Using Technology to Bring ATM to the Desktop

Richard L. Townsend, Jr.
Jean-Jacques Werner
Mai-Huong Nguyen

One of the newest technologies to impact both network transmission systems and local area networks is asynchronous transfer mode (ATM), due, in part, to the fact that it is a key technology in offering multimedia services. One issue is how to make the transition from the current embedded network switching systems to ATM-based systems, while minimizing the cost of developing a new infrastructure. This paper discusses how ATM capabilities can be extended to the desktop by applying the latest modem technology to send medium-to-high bit rates (51.84 Mb/s) on twisted-pair copper wiring over distances of up to 100 meters.

Market Need

For asynchronous transfer mode (ATM) to become the ubiquitous service envisioned by many, it must include not only the high-volume traffic in backbone networks and in the public switched network, but it must connect directly to the desktop as well. While much of the development effort for ATM has been focused on fiber transmission systems and high-capacity backbone links, advances in integrated circuitry now make it possible to have end-to-end ATM—to the desktop—without requiring new cabling to support multimedia services.

Whereas the costs of shared-use backbone and public switched systems can be amortized across a large number of users, desktop connections on a customer's premises are either individual or shared by several users, making the per-user cost more sensitive. However, in just the past few years, great strides have been made with transmission techniques for sending high-speed analog signals along ordinary twisted pair copper wires, due to the low cost of highly complex silicon chips. Taking advantage of this progress, it is possible that ATM signals can now be transmitted on twisted pairs by using digital signal processing in the receiver and by transforming the ATM digital data signals to *analog-like* components.

Thus, market need and technology are coming together to provide an end-to-end ATM product driven by the increasing needs for more bandwidth for data exchange, such as video and multimedia, and by the availability of technology to transmit high-speed data on existing twisted pair wires.

Key Criteria

What are the key criteria driving this development from the marketplace?

AT&T studies have shown that 97 percent of all desks in commercial buildings can be reached with cable runs of 100 meters or less from hubs located in closets (see Figure 1). The use of existing unshielded twisted-pair cabling, the type of telephone voice-grade wiring that dominates the installed base in the current indoor customer premises environment, provides an opportunity to save on installation costs by not requiring the deployment of a new transmission medium. Thus, the length of existing twisted pairs is a factor.

Using existing two-pair telephone wiring provides a beneficial compromise between the cost of an expensive echo canceler required for a single-pair wiring solution and multiple pairs operating in a ping-pong fashion, where several inexpensive parallel links work in concert unidirectionally—with the direction of all the links revers-

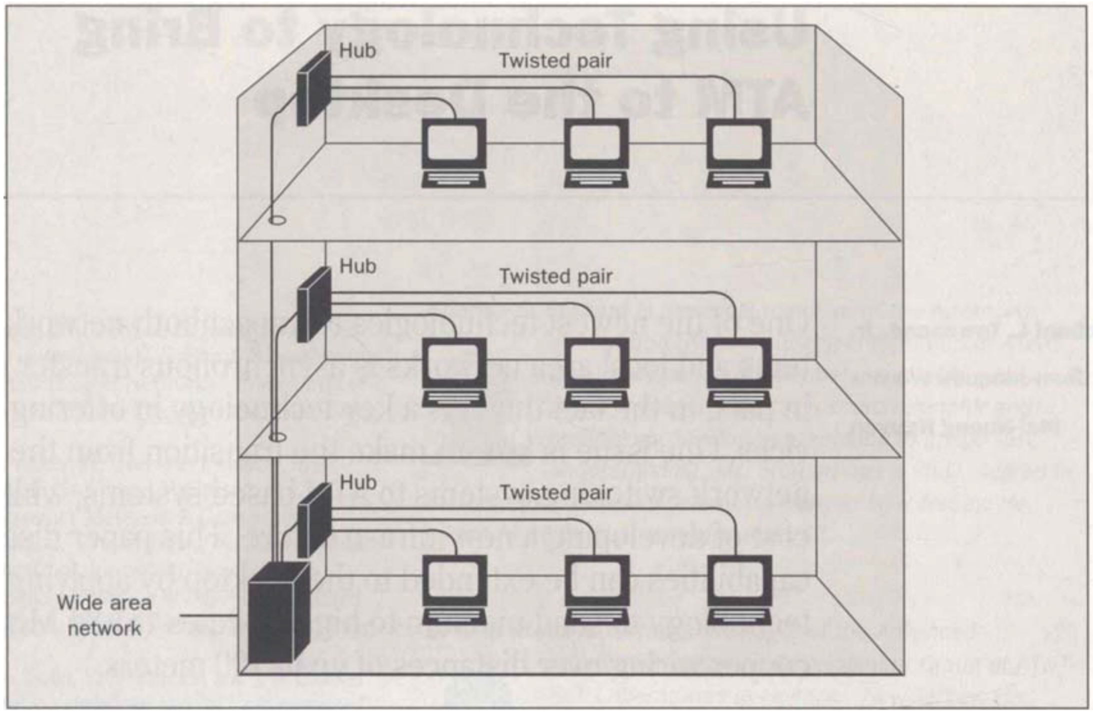

Figure 1. The use of existing unshielded twisted-pair cabling, the type of telephone voice-grade wiring that dominates the installed base in the current indoor customer premises environment, permits ATM to be delivered to the desktop without the cost of rewiring an office building.

ing together—at the loss of simultaneous two-way communication. Thus, minimizing cost is a factor.

About 10 years ago, AT&T pushed for the standardization of indoor wiring comparable to its SYSTIMAX® structured indoor wiring product line, which bundles four twisted-pairs of wires into the cabling that runs to the telephone outlet on the wall. This scheme supported analog and integrated service digital network (ISDN) telephones, as well as low-speed data connections to both circuit-switched and packet-switched networks. Thus, the use of existing standards is a factor.

In addition, the plan supported AT&T's StarLAN network, first at 1 Mb/s and then at 10 Mb/s transmission, at a time when most LANs were still using coaxial cable. High-speed data transmission is now approaching 100 Mb/s on twisted pair, making ATM transmission a practical, cost-effective possibility. Thus, technological advances are a factor.

In addition, the system to the desktop must operate within the legal regulatory bounds, principally those imposed by the Federal Communications Commission (FCC), and be robust against interference from outside sources, including meeting the worst case bit error rate (BER) for ATM systems of 10^{-10}.[1]

Lastly, any system must provide maintenance and management capabilities that will attract customers, lower costs, and provide an attractive alternative to competing systems.

The key criteria for such a system can be summarized as:

- Ubiquitous connections to the desktop;
- Connections up to 100 meters in length;
- Use of existing twisted pair cabling;
- BER of 10^{-10} or better
- Compliance with various mandatory standards, such as FCC regulations;
- Robustness and lack of susceptibility to surrounding electrical noise;
- Resource management; and
- Network compatibility.

This paper deals with the implementation of a transmission system that eases the transition to ATM by using the existing copper wiring plant already in place in most business establishments.

Existing Cable Environment

This section discusses the existing building cable plant, and its implications for ATM transport.

Existing Building Cable Plant. Most business establishments have Category 3 unshielded twisted-pair (UTP-3) in their walls and often times use this type of cabling for data communications.[2,3] A better grade of cable, Category 5 unshielded twisted-pair (UTP-5), also known as *data grade cable*, is used for upgrading existing installations and for new installations where premium quality is required for high-speed data communications. While most of the new sales for twisted pair are UTP-5, most of the installed base in existing commercial buildings is UTP-3 and will remain in place for a number of years—if only because of the high cost of installing new cabling in existing structures.

While the "cost per desktop" is often quoted as being the critical figure for deciding on whether or not to upgrade indoor cable, that cost often includes only the adapter boards and other equipment at each end of the cable link. Often overlooked or ignored is the largest cost of installing a new cable system to provide connectivity to a multiplicity of desktops in a business establishment. Obviously a system that does not require new cable to be installed can offer significant savings in costs.

Therefore, facilitating the early adoption of ATM by using the existing UTP-3 cabling makes good economic sense. Though not aimed at the highest "power users," those with needs of 155 Mb/s and up, but at those multimedia users where 52 Mb/s currently is quite sufficient, this strategy provides opportunities for a large percentage of that market share. Such applications include groupware activities, medical imaging, physical chemistry, and pharmaceutical design.

Many of these users could be served with existing high-speed LAN products, but as will be seen, there are important advantages to using ATM technology. The competition for inexpensive data transfers in the local area are switched Ethernet systems at 100 Mb/s and a vast array of LAN products at bit rates as low as 10 Mb/s. ATM is not principally a LAN service; it originally was conceived as a global interconnection service between common-carrier and privately owned switching points. But the concept of combining global interconnection and LANs, and extending ATM capabilities into these LANs, adds significant new user and monitoring features to the current LAN-based applications, particularly if it can be provided at the same low cost as existing LANs.

The advantages of an ATM desktop solution will be developed later in this paper. Here it will be noted only that an ATM system operating at a given bit rate is equivalent to a shared media system operating at more than *twice* that rate, because the ATM system is available bidirectionally at all times and the data rate is dedicated to a single user.

Twisted-Pair Cable Configuration. Two pairs of twisted wires, that is, four wires, are used to link each user on a point-to-point ATM desktop connection. For each user, one twisted pair is connected to the transmitter in the outgoing direction and the other twisted pair is connected to the receiver in the in-coming direction. This configuration has advantages over using a single twisted pair for both transmit and receive because the echo cancellers required on a single pair for 51.84 Mb/s rates are difficult and expensive to build.

As noted, twisted pairs rarely are run except in bundles; UTP-3 cable has four twisted pairs and cables with 25 pairs are commonly found in older buildings. Other users within a large cable bundle would each use their own pair of twisted-wire pairs.

Losses in the Cable Medium. Figure 2 shows plots for a model of the two predominate types of losses in a

Panel 1. Abbreviations, Acronyms, and Terms

2-CAP — carrierless amplitude modulation/phase modulation with 2 constellation points
4-CAP — carrierless amplitude modulation/phase modulation with 4 constellation points
16-CAP — carrierless amplitude modulation/phase modulation with 16 constellation points
A/D — analog/digital
ANSI — American National Standards Institute
ATM — asynchronous transfer mode
B-ISDN — broadband integrated services digital network
BER — bit error rate
CAP — carrierless amplitude modulation/phase modulation
D/A — digital/analog
DFE — decision feedback equalizer
EIA — Electronic Industries Association
FCC — Federal Communications Commission
FEXT — far-end cross talk
FSLE — fractionally spaced linear equalizer
ISDN — integrated services digital network
ISO — International Organization for Standardization
ITU-T — International Telecommunications Union-Telecommunications Standardization Sector
LAN — local area network
NEXT — near-end cross talk
NRZ — non-return to zero
OA&M — operations, administration, and maintenance
PBX — private branch exchange
SONET — synchronous optical network
STS — synchronous transport signal
TIA — Telecommunications Industry Association
UTP — unshielded twisted pair

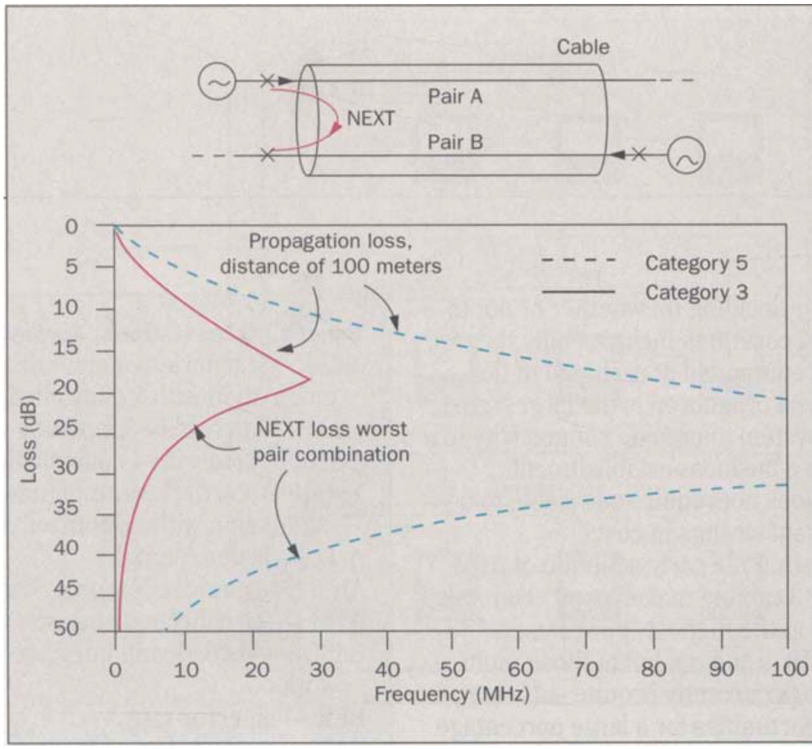

Figure 2. This figure shows plots for a model of the two predominate types of losses in a copper twisted-pair cable—propagation loss (top two curves) and crosstalk (bottom two curves). Parenthetically, this figure shows the effect of higher quality UTP-5 cable (dashed lines). The crosstalk coupling mechanism is illustrated in the insert. The wires in a cable both radiate signals into the environment and receive signals from the environment.

copper twisted-pair cable—propagation loss and crosstalk. The model comes from the Telecommunications Industry Association (TIA) and is published as standard TIA/EIA-568A.[2,3] Parenthetically, Figure 2 shows the effect of higher quality UTP-5 cable, where better quality insulation material and a more controlled mechanical configuration within the cable produce lower attenuation and lower crosstalk.

Attenuation. Propagation loss is straightforward attenuation—the farther you go down the cable, the less signal you have. Attenuation losses also increase with increasing frequency, so that high-frequency components like sharp edges and fast rise times do not propagate as far down the cable. Cables, in effect, act like low-pass filters.

Crosstalk. The major interaction among the wires in the cable bundle is *crosstalk*,[4] the coupling of a signal from one pair of wires to another wire pair. The crosstalk coupling mechanism is illustrated in the insert to Figure 2. The wires in a cable act to couple signals from one wire to another, each both radiating and receiving signals to and from the other wires in the bundle. Since all active wires terminate at receivers at one end of the cable or the other, the signal on each wire pair contains noise picked up from all the other transmitting wires in the cable, in addition to the intended signal propagating down the wire. This *combined* signal is then presented to the input of the receiver attached to that wire pair.

The crosstalk effect is strongest for wires that are nearest to one another geometrically, decreasing rapidly with physical separation. We ignore for the moment the fact that these wires can also couple with signals from external sources.

The effect of crosstalk is described in terms of loss, so its effect can be compared to the effect of attenuation. Little crosstalk is high loss, that is, a small amount of induced signal, and, correspondingly, more crosstalk is less loss.

Crosstalk is produced from two sources relative to the receiver, *near-end crosstalk (NEXT)* and *far-end crosstalk (FEXT)*. Consider a cross section of cable, the simplest having four twisted pairs of copper wire bundled together. Crosstalk into any one receiver is generated by all of the other wires in the bundle. The crosstalk produced from a transmitter associated with a particular receiver is the known outgoing signal and is, therefore, *deterministic*, making it predictable and, therefore, correctable; whereas crosstalk produced from the other transmitters in the bundle can be considered *random* and, therefore, requires more sophistication to handle.

This distinction in the type of noise from the two types of crosstalk has implications for the technique chosen to mitigate the effects of crosstalk. In particular, since the noise at the receiver is composed of both types and, therefore, cannot be known *a priori*, statistical methods must be used.

Since FEXT is generated from the transmitters at the far end of the link, it experiences propagation loss for the whole length of the cable. As a result, the effects of FEXT are significantly less than that of NEXT in the LAN environment.

A typical NEXT is plotted as the lower curve in Figure 2. High-frequency components have higher coupling than low-frequency components. Consequently, signals with sharp edges, that is, high-frequency components, produce a worse environment within the cable than signals with rounded edges.

Note that above approximately 30 MHz at 100 meters the propagation loss is greater than the NEXT loss. Above that frequency, the amplitude of the induced NEXT becomes more significant than the amplitude of the signal propagating down the wire, and the received signal-to-noise ratio is decreased to the point that the link is not usable above this frequency.

Suppression Technique in a Distorted Medium. Typically, NEXT from the transmitter co-located with the receiver offers the strongest source of noise; but since the pairs of wires in UTP-3 cable are placed loosely within the bundle without regard for their geometrical placement, NEXT from another transmitter or the composite NEXT from other transmitters may be greater. This uncertainty concerning the amplitude of the NEXT interaction, particularly for cables from different vendors, means that a straightforward *cancellation technique* is not appropriate.

Cancellation techniques can only eliminate the crosstalk from a particular receiver's transmitter, but can do nothing with the remaining sources of crosstalk. Since one does not know *a priori* what the source of the dominant crosstalk will be because of the variability of crosstalk among the wire pairs in the cable, cancellation is not a generally effective technique in this application. Instead, *suppression techniques* are the more viable choice. Suppression techniques minimize the composite crosstalk from all sources, not just the worst one. Another advantage is that the receiver can notch out narrowband external disturbances such as signals from nearby radio stations.

As an additional complexity, attenuation and crosstalk characteristics are not as well behaved as the diagram in Figure 2 indicates but are, in fact, much more variable. In the standards documents, the cable is only specified in the range of 0-16 MHz;[2,3] the plot in Figure 2 is only an extrapolation for the range above 16 MHz, and it can vary greatly among cable manufacturers. Even in the range 0-16 MHz, while attenuation is a slowly varying function of distance, crosstalk is not and may vary greatly in amplitude across the frequency spectrum. The use of digital signal processing is necessary as part of the receiving function because:

- It is not possible to anticipate the characteristics of the signal through an unpredictable cable, and
- It is necessity to deal with NEXT by suppression techniques.

Emissions. In addition to attenuation and crosstalk losses in the cable, limits on emissions of electromagnetic radiation impose critical constraints on the frequency spectrum of the signal on the cable. Equipment sold in this country must meet FCC Class A regulations for business and Class B regulations for residential use. Currently, federal regulations place limits only on emissions above 30 MHz. Internationally, constraints also are placed on emissions above 30 MHz. Additionally, cabling is susceptible to electromagnetic radiation from external sources and the implementation of the technology must deal with these additional interferences.

Since propagation loss exceeds crosstalk loss above 30 MHz and since emissions must be restricted above 30 MHz, the power spectrum of the signal must be significantly restricted above 30 MHz. One way to meet these restrictions is to design the system so that nearly all of the frequency spectrum power is below 30 MHz. The next section explains the technology for building a system that meets these constraints.

Physical Layer Transmission Technology

The high speeds now attainable over twisted-pair cables are due to a recent development, the implementation of carrierless amplitude modulation/phase modulation (CAP) technology.[4] Although CAP is based on established modem technology, it opens up new opportunities, since it provides transmission along twisted-pair copper wires at high bit rates that were considered impossible just a few years ago. Since the technology allows such bit rates, two questions naturally arise:

- How should these high bit rates be organized, and
- What should the actual bit rates be?

A number of factors other than the technology of the physical medium enter into those decisions. First, we shall discuss the organization question, and then in the next section, "Physical Media Dependent Sublayer," we will see how an appropriate bit rate was chosen.

Transmission Convergence Sublayer. Aside from the aspects of physical technology, another issue is how the transmitted bits should be organized, since that decision impacts not only how information is transmitted, but also how maintenance and performance monitoring are to be accomplished.

In an OSI model of a transmission system, the physical layer is divided into two parts, the physical media dependent sublayer and the transmission convergence sublayer. The transmission convergence sublayer addresses the format of how the transmitted bits are organized and determines how bits of user information and functions, such as maintenance and performance monitoring, are transmitted between pieces of equipment.

Transmission Bit Rate and Frame Structure. Some people have argued that a small increase from the currently available bit rates in the neighborhood of 16 Mb/s is sufficient for desktop connectivity, based on applications currently being used. Others have argued that telecommunications bandwidth is analogous to memory and speed in personal computers—plenty is barely enough. This is a classical "chicken and egg" situation—high bandwidth is not yet readily available for ATM, so applications have not been developed and, on the other hand, since no applications have been developed, the industry cannot afford to install extra bandwidth. The answer, of course, will be determined by the direction video, multimedia, and data communications will take, where, historically, more is always better.

The use of ATM in the local area network will be competing with many existing LAN systems. The real advantage of ATM in the marketplace is its capability to connect a user in one local area to a user in another local area across a wide area network using the switched public telephone network. Being able to satisfy both the local area network needs, as well as transmitting information across the public switched network, is the key to differentiating ATM from its competitors and gives it a unique advantage in the world of voice/data/multimedia communications.

SONET-based Transmission. The switched public telephone network has, for the past few years, been developing and implementing a network based on a frame structure—that is, the organization of the bits as they are transmitted through the transmission system—known as synchronous optical network (SONET).[5] SONET supports optical transmission in multiples of 51.84 Mb/s up to 6.448 Gb/s. The SONET structure, as modified to carry ATM cells, is known as broadband integrated services digital network (B-ISDN). Being able to apply the attributes of an optical SONET-based frame to the copper-wire desktop connection offers a number of pertinent capabilities.

Use of a SONET-based format allows for the reuse of both the software and hardware developed for this technology, saving in development costs. Systems can be built having significant commonality for software and hardware modules across several bit rates. Shared design expertise also is a factor in producing equipment that is both better and less expensive, with higher degrees of interoperability. The use of the same silicon across several interfaces also will produce cost savings because of economies of scale, particularly since the link to the desktop will be the full 51.84 Mb/s bit rate.

The primary rate for B-ISDN is 51.84 Mb/s,[1] and so this bit rate was chosen, along with the accepted SONET-based frame format known as synchronous transport signal-1 (STS-1). The STS-1 frame has a structure composed of two parts, a *payload part* of 48.386 Mb/s carrying the user information and an *overhead part* carrying information necessary for management of the connection. For this application, the payload is simply filled with ATM cells. The interface is symmetric, that is, it has the same bit rate is both directions.

Connection with the Telephone Network. One of the emerging thrusts in the marketplace is the concept of end-to-end management of the users' connection, that is, being able to manage both the public switched network as well as the network on the customer's premises. Because the format chosen here is based on the SONET/B-ISDN

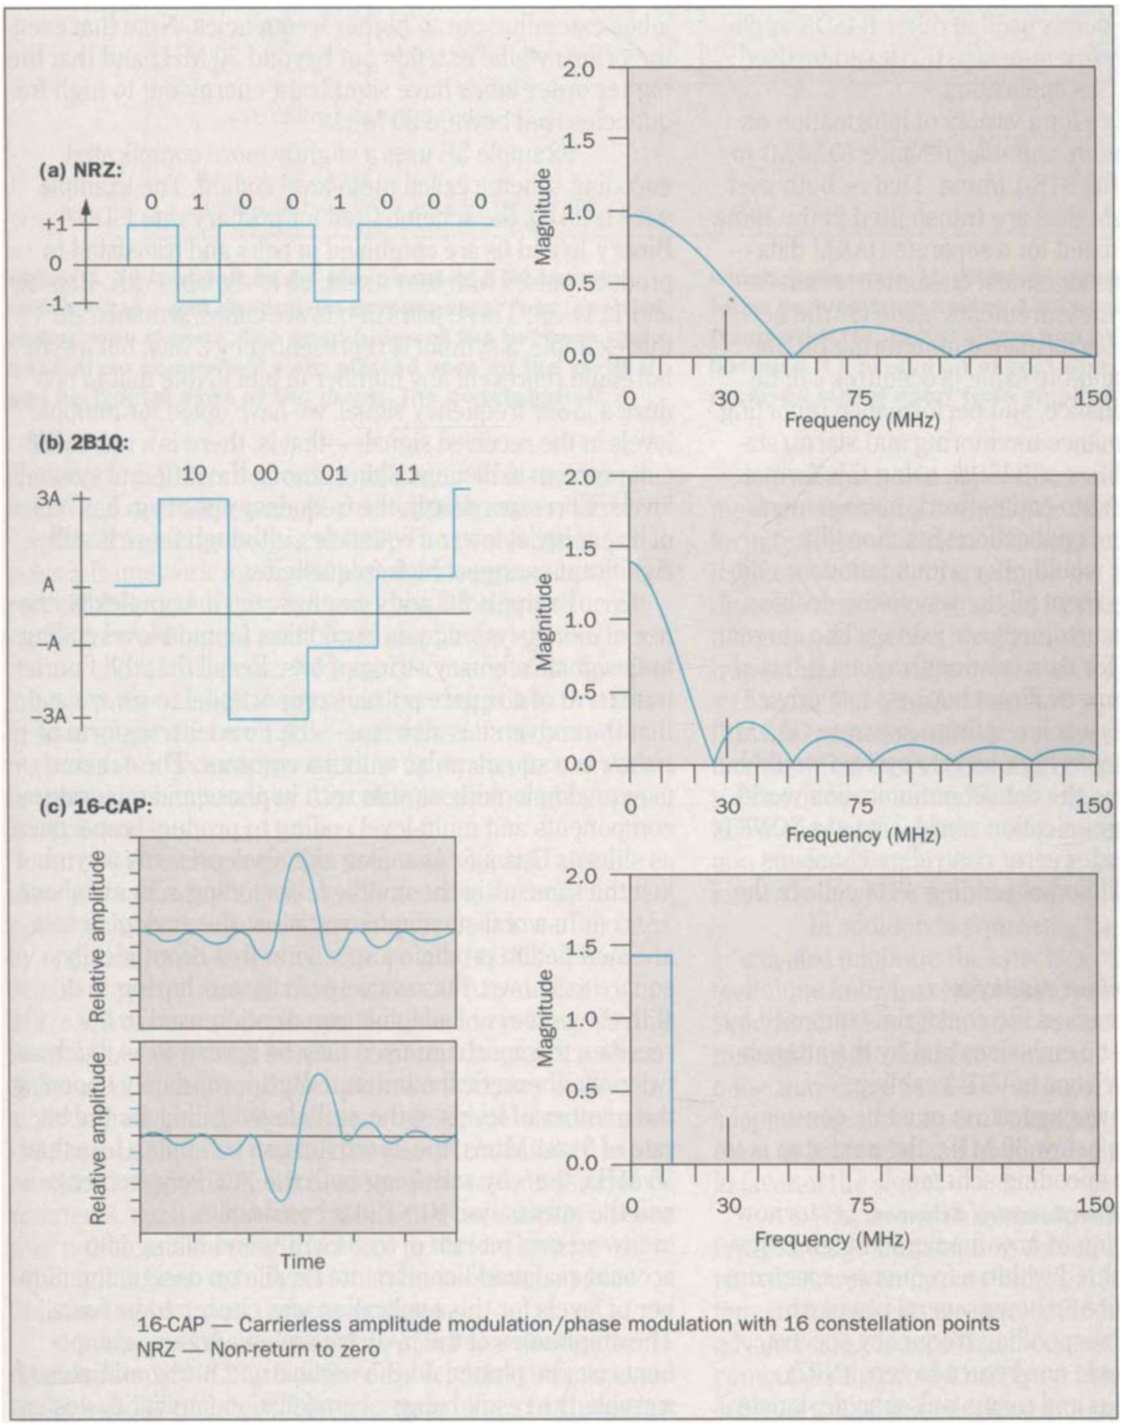

Figure 3. This figure shows several waveform examples and their corresponding frequency spectra. In 3A, a simple non-return to zero (NRZ) encoding scheme, the primary lobe extends out beyond 30 Mhz. 3B shows a slightly more complicated encoding scheme called multi-level coding, which has more of its energy at lower frequencies, although there is still significant energy at frequencies above 30 MHz. In 3C, analog-looking signals with in-phase and quadrature components and multi-level coding produce a frequency spectrum well below 30 Mhz.

STS-1 frame, the same octets used in other B-ISDN applications at the user-network interface (UNI) can be used in this customer premises application.

SONET provides for a variety of information on operations, administration, and maintenance (OA&M) to be transmitted within the STS-1 frame. That is, both user data and network OA&M data are transmitted in the same frame, eliminating the need for a separate OA&M data link now required to manage most customer premises equipment. The same measurements made for the network can be made for performance monitoring on the customer's premises, and the same procedures can be used for alarm, maintenance, and performance reporting. Because of the performance monitoring and alarm/status monitoring capabilities of B-ISDN, using this format lays a foundation for end-to-end network management, enabling true end-to-end applications functionality.

Such a system would offer a foundation for end-to-end network management all the way to the desktop, a concept for which our customers are asking. The current desktop configuration for data communications is based either on a LAN structure or direct hook-up to a private branch exchange (PBX), each requiring separate OA&M. A SONET-based medium-to-high bit rate system would be a step forward in uniting the data communication world with the voice telecommunication world. Use of a SONET-based format also includes error control mechanisms that minimize the likelihood of sending ATM cells to the wrong address.

Physical Media Dependent Sublayer

Earlier we discussed the constraints imposed by limits on electromagnetic emissions and by the attenuation-crosstalk characteristics of UTP-3 cabling. Concluding that the power spectrum must be constrained to frequencies below 30 MHz, the next step is to decide on a bandwidth encoding scheme.

Bandwidth Efficient Encoding Schemes. Let us now develop an understanding of how medium-to-high range bit rates can be transmitted within a frequency spectrum less than 30 MHz. Figure 3 shows several waveform examples and their corresponding frequency spectra. Example 3A uses a simple non-return to zero (NRZ) encoding scheme of plus and minus pulses to designate binary 1s and 0s. Its spectrum is the classical *sin x/x* with a primary lobe at lower frequencies and higher order lobes extending out to higher frequencies. Note that even the primary lobe extends out beyond 30 MHz and that the higher order lobes have significant energy out to high frequencies well beyond 30 Mhz.

Example 3B uses a slightly more complicated encoding scheme called multi-level coding. The example used is 2B1Q, the scheme used for primary rate ISDN.[6] Binary 1s and 0s are combined in pairs and translated to produce pulses with four levels: 00 is +A, 01 is +3A, 10 is -A, and 11 is -3A. These pairs of bits are called symbols. (In this example, a symbol is represented by 2 bits, but a symbol could represent any number of bits.) Note that to produce a lower frequency signal, we have opted for multiple levels in the received signals—that is, there is a more difficult problem in distinguishing among the different symbol levels. Correspondingly, the frequency spectrum has more of its energy at lower frequencies, although there is still significant energy at high frequencies.

Example 3C adds another step in complexity, the use of *analog-type* signals in addition to multi-level coding to designate a binary string of bits. Recall that the Fourier transform of a square pulse is proportional to *sin x/x* and that the converse is also true—the Fourier transform of *sin x/x* is a square pulse within a constant. The scheme uses analog-looking signals with in-phase and quadrature components and multi-level coding to produce a spectrum as shown. This pair of analog signals represents a symbol just the same as in the multilevel encoding scheme above.

In a realistic implementation, the analog signals are modified to produce a spectrum that is not ideally square as shown. For a variety of reasons having to do with the process of adaptive equalization used in the receiver, the spectrum used may be spread to as much as twice the theoretical minimum. By appropriately choosing the number of levels of the multi-level coding for the bit rate of 51.84 Mb/s, the spectrum can be kept to less than 30 MHz, thereby satisfying both the FCC requirements and the attenuation-NEXT loss constraints.

For a bit rate of 51.84 Mb/s and taking into account real world constraints for silicon design, the number of levels for this application was chosen to be four. The amplitudes of the in-phase and quadrature components can be plotted on the vertical and horizontal axes of a graph, the result being a constellation familiar to designers of modems (see Figure 4). What the constellation represents is a diagram showing the 16 different combina-

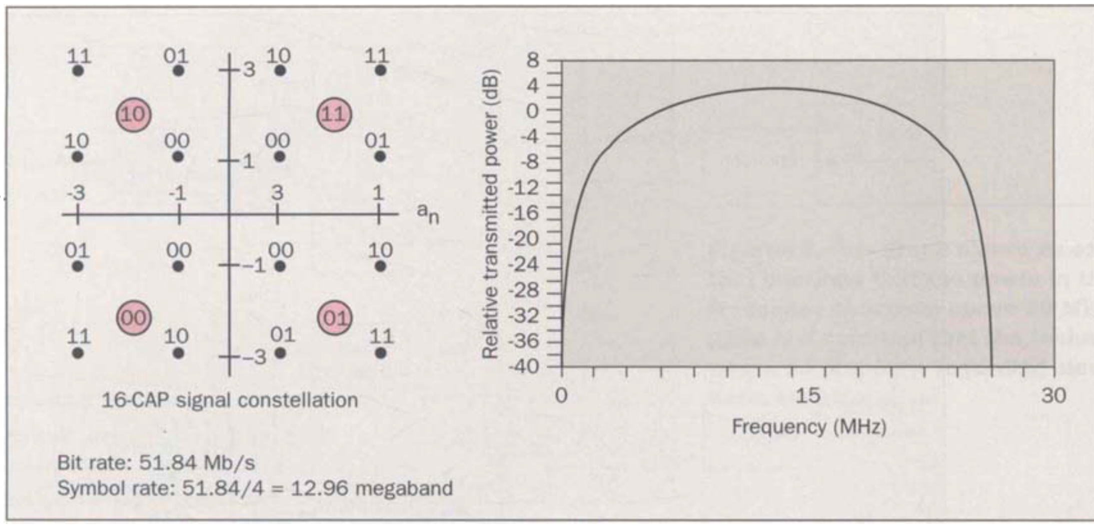

Figure 4. To transmit 51.84 Mb/s analog ATM signals with in-phase and quadrature components, four levels of coding was chosen. The amplitudes of the in-phase and quadrature components are plotted here on the vertical and horizontal axes of the graph. The constellation, which represents 16 different combinations of bit patterns derived from having 4 bits per symbol, provides a framework for digital signal processing in the receiver because it conveys an expectation of the format of the received signal apart from amplitude.

tions of bit patterns derived from having 4 bits per symbol. When operating with 16 constellation points, the name of this technology is 16-CAP. The constellation provides a framework for digital signal processing in the receiver because it conveys an expectation of the format of the received signal apart from amplitude.

So, at this point, we have discovered that, by taking advantage of digital signal processing power that can fit on a single silicon chip, we can transmit medium-to-high bit rates over copper wires in an analog format within constraints imposed by the cable medium and by regulatory mandates.

CAP technology also can easily operate at subrates of its primary bit rate by decreasing the number of points in its constellation[4] and drop back to lower bit rates by operating with 4 constellation points (4-CAP) at 25.92 Mb/s or even 2 constellation points (2-CAP) at 12.96 Mb/s. Thus, an end point wouldn't have to receive the full 51.84 Mb/s signal if it weren't required. In practice, the system can operate at these two lower bit rates without any additional hardware because the symbol rate remains the same in all these cases. This also permits lower bit rates for longer-than-normal extensions or for nosier environments, such as in older buildings where wiring closets may not be as close as 100 meters to the desktop or where cabling has been optimized for audio telephony applications and not for data communications.

Adaptive Equalization. As we just saw above, the combination of NEXT, undefined cable characteristics, and the possibility of external interference led us to choose an adaptive equalization technique to implement the suppression in the receiver.

Earlier, this paper discussed the harsh environment presented by twisted-pair cabling due to non-uniformity of the electrical characteristics of the cabling itself, as well as by the complexity of the crosstalk conditions. At the receiving end of the link, the receiver is not going to see the idealized constellation sent by the transmitter, but will see, instead, a distorted, noisy constellation. The adaptive equalizer makes use of the fact that the constellation at the receiver has a known structure modified by the effects of attenuation and phase distortion from the cable. The equalizer makes use of the knowledge of this pattern to compensate for the distortion imposed by the cable, a compensation technique known as *adaptive equalization*.

In addition to equalizing the channel, the equalizer can also minimize the effects of NEXT interference, a technique known as NEXT *suppression*. This technique is superior to other techniques such as cancellation for maximizing the signal-to-noise ratio in cases where the noise may come from several sources of crosstalk. An adaptive equalizer basically looks at the entire environment as one entity and minimizes the totality of the noise in favor of the signal.

Why is such a system possible now when only a few years ago transmitting these high bit rates over copper wires was an impossibility? The enabler for this system, as it has been for much of today's technological advances, is advances in silicon technology. Complex computers that could barely be imagined a decade ago are now available in desktop and portable units. So it is with transmission over copper wire pairs. The processing needed to compensate for the electrical characteris-

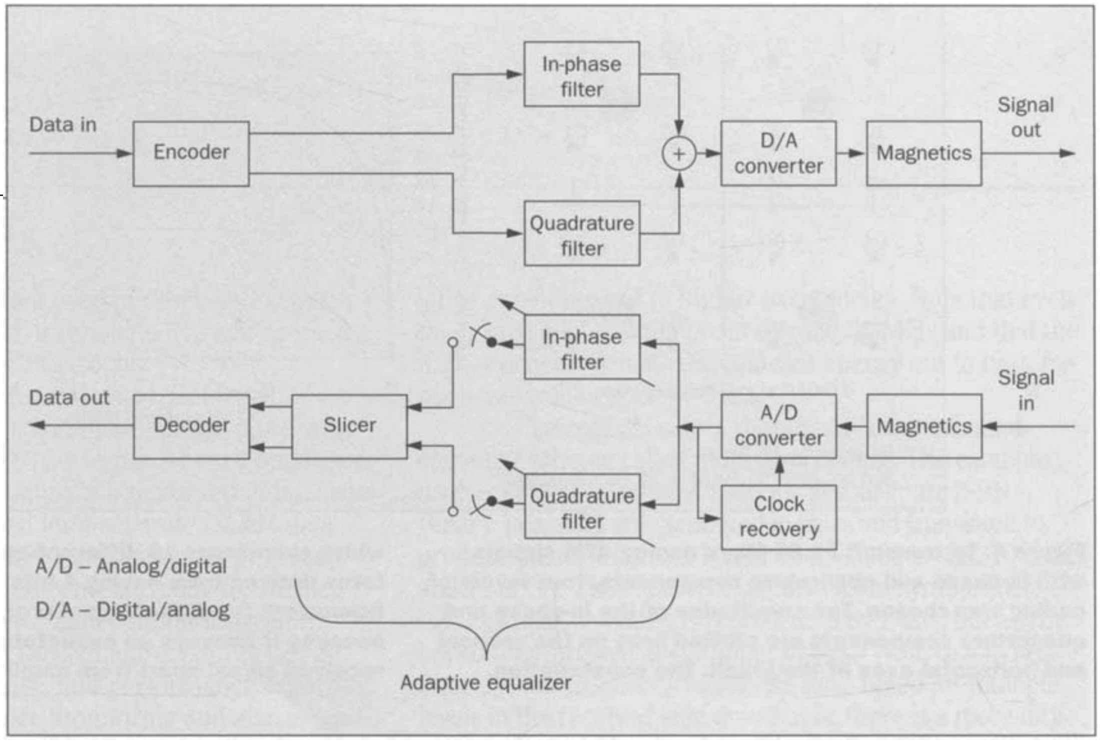

Figure 5. This illustration shows a block diagram of a digital implementation for a transmitter and a receiver, called a transceiver. On the transmitter side, bits are encoded four at a time into a symbol and are passed to the transmitter filters. Those filters produce in-phase and quadrature representations of the symbol, which are summed, passed through a digital-to-analog (D/A) converter, and presented to the transformer magnetics, the elements that perform the matching function for the signal so it may be placed on the cable.

tics of the medium is complex—analogous to the equalizer that takes up a whole shelf in your stereo at home—but today can fit easily on a single silicon chip.

An equalizer based on digital signal processing principles performs the adaptive function. As in the transmitter side, the receiver is using fairly sophisticated digital processing, supported by the higher speed and complexities of silicon chips, to carry out a function unheard of just a few years ago.

Implementation. Figure 5 illustrates a block diagram of a digital implementation for a transmitter and a receiver; the term transceiver is used to describe this functional pair.

On the transmitter side, bits are encoded four at a time into a symbol and are passed to the transmitter filters. Those filters produce in-phase and quadrature representations of the symbol, which are summed, passed through a digital-to-analog (D/A) converter, and presented to the transformer magnetics, the elements that perform the matching function for the signal so it may be placed on the cable.

At the receiver, the analog signal is recovered, converted to digital in an analog-to-digital (A/D) converter, and passed to an adapter equalizer. Given the unknown electrical characteristics and crosstalk environment at the receiver, the receiver needs a mechanism for developing a uniform presentation so it can extract the correct binary signal.

As a brief description, the receiver compensates the incoming signal much like a person compensates the sound from an audio system to get the best balance. The equalizer does this in an adaptive fashion by fitting the incoming constellation to the expected constellation using a mean square fit. No matter what the cable characteristics are and no matter what the details of the data being transmitted, the noise in the environment can be considered slowly varying which, therefore, can be tracked by the adaptive equalizer. The technique works by emphasizing a signal at one part of the spectrum and minimizing it in others.

There are several types of adaptive equalizer. At the high-quality end of the spectrum is the decision feedback equalizer (DFE) that can do a superb job of recovering the digital signal. Less sophisticated and, therefore, less costly is the fractionally spaced linear equalizer (FSLE), which gives sufficient margins for any of the normal systems encountered during experiments and trials. The technique of adaptive equalization using a FSLE is a series of multiply-and-add operations. Using some clever

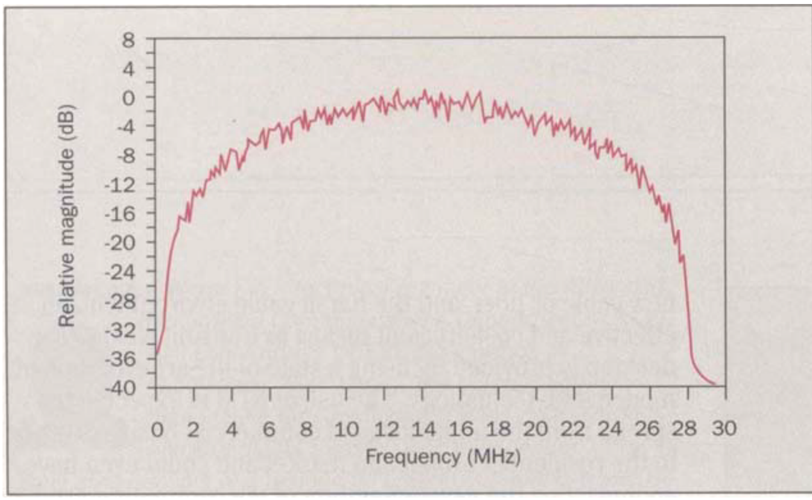

Figures 6. This graph shows an experiment that confirms that the power in the CAP frequency spectrum above 30 MHz is negligible and confirms that the technology poses no problems regarding electromagnetic emissions.

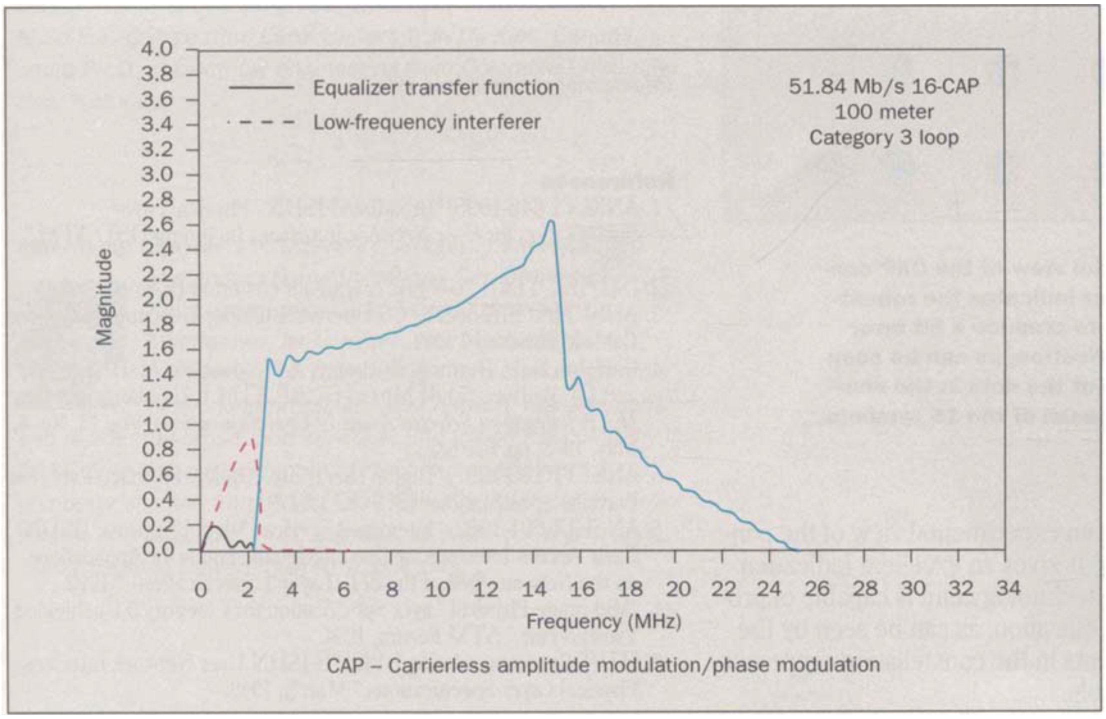

Figure 7. This simulation illustrates the effect on the CAP power spectrum when an interferer (dashed line) is induced at low frequencies. The adaptive equalizer compensates to maintain signal quality. Power in the spectrum at low frequencies, where interference is present, is moved to high frequencies to avoid the interference and the bit error rate is maintained within specifications.

layout schemes, the multiply-and-add circuits can be reused in a folding technique in which the result of a multiply-and-add is saved and fed back into the input for the next iteration.

Tradeoffs between the speed of computation and the cost of the silicon, as represented by size of the silicon chip, are made by choosing the size, correspondingly the complexity, of the FSLE. By choosing an FSLE of the appropriate complexity, a proper balance can be maintained between speed and cost.

Because the technology provides such a clear cutoff to the power spectrum at 30 MHz, robust margins can be used for signal amplitudes to avoid problems with susceptibility from external interference.

Results of a CAP System

Figures 6–8 show the resulting signal from a CAP system. Figure 6 is an experimental result confirming that the power in the frequency spectrum above 30 MHz is negligible and confirming that the technology poses no problems regarding electromagnetic emissions.

Figure 7 shows a simulation illustrating the effect on the power spectrum of having an interferer induced at low frequencies and illustrating how the adaptive equalizer compensates to maintain signal quality. Power in the spectrum at low frequencies, where interference is present, is attenuated and power of higher frequencies is amplified to avoid the interference and the bit error rate is maintained within specifications.

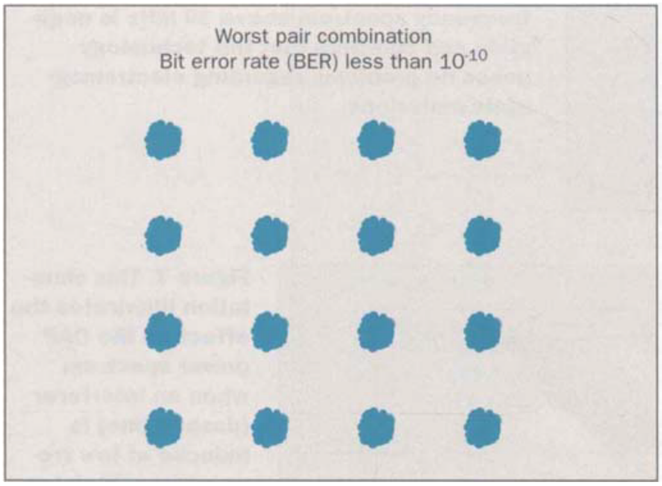

Figure 8. An experimental view of the CAP constellation at the receiver indicates the robustness of the technology to produce a bit error rate (BER) within specification, as can be seen by the clear separation of the dots in the constellation representing each of the 16 symbols.

Figure 8 shows an experimental view of the constellation at the receiver. It gives an excellent indication of the robustness of the technology and is capable of producing a BER within specification, as can be seen by the clear separation of the dots in the constellation representing each of the 16 symbols.

Standards Status. The specification for a 16-CAP physical media dependent sublayer and a transmission convergence sublayer based on STS-1 has been approved by the ATM Forum[7] and is under consideration by both the American National Standards Institute (ANSI) Committee T1 (T1E1.2) and the International Telecommunications Union-Telecommunications Standardization Sector (ITU-T) SG 13 (as part of Question 13 on B-ISDN, Recommendation I.432.[8]).

Conclusions

Given the market need for medium-to-high bit rate transmission, the economics of not having to install new cable or fiber, and the harsh cable environment, an effective and cost-efficient means to transmit ATM to the desktop is provided by using a state-of-the-art evolution of modem-like technology. The use of ATM is expected to spread from the initial focus of commercial business traffic to the residential broadband market and could even have an impact on the new generation of telecommuting products and services.

(Manuscript approved June 1995)

References

1. ANSI T1.646-1995, "Broadband ISDN - Physical Layer Specification for User-Network Interfaces Including DS1/ATM," 1995.
2. ISO/IEC 11801 "Generic cabling for customer premises," 1995.
3. ANSI TIA/EIA-568A, "Commercial Building Telecommunications Cabling Standard," 1994.
4. G. H. Im, D. D. Harmon, G. Huang, A. V. Mandzik, M.-H. Nguyen, and J.-J. Werner, "51.84 Mb/s 16-CAP ATM LAN Standard," *IEEE Journal on Selected Areas In Communications*, Vol. 13, No. 4, May, 1995, pp. 620-632.
5. ANSI T1.105-1991, "Digital Hierarchy - Optical Interface Rates and Formats Specifications (SONET)," 1991.
6. ANSI T1.601-1992, "Integrated Services Digital Network (ISDN) Basic Access Interface for Use on Metallic Loops for Applications on the Network Side of the NT (Layer 1 Specification)," 1992.
7. "Mid-range Physical Layer Specification for Category 3 Unshielded Twisted-Pair," ATM Forum, 1994.
8. ITU-T Recommendation I.482, "B-ISDN User Network Interface - Physical Layer Specifications," March, 1993.

Richard L. Townsend, Jr., *is a distinguished member of technical staff in the Standards Planning Department in the AT&T Architecture Area of AT&T Bell Laboratories in Holmdel, New Jersey. He represents AT&T on various standards bodies and industry consortia dealing with physical layer aspects for ATM-based telecommunications issues. He joined the company in 1969. He has S.B. and S.M. degrees in electrical engineering from the Massachusetts Institute of Technology in Cambridge, and a Ph.D. in electrical engineering from Stanford University in Palo Alto, California.*

Jean-Jacques Werner *is a technical manager in the Advanced*

Multimedia Communications Department of AT&T Bell Laboratories QUEST organization in Middletown, New Jersey. His organization is responsible for conducting studies on bandwidth-efficient digital transmission over copper wiring for local area network and residential broadband applications. He joined the company in 1973. He has a B.Sc.E.E. degree from Institut National des Sciences Appliquèes in Lyon, France; a M.Sc.E.E. degree from Laval University in Quebec, Canada; and a Ph.D. in electrical engineering from Columbia University, New York City.

Mai-Huong Nguyen *is a technical manager in the Advanced*

Multimedia Communications Department of AT&T Bell Laboratories QUEST organization in Middletown, New Jersey. Her organization is responsible for developing distributed multimedia networking protocols, ATM network management, and residential broadband services. She joined the company in 1985. She has a B.S. degree in electrical engineering and computer science from Princeton University in Princeton, New Jersey, and M.S. and Ph.D. degrees, in electrical engineering, from Columbia University in New York City.