Abstract

Regular analog telephone lines can be directly connected through the telephone network to ISDN terminals. By using appropriate modulation formats, they can carry digital traffic at rates approaching those of ISDN, without extra recurring costs for the customers or new investments in the network. Such a system thus provides an affordable and universal "driveway" to information services. This article explains the basic concepts and reports on some measurements that show their feasibility.

The Information Driveway

Pierre A. Humblet and Markos G. Troulis, Institut Eurecom

"The fundamental problem of communication is that of reproducing at one point ... a message selected at another point."
Claude E. Shannon

The coming of age of a society based on the processing of information necessitates wide access to information sources through communication networks, as evidenced by the success of services such as the World Wide Web. The current situation is satisfactory for larger organizations that can connect their local area networks together and to international backbones through high-speed links. Even if the available throughput is not always sufficient, realistic migration paths based on ATM technology are in sight. Unfortunately the cost of these "information highways" is too high for smaller companies, schools and individuals at home.

There is no doubt that the situation will evolve. The need to distribute high-bandwidth entertainment material to homes will be met by a future broadband network. "Fiber to the home" was thought to be the solution, but new scenarios based on asymmetric digital subscriber loop (ADSL) on twisted pairs, or on an extension of cable TV distribution using fiber and coax cable technologies, now appear to be more cost effective. This evolution to shared media is paradoxical: In the 1970s, the computer industry discovered the merits of broadcast media, such as Ethernet, but new high-speed local networks are hub- and switch-based. On the contrary telecommunication companies, which traditionally use point-to-point facilities, are now investigating the use of shared broadcast media to distribute high-rate traffic!

Whatever its structure, it is clear that such a network will not be universally deployed for decades, and that its cost will be tremendous. While they wait for a future broadband network, individuals, schools, and smaller companies must transmit data on the telephone network. They have a choice between two technologies.

The most popular solution is to use modems, which commonly provide communication at 14.4 kb/s, although newer standards such as V.34 allow transmission at up to 33.6 kb/s. To support that option, service providers must install and maintain modem banks for local access, and multiplexers to carry the long-haul traffic.

The other solution is to use basic rate integrated services digital network (ISDN). It provides two 64 kb/s transparent channels and a 16 kb/s signaling channel. In conjunction with current advances in compression technology, this solution provides more comfortable data rates for accessing data services, telecommuting, carrying forms of personal video communications (more so with the advent of Motion Picture Experts Group version 4, MPEG-4), and providing distance teaching. However, it is not available everywhere and requires a significant investment for the customers and for the telephone companies. It usually also results in significantly higher monthly fees. Depending on the protocols and on the statistics of the traffic, service providers can either simply transport the 64 kb/s streams on the long-haul network or install statistical multiplexing equipment to reduce the cost of long distance transmission.

Our purpose in this article is to point out the existence of a third solution that relies on ISDN while requiring little investment and no recurring costs for the customers. This "information driveway" will provide greater access to information and education services for the vast majority of the population.

In the next section we will explain the nature of our proposal and investigate some of its general characteristics. We will then report on some measurements that show the feasibility of the concept.

The Hybrid Communication Network

We recognize that information servers can usually afford a digital connection to the telephone network, often at the rates of 1.5 Mb/s or 2 Mb/s. Many private customers, on the other hand, are reluctant to engage the expense and continue to rely on plain old (analog) telephone service (POTS). It is possible to take advantage of this situation as illustrated in Fig. 1 where the ISDN and POTS networks are directly connected.

ISDN and analog terminals can dial each other and exchange information. This is commonly done today (e.g., between a digital phone and an analog phone). The digital signal is simply a sampled and quantized representation of speech.

It is possible to do much better, as will be shown shortly. By analogy with satellite and mobile systems, we call the direction from the ISDN terminal to the analog line the "downlink" direction, while the reverse direction is the "uplink." We discuss them separately.

The Downlink Channel

Consider what happens after an ISDN terminal dials an analog phone number and starts transmitting octets. The backbone telephone network is almost all digital, and the octets will arrive with a very low error rate at the digital-to-analog converter (D/A) in the line card driving the analog telephone line.

There they will be interpreted as octets representing speech encoded with the A-law or μ-law, depending on the region of the world, and the D/A will produce a continuous-time waveform

0163-6804/96/$05.00 © 1996 IEEE

$$x(t) = \sum_k g(t - kT)s(u_k) \quad (1)$$

where u_k represents the octets; $s(u_k)$ is the level specified by the appropriate quantization law, $T = 125$ µs is the sampling interval, and $g(t)$ is an interpolation function approximately bandlimited below $1/(2T)$, or about 4000 Hz, to meet Nyquist's reconstruction theorem.

Normally $x(t)$ represents speech, but if the u_k are data octets it will sound like noise. To a communication engineer, however, $x(t)$ in Eq. (1) has another interpretation. It can be seen as a pulse amplitude modulated signal. A digital receiver located at the customer site should be able to process the signal and recover the octets u_k. We will see that this method of communication can support about twice the data rate of voiceband modems.

Viewing a D/A as a pulse amplitude modulator is unusual, and it is important to be aware of the characteristics of the quantizing law to understand the capabilities of the system. More details are available in standard references such as [1].

In principle the 256 reconstruction levels (128 positive, 128 negative) follow a logarithmic law, so the ratio of the worst quantization error to the original level is independent of the level. Practical laws, such as the A- and µ-laws, are piecewise linear and can be decomposed in 8 positive segments, labeled from 0 to 7, and 8 negative segments forming a mirror image. There are 16 points in each segment.

In the A-law, segment i, $i \geq 1$, covers the interval $16\ 2^i$ to $16\ 2^{i+1}$. Its points are separated by 2^i, the first point being at $16\ 2^i + 2^{i-1}$. The interval between the last point of segment i and the first point of segment $i + 1$ is thus $1.5\ 2^i$. Segment 0 ensures a linear transition between the positive and negative regions. It extends by steps of 2 between 1 and 31 and is colinear with segment 1.

Points in segment $i \geq 1$ of the µ-law have values $2a - 33$, where a is the corresponding A-law level. Segment 0 extends by steps of 2 between 0 and 30. Thus, there are only 255 distinct levels, 0 being duplicated in the positive and negative regions.

In the A-law, the last segment extends between 2112 and 4032, while in the µ-law it goes from 4191 to 8031. The equipment is designed so that the highest levels of the A- and µ-laws are mapped to about the same analog voltage levels.

In some µ-law systems, the least significant bit of an octet is sometimes used for signaling("robbed bit signaling"). This halves the number of levels available for data transmission, since two distinct received levels can correspond to the same input.

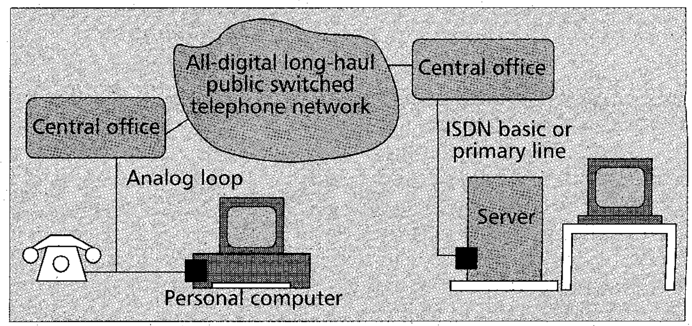

■ **Figure 1.** *Illustration of a hybrid system.*

Rate vs. Noise Resistance Tradeoff — If one wishes to transmit at 64 kb/s all 256 levels must be used. For transmission in additive Gaussian noise, the error probability is largely determined by the minimum distance $d_{\min}$ between adjacent levels. The symbol error probability is bounded by

$$P_e \leq 2Q\left(\frac{d_{\min}}{2\sigma}\right) \quad (2)$$

where we have assumed that the noise is Gaussian with variance σ and we have neglected possible intersymbol interference. There the Q function is the tail probability of a normalized Gaussian distribution. Equation 2 is also an excellent approximation for the bit error probability when Gray coding is employed.

One can trade the bit rate for a higher resistance to noise by only using a subset of the available levels. In normal PAM with equidistant levels, the number of available levels is halved each time the minimum distance is doubled. The situation is much better with the A-law: one can see that when the minimum distance is increased from 2 to 4, only 33 levels out of 256 are eliminated.

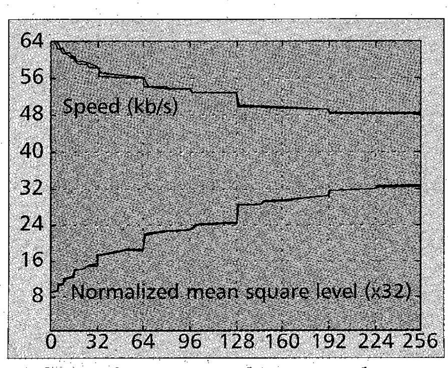

■ **Figure 2.** *Maximum bit rates and average powers for the A- and µ-laws, as a function of the minimum distance d_{min} (d doubled for A-law.)*

The relationship between the maximum bit rate (the logarithm base 2 of the number of levels times the symbol rate) and the minimum distance $d_{\min}$ are illustrated in Fig. 2. The curves for the A-and µ-laws are almost superimposed at the top. The separation of the A–law has been magnified by 2, to make it directly comparable to the µ–law. One can transmit at a rate close to 48 kb/s by using only levels separated by 128 (*A-law*) or 256 (µ–law).

Other curves on the same figure gives the mean square level, normalized with respect to the mean square level of a signal set uniformly distributed between ± 4096, i.e. $4096^2/3$, for the A-law, and between ± 8159 for the µ–laws. Note that the mean power level increases with the minimum level separation, and that its order of magnitude is 10^6.

The power levels considered here exceed those allowed under current regulations. The motivation of these regulations is to make data signals have the same characteristics as voice signals (e.g., to limit saturation and crosstalk in the analog backbone network). These issues disappear in the configuration of Fig. 1, and the regulations could be modified.

The Capacity of the Idealized Channel — A more fundamental understanding is offered by Shannon's information theory. Assume that the convolution of $g(t)$ and of the channel response has a flat Fourier transform bandlimited to 4000 Hz, and that the channel introduces additive white Gaussian noise with two-sided spectral density $N_0/2$. An optimal receiver will consist of a lowpass front-end filter with a phase matched to that of the received signal, followed by a sampler. The decisions are made by finding the legal transmitted level closest to the received level.

The equivalent discrete-time channel for this situation is memoryless. It has a unit gain and additive Gaussian noise with variance $\sigma^2 = N_0/2$. Its inputs are discrete, but its output is continuous. We have evaluated the capacity of such a channel, assuming that a 16-bit linear analog-to-digital converter follows the sampler. The capacity, translated in bits per second, is displayed in Fig. 3 as a function of the "standard signal to noise ratio" (i.e., $1/3\ 4096^2/\sigma^2$, the signal to noise ratio of the uniformly spaced signal set). The curves are parametrized by $d_{\min}$. The asterisks in Fig. 3 represents operating points with uncoded modulation when the upperbound Eq. (2) is equal to 10^{-6}.

The key unknown at this point is the value of σ^2. Measurement of telephone loops [2, 3] yield only a partial answer because they are geared toward voice transmission, using weighted noise measurements or studying end-to-end characteristics. We can try another approach. The designers of digital transmission systems made sure that the digitization of speech was not perceptible, without actually overdesigning the system. Reference [1] indicates that the quantization noise of the A- and μ-laws is about 10 dB lower than the maximum background noise. In fact, in some systems the least significant bit can be robbed for signaling purposes without being noticed, which leads to an error varying between 2 and 256 depending on the segment where it occurs. It appears reasonable to expect a noise variance greater than 1, but probably not exceeding 10 or 20, dBs, corresponding to a standard SNR between 47 and 57 dBs. In this range, uncoded transmission rates at $P_e = 10^{-6}$ vary between 48 kb/s and 56 kb/s, while coding would allow us to ultimately approach 60 kb/s.

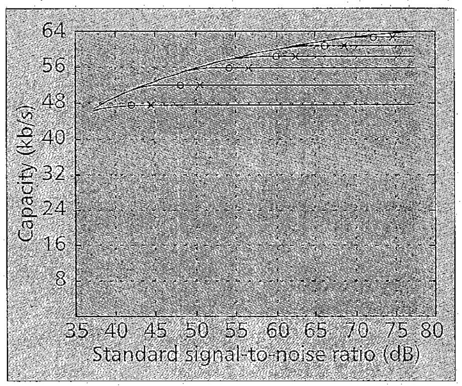

■ **Figure 3.** *Capacity of the ideal downlink channel with A-law levels, for d_{min} = 2, 4, 8, 16, 32, 64, and 128. Uncoded modulation: *, coded modulation: ○.*

The "capacity gap" is about 7 dB in the horizontal scale. Trellis coded modulation [4–7] can greatly improve performance of communication systems, and our situation is no exception. Circles in Fig. 3 indicate the performance achievable by a simple four-state trellis code, still for $P_b = 10^{-6}$. It reduces the capacity gap by about 3 dB.

With full-duplex communication, the equipment on the analog side will need to cancel the echo of the signal it originates. This problem is similar to echo cancellation in voiceband modems, although it is simpler because only near-end echo will be present.

The Uplink Channel

We consider now the communication from the customer to the server. The equipment at the customer location must send an analog signal in the usual telephone band. At the central office the signal will go through an anti-aliasing filter, be sampled at 8000 samples/s, and be represented by octets following the A or μ quantization law. The octets will be transported by the digital long-haul phone network to the ISDN terminal.

A simple-minded approach to communicating on the uplink is to use a standard analog modulation, and to have equipment on the ISDN side demodulate it. This would be suitable in situations where traffic is asymmetric.

If faster communication is desired on the uplink, one should reduce the quantization noise in the central office. This is possible in principle, as we presently outline, ignoring the presence of an echo.

By making use of the timing information obtained from the downlink, and by measuring the combined impulse response of the line and the anti-aliasing filter, the equipment on the customer side can predistort its output and ensure that the signal at the analog-to-digital converter has a predetermined level at the sampling time, depending on the information being transmitted. The previous discussion relating to the bit rate and required precision is also applicable here. In addition to allowing higher speeds, this approach probably requires less processing in the ISDN receiver than demodulating a standard V.34 modem signal.

Again we can gain some insight by taking a fundamental look at the problem. Assume that the impulse response between the customer and the sampler in the central office is $g(t) = 1/\sqrt{T}\,\mathrm{sinc}(t/T)$ with a 4000 Hz bandwidth, and that the central office filter limits the noise to that band. By Nyquist's sampling theorem, any analog signal at the A/D input can be written as $x(t) = \Sigma_k a_k g(t - kT)$, and it will be received in lowpass Gaussian noise. We ignore at this point the presence of an echo. This gives rise to the same memoryless discrete-time channel as for the downlink, with one important difference: the central office quantizes the received signal, following the A or μ-law.

Taking quantization into account leads to a continuous-input, discrete-output memoryless channel model. The transition probabilities can easily be computed in terms of the Q function. We can find its capacity by maximizing over the input distribution. It is known that it will also be discrete and that the number of input values with positive probabilities does not exceed the number of output values [8]. We have assumed that the optimal input levels are equal to the nominal output levels (we do not expect this to lead to a significant reduction) and have carried out the computation. The results appear in Fig. 4, which shows the capacities of both the uplink and the downlink channel with the A-law. The quantization in the central office reduces capacity by a few dBs, as expected. This capacity also bounds the "capacity of the telephone line," an elusive value that keeps increasing with time.

The bit error rate in the absence of coding is identical to that on the downlink, and the corresponding points are shown by asterisks. It is self-evident that coding will lead to gains slightly smaller than those on the downlink channel.

With full-duplex communication, the equipment on the ISDN side will need to compensate the far-end echo of the signal it originates. This echo will reduce the capacity, since it cannot be simply canceled.

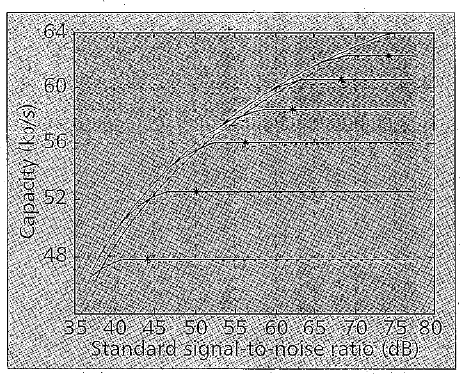

■ **Figure 4.** *Compariosn of the uplink (dashed) and downlink (solid) ideal chanel, A-law quantized, for d_{min} = 2, 4, 8, 16, 32, 64, and 128. Uncoded modulation: *.*

The Channel Characteristics

To get some understanding of the real behavior of the channel, we have taken some measurements on the telephone lines leading to our laboratory and our homes in the Sophia-Antipolis area. They should not been seen as a definitive study!

Our equipment consists of two personal computers, one with an ISDN card, the other with a digital signal processing (DSP) card. From the PC equipped with the ISDN card we placed local calls through the regular French telephone system, transmitted test sequences in transparent mode, and recorded the signal echo. The other PC answered the phone and sampled the signal at 16 kb/s, saving the samples on disk for further processing in a workstation environment using Matlab. We have investigated three parameters: the response of the system, the noise spectral density, and the infidelity of the D/A in the central office, that is, deviations of the transmitted levels from the A-law (we prefer the word "infidelity" to "nonlinearity", as an A-law quantizer is nonlinear by definition). We report on these three measurements.

The channel impulse response, which comprises the response of the interpolating filter in the D/A, that of the actual telephone line, and that of the front-end filter on the DSP card, was estimated by sending 125-symbol pseudo-random sequences of ± 2112 (preceded and followed by long periods with a constant level of 1). A least squares estimation algorithm, assuming the length of the response was 500 samples (31.25 ms), yielded the response (normalized to unit energy) displayed in Figs. 5, 6, and 7 on a logarithmic scale. The noise in the estimate can be observed on the graph.

The Fourier transforms of the responses also appear in the figures. They follow the specifications given by data sheets from Personal Computer Manufactueres (PCM) encoder/decoder (codec) and filter manufacturers. We note in passing that these data sheets show that in the uplink direction, the anti-aliasing filters are designed to reject noise at 50 or 60 Hz.

We now turn to the additive noise measurement. It was obtained by transmitting 10^5 constant samples and measuring the output. The Elf card is equipped with a 16-bit analog-to-digital converter, but the dynamic range of its telephone line input is only dB. Concretely that means that when we adjusted the gain to obtain signals with a peak amplitude of about 20,000 (on a full scale of 32,767), we observed noise with a variance of X even in the absence of an input signal (instead of the expected variance of 1/12). The A/D exhibits a slowly varying bias which was removed. This introduces some uncertainty about the behavior of the noise spectral density near $f = 0$. The power spectral density was estimated by a variant of the periodogram method with a 10 Hz resolution. It is displayed in the figures, normalized by the same factor as the impulse response. We note that power line harmonics are the dominant source of noise. We believe a large part of the noise "floor" originates in our A/D. The average noise level between 0 and 4000 Hz is shown by the horizontal dashed line.

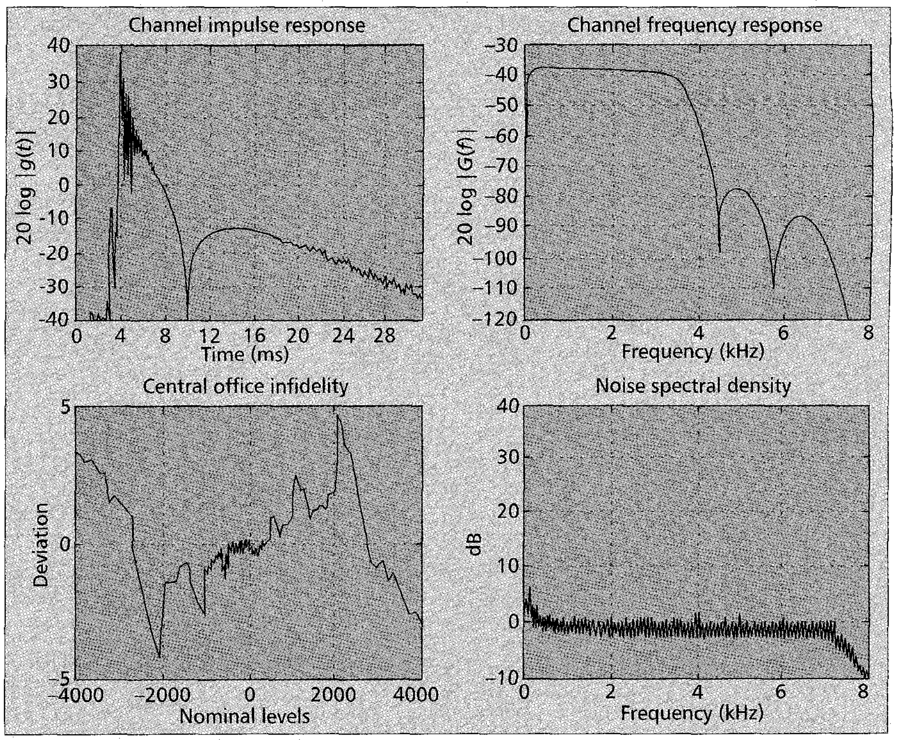

■ **Figure 5.** *Measurements of the LAB line.*

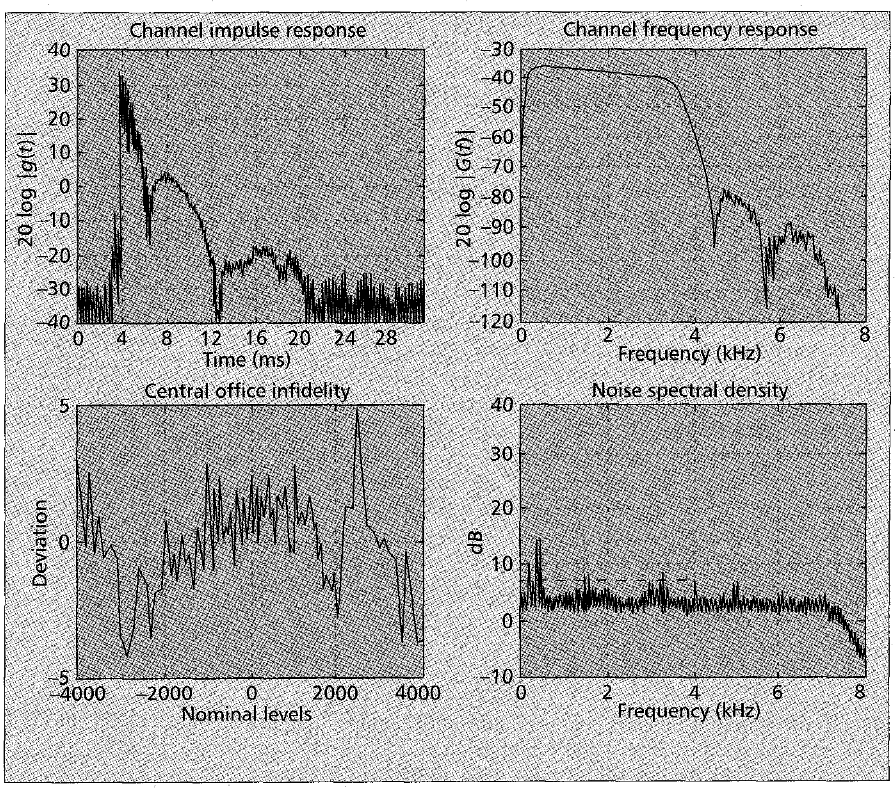

■ **Figure 6.** *Measurements of the MGT line.*

Finally, the infidelity of the D/A was obtained by transmitting a random 20,000-point sequence, subtracting the expected signal (using the estimated response) to measure the error, and using a least squares estimator. We see that the infidelity of the D/A can be very significant, especially in the extreme segments. It dwarfs the effect of the additive noise. As each level occurs only about 160 times in our experiment, the infidelity estimates are still noisy.

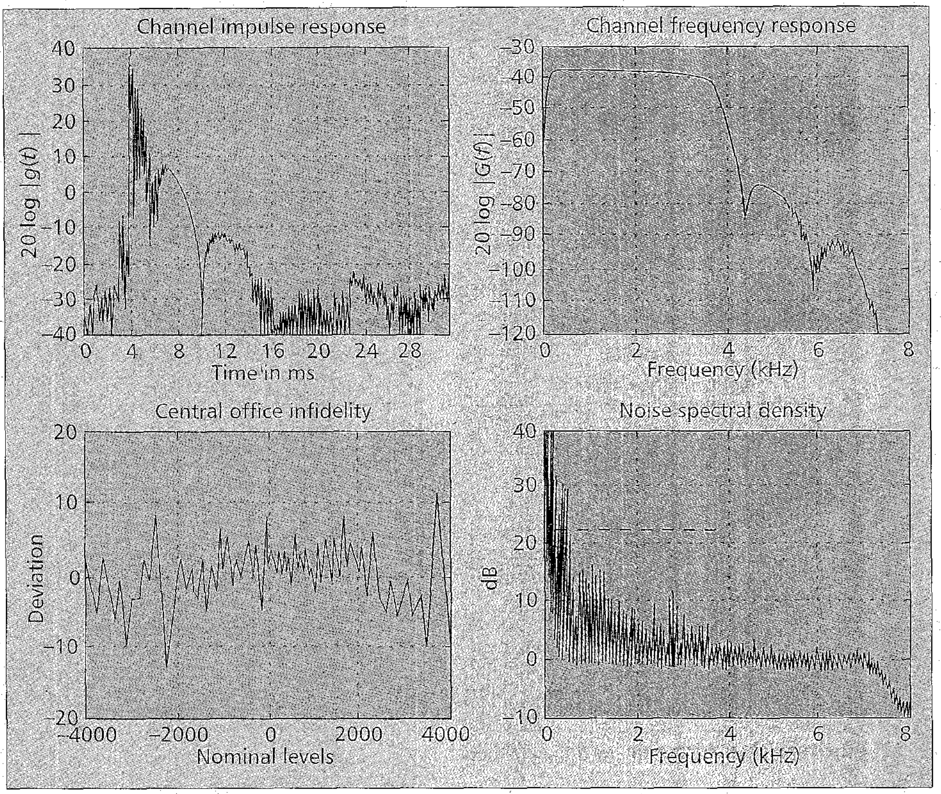

■ **Figure 7.** *Measurements of the PAH line.*

The Eurecom System

In the previous sections we have outlined the basic transmission capabilities of the downlink and uplink telephone channel, and we have reported some measurements.

We now face the task of outlining the structure of a receiver that can meet our goals. We developed a system combining in an original manner traditional detection and estimation techniques; its details will be reported elsewhere.

On the downlink there are two key problems. One is controlling intersymbol interference; the other is acquiring precise timing and interpolating the received signal. Real signals were transmitted from the ISDN side, captured by the DSP board, and processed off-line in Matlab. Our system could transmit at 48 kb/s on the LAB and MGT lines, but not on the PAH line.

Similarly, on the uplink one must predistort the signal on the analog side and synchronize it to the central office A/D clock. We were unable to experiment in that direction, because synchronization to the central office clock requires real-time operation.

Conclusion

We have shown that using the existing analog telephone loop, in conjunction with ISDN, provides a cost effective "information driveway" to information and education services for consumers, at double the speed of voiceband modems. Although the project will benefit telecommunication companies as traffic increases, no new investments or modifications to the digital transmission and switching network will be required.

Through some measurements and the sketch of a system architecture we have shown that our approach is feasible on some telephone loops, and that its range of applicability can be increased with some effort. We believe that data transmission at rates approaching 56 kb/s on the downlink, and 48 kb/s on the uplink should be possible on a fair number of telephone lines.

The idea of explicitly considering the presence of the quantization process in modem design has apparently occurred to many familiar with Shannon's quote [9] which prefaces this article. The idea of avoiding the quantization noise by pre-adjusting the transmitted analog signal was first discussed publicly by Kalet *et al.* in [10], after the initiation of our research. However, this approach did not take advantage of a digital connection in one of the modems.

Recently, several companies announced intentions to market high-speed modems which appear to be based on the concept described in this article. We believe they can be realized on standard DSPs at a cost comparable to V.34.

Acknowledgments

We are grateful to Motorola ISG, in particular to Dr. Vedat Eyuboglu, for their support of this research. We also thank Dr. Jim Mazo of Lucent Technologies for his comments on an earlier version of this article.

References

[1] J. Bellamy, *Digital Telephony*, 2nd Ed., Wiley, 1991.
[2] D. V. Batorsky and M. Burke, "1980 Bell System Noise Survey of the Loop Plant," *AT&T Bell Labs Tech. J.*, vol. l63, May-June 1984, pp. 775–818.
[3] M. Carey *et al.*, "1982/1983 End Office Connection Study: Analog Voice and Voice Band Data Transmission Performance Characterization of the Public Switched Network," *AT&T Bell Labs Tech. J.*, vol. 63, Nov. 1984, pp. 2059–2119.
[4] G. Ungerboeck, "Channel Coding with Multi-Level/Phase Signals," *IEEE Trans. Info. Theory*, vol. IT-28, Jan. 1982, pp. 55–67.
[5] G. Ungerboeck, "Trellis-Coded Modulation with Redundant Signal Sets, Part I: Introduction," *IEEE Commun. Mag.*, vol. 25, Feb. 1987, pp. 5–12.
[6] G. Ungerboeck, "Trellis-Coded Modulation with Redundant Signal Sets, Part II: State of the Art," *IEEE Commun. Mag.*, vol. 25, Feb. 1987, pp. 12–22.
[7] E. Biglieri *et al.*, *Introduction to Telllis-Coded Modulation with Applications*, Macmillan, 1991.
[8] R. G. Gallager, *Information Theory and Reliable Communication*, New York: John Wiley & Sons, 1968.
[9] C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication*, University of Illinois Press, 1949.
[10] I. Kalet, J. E. Mazo, and B. R. Saltzberg, "The capacity of PCM voice band channels," *Proc. ICC '93*, vol. 1, Geneva, Switzerland, May 1993, pp. 507–11.

Biographies

PIERRE A. HUMBLET [F '93] received the Electrical Engineer degree from the University of Louvain, Belgium, in 1973, and the M.S.E.E. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, Massachusetts. After graduating in1978 he remained at MIT where he became a professor of electrical engineering. In 1993 he joined the Eurecom Institute, in France, where he heads the Mobile Communications Department. His teaching and research interests are in the area of communication systems, particularly mobile digital networks and optical networks.

MARKOS TROULIS [StM] received the diploma in electrical engineering from the Aristotle University of Thessaloniki, Greece in 1991 and the Diplôme d'Ingénieur in telecommunications from the Ecole Nationale Supérieure des Télécommunications de Paris, France, in 1993. In September 1993 he joined the Eurecom Institut, Sophia Antipolis, France, where he is working toward the Ph.D. degree. His thesis deals with signal processing applied to digital receivers. His interests include statistical communication theory, signal processing for communication applications, synchronization algorithms, and information theory. He is a member of the Technical Chamber of Greece.