

ISSUES IN HANDLING CABLE SIGNALS WITHIN DIGITAL OPTICAL INTERCONNECT NETWORKS

James O. Farmer
Electronic System Products
An Antec Corporation Business

Edward J. Callahan
Antec Corporation

INTRODUCTION

As cable systems seek to add telephony services and interconnect headends for operational efficiencies, they will migrate toward use of digital fiber optics for the interconnections. This is true even of systems which expect to continue to deliver analog programming to the home for the immediate future. When developing a digital optical interconnection system, one may choose either proprietary systems, or may select compatible standard systems available from a number of manufacturers. Here-in the case is made for the latter approach, in which standard systems, originally developed for telephony applications, are used to transport video and ancillary signals from one headend location to another.

Interface issues include not only video, but also audio services, addressable data, control and new data services. The versatility of the network is enhanced by integrating cable television's needs with standard multiplexing systems. Multiplexing is explained, and the hierarchies in common usage are introduced. Finally, a short space is devoted to describing some of the quality issues in hybrid analog/digital networks.

INTERCONNECTING HEADENDS

Headend interconnection is becoming popular today to allow efficiencies of operation, more controllable advertising insertion and improved reliability. Figure 1 shows how

headends may be connected. Two architectures are shown. Figure 1 (A) shows a single master headend connecting with three hubs, or sub headends. Fiber cables interconnecting the headends are routed two ways, so if one path is interrupted, for example by a car hitting a pole, the other path can be used. Figure 1 (B) shows a modified concept in which the headends are connected in rings. Signals travel around the ring in both clockwise and counterclockwise directions, passing through each headend on the way to the next. In this case two master headends are used. Signals are routed completely around the ring in two directions. As in (A), if one path is interrupted, the path around

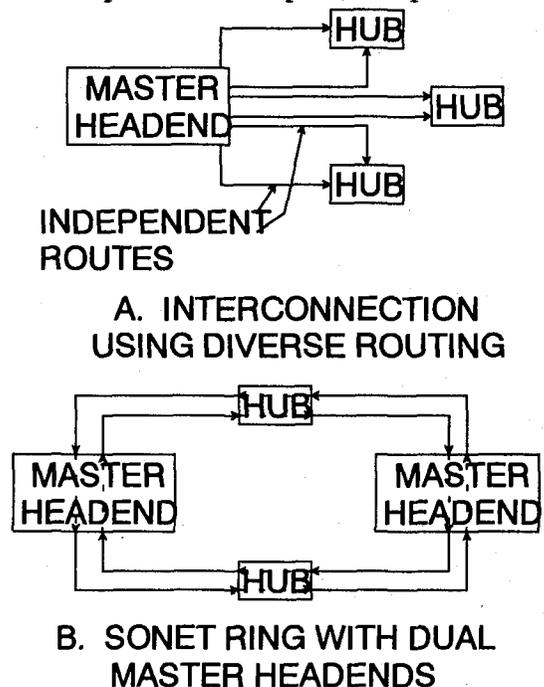


Figure 1. Options in Interconnection

the ring in the other direction is available to carry signals. Further, if something happens at one headend, there is a second master headend to take over, so that service is interrupted, if at all, to only one node.

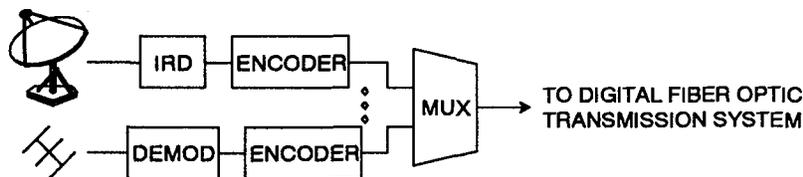


Figure 2. Simplified Master Headend

The preferred method used to transmit the signals is the SONET Synchronous Optical Network. It is a fiber optic based system with digital transmission of all signals. Figure 2 shows the basic conditioning of a video signal at the master headend. Signals are converted to baseband and supplied to *encoders*, which convert to digital and compress the signal. At this time MPEG compression is not being used due to the cost of the encoders, though this may change in time. The digital signals are multiplexed (described below) onto a single data stream for transmission throughout the network. At each node, the signals are passed through with provision to drop and add signals if necessary. The signals are *demultiplexed* and converted to NTSC for modulation and conventional transmission to subscribers.

WHY DIGITAL SIGNAL TRANSMISSION?

Within the network, there are excellent reasons for transmitting TV signals digitally. Digital transmission can be less costly *when applied intelligently*. Because fiber transmitters for baseband digital signals don't have to control distortion (be linear), they are potentially lower in cost. Receivers are also potentially less costly. Of course realization of the cost savings will depend on volume, which cable television may drive in the future.

Another reason for digital transmission is that, within limits, no video degradation

occurs during transmission or switching.ⁱ Of course we can and do pick up distortion in the process of digitizing (and compressing) a signal. That distortion is a trade-off in how few bits we use to represent the signal, the quality of the signal and our ingenuity in compressing the video.

In the future the industry is likely to begin switching signals more, as we subdivide systems into fiber nodes to which we can route different

signals and as we provide more customized cable features and transmission path back-up. These will require more switching and routing of signals. These operations can be done much better if the signals are in digital format, since we don't have to worry about signal degradation.

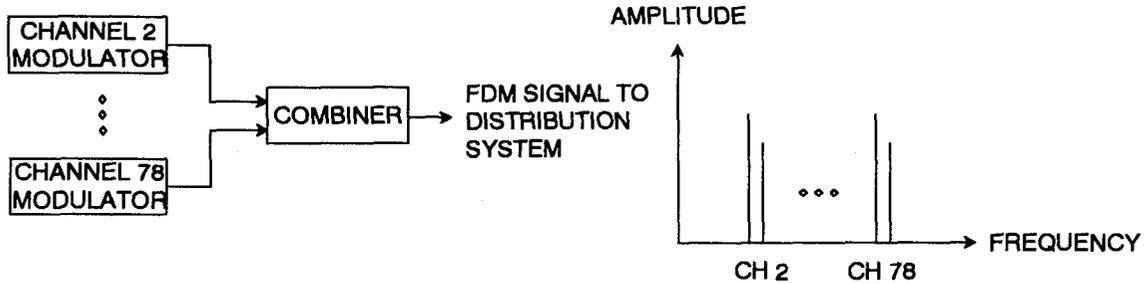
Video file servers, big brothers to the file servers used in local area networks (LANs) today, will, in future years, become commonly used for storage of programs. The programs will then be transmitted to subscribers, either on a subscription basis or on demand. File servers store video digitally, so in many cases it will make sense to transmit the signals in digital form, at least part way to the subscriber. The trick will be to balance the cost of transmission against the cost of converting the signal back to NTSC.

Once we have a signal digitized and compressed, the next effort is route the signal to a given subscriber, along with other signals. Simultaneously other signals will be routed to other subscribers. The process of combining all of the signals for transmission is the subject of this paper. We call this combining process "multiplexing."

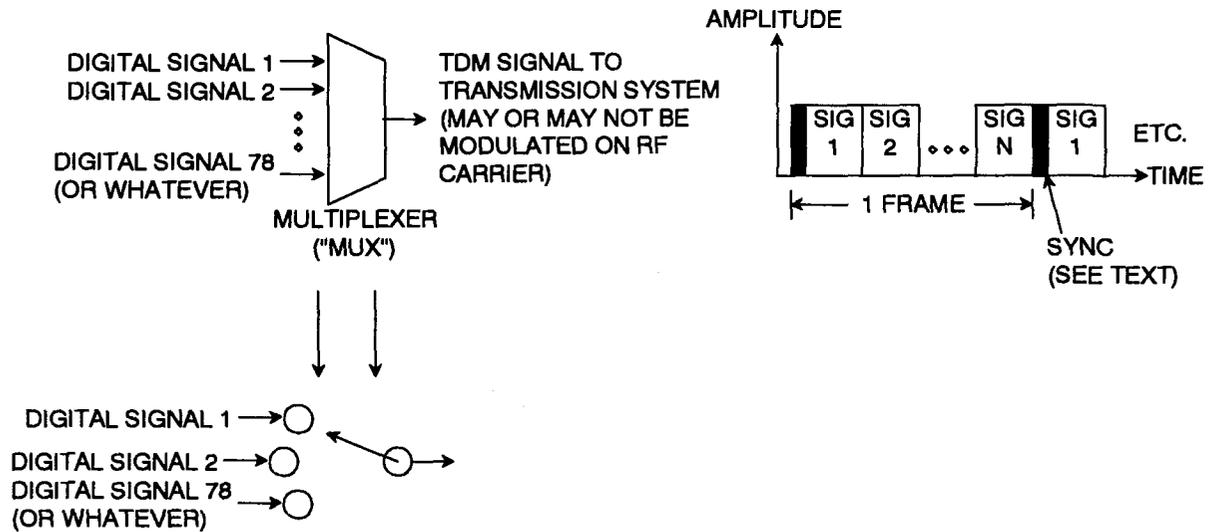
INTRODUCTION TO DIGITAL MULTIPLEXING

As we move into a world in which video is carried in digital form, we encounter the term *multiplexing*. This is another way of saying that we are combining signals. In the cable

more than a different way of putting signals together on a single cable. Rather than putting each signal on its own unique RF carrier **frequency** and multiplexing by combining the signals on a cable, we assign each signal to its own unique **time** slot, switching from one signal to another at the appropriate time.



A. FREQUENCY DIVISION MULTIPLEXING (COMBINING)



B. TIME DIVISION MULTIPLEXING

Figure 3. Concept of Multiplexing

television world, we are accustomed to combining RF signals in a *combiner*, which puts many signals together on a single cable. In traditional cable television, these signals are put together using a technique called *frequency division multiplexing* (FDM). In FDM, we put several signals together on one cable by assigning each its unique band of frequencies. The same idea is practiced in the digital world, except instead of employing FDM we normally use *time division multiplexing* (TDM). TDM is nothing

Figure 3 compares time and frequency division multiplexing. In (A) we are doing frequency division multiplexing at a headend, by sending each signal (on its own frequency) to a combiner. The spectrum diagram shows what the resulting signal looks like. We have the familiar diagram of each channel in its place on the cable. In order to recover one of the signals, we tune it with circuits that can differentiate one frequency band from another.

In contrast, (B) shows time division multiplexing, as practiced in the types of systems discussed here. Each digital signal (a series of ones and zeros coming one after the other) is applied to a multiplexer (commonly called a "mux"), which serves an analogous function to the combiner in a conventional headend. The amplitude vs. time diagram shows the result. A short segment (one or more bits) of Signal 1 is transmitted, followed by a short segment of Signal 2 and so on. Notice that, as in all digital systems, everything has the same amplitude: amplitude data is carried as a series of bits which represent the numerical amplitude. In order to recover the individual signals, we apply the data stream to a demultiplexer ("demux"), which takes the signal back apart, into its individual data streams. The demux does this by switching Signal 1 to its "bin," then switching Signal 2 to its bin and so on. In order to let the demux know when a particular signal is being sent, we must add synchronization (analogous to vertical and horizontal sync in TV) to the data stream. The sync is shown in the diagram. Depending on the system, sync may be part of each data stream, or as shown here, may be added to the composite data stream.

In order to get all of the incoming digital signals into the same data stream, we must speed up the data in the composite stream. (Think of a river having several tributaries: streams that feed it. As more water is added to the river, it must either flow faster, wider or deeper.) For example, suppose we had 24 streams of 64 Kb/s (kilobits per second - often informally called, simply, "kilobits"). The data rate required to transmit all of the data streams on one cable is at least $24 \times 64 \text{ Kb/s} = 1.536 \text{ Mb/s}$. The actual data rate may be higher still if sync is added external to the data, as in Figure 3. For example, we may add enough sync to require an additional eight Kb of data to be transmitted each second. In this case, the resultant data stream would be $24 \times 64 \text{ Kb/s} + 8 \text{ Kb/s} = 1.544 \text{ Mb/s}$. By the way, these are not

arbitrary numbers: we shall encounter them again.

SPECTRUM

The spectrum of a digital signal extends to at least one half the bit rate of the data stream: if we transmitted a one followed by a zero, we would have one cycle of a square wave. In some cases, we want to filter the signal so we have nothing above the minimum frequency required to get data through. In other cases, we would handle the signal as a square wave, with a spectrum much wider than required.

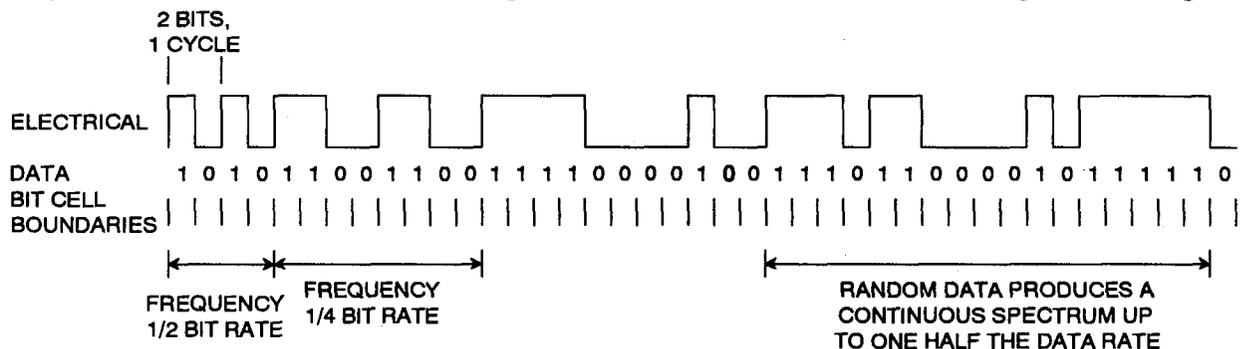
Figure 4 shows a sample sequence of a baseband digital signal and the resulting spectrum. In (A) we see a sequence of bits that starts out with the sequence 1010, followed by a sequence 11001100, then a random sequence. The 1010 sequence represents the highest frequency the system will be called on to transmit. As shown, two bits, a one followed by a zero (or vice-versa), make up one cycle of a periodic waveform. Thus, we have an inherent "efficiency" of two bits per Hertz with this form of data transmission.ⁱⁱ In (B) we show the spectrum of this square wave (1010...). The fundamental is shown as a solid line at one half the bit rate. Since we are dealing with a square wave we have harmonics of this signal, and these are shown as dotted lines at two and three times the fundamental frequency. (Granted, a square wave doesn't have even harmonics, but we will have them by the time we get differing duty cycles caused by different data patterns.)

The sequence 11001100... has a rate one half that of the 1010... sequence, so the fundamental of its waveform is one half the fundamental of the 1010... sequence. This is shown as the long-dashed dashed line in (B). The harmonics are not shown. As we go to waveforms having longer sequences of ones and zeros we have even lower frequency components. If we have sequences of random lengths of ones and zeros, such as suggested at the right of (A), the spectrum is filled in from the one half bit rate point of (B) down to nearly zero frequency. Harmonics exist, and may or may not be removed, depending on the following transmission system. This is illustrated in (C), which shows the frequencies occupied by the fundamental and harmonics of the random data waveform. If we didn't do something to limit the number of ones or zeros that could occur in succession, the fundamental would occupy frequencies down to zero Hz. In most systems, this would cause problems with clock recovery and modulation, so something is

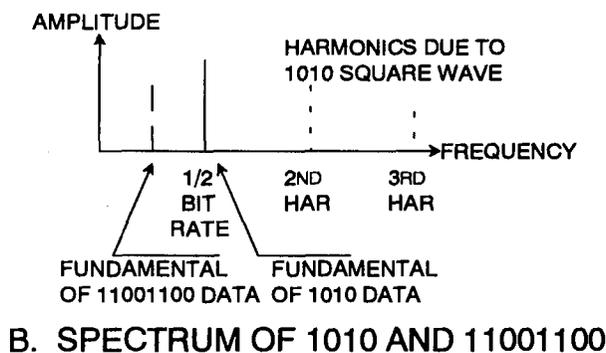
added to the data to prevent excessively long strings of ones or zeros. This limits the low frequency end of the spectrum, as shown in (C).

The limiting of the longest sequence of ones or zeros is normally done by exclusively OR-ing the data stream with a pseudo-random data sequence known to both the transmitter and receiver. This process, unfortunately, is called "scrambling." It has nothing to do with hiding a signal from an unauthorized viewer (the definition of scrambling traditionally used in our industry).

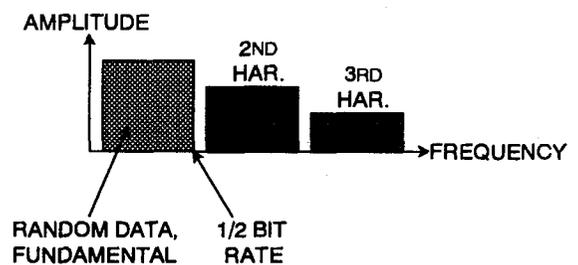
Of course a multiplexed digital signal can be modulated onto an RF carrier, and this is frequently done. Examples are the digital music services carried on many cable systems today. In both systems the left and right channels, along with certain data, are muxed onto a single RF carrier. In one of the systems, the RF carrier consists of only the left and right channels and data, for one stereo pair. The signals



A. SAMPLE DATA STREAM SHOWING 2 BITS IN 1 CYCLE



B. SPECTRUM OF 1010 AND 11001100



C. SPECTRUM OF UNFILTERED RANDOM DATA

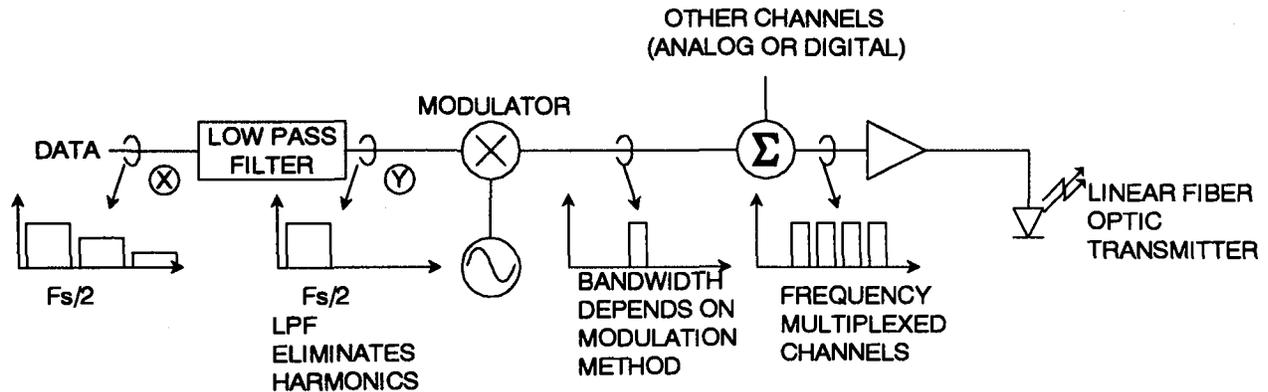
Figure 4. Spectrum of a Baseband Digital Signal

for various programs are then combined (frequency division muxed) onto the cable, in channels about 600 KHz wide. In the other system, some additional TDM is done, in that five stereo pair are TDM'd and modulated onto a single carrier about three MHz wide. The arguments concerning the merits of the two approaches are far beyond the scope of this paper.

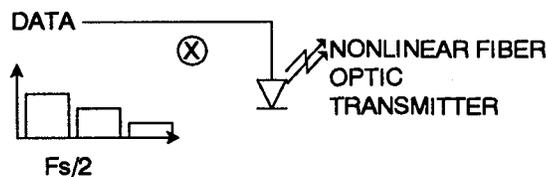
COMPARISON OF FREQUENCY AND TIME DIVISION MULTIPLEXING

It is quite instructive to consider further the difference between a frequency multiplexed system and one that is time division multi-

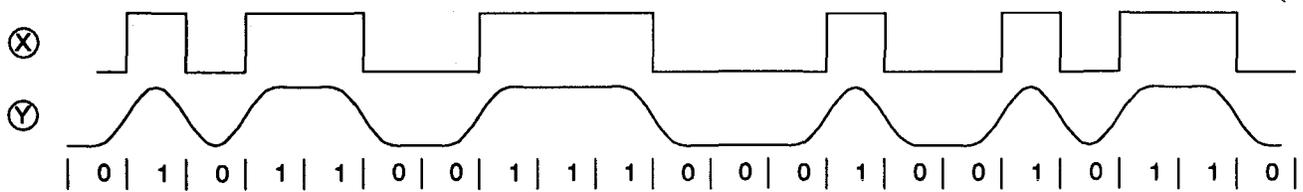
plexed. This will help illustrate the advantages of handling baseband digital signals where practical. Figure 5 shows a system for transmitting analog or digital data using modulation and frequency division multiplexing in (A), and using digital baseband transmission in (B). (C) shows waveforms useful in understanding (A) and (B). In (A) we have taken in digital data at the left. A sample waveform is shown in (C), as circle X. This is the way data would look coming out of any normal digital process, such as a codec that converts video to compressed digital data. The spectrum is shown in (A), as explained above. In order to reduce the transmitted spectrum as much as possible, the data is supplied to a low pass filter having a cut-off nominally at one half the bit rate, to



A. ANALOG TRANSMISSION



B. DIGITAL TRANSMISSION



C. WAVEFORMS

Figure 5. Linear vs. Non Linear Transmission

eliminate harmonics. If transmitted, the harmonics would require power, and would spread out the bandwidth required for transmission, without adding any useful intelligence to the signal. The spectrum at the output of the filter is shown. The resulting waveform is shown as circle Y.

From the low pass filter, the signal is applied to a modulator, which places it on a suitable frequency for carriage on the cable system. At the output of the modulator, we have a spectrum at some RF frequency, depending on the frequency of the modulator. The occupied bandwidth of the signal is a function of the low pass filter, and also of the type of modulation used: higher orders of modulation take up less spectrum, but are more expensive to decode and are more susceptible to transmission errors.

The output of the modulator corresponds to the output of a conventional cable television modulator, which is supplied to the headend combiner. In (A) Σ (sigma) represents the combiner, which combines this signal with other digital and analog signals. The composite spectrum is shown at the output of the combiner, and is amplified and supplied to a laser transmitter in this example. So far, we have a conventional cable television system with a mix of analog and digital signals. Because of the multiplicity of signals, the fiber optic transmitter must be linear. That is, it must handle any instantaneous signal amplitude in the correct proportion, a requirement to which we are accustomed in cable television. The transmitter is relatively expensive due to this need for linearity.

By way of contrast, section (B) of the figure shows what is required for time division multiplexed digital transmission. The waveform of circle X is applied directly to the laser. Of course, the data rate is much faster than is the data rate of a signal in (A). Rates up to 2.4

Gb/s are in use today. Note that the laser transmitter only needs to be on or off, according to whether a one or zero is being transmitted. This allows a transmitter which potentially costs less than a linear transmitter (as in (A)) having similar information capacity. Of course the cost of getting data into and out of the digital format is high, and this must be considered in developing a system. Fast multiplexed data may not be a suitable format for transmission to subscribers, due to the recovery cost, but is suitable for headend interconnect or for use anywhere in the network where costs are shared among a number of subscribers. (We may see some applications for data multiplexed at lower rates and transmitted directly to consumers, though.)

WHY MULTIPLEXING STANDARDS

We need a standard way of muxing in the TDM world. Granted, there are manufacturer-specific (non-standard) muxing schemes in use today, and they work well within the confines of what they were designed to do. However, they have the disadvantage that one can only obtain equipment from that specific manufacturer. If he doesn't have the particular function you need or if he leaves the business, you are stuck. If you want to interchange programming with another cable system (in a regional headend interconnect for example) and that system uses different equipment, you will have to bear the cost of some sort of conversion, if you can make the conversion at all. Maintenance is a problem because personnel must become familiar with specific techniques that won't apply if the vendor is changed. A proprietary multiplexing system will not allow a system to add telephony services easily. Finally, we can expect limited cost reductions because the vendor has no competition once an initial purchase is made.

The telephone industry has had these same problems and has evolved a standard way of handling data. We may as well take advan-

tage of what they have developed. Not that the telephone industry has implemented their data standard perfectly: only now are standards being implemented that will let a North American signal interconnect gracefully with a European signal. Further, despite efforts to the contrary, just because a piece of equipment conforms to a particular standard (such as DS3), that doesn't mean that one manufacturer's piece of hardware will interconnect with another. However, the telephone industry had a big head start on the cable television industry in solving these problems, so we may as well take advantage of what they have developed. That way, we don't have to reinvent the wheel ourselves, and if we want to interconnect with them in the future (a situation almost guaranteed by recent events), we will be ahead of the game if we use the same standards.

Finally, the cable television industry is likely to begin handling telephone signals in the next few years, with networks we are installing now. In order to interface with other telephone vendors, we will have to use the same set of standards that they use.

In this spirit, we shall attempt here-in to describe how multiplexing is done in the telephone industry, and relate it to our needs in cable.

The multiplexing standards allow for various levels of interconnect, mostly differentiated by data rate. In order to facilitate the development of standardized equipment, only certain data rates are allowed. These can be combined in prescribed manners. The process of defining the particular ways data can be combined is called a "hierarchy." Consider relevance to the commonly understood definition of the word: "a group of persons or things arranged in order of rank, grade, class, etc."ⁱⁱⁱ In this case, the *group* is a set of standard data rates, and the *rank* is the way in which they are ordered. There are at least two hierarchies involved, an older asynchronous digital hierar-

chy (ADH) and a newer synchronous digital hierarchy (SDH). The SDH in North America is the still evolving Synchronous Optical Network (SONET). The words "asynchronous" and "synchronous" refer to timing and synchronization between individual data streams. While important, this is not a topic that can be covered within the present scope.

In connection with a discussion of the hierarchies, we shall introduce a number of relevant terms and the data rates associated with them. Everything builds on what came before, though the manner in which this is done is a bit inconsistent for historical reasons. We shall not go into the history of how things came to be as they are, though a paper^{iv} published not long ago contains some interesting insights into the topic, as well as other useful information.

ASYNCHRONOUS DIGITAL HIERARCHY

The older asynchronous hierarchy developed during the 1960s, as telephone companies attempted to digitize voice signals for efficient switching and transmission. A suitable data rate for carrying voice traffic is 64 Kb/s.^v Thus, everything is developed around this data rate. Figure 6 shows the hierarchy that exists. The voice signal to be transmitted is digitized in a device called a "codec" (for coder-decoder: in the telephone industry everything is symmetrical, and where a signal must be coded, a related signal must be decoded). The output of the codec (really the coder portion) is a 64 Kb/s data stream, a *DS0* (spoken "DS zero") data stream. It is muxed with 23 other 64 Kb/s data streams, plus eight Kb/s of sync, into a 1.544 Mb/s data stream called a *DS1* signal. (Notice the data rates used here, compared with the example of multiplexing given earlier.)

The reader may recognize this as the so-called T1 rate that has been used in earlier stand-alone cable television data applications. Technically, calling this T1 is not correct. The

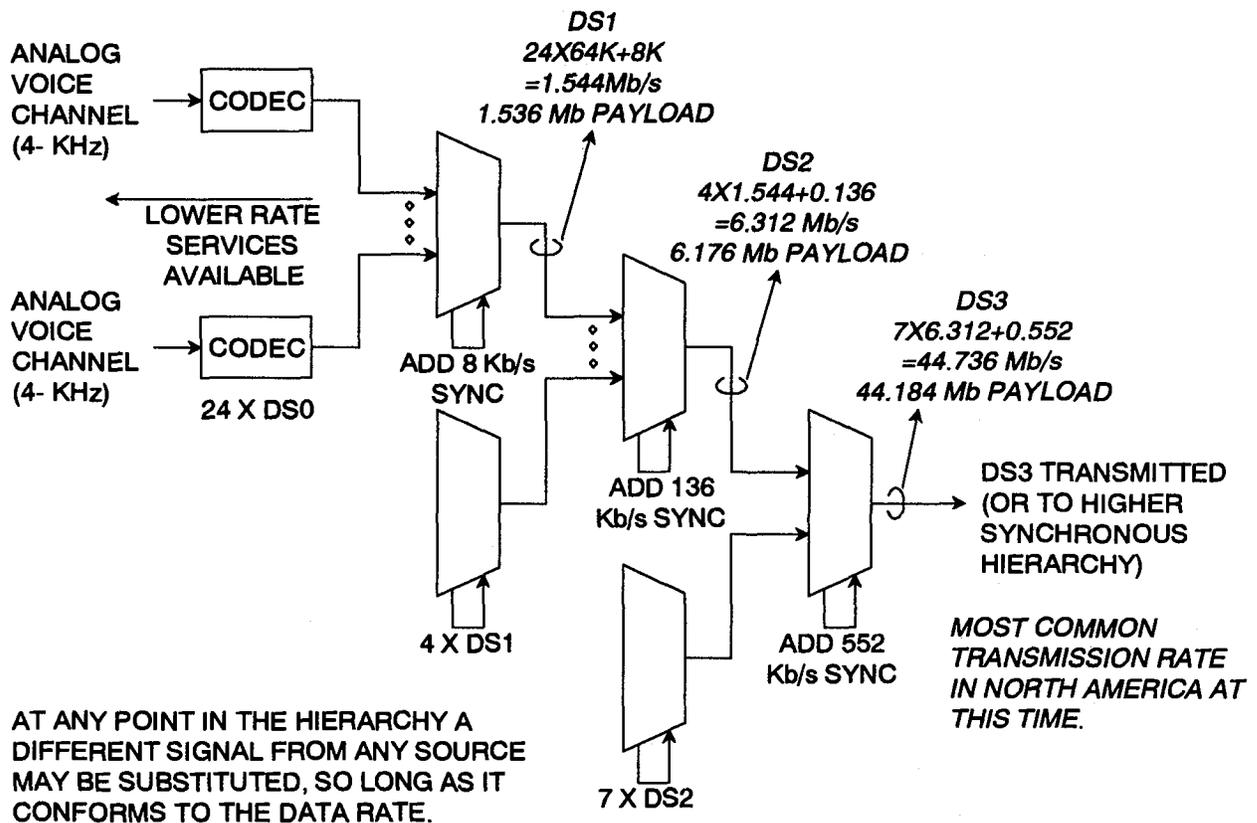


Figure 6. Asynchronous Digital Hierarchy

term T1 is reserved for a DS1 data stream carried on twisted pair (conventional telephone) cable. The definition of the data stream, which is what cable television has used before, is contained in the DS1 specification. T1 refers to a particular type of modulation used to get the data onto a twisted pair.

Note that data rates lower than 64 Kb/s (the DS0 rate) exist and can be muxed into a data stream.

Four DS1 data streams can be muxed into a DS2 data stream, and seven DS2s can be muxed into a DS3, which operates at 44.736 Mb/s. At each stage in the multiplexing, sync is added as overhead, as illustrated in Figure 6 above. The overhead is necessary in order to recover the data, but does not carry data itself. This is one of those necessary evils of life: we have to tolerate the overhead, even though it doesn't help us by carrying useful (to the user)

data. Later we shall see hierarchies that depend on imbedded sync and don't utilize added overhead.

Figure 6 shows the computation of the amount of payload (useful data) that exists at each point in the hierarchy. The payload numbers shown consider sync from lower levels (to the left), to be payload, not overhead. Whether or not this is true depends on the application. If the application is voice traffic, where everything is transmitted in 64 Kb/s channels, then the sync overhead from previous levels of the hierarchy is clearly still overhead, not payload. However, if the data from a previous level is something else, it may all be useful payload data. A simple computation from information in Figure 6 shows that a DS3 channel bears up to 672 voice channels of 64 Kb/s each, or a net payload of voice data, of 43.008 Mb/s. The difference between this and the 44.184 Mb/s payload suggested in the figure, is the overhead added at the DS1 and DS2 levels.

To get a feel for transmitting video, consider that the professional D1^{vi} data rate, a method of recording and transmitting video used in some broadcast and production studios, is 140 Mb/s. This is uncompressed video, with no attempt to minimize the bandwidth. Compression systems that reduce the data rate to permit a video channel to fit in a DS3 data stream, have been used for specialized transmission of professional video for a while. The motivation to squeeze the video into one DS3 is that this is currently the most commonly available data rate interface in North America.^{vii} This year equipment has come on the market

that MPEG is an asymmetrical process: the cost of compression is very high, while the cost of decompression is low. This works in our favor for transmission to the subscriber, but works against us in the network.

Signals may be muxed into the hierarchy at any point. For example, at the DS2 level, one could mux together some combination of digital audio, voice and data signals as required, so long as each existed at a data rate supported by the commercially available multiplexers. Out of band addressable data from a headend computer to set top converters is normally transmit-

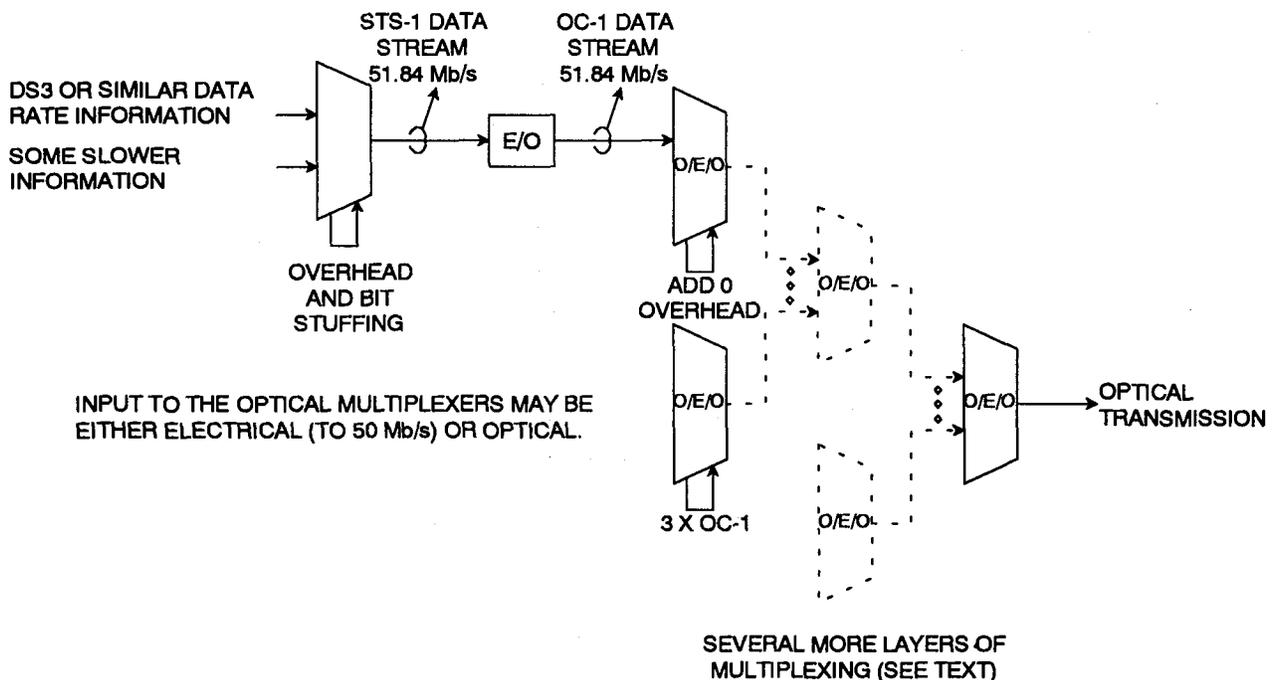


Figure 7. Synchronous Digital Hierarchy

for cable television use, that puts two NTSC signals in one DS3. Each signal includes three audio channels (stereo and SAP) and an 18 Kb/s data stream. The application is headend interconnect.

Depending on what one believes MPEG compression will ultimately deliver, one may be able to put up to about ten compressed NTSC signals in one DS3, though the hardware to do so is not quite available yet. A disadvantage of MPEG compression systems for network use is

ted at either 9.6 or 19.2 Kb/s, depending on the system. These rates are standard interface rates and can be interfaced to the multiplexer at any level where a suitable bandwidth channel is available.

On the other hand, digital audio systems don't operate at standard data rates, so some sort of converter must be supplied by the manufacturer or a third party, to raise the data rate to the next highest standard so it can be muxed in. The data rate is raised by "stuffing"

bits. That is, extra, or "junk" bits are added to the data stream to bring the total rate up to the next standard rate that can be muxed. These extra bits are thrown away at the receive end.

SYNCHRONOUS DIGITAL HIERARCHY

The newer synchronous digital hierarchy (SDH) is currently in an advanced state of development and in initial stages of deployment. It is fiber optic based. A significant motivation for SDH is that we cease "burning" bandwidth for synchronization: the sync information already in the data is used for synchronization, so all we have to add is payload. The synchronous hierarchy begins at the DS3 level. Figure 7 shows this hierarchy.

Data at the DS3 rate (or lower) is supplied to a multiplexer whose output is at the synchronous base rate of 51.84 Mb/s. This rate is chosen to allow multiplexing in either a DS3 rate or the European equivalent. This is the point at which an attempt is made to reconcile the North American and European rate standards. (Full compatibility is not attained, however, until the next level of multiplexing.) The electrical output at 51.84 Mb/s is known as an STS-1 data stream. In the past, this signal has not been utilized for external interface with multiplexers, though one multiplexer vendor has announced product availability about mid 1994. A more common interface point is after the data stream has been converted to optical form in the electronic to optical conversion process shown. From this point on in the multiplexing hierarchy the signal interface is optical, even within the confines of one headend or other center at which multiplexing is done ("office" in telephone terminology). Multiplexers at later points in the hierarchy include conversion of the input signals from optical to electronic form, internally muxing the signals in electronic form, then re-conversion to optical on the output. Optical interfaces are much easier to handle, even in the space of one room, than are electronic interfaces. This fact

is painfully apparent to anyone who has spent late nights fighting ground loop problems!

The hierarchy is a bit more consistent in the synchronous world than it was in the asynchronous. At each level, three channels of the previous level are muxed. All synchronization is built into the data streams from the DS3 level or equivalent, and this synchronization is interpreted by the multiplexers and demultiplexers. Consequently no external synchronization is needed, and no additional overhead is added in the process of muxing to higher levels.

From the OC-1 level, three data streams are muxed to form an OC-3 stream at three times 51.84 Mb/s, or 155.52 Mb/s. The process continues in groups of three to the maximum currently defined rate of OC-48, 2.48832 Gb/s. This rate can actually be carried on a single fiber optic cable today. To compute the data rate at any level of multiplexing, multiply the data rate at OC-1, 51.84 Mb/s, by the level of multiplexing. For example, the OC-48 data rate is 48 times 51.84 Mb/s. All of these frequencies are exact and are shown with no rounding error.

REGIONAL INTERCONNECTS

Now that we see the why and the how of multiplexing signals, let us examine the use of digital SONET rings to accomplish cable television headend regional interconnects. There are several reasons why one or more cable operations might want to interconnect headends.

1. Ad insertion. A regional interconnect allows ad insertion equipment to be centralized at the master headend site, with regionalized ads inserted in the appropriate channels for transmission to all other hubs on the ring. Localized ads specific to one or more hubs on the ring could be stored at the master hub site and then

transmitted over the ring to specific hub locations at the appropriate time.

2. Headend interconnect and/or consolidation. Off air satellite-delivered signals can be received at the master headend, encoded as necessary, then transmitted on the ring to all other hubs. Depending on the situation in a particular market, it might be wise to designate one of the other hubs on the ring as a back-up headend site with appropriate off-air and satellite reception capabilities along with encoding equipment necessary to interface those backup signals with the digital ring.

Satellite delivered digitally compressed video signals (such as MPEG-2) for delivery to a digital set-top box at a subscriber location will be interfaced with the ring at the master headend with devices called "variable rate multiplexers". As their name implies, those devices can flexibly package various lower speed, non-standard rate data streams, into virtual containers (called "virtual tributaries") which fit nicely into higher speed, standard rate data streams. The input to a variable rate multiplexer connects to a line card tailored to the application. The output of this line card conforms to the appropriate higher level standard rate, making it easy to fit into a standard hierarchical structure.

Other digital sources which operate at non-standard data rates, such as digital music services, electronic program guides, and some downloadable game channels, can also be handled with variable rate multiplexers.

At the remote hub sites, video sources which were digitized, compressed, and multiplexed into higher rate data streams at the master headend, are now demultiplexed, decompressed, and converted back to analog format, and are then modulated to the appropriate cable channel.

Compressed digital video services which had been satellite delivered to the master (or backup) headend, are received from the high speed digital ring, demultiplexed to the correct rate, and then modulated onto an RF carrier for transmission to the subscriber location. Other digital data sources (digital music services, game channels, electronic program guides, etc.) are received, demultiplexed, and presented to an appropriate digital modulator for transmission over the broadband hybrid fiber/coax cable television system to the subscriber.

MULTIPLEXING INTO THE STANDARD

Equipment is available to multiplex in all of the standard data rates used today. These start at low rates we know from the EIA 232 (formerly RS-232) standard: 1200 b/s, 2.4 Kb/s and so on through at least 56 Kb/s. Any 64 Kb/s data stream can be multiplexed into a higher level data stream using off the shelf equipment, as can any higher level of multiplexing shown in the figures above. 1.544 Mb/s (DS1) data is commonly encountered. The 10 Mb/s data rate of Ethernet local area networks (LANs) can be accommodated with commercially available muxing equipment. We mentioned above that DS3 data (44.736 Mb/s) is the most common interface currently available in North America. Any signal at any of these data rates may be multiplexed into a higher level using equipment available from a number of manufacturers.

The table shows the interfaces that are currently available for OC-48 multiplexers. Also shown are some lower data rate levels not generally supported in OC-48 multiplexers. These rates may be multiplexed in at lower asynchronous rates, before one gets to the optical multiplex levels.

To our knowledge the data rates used in addressable converter systems are standard, so one can get addressing data from one point in a network to another with no difficulty. Of

Table 1. Interfacing Optical Multiplexers

INTERFACE	COMMENTS
DS0, DS1	Not likely to be interfaced directly in OC48 mux due to cost. Can be interfaced at lower level.
DS3	44.736 Mb/s. Most common North American interface today.
STS-1	51.84 Mb/s. Electrical equivalent of OC-1 optical rate.
STS-3	155.52 Mb/s Electrical.
STS-3c	155.52 Mb/s. Not likely due to cost, interface issues. (The "c" indicates a concatenated signal - made up of lower level signals transported as a single signal.)
OC-3, 3c	155.52 Mb/s Optical.
OC-12, -12c	622.08 Mb/s Optical.
OC-24	Uncommon.

course, there may be issues regarding how headends are logically configured. One data stream inserted at a master site will serve all hubs identically.

The situation is not as clear today concerning some of the non standard data rates used in cable television today. Various signals in use in headends today include 3.3 Mb/s and 5.33 Mb/s. These non standard rates were adopted for good reasons, but they cannot be accommodated with existing multiplexers. As cable television interests begin to use the standard multiplexing system, one assumes that equipment will become available to accommodate these data rates. One company has announced such. These data rates can be accommodated by adding extra bits to raise the rate to one of the standards that can be handled with off the shelf multiplexing equipment. For example, if a data stream operates at 5.33 Mb/s, an extra 982 Kb/s can be added to the data stream to produce a DS2 payload. That is, a DS2 data stream has a total data rate (payload plus sync added at the DS2 level) of 6.312 Mb/s. The difference between this and 5.33 Mb/s is the data rate that needs to be added, or "stuffed." The extra bits may be used for another data stream of less than 982 Kb/s rate, or if the data "bandwidth" is not needed, random data bits can be inserted.

WHAT TO EXPECT FROM DIGITAL TRANSMISSION

Arguably, this is off the subject a bit, but due to confusion we wish to briefly address the concept of quality of a digital signal in a cable television environment. This is an area subject to considerable mis-interpretation today. Figure 8 illustrates qualitatively the difference between digital and analog transmission. We plot noise level (roughly equivalent to carrier to noise ratio) vs. video quality, measured by any convenient metric. Analog transmission quality vs. noise level is shown as a heavier curve. It follows the well-known rule of gradual degradation with increasing noise. At very low noise levels the quality of the video is excellent. As the noise level increases, for a while a viewer cannot perceive the noise, so the picture quality appears to remain unchanged. At some noise level (perhaps 50 dB C/N) the picture quality begins to degrade with increasing noise. The degradation is gradual at first but as the noise gets worse the picture quality degrades until at some moderately arbitrary level we say the picture is unusable.

The lighter lines describe two different scenarios for digital video. One scenario is without error correction to the digital signal, and the other is for error correction. Error

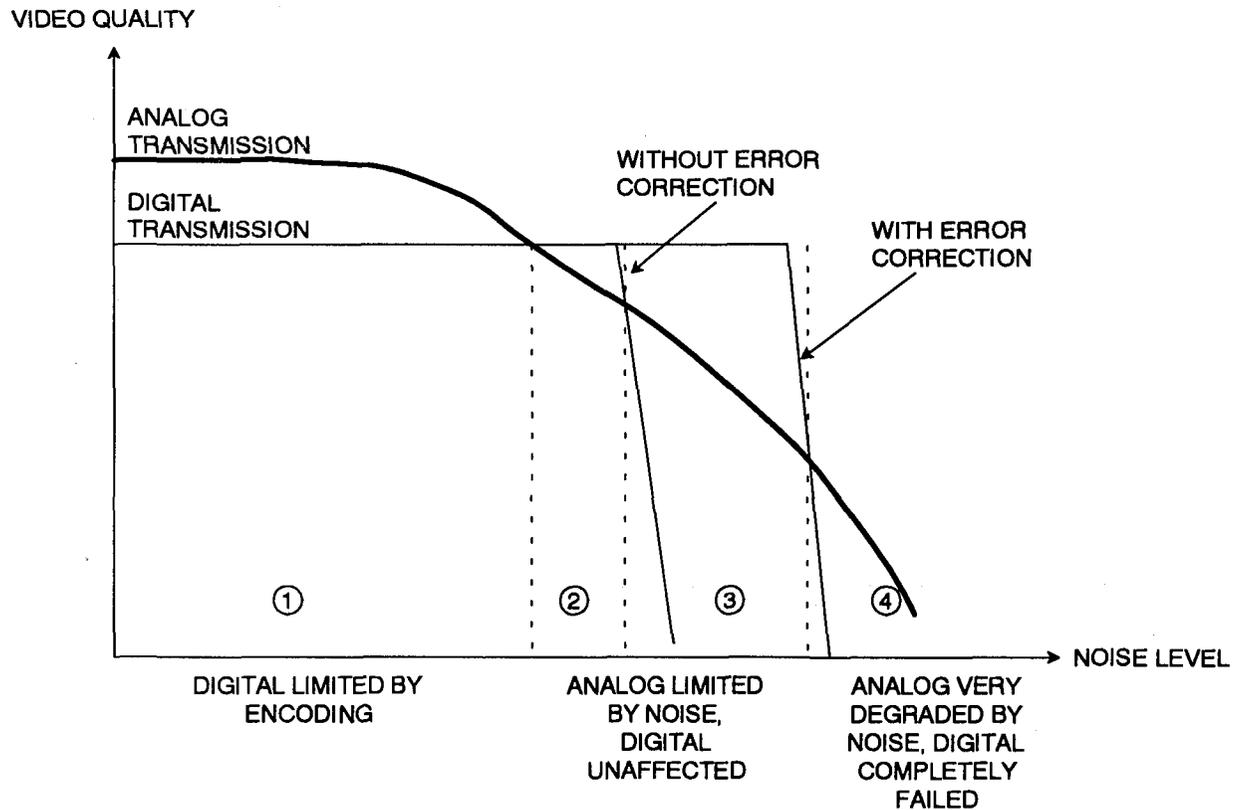


Figure 8. Degradation of Analog vs. Digital Transmission

correction is a way of transmitting extra bits with the digital signal, such that if some bits are received in error, they can be identified and corrected. Without error correction digital transmission would hold few if any advantages over analog transmission.

In order to compare analog with digital transmission, we have divided the graph into four regions identified by circled numbers. In region 1 the analog signal is better than the digital signal. This is because we have inevitably lost something when we digitized the signal, an added step. In some cases, one can argue that the digital picture is better than the analog, but this will only be true for non compressed signals that were never in NTSC format before being digitized. Compressed pictures will certainly be degraded over the corresponding analog signal, by definition (compression is a lossy process, in which some information is removed). The trick in compression is to remove information from the picture in such a

way that the removal is undetectable by the viewer. In region 1, both the analog and digital pictures would be considered excellent.

In region 2 noise is beginning to affect the analog picture noticeably. The digital picture, with or without error correction, is not affected by the noise, because the noise is not yet sufficient to cause errors in recovering the bits that make up the digital signal. Above a certain noise level the non error corrected signal degrades fairly quickly, going from "perfect" to unusable with just a few decibels increase in noise (region 3). The error corrected signal in region 3 is surviving with no noticeable degradation because error correction is masking the transmission errors. The analog signal is continuing to degrade significantly.

Finally in region 4 the noise level is so great that even the error corrected signal fails. When it fails, it does so very quickly, going from a near perfect picture to nothing with very

little increase in noise. The analog signal is very degraded but may remain recognizable.

The exact shape of the curve and the relative position of the analog and digital curves is a function of a lot of variables including the type of system through which the signals have passed, the type of digitization and compression used and the error correction used. If a digital transmission link is cascaded with an analog link, as in a regional interconnect done with SONET digital transmission between headends, followed by analog distribution, the signal to noise ratio of the digital link will roughly add to the carrier to noise ratio of the analog portion. The reason the addition is not exact is that the digital noise is composed exclusively of quantizing noise related to the number of bits, the compression algorithm used. This noise probably does not have the same spectral characteristic and peak to RMS value that thermal noise has.

Most certainly the signal to noise contribution of the digital link will NOT add to the carrier to noise contribution of the analog link. This can be seen from Figure 8, which shows the digital link signal to noise ratio (video quality) not changing as link noise increases in the way noise adds in the analog world. The signal to noise ratio measured at the originating point is the same as the signal to noise ratio at the receiving point.

We have been a bit sloppy in mixing baseband signal to noise and RF carrier to noise above. If the baseband signal to noise ratio is measured using the CCIR unified weighting network, using the proper definition of "signal," then it is numerically within a couple of tenths of a decibel of the RF carrier to noise ratio as defined by the NCTA. This is the way many instruments make signal to noise measurements today. However, it is a source of considerable confusion, and addition of noise on combined digital and analog links must be done very carefully.

ACKNOWLEDGMENTS

K. Lynch and R. Reynard contributed valuable research material. T. Engdahl contributed a vast amount of information concerning telephone practices. M. Dionne was very rough on the manuscript, but her suggestions improved the paper immeasurably.

END NOTES

ⁱThere are limits to what we can do with the signal: digital transmission is generally characterized by an all or nothing situation. If the signal is above a threshold, the signal out is just as good as the signal in. Below that threshold, the signal is lost completely.

ⁱⁱIn some transmission systems we would have two transitions per bit, but this is usually done when economy is at a higher premium than is bandwidth efficiency.

ⁱⁱⁱ*Webster's New World Dictionary of the American Language*, World Publishing Company.

^{iv}McGrath, C. J., *Digital Delivery Technology for CATV Networks*, Technical Papers of the SCTE Fiber Optics Conference, 1991.

^vIn telephone terminology, the voice transmission quality needed is called "toll quality." As described in the McGrath paper, achieving this quality at 64 KB/s required compression and expansion respectively at the transmit and receive ends of a circuit. Compare this rate for transmitting voice with the approximately 704 KB/s data rate to transmit one channel of CD quality audio!

^{vi}Perhaps the use of the term "D1" here is confusing: D1 stands for the *first digital* video recording standard to be used. It has nothing to do with the DS1 multiplexing level of telephone company usage.

^{vii}The corresponding European rate is 2.048 Mb/s, known as E1.