, rate renegotiation, and cleardown. Half-duplex operating procedures include both primary and control channel startup, resynchronization, and retrains.

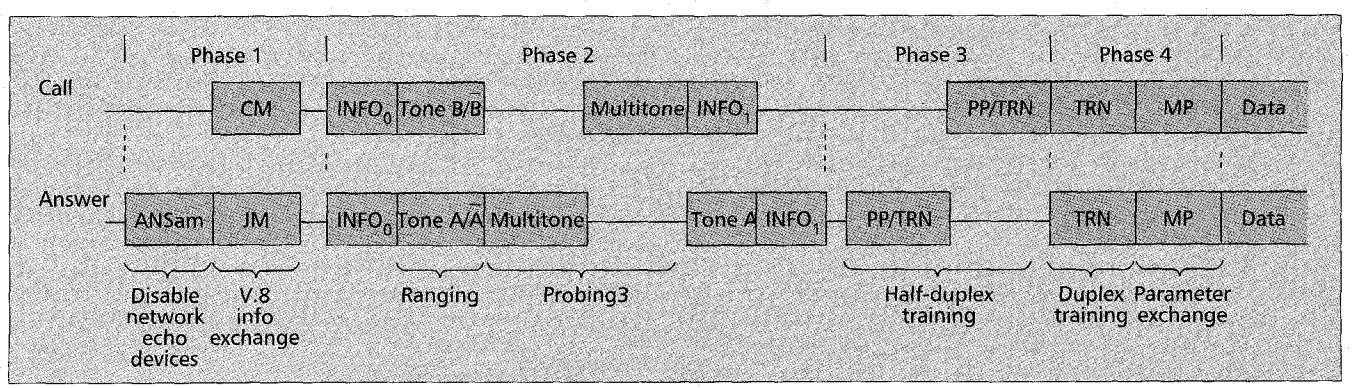
We shall discuss only duplex startup procedures in detail in this article. The half-duplex startup procedures are similar.

All operating procedures are designed to work well both under error-free conditions and in the presence of errors. The goal is to minimize startup time in the absence of errors, while providing slower but robust error recovery mechanisms to cope with errors.

V.34 duplex startup consists of four phases. Figure 3 is a simplified illustration of this sequence.

Phase 1, the network interaction phase, is based on Recommendation V.8 [23]. Its purposes are to disable network echo suppressors and cancellers, to provide for limited terminal selection between modem, fax, videotex, and text telephony, to provide for an improved "automoding" procedure over that in V.32bis (i.e., to determine the V-series modem recommendation used by the remote modem), and to allow network circuit multiplication equipment (CME) to switch in demodulation-remodulation (demod/remod) facilities; such facilities are used today for facsimile on international circuits in order to save bandwidth.

In order to allow for a V.8-capable calling modem to be able to differentiate between a V.8-capable answering modem and an older modem, while not causing existing equipment to malfunction, the normal 2100 Hz answer tone is modulated with a distinctive 15 Hz sine wave with a 20 percent modulation index.



■ **Figure 3.** *Simplified duplex startup sequence.*

Recommendation V.8bis, recently approved by the ITU, improves over V.8, particularly in multimedia applications, by allowing a call to start in the regular voice mode and then smoothly transition into a voice-and-data or video mode.

Phase 2, the ranging and probing phase, consists of an initial information exchange ($INFO_0$), ranging and probing sequences, and a second information exchange ($INFO_1$). The information exchanges use 600 bit/s frequency-division multiplexed (FDM) differential phase shift keying (DPSK) modulation at carrier frequencies of 1200 Hz and 2400 Hz.

$INFO_0$ is used to convey capability information such as which optional symbol rates are supported, regulatory bandwidth restrictions, and maximum allowed asymmetry of the transmit and receive symbol rates.

Ranging uses phase reversals during tone transmission to determine the round-trip delay of the connection for proper placement of the modem's far-end echo canceller taps.

Probing is used to determine the channel characteristics. The probing signal consists of a set of tones of equal amplitude spaced 150 Hz apart at frequencies from 150 Hz to 3750 Hz. This signal may be used to measure the amount of amplitude distortion across the band, the SNR across the band, and the frequency offset for a given channel. Tones at 900, 1200, 1800, and 2400 Hz are omitted to allow for the measurement of the level of intermodulation distortion products. The probing signal is transmitted first at 6 dB above the nominal power level and then at the nominal power level. This permits the measurement of overall nonlinearity, which in turn can be used to determine the amount of power reduction that a modem requests of a far-end transmitter.

$INFO_1$ is used to convey the results of probing measurements in terms of projected maximum bit rate, and to select the symbol rate, carrier frequency, pre-emphasis filter, and a range of power reduction to be used by each modem. $INFO_1$ also indicates the duration of the optional Phase 3 echo-canceller training sequence.

Phase 3, the equalizer and echo-canceller half-duplex training phase, consists of a series of signals transmitted first by the answering modem and then by the calling modem. Each series consists of an optional manufacturer-defined echo-canceller training signal, a short periodic sequence for fast equalizer training, a sequence of scrambled binary 1s for fine tuning of the equalizer and echo canceller, and a repeating 16-bit scrambled sequence indicating the constellation size that will be used during Phase 4 of the startup procedure. These scrambled sequences are transmitted using a four-point constellation. The repeating 16-bit scrambled sequence continues until interrupted by reception of a tone from the remote modem indicating that its receiver has completed equalizer training.

Phase 4, the final duplex training phase, consists of a sequence of scrambled binary 1s and a modulation parameter exchange, following which the selected modulation features and options are enabled. The sequence of scrambled binary 1s uses either a 4- or 16-point QAM constellation and is used to train precoder coefficients and fine-tune the equalizer and echo canceller. If a 16-point constellation is used, measurements of nonlinear distortion may also be made, which may be used to decide whether or not to invoke the nonlinear encoding option.

The modulation parameter exchange is used to indicate the final modem bit rates (based on duplex performance measurements), exchange the precoder coefficients to be used by each modem, select the trellis encoder and the degree of nonlinear encoding and shaping to be used by each modem, and indicate the capability to use a 200 b/s auxiliary channel. The modulation parameter exchange is terminated by a well-defined marker, after which all of the selected modulation features are immediately enabled and the transmission of user data commences.

The total startup time for duplex operation is nominally about 10 s, but may vary from 4 s plus 9 round-trip delays to 13 s plus 12 round-trip delays, depending on implementation.

Either modem may initiate a retrain during duplex data mode. There are two types : long retrain, starting at Phase 2 of duplex startup, and short retrain/rate renegotiation, starting at the beginning of Phase 4 of duplex startup. The short retrain is typically used to change bit rates (up or down) or resynchronize the receiver without performing a full retrain. The long retrain is typically used when a short retrain has failed, or a modem wishes to change its symbol rate or carrier frequency, or needs to retrain its echo canceller.

PERFORMANCE AND TESTING

The technologies used in V.34 were selected on the basis of their ability to expand the percentage of connections over which high bit rates could be achieved on the actual GSTN.

It was understood that traditional measurements of bit error probability as a function of SNR on a linear Gaussian channel could not be the whole story of predicting actual GSTN performance. Therefore, a statistical model of actual transmission facilities and loops was developed to evaluate V.34 technologies [24]. This novel approach was based in large part on available GSTN survey data in the United States for the likelihood of occurrence of certain impairments as a function of transmission facility type [25] and the likelihood of occurrence of local loops [26]. This 1984 data base was updated using 1989 Federal Communications Commission (FCC) petition filings [27] detailing the percentages of types of switches installed in the U.S. GSTN and projections of future installations up to 1994.

The performance evaluation criterion was reduced to a single percentage score representing the network model coverage (NMC). This criterion has proved to be a reasonably accurate predictor of how a given V.34 modem will operate in the real GSTN, and has been widely adopted by laboratories that perform V.34 modem evaluations. It has also been embodied in the U.S. standard Telecommunications Industry Association (TIA) TSB-37A in the United States and the international ITU-T Recommendation V.56bis [28].

The V.56bis test suite is a network model comprising 168 connection types in total, based on 7 local loop combinations and 24 end-office-to-end-office combinations. All combinations are weighted based on their likelihood of occurrence (LOO) in the network. The NMC is the total LOO for which a specified block error rate (BLER) can be achieved. The behavior of NMC as a function of BLER can be used as a more realistic replacement for the traditional waterfall curves of BLER versus SNR for ideal Gaussian channels, although of course the amount of testing required to develop these numbers is much greater.

In practice, different V.34 implementations achieve different bit rates on the same connection, depending on which options are implemented and how they are implemented. It thus becomes rather difficult to give a precise answer to the question, "How fast do V.34 modems really go?" Experience indicates a rate of 28.8 kb/s can be achieved over the majority of lines in North America, Europe, and Japan, and 24 kb/s over practically all lines except for intercontinental links with ADPCM, where 16.8 or 19.2 kb/s is often the practical limit.

WHAT'S NEXT FOR MODEMS?

Some say that "V.fast = V.last" — V.34 is the ultimate modem standard. History tells us this is not likely to be the case. Some modem manufacturers have recently announced plans to develop super-high-speed modems that can operate at rates well above 33.6 kb/s. The technical concept behind these new modems is described by Humblet and Troulis in this issue [29]. With the growing popularity of multimedia applications and the Internet, modem engineers will continue to seek ways to extend the life of voiceband modems.

A project to develop an enhanced version of V.34 is now underway in the ITU-T. This work could lead to faster startup and the possibility of seamless rate switching, which are important enhancements for the multimedia applications discussed in this issue. Other features will surely be added over time.

ACKNOWLEDGMENTS

The V.34 standard was created through the efforts of a great many people in many companies over several years, including a number of our colleagues in Motorola, as well as other participants in the TIA 30.1 and ITU-T Study Group 14 standards meetings. We cannot possibly acknowledge all these contributors individually; we can only repeat that we regard V.34 as a major triumph of the standards process. We would also like to acknowledge the helpful suggestions of the reviewers.

REFERENCES

[1] ITU-T Rec. V.34, "A Modem Operating at Data Signalling Rates of Up to 28,800 bit/s for Use on the General Switched Telephone Network and on Leased Point-to-Point 2-Wire Telephone-Type Circuits," 1994.
[2] CCITT Rec. V.32bis, "A Duplex Modem Operating at Data Signalling Rates of Up to 14 400 bit/s for Use on the General Switched Telephone Network and on Leased Point-to-Point 2-wire Telephone-Type Circuits," 1991.
[3] L.-F. Wei, "Rotationally Invariant Convolutional Channel Encoding with Expanded Signal Space, Part II: Nonlinear Codes," *IEEE JSAC*, vol. SAC-2, Sept. 1984, pp. 672–86.
[4] L.-F. Wei, "Trellis-Coded Modulation Using Multidimensional Constellations," *IEEE Trans. Info. Theory*, vol. IT-33, July 1987, pp. 483–501.
[5] R. G. C. Williams, "A Trellis Code for V.fast," CCITT V.fast Rapporteur meeting, Bath, U.K., Sept. 1992.
[6] L.-F. Wei, "A New 4D 64-State Rate-4/5 Trellis Code," Cont. D19, ITU-T SG 14, Geneva, Switzerland, Sept. 1993.
[7] E. J. Rossin and C. D. Heegard, private communication, Mar. 1993.
[8] G. R. Lang and F. M. Longstaff, "A Leech Lattice Modem," *IEEE JSAC*, vol. 7, Aug. 1989, pp. 968–73.
[9] P. Fortier, A. Ruiz, and J. M. Cioffi, "Multidimensional Signal Sets through the Shell Construction for Parallel Channels," *IEEE Trans. Commun.*, vol. 40, Mar. 1992, pp. 500–12.
[10] F. R. Kschischang and S. Pasupathy, "Optimal Nonuniform Signaling for Gaussian Channels," *IEEE Trans. Info. Theory*, vol. 39, May 1993, pp. 913–29.
[11] A. K. Khandani and P. Kabal, "Shaping Multidimensional Signal Spaces — Part I: Optimum Shaping, Shell Mapping," *IEEE Trans. Info. Theory*, vol. 39, Nov. 1993, pp. 1799-1808.
[12] R. Laroia, N. Farvardin, and S. Tretter, "On Optimal Shaping of Multi-dimensional Constellations," *IEEE Trans. Info. Theory*, vol. 40, July 1994, pp. 1044–56.
[13] W. Betts, A. R. Calderbank, and R. Laroia, "Performance of Nonuniform Constellations on the Gaussian Channel," *IEEE Trans. Info. Theory*, vol. 40, Sept. 1994, pp. 1633–38.
[14] M. Tomlinson, "New Automatic Equalizer Employing Modulo Arithmetic," *Elect. Lett.*, vol. 7, Mar. 1971, pp. 138–39.
[15] H. Harashima and H. Miyakawa, "Matched-Transmission Technique for Channels with Intersymbol Interference," *IEEE Trans. Commun.*, vol. COM-20, Aug. 1972, pp. 774–80.
[16] G. D. Forney, Jr. and M. V. Eyuboglu, "Combined Equalization and Coding Using Precoding," *IEEE Commun. Mag.*, vol. 29, no. 12, Dec. 1991, pp. 25–34.
[17] M. V. Eyuboglu and G. D. Forney, Jr., "Trellis Precoding: Combined Coding, Precoding and Shaping for Intersymbol Interference Channels," *IEEE Trans. Info. Theory*, vol. 36, Mar. 1992, pp. 301–14.
[18] Motorola Information Systems Group, "A Flexible Form of Precoding for V.fast," Cont. D194, CCITT SG XVII, Geneva, Switzerland, June 1992.
[19] General Datacomm, Inc., "Distribution-Preserving Tomlinson Algorithm," Cont. D189, CCITT SG XVII, Geneva, Switzerland, June 1992.
[20] R. Laroia, S. Tretter, and N. Farvardin, "A Simple and Effective Precoding Scheme for Noise Whitening on Intersymbol Interference Channels," *IEEE Trans. Commun.*, vol. 41, Oct. 1993, pp. 1460–63.
[21] AT&T, "ISI Coder – Combined Coding & Precoding," Cont. D24, ITU-T SG 14, Geneva, Switzerland, Sept. 1993.
[22] United States of America, "Recommendation for V.fast — Precoding," Cont. D17, ITU-T SG 14, Geneva, Switzerland, Sept. 1993.
[23] ITU-T Rec. V.8, "Procedures for Starting Sessions of Data Transmission over the General Switched Telephone Network," 1994.
[24] AT&T and Motorola Information Systems Group, "Proposed update of V.fast test suite," Cont. D229, CCITT SG XVII, Geneva, Switzerland, Jan. 1993.
[25] Bellcore, "1983 Exchange Access Study: Analog Voice and Voiceband Data Transmission Performance Characterization of the Exchange Access Plant," TR-NPL-000037, issue 1, 1984.
[26] Bellcore, "Characterization of Subscriber Loops for Voice and ISDN Services (1983 Subscriber Loop Survey Results)," ST-TSY-000041, 1984.
[27] FCC, "FCC (Common Carrier Bureau) Petition Filings by the Regional Bell Telephone Operating Companies (RBOCs) for Exchange Carrier Facility Equipment," Docket no. 89-624, 1989.
[28] ITU-T Rec. V.56bis, "Network Transmission Model for Evaluating Modem Performance over 2-Wire Voice Grade Connections," 1995.
[29] P. Humblet and M. Troulis, "The Information Driveway," *IEEE Commun. Mag.*, this issue.

BIOGRAPHIES

G. DAVID FORNEY, JR. [F '73] received the B.S.E. degree in electrical engineering from Princeton University, New Jersey, in 1961, and the M.S. and Sc.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1963 and 1965, respectively. In 1965 he joined the Codex Corporation, and became a vice president and director in 1970. From 1982 to 1986 he was vice president and director of technology and planning of the Motorola Information Systems Group, Mansfield, Massachusetts. He is currently a vice president of technical staff at Motorola, Inc. He has been an adjunct professor at MIT from 1978 to 1980 and 1996, as well as a visiting scientist in 1991 and 1994–1996, and was also a visiting scientist at Stanford University, Stanford, California, in 1972–1973 and 1990. Dr. Forney was Editor of *IEEE Transactions on Information Theory* from 1970 to 1973. He was a member of the Board of Governors of the IEEE Information Theory Society during 1970–1976 and 1986–1994, and was President in 1992. He has been awarded the 1970 IEEE Information Theory Group Prize Paper Award, the 1972 IEEE Browder J. Thompson Memorial Prize Paper Award, the 1990 IEEE Donald G. Fink Prize Paper Award, the 1992 IEEE Edison Medal, the 1995 IEEE Information Theory Society Claude E. Shannon Award, the 1996 Christopher Columbus International Communications Award, and the 1997 Marconi International Fellowship. He was elected a member of the National Academy of Engineering (U.S.) in 1983, and an honorary member of the Popov Society (Russia) in 1994.

LES BROWN was born in Vancouver, Canada, on December 8, 1949. He received the B.A.Sc. degree in electrical engineering from the University of Toronto in 1972, and the MASc. degree in electrical engineering, specializing in active filter design, from the University of Toronto in 1977. In 1981 he joined ESE Ltd., then a subsidiary of Codex Corporation (now Motorola ISG), and currently holds the position of principal staff engineer. Mr. Brown has been involved in all facets of modem development over the years. He is also the Motorola ISG representative in modem-related standards committees, holding several positions of leadership, including the chairmanship of TIA TR-30.1 and several rapporteurships in ITU-T SG14.

VEDAT EYUBOGLU [F] received his B.S. degree from Bogazici University, Istanbul, Turkey, and M.S. and Ph.D. degrees from Rensselaer Polytechnic Institute in Troy, New York, in 1980 and 1984, respectively. He is a vice president at Motorola's Information Systems Group and is the manager of the Signal Processing R&AD Department in Mansfield, Mass. His group is involved in research and development on high-speed modems, multimedia and broadband access technologies, in particular for Hybrid Fiber/Coax Networks. Dr. Eyuboglu was a key contributor to the development of the V.34 high-speed modem standard and the H.223 multimedia multiplexing protocol of H.324. Dr. Eyuboglu is a Dan Noble Fellow and Distinguished Innovator of Motorola. He holds over 25 issued/pending patents.

JACK L. MORAN is a principal staff engineer at Motorola Information Systems Group in Mansfield, Massachusetts. He has worked as an electrical engineer in the communications industry for over 19 years. He has been with Motorola for 15 years where he worked primarily on analog hardware design. Over the past 12 years he has conducted numerous surveys to characterize the public telecommunications network throughout the world and has applied this knowledge to develop network models to evaluate modem performance. Mr. Moran has been an active participant in IEEE Committee P 743, TIA TR-30, T1A1, and ITU-T SG 14. He was Editor and chief sponsor for TIA Telecommunications System Bulletin TSB-37-A and ITU-T Recommendation V.56 bis. More recently, Mr. Moran has been involved in the design and testing of cable modems. He holds several patents in the field of data communications.