INTRODUCTION BY EDITOR

As noted in the following paper by Fred Andrews on the genesis of the PCM-based T1 system at Bell Labs, pulse code modulation had been invented by the British engineer Alex Reeves in 1937. But, as Reeves himself noted much later, in 1964, this was an invention that came before the technology required to implement it commercially was available. It wasn't until the transistor revolution of the late 1940s and early 1950s that PCM, with all the advantages it offered of fully digital, relatively low-noise transmission and regeneration, could be developed into the hugely successful T1 system. Andrews, himself a leading engineer on the T1 project, tells the story of its development at Bell Labs with authenticity and gusto. He describes the efforts of the young engineers like himself working so enthusiastically on the project, solving problems of implementation as they came along. He then describes the follow-up digital transmission hierarchy carrying ever higher transmission bit rates, and the digital telephone switching systems this made possible. These developments led for the first time to all-digital telephone communication systems worldwide, first carrying voice and then data as well. This history of T1 and transmission systems makes for a fascinating story. I commend it to you for your reading and study.

—Mischa Schwartz

EARLY T-CARRIER HISTORY

FREDERICK T. ANDREWS

ABSTRACT

In the 1950s AT&T faced the challenge of providing for the growth of traffic in the local telephone plant at lower costs than by simply continuing to expand the number of voice cable pairs between switching offices. Engineers identified pulse code modulation, PCM, as offering a robust solution to the problem and launched a system development effort in 1956. Some difficult technical problems were identified; the solutions were found and incorporated into a new system design. The product was T-Carrier, which was introduced into service in 1962. This new application of PCM was highly successful and the technology was licensed to the independent telephone industry. The race toward an all-digital network of digital transmission and switching was off and running.

INTRODUCTION

The year 1956 marked the beginnings of an exploratory development project at Bell Laboratories that was to have major impact on the future of communications. This project introduced a transmission system architecture fundamentally different from what had been used in telecommunications networks in the past. It had far reaching consequences, both beneficial and unanticipated at the time. Before getting into the specifics of the project we should review the state of the evolution of telecommunications networks in the 1950s. What was driving research and development, and what were the technological enablers of future progress?

A decade after the end of World War II, global economies were booming and the rapid growth of telephone voice traffic demanded ever more capable telephone networks. The telephone carriers and their equipment suppliers responded by deploying the best technologies available to support that growth in traffic capacity demand at reasonable cost. AT&T, the dominant carrier at the time, had led the way with a series of frequency-division multiplex, FDM, carrier systems optimized for specific applications within the telephone network.

The year 1956 marked the beginnings of an exploratory development project at Bell Laboratories that was to have major impact on the future of communications.

One of the simplest of these FDM systems was N1 Carrier, an example of the state of the art at that time. Twelve voice channels were stacked much like double-sideband radio signals in a 96 kHz wide band from 172 to 268 kHz. This multiplex of 12 channels was transmitted over the cable pairs heretofore used for a single voice channel. The 36 to 132 kHz band was used for the opposite direction on a second cable pair, thereby avoiding cross coupling within a cable sheath. Most of the terminal electronics — carrier generators, modulators, and band-limiting filters — was on a per-channel basis. It was the sharing over multiple channels of the line facilities — cable pairs and line repeaters — that led to cost savings.

Later versions of N-carrier used single-sideband modulation to double the capacity to 24. The N-carrier family of systems found extensive applications on 24- to 16-gauge cable pairs first in urban areas and later over distances as great as 250 miles. The contemporary L-carrier series used similar FDM technology to stack thousands of single sideband voice circuits on lower-loss coaxial cable for spanning cross-continental distances. The first transatlantic undersea coaxial cable system for voice, TAT-1, went into service in 1956.

The first versions of N carrier and L carrier were based on well-proven electron tube technology. The work to transistorize the FDM architectures for short-haul applications was already gaining headway in 1956. The advantage of the transistor was not so much lower cost in these applications but lower power consumption, reliability, and size of the electronic equipment. Savings in the shared medium and the repeaters exceeded the costs of the per channel elements — modulators, demodulators, and filters — required to implement FDM. But the savings were not enough to make FDM systems economic for the shorter distances between local switching offices.

What to do? The need for more system capacity was definitely there, propelled by an ever increasing quality of transmission and in the ease of placing calls using direct distance dialing instead of manual operators.

To better understand the problem we need a few words about how telephone networks were configured at that time. The basic network node was (and is) the local central office to which residential and business telephones and private branch exchanges are connected. Chances are that much of the traffic

(Continued on page 14)

 0163-6804/10/$25.00 © 2010 IEEE

(Continued from page 12)

originated by individual callers were to other local central offices located in the same or adjacent towns, or the same or adjacent counties, and so on.

The plan for interconnection of central offices is by a combination of direct circuits and circuits switched through intermediate or tandem nodes. This plan is far more economic than having all offices directly interconnected. This principle of connection through multiple links is followed through a five-layer hierarchy, with the bottom layer, consisting of "class 5" offices, being by far the most numerous.

In urban population centers class-five offices and those serving as tandems are typically 10–20 miles apart and are interconnected with groups of several hundreds circuits. There were at the time some 10,000 class 5 offices, and the large number of interconnecting circuits had not yielded to electronically based cost reductions. On the contrary, they were still being provided by individual pairs of wires, as visualized by Alexander Graham Bell in the plan he devised on his honeymoon trip in 1876.

Of course, there had been a myriad of improvements in cable construction affecting both transmission performance and reliability. For example, within a lead or polyethylene cable sheath, the individual wire pairs were bound together with other similar pairs, each with a different twist length. This feature minimized inductive and capacitive coupling and "crosstalk" from one pair to another, and made possible the use of these cables for the N-carrier family of systems discussed earlier. As said before, the cost of the multiplexing gear was higher than the cost of the cable pairs that could be saved over the typical distances between class 5 offices.

Is there a better way to exploit the still higher speed capabilities of ordinary cable pairs and save in the costs of multiplexing terminal equipment? As it turns out, yes, there is. The answer lay in a combination of time-division multiplexing (TDM) and pulse code modulation (PCM). TDM offered the possibility of lower-cost multiplexing and PCM the improvement in tolerance of the noise found on cable pairs in the class-five office environment [1]. With the use of fully regenerative repeaters along the line facility, the exponential accumulation of noise could be totally avoided. In fact, the system noise would be the quantizing errors introduced by the coding-decoding function.

THE T-CARRIER CONCEPTS

At the time, research people in the Bell Labs transmission research organization — Robert Carbrey and Fred Kammerer — were building an experimental system that showed promise for just the application the systems engineers and developers had in mind. It was the specific application, not the TDM/PCM principles involved, that was particularly novel. Alec H. Reeves, a British engineer working for Standard Communications Laboratories of the International Telephone and Telegraph Corporation, invented PCM in May 1937 while at the ITT laboratories in Paris [2]. It soon was apparent that PCM was the ideal digital format to use in this difficult, very noisy local cable environment.

The Bell Labs researchers had put together a skeleton system they were demonstrating over a reel of cable stored in the basement of Bell Labs at Murray Hill, New Jersey. The reel of cable replicated the transmission characteristics of cables found in the telephone plant, but not the typically noisy environment.

Eric Sumner headed the exploratory development department (of which I was a member), and Dan Hoth headed the system engineering department that were asked to examine these concepts and take them a step further toward actual application in the real telephone plant.

A number of desirable features combined to make TDM/PCM appear practical and perhaps even economically attractive for use on the cable pairs. Here are the principles on which the research experiment was based:

1 Most fundamental was the sampling theorem, often attributed to Nyquist. Actually, the theorem was first demonstrated in a 1915 publication of the Royal Society of Edinburgh [3]. It showed that one can fully reconstruct a signal from samples of its amplitude taken at a rate at least twice that of the highest frequency component in the signal. That meant an 8 kHz sampling rate for a telephone voice signal having an upper frequency limit of 4 kHz.
2 The voice signal would be kept within a 4 kHz limit by low pass filters of simple design. (Remember that an FDM system required a separate, more complicated, bandpass filter design for each channel.) This is a key cost saving of TDM over FDM.
3 The amplitude of the voice samples would be adequately represented by a 7-bit binary code with 128 code levels, provided that the levels were assigned nonlinearly over the amplitude range of the signal sample. The codes are an approximation of the actual sample level, the error resulting in what is called quantizing noise. Using more coding levels at lower sample amplitudes preserves a relatively constant quantizing noise ratio over a wide range of input levels. This nonlinear coding is similar in results to the syllabic compression-expansion (companding) used in FDM systems. Here it is called instantaneous companding. The combination of 7-digit encoding and instantaneous companding would keep the overall quantizing noise within acceptable limits for circuits in class 5 office environments. The instantaneous companding and encoding/decoding circuits would be shared over all channels, another key cost saving over FDM systems.
4 Digital regenerators, installed at approximately one mile intervals and powered over the line, would recover a timing signal from the incoming pulse stream, detect the presence or absence of a pulse in a given time slot, and faithfully regenerate the incoming pulse pattern. There would be no accumulation of pulse distortion or noise.

Given these basic concepts from research, there were still many issues to be addressed before we had a practical system for the local circuit application we had in mind. Identifying and resolving these issues was a challenging job for our systems engineering and systems development groups. But that is what is expected in the handoff from research to development.

THE EXPERIMENTAL DEVELOPMENT PROJECT

Our goal was not to duplicate the research experiment but to construct a further experiment suitable for testing in a real world environment. The development and trial was to be accomplished in roughly a two-year timeframe. As you can imagine, meeting this schedule required many new design and development tasks going on in parallel, much of the work done by a cadre of young inexperienced engineers. What these engineers lacked in know-how they made up for with their enthusiasm for developing a brand new system. These were exciting times for us members of the team. And remember that while the transistor had been invented eight years earlier, its most notable commercial application up until that time had been the transistor radio. The device speeds required to implement the common circuit functions challenged the state of the transistor art in 1956.

One of the first questions to be addressed was how much channel sharing of common equipment made sense for this local circuit application. Systems engineers and developers agreed that 24 channels seemed about right, from the standpoint of the application and the current device capabilities. We also decided to devote a whole bit per channel sample solely for the out-of-band signaling required for call control. This was a lavish use of bits for a slow signaling function, but the excess capacity proved very useful for other purposes in later systems. As a result of these design decisions, the overall bit rate for a 24-channel system was 1.536 Mb/s (24 channels × 8 b/channel × 8 kHz sampling rate). But wait, a bit more was required for yet another function.

A feature not inherited from research was a means of synchronizing the time-division multiplexers at opposite ends of the system. This synchronization is needed to prevent the signal sent on channel x from being demultiplexed onto channel y. This time we decided to avoid adding a lot more bits to what now seemed a very large total. Henry Mann and I were agonizing over this problem in his office one day when we realized that none of the 192 bits in the frame that had been defined would be alternately on and off except briefly by chance. We simply needed to add a forced on/off bit to the 192-bit frame and a suitable detector to find the starting point of a frame. That is how the frame got to be 193 bits instead of 192, and the bit rate became 1.544 Mb/s. (Henry and I later took a lot of kidding for giving the frame an inconvenient prime number of bits, but the idea survived.).

A key requirement of this application was that the system could be applied to the cable pairs already in place for handling voice traffic. The only conditioning of cable pairs that was contemplated before the installation of the digital repeaters was the removal of any loading coils.

For those of you not familiar with such ancient art, loading coils have magnetic cores with balanced windings for the two conductors in a pair. The coils were typically used to insert 88 mH of inductance at 6000 ft intervals, which combined with the capacitance between the conductors of the pair to form a distributed low pass filter. This filter had less loss in its passband than untreated pairs, thereby extending the range of voice transmission. Of course, it was necessary to remove the load coils in order to transmit 1.544 Mb/s signals.

Is there a better way to exploit the still higher speed capabilities of ordinary cable pairs and save in the costs of multiplexing terminal equipment? As it turns out, yes, there is.

Very little was known at the time about the transmission characteristics of cable pairs at these frequencies. In particular, how much coupling could one expect between cable pairs in the same binder group? How much between pairs in adjacent binder groups? This coupling would determine the rules for assigning digital lines to cable pairs and whether the system could work at all. More about that later.

But first, let us consider what may appear to be a mundane problem. A string of on/off pulses of the kind generated by our system has a slowly moving low-frequency component, which depends on the short-term average density of 1s and 0s. This leads to a wandering up and down relative to the decision threshold in a digital repeater. This is called baseline wander and erodes the noise margin in a major way. The research repeater used a DC restoration scheme that was far from perfect in minimizing the effects of this baseline wander.

One day the proverbial light bulb went on and we had the idea that solved this and, as it turns out, several other problems. If 1s were transmitted by alternate positive and negative pulses and a 0 by no pulse, there would be virtually no wander. It was far easier to build what we dubbed a bipolar repeater than to implement some fancy baseline wander correction scheme.

Bob Aaron analyzed the new bipolar proposal, and found an unintended consequence that literally saved the day. At the first field experiment in South Orange, already underway in 1956, the original line format was causing higher than expected crosstalk between pairs. Eric Sumner did some quick extrapolating to a cable full of PCM systems and concluded that the whole system proposal would probably not work in the real world. Bob was quick to show that the new bipolar line signal's frequency components extended to only half the frequency of the original line signal. The resulting lower crosstalk levels saved the day. Lower crosstalk became the most compelling reason for changing to the new bipolar format.

Here is a case of the unintended consequence of a change being very good, not bad. Adam Lender, then at Lenkurt, went on to develop more sophisticated versions of what was really a three-level code. To the best of my knowledge, these variations were never incorporated into a working system.

One other advantage of the bipolar line format soon emerged. By counting violations of the alternating pulse-polarity rule, it was possible to tell when errors were being caused by noise or interference. This feature provided a useful tool in monitoring and maintaining the quality of digital lines. Some form of monitoring is essential so that bad lines can be swapped out and repaired before customer complaints are generated.

The implementation of the companding function turned out to be more difficult than anticipated. Remember that the multiplexed series of samples from 24 channels had to be instantaneously and accurately compressed ahead of the coder function and expanded after the decoder function. A mathematical function, known as the "mu" law, was selected as the desired nonlinearity for minimizing the subjective effect of the quantizing noise inherent in PCM. The trick was to find a nonlinear solid-state circuit element that, suitably biased and temperature controlled, matched the desired mu curve.

By painstaking effort Harold Straub was able to do that. But his design had to include a small temperature-controlled oven to stabilize the nonlinear diode characteristic. A coder/decoder was later invented that implemented the companding curve in a piece-wise linear fashion in the coder/decoder itself. This more robust approach was used in later, more advanced system designs. The engineers who carried the design through final development and the later improved systems did an excellent job of delivering better and better products.

What follows is the chronology of the trials leading to the commercial deployment of this new digital carrier technology [4]:

1956: An initial field experiment in South Orange, New Jersey, leading to the unhappy conclusion that the original unipolar line format would not support the deployment of multiple systems within the same cable sheath. There was too much crosstalk between pairs at the 1.544 Mb/s line rate.

1959: A testbed on a Summit-South Orange route examining all aspects of the repeatered line and confirming the satisfactory performance of the new bipolar system.

1961–1962: Bell Labs tested a prototype system, manufactured by Western Electric, on a route between Newark and Passaic, New Jersey. This was a full commercial service trial.

1962: An aerial cable trial to determine the ability of the system to survive lightning surges and adapt to cable temperature variations. This trial resulted in some changes in the engineering rules for installation.

1962 (Fall): Installation of a full production system on a 13.6 mile route between Chicago and Skokie, Illinois. This successful application led Illinois Bell to order 300 systems for installation in 1963.

I mentioned earlier that the overall goal was to push the economic application of carrier systems down into the low end of the range of circuit lengths — from roughly 20–25 miles to the order of 10 miles. As a very rough rule of thumb, this involved cutting the cost of carrier-derived circuits by about 2:1. Because of concurrent work on reducing the cost of N-carrier, this goal was a moving target. The first credible cost comparison found a cost saving of just 15 percent, but also an expectation of much greater savings from digital technology in the future [5].

In fact, these PCM systems were so successful that designs for other specific applications became very attractive. To accommodate these additions to the product line, the original PCM terminals were designated as D1 banks, and the original digital line was designated as the T1 line. D banks and T lines henceforth evolved more or less independently for different applications. By 1981 T-carrier with several successive generations of terminal and repeater designs was providing over 100,000,000 voice circuit miles in the Bell System, far exceeding initial expectations [6].

This article would not be complete without mentioning the contributions of John Mayo, who took over from me when I assumed broader responsibilities for transmission systems engineering. John played a major role in bringing the system work to practical fruition and was co-recipient with Sumner and Aaron of the 1978 IEEE Alexander Graham Bell medal for their contributions to digital communications. Those thirsting for greater technical depth on the T-carrier system than presented here, including systems diagrams and pictures, should refer to the January 1962, Vol. 41 issue of the Bell System Technical Journal. It contains a series of five articles by the major contributors to the experimental systems work, including Claude Davis, John Mayo, Bob Aaron, Bob Shennum, Jim Gray, Henry Mann, Harold Straube, and Claude Villars. The *Bell System Technical Journal* archive of these articles is now available on the Internet at http:/bstj.bell-labs.com. (For some reason I was unable to find this trove of *BSTJ* information via either a Google or Bing search!)

DIGITAL LINES AND TERMINALS

For your further information, here are the characteristics of the digital lines that were developed for specific applications, some at higher bit rates and for improved cable pairs [7].

T1, 1.544 Mb/s: This, the original line, operates at what is called the DS1 line rate in the North American digital hierarchy. (DS0 is the 64 kb/s rate of an individual channel.) The digital repeaters are typically installed on 22 gauge pulp-insulated cable pairs of the kind found between local switching offices for voice transmission. The pulse format is bipolar, as discussed above, and the two directions of transmission are in non-adjacent binder groups within the same cable sheath. The digital repeaters, installed in waterproof containers, much like inverted stainless steel lobster pots, have a nominal spacing of 6000 ft. End sections are about half that length, because they are located next to the major sources of impulse noise, e.g. switching equipment.

To simplify engineering and installation, all repeaters have equalizers for the longest repeater spacing, and line build-out networks are selected to make all sections look approximately the same length. The latest versions of T1 repeaters include automatic line build-out that is actuated by the incoming pulse stream. Build-out selection is not required, and variations of loss with temperature are compensated.

T1C, 3.152 Mb/s: Favorable engineering and performance experience resulted in a new repeater design operating with the same bipolar format but at twice the bit rate of the original T1. This is accomplished using the same pairs and spacing as in T1. New engineering rules provided greater control of crosstalk and interference, making possible the higher frequencies involved.

Note that the pulse rate of T1C is actually 64 kb/s higher than twice the T1 rate of 1.544 Mb/s. These overhead bits are added for maintaining the multiplex equipment used to combine the signals from two 24-channel D1 banks. The resulting 48-channel system is a more economical way of providing point-to-point circuits, and is possible with improved regenerators and better engineering methods. It is not a defined step in the digital hierarchy.

T2, 6.312 Mb/s: This next offering in the T line series transmits at the DS2 rate, the next step in the formal North American digital hierarchy. The line has the capability of carrying four DS1 signals plus 136 kb/s for multiplexer overhead. Designed for use on special low capacitance 22 gauge pairs, the line uses a separate cable for each direction of transmission. Because of lower loss and better controlled crosstalk, the nominal repeater spacing is 14.8 kft, with closer spacing adjacent to powering stations, switching offices, and aerial cable sections. Up to 250 regenerators, powered over the cable pairs, can be used in tandem.

The regenerators perform the same functions of those on a T1 line. The pulse format is bipolar with an added twist. Whenever six consecutive 0s are sensed, they are replaced with a special code word. This avoids excessive loss of energy in the timing extraction circuit. This is known as the B6ZS format, is a defined feature of the DS2 signal, and is required from all terminals connected to a T2 line.

Now let us turn to characteristics of the family of terminal equipment that evolved, beginning with the first product [8]:

D1A: This is the original bank, aimed at providing 24-channel systems for the so-called exchange trunk environment. The first design went into service in 1962. It achieved its objectives so well that it spawned variants for special services applications with different signaling features. These were designated as D1B and D1C. D1B has four signaling states, still derived from the 8th channel bit, and D1C has the features required for the operation of remote traffic position systems.

D2: Introduced in 1969, D2 had the higher quality required for toll circuit applications. Higher quality was accomplished by using all eight bits assigned to voice channels for encoding the voice samples except for every sixth frame. Then the least significant bit in the code was "borrowed" for transmission of the signaling states. Also, the mu companding curve was approximated by a controlled segment linear companding characteristic rather than by the nonlinear semiconductor diode used in D1. Finally, the encoder and decoder were shared by 96 rather than 24 channels, making it a 96-channel system. It can be connected to either four T1 lines or one T2 line.

D3: The D3 design was motivated by advances in integrated circuit technology and the desire to achieve compatibility with the new 4ESS digital toll switching system. It is a lower-cost better performing replacement for D1 banks in the digital system product line and incorporates the same improvements achieved in D2. Hybrid integrated circuits were used, and the sampling gates and channel filters were moved to the channel unit plug-ins. D3 banks together with T1 lines became more economical than voice frequency cable pairs at distances of 5–8 miles depending on the particular circumstances. When one end terminated on a 4ESS, these systems proved in at any length.

CONCLUDING REMARKS

The introduction of digital switching in the Bell System focused first on toll switching applications with the 4ESS as the vehicle. Meanwhile, the independent telephone industry was moving rapidly toward the use of digital switching in local central offices. The combination of these initiatives created very fertile ground for the growth of digital transmission.

Most intriguing was the possibility of terminating digital lines on both local and toll digital switches, thereby achieving digital transmission from end to end of a connection. Frankly, we hadn't given much thought to that possibility in 1956 when the foundations for T1/D1 were laid. As it turns out, the system choices we made at that time were not too bad. But it was important that the systems designed by independent equipment manufacturers and deployed by the independent carriers mesh with our designs if the goal of end-to-end digital transmission in a national network were to be achieved.

There is a story to be told, from a historical perspective, about how new Internet services have been overlaid on the digital transmission infrastructure that has grown out of T-carrier. I leave this task to future historians.

The AT&T/Justice Department Consent Decree of 1956 required the licensing of our telephone network technology to all comers. In a Letter to the Editor commenting on a history column by Irwin Dorros entitled "Retrospective — 25 Years Later," Jack McDonald mentions how this licensing caused a whole new industry to be formed using the latest technology from Bell Labs [9].

Mr. McDonald goes on to say that Vidar Corporation, where he was vice president for research, licensed T Carrier and sold it to the independent telephone companies. Vidar built improved terminals 1/4th the size of the D1 channel bank. When he asked for Bell System help in creating an end-to-end digital network, he first encountered a lack of enthusiasm for this as a near-term goal, but later received elegant guidance on how the industry might proceed in working together. He comments that "these standards are now the basis for today's seamless network both for wire-line, wireless and the Internet."

Yes, T Carrier had a major impact. My supporting role began with a visit to a research laboratory at Murray Hill, New Jersey in1956. I will always be grateful for having had the opportunity to play a brief part as a hands-on systems designer in this industry-changing development.

As a final note, I should remark that the European telephone industry followed along with its own common standard that turned out to be quite different from ours. For example, their equivalent of the T1/D1 had 32, not 24, channels, and the instantaneous companding curve was "A" law, not "mu" law. Two full channels (16 bits) were used for a common channel signaling systems. This provides greater flexibility than "robbing" the eighth bit and preserves 8-bit words through tandem digital circuit switches.

These changes, made for good reasons in the European environment, complicated achieving end-to-end global connectivity. In the end, the incompatibility problems were resolved, and circuit-switched global digital networks have been successfully implemented. However, packet switching of voice and data via the Internet now appears to be the way of the future. There is a story to be told, from a historical perspective, about how new Internet services have been overlaid on the digital transmission infrastructure that has grown out of T-carrier. I leave this task to future historians.

REFERENCES

[1] M. Schwartz, "Improving the Noise Performance of Communications Systems: 1930s to Early 1940s," *IEEE Commun. Mag.*, vol. 48, no. 5, May 2010, p. 18.
[2] Op cit, M. Schwartz, p. 19
[3] E. T. Whittaker, "Expansion of the Interpolation Theorem," *Proc. Royal Soc. Edinburgh*, 1915, pp. 181–94.
[4] *A History of Engineering and Science in the Bell System: Transmission Technology*, AT&T Bell Laboratories, E. F. O'Neill, Ed., 1985, pp. 559–60.
[5] Op cit, *A History of Engineering and Science in the Bell System*, p. 560.
[6] Op cit, *A History of Engineering and Science in the Bell System*, p. 562.
[7] *Engineering and Operations in the Bell System*, Bell Telephone Laboratories, P. V. Dimock, Ed., 1977, pp. 320–22.
[8] Op cit, *Engineering and Operations in the Bell System*, pp. 358–60.
[9] J. McDonald, Letter to the Editor, Comments on "Retrospective — 25 Years Later" by I. Dorros, *IEEE Commun. Mag.*, vol. 48, no. 1, Jan. 2010, p. 18.

BIOGRAPHY

FREDERICK T. ANDREWS, JR. retired from what is now Telcordia Technologies in 1990 as vice president, technology systems. He began his career at Bell Laboratories in 1948 upon graduation from the Pennsylvania State University with a B.S.E.E. degree. Post-graduate work followed in a program of studies taught by Bell Lab's own expert staff. His early research was in the application of the electronics of the day to the reduction of costs in the local telephone plant, a subject that remained close to his heart throughout his career. In 1956 he assumed leadership of a group charged with the application of digital carrier principles as an alternative to local voice frequency circuits, which is the subject of this column. Broader responsibilities followed from 1958 to 1982, and included transmission planning, military communications engineering, loop electronics development, and switching system engineering. The AT&T divestiture plan of 1982 included a new research and engineering organization, now known as Telcordia Technologies. It provided shared technical support on behalf of the seven spun-off local telephone companies. He headed the vice presidential area responsible for generic requirements and technical analysis of equipment for building telephone networks until his retirement. He has been President of the Communications Society, a member of the IEEE Board, and Chairman of its Strategic Planning Committee. He was elected to the National Academy of Engineering in 1988.