DEPLOYING TELEPHONY SERVICES OVER CATV SYSTEMS: SYSTEMS AND ARCHITECTURAL CONSIDERATIONS

Gaylord A. Hart, XEL Communications
17600 East Exposition Drive,
Aurora, CO 80017
(303) 369-7000

Richard Lyford, ANTEC
8101 East Prentice Avenue, Suite 210
Englewood, CO 80111
1-800-FIBER-ME

ABSTRACT

Today's cable systems are in a unique position to deliver telephony services via the two-way transmission of digital signals. Several methods exist for combining telephony and entertainment video signals within the distribution plant and for separating those signals at the subscriber end of the system. Each of these methods has unique technical and economic advantages and disadvantages.

This paper discusses the architectural and design considerations when developing a network capable of delivering telephony services. Architectural options are outlined, and the tradeoffs associated with design choices are discussed. Node sizing, optimization, and evolution are considered, as well as fiber counts, bandwidth utilization, reliability and redundancy, network management, and subscriber terminal deployment.

INTRODUCTION

The ability to deliver a multitude of new services to the home (both analog and digital; both interactive and non-interactive) via cable systems has become feasible in the last five years due to the rapid deployment of fiber optic technologies. Indeed, most experts now agree that hybrid fiber/coax networks provide the most economical means of delivering these new services. One such service is POTS, or Plain Old Telephone Service.

Deploying telephony over cable systems is only a short time away. Many manufacturers have announced cable telephony products or are already in production. Effective deployment of these systems requires a good understanding of spectrum utilization, network architectures, drop architectures, system evolution, and the tradeoffs surrounding decisions made today for tomorrow's services.

SPECTRUM UTILIZATION

CATV Spectrum Allocation

New technologies and services are bound to place new demands on future allocations of CATV spectrum, requiring much more bandwidth than has been available in the past, including a much larger or more densely utilized return path allocation. Services such as telephony, HDTV, data communications, video-on-demand (VOD), and multi-channel compressed NTSC video delivery are already staking out portions of the CATV spectrum, and as it usually turns out, services expand to consume the maximum amount of bandwidth available. Of course, new digital compression technologies will help operators efficiently utilize this spectrum.

Figure 1 shows a typical frequency allocation plan for a 1 GHz cable system. Below 550 MHz, the system looks exactly like current systems, with analog TV channels occupying 50-550 MHz and the return path occupying 5-30 MHz.

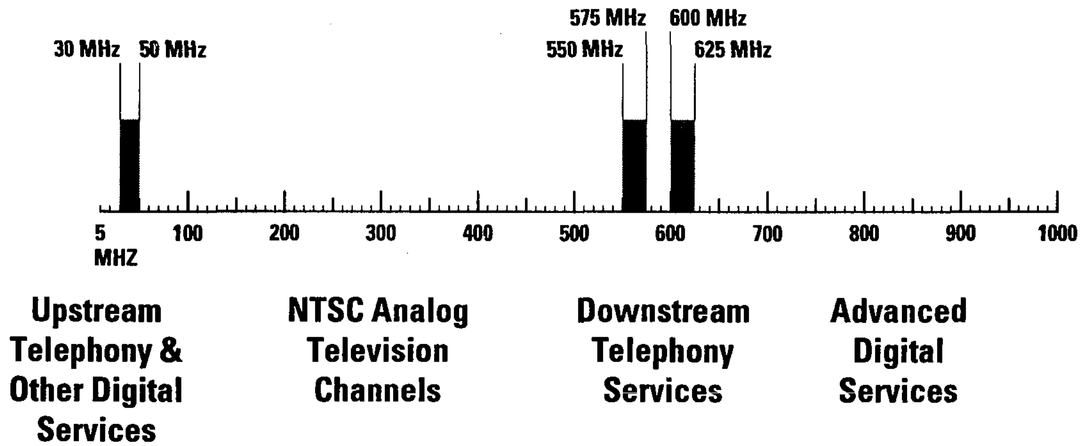

Figure 1. CATV Spectrum Allocation with Telephony and Advanced Digital Services

It is in the area above 550 MHz that the greatest changes are taking place. In this example, 25 MHz of bandwidth has been allocated for telephony in the reverse path, but in practice this may not be possible (or at least not without forward error correction or other means of maintaining a low bit error rate). System ingress originating in the subscriber's home wiring is not uncommon, and this ingress may come from several sources. It is also in the 5-30 MHz band where many amateur radio and CB transmitters are active in urban neighborhoods. In this example, another 25 MHz has been allocated for telephony in the forward path (575-600 MHz, with guard bands shown on either side; depending on system hardware, these guard bands may not be required). It is important to point out that the telephony signals in the forward path could be placed anywhere (even below 550 MHz), and that the choice of frequency location is left to the operator to maximize the particular application.

Typically, the spectrum above the telephony channels would be used for delivering advanced services such as compressed digital NTSC or VOD. As the deployment of interactive services and digital communications becomes ubiquitous, it is likely that the 5-30 MHz return path used today will not be adequate to support these new services.

Additional return spectrum can be gained by converting a system to a mid-split return (e.g., 5-112 MHz return, 150-1 GHz forward), but this is unlikely due to the analog spectrum currently in use for TV services and the vast number of consumer products supporting this spectrum. It is far more likely that additional return spectrum will be carved out of the higher end of the spectrum, likely from 850-1000 MHz.

This raises important questions about the 5-30 MHz spectrum. If additional return spectrum is provided at the top of the CATV spectrum, the 5-30 MHz spectrum may not be needed for return, thus eliminating the need for two sets of diplex filters for the two return paths. In this case, this spectrum could be allocated for forward use. Or perhaps more interesting, the 5-30 MHz spectrum could be utilized as a fully bi-directional path operating completely passive, without any intervening amplifiers. The cable losses at these frequencies are very low, and if the node sizes are small enough reliable transmission should be achievable. Such a passive path would be highly reliable since transmission could be accommodated even if power failed in the amplifiers. This has positive implications for services such as telephony and network management which need to be highly reliable.

Telephony Spectrum Allocation

Fortunately, only a relatively small allocation of bandwidth is required to deliver telephony services over cable systems. This is because not all frequencies discernible by the human ear (typically >15 kHz) need be transmitted for intelligible, natural conversation. In fact, the actual bandwidth occupied by a voice signal is limited to 4 kHz in telephony to conserve spectrum. Of course, many telephony channels have to be provided to maintain reliable service, and collectively these can occupy significant bandwidth.

Fortunately again, the coming of age of fiber technology in cable networks provides all the bandwidth necessary for adding telephony services to existing cable systems. This is accomplished through fiber division multiplexing several voice channels on several fibers. That is, for a given number of voice circuits, telephony bandwidth on an individual fiber is conserved by spreading out the voice circuits over multiple fibers, each fiber serving a different geographical area. The fiber counts and bandwidths required for economical telephony delivery on cable systems coincide very well with bandwidth requirements and fiber counts required for existing services, and in fact provide additional capabilities and leave room for contemplated new services.

Universally, digitized voice signals are created as 64 kb/s digital channels (commonly referred to as a DS0). These individual channels may then be multiplexed together in a number of ways and at varying bit rates. Higher rates, of course, support more voice channels. In North America, the most commonly used next order of multiplexing is the T1, which consists of 24 digitized voice channels (24 DS0's, for a total of 1.544 Mb/s, including framing overhead). Although many different transmission rates could be used for delivering telephony over cable systems, this is a logical rate supported by much existing hardware and allowing relatively flexible use of existing CATV spectrum and smooth migration as service penetration increases.

Channel overhead is also required in the bit stream for control purposes within the cable telephony delivery system itself (e.g., for remote provisioning of the Subscriber Terminal Unit). This overhead may come at the expense of capacity taken from the T1, or by using the T1's embedded extended data facility (Extended Super Frame format only), or by providing additional channel overhead by transmitting at a higher data rate. For the purpose of discussion, this paper assumes T1 transmission rates will be used.

Given a T1 transmission rate for the bit stream, many choices exist for modulating the T1 onto an RF carrier: FSK, BPSK, QPSK, QAM, and others. The choice is typically a tradeoff between bit rate vs. bandwidth used and circuit complexity. Thus the modulation scheme chosen directly effects the overall system traffic capacity. QPSK presents a good balance here (providing a transmission efficiency of two bits per Hertz with relatively inexpensive circuitry and good noise performance), and many proposed cable telephony systems employ QPSK. For the purposes of discussion, this paper assumes QPSK data transmission in both the forward and return transmission paths.

Given the same data rates and modulation methods, the transmission bandwidth required for the system will be the same for upstream and downstream. The upstream path is the limiting factor, providing 25 MHz of usable bandwidth from 5-30 MHz (there is currently active interest in extending this to 40 MHz, thus providing 35 MHz of usable bandwidth). Referring to Figure 2, 25 MHz of bandwidth allows 24 QPSK signals to be transmitted, each carrying one T1 (1.544 Mb/s divided by 2 b/Hz = .772 MHz bandwidth; allowing 30 percent more for adjacent channel guard bands gives 1 MHz of bandwidth required for actual transmission). Given that each T1 supports 24 voice channels, 576 voice circuits are available in 25 MHz.

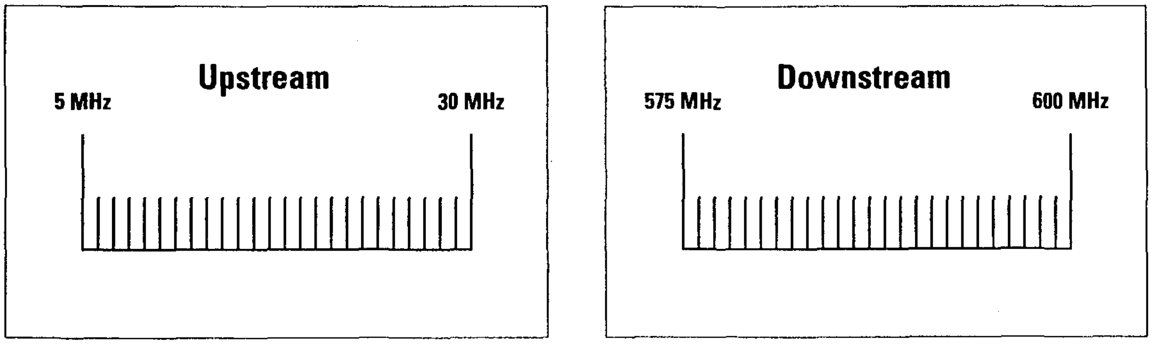

24 QPSK Signals
= 576 Voice Channels (DS0s)

Each QPSK Signal
= 1 T1
= 24 Voice Channels
= 1 MHz Bandwidth

24 QPSK Signals
= 576 Voice Channels (DS0s)

Each QPSK Signal
= 1 T1
= 24 Voice Channels
= 1 MHz Bandwidth

Figure 2. Typical Telephony Spectrum Allocation

One great advantage of transmitting telephony over several different RF data channels is the great flexibility provided. RF channels need only be added as telephony service penetrations increase (capital expenditures are tied directly to revenue generating services). RF channels may be flexibly assigned frequencies around existing or contemplated services (particularly important in the return path where service is often shared with network management carriers and PPV authorization channels utilized by addressable converters). Finally, frequency agility of the return and forward paths allows efficient use of CATV spectrum and the ability to shift spectrum usage in the future as new services and technologies are introduced.

Contention and Service Penetration

A discussion of spectrum utilization would not be adequate without some consideration of contention and penetration. Two basic approaches exist for handling telephony traffic over cable systems: dedicated circuits and contention based channel assignment. A dedicated voice circuit may be assigned to a subscriber (actually a virtual circuit, consisting of an RF channel assignment and a time slot within the T1 carried on that channel). Whenever that subscriber uses his phone, that carrier and time slot are accessed. When that subscriber is not using the phone, those resources are idle and cannot be used by other subscribers. A voice channel must be dedicated for each telephony subscriber, thus potentially requiring a large number of dedicated voice channels. For spectrum efficient telephony systems, this is not a problem. Dedicated channel assignment has the advantage that a subscriber always has guaranteed access to his voice channel and the channel assignment process is simplified.

Contention based channel assignment leaves idle voice circuits free in a pool. When service is requested either by an incoming call or by a subscriber trying to place an outgoing call, a free channel is taken from the pool and assigned to a subscriber for the duration of the call. When the call is finished, the channel is returned to the pool for use by other subscribers. Since it is highly unlikely that all subscribers would attempt to make (or would receive) calls simultaneously, this approach allows a finite number of voice circuits to be used by a much larger number of subscribers, thus increasing hardware utilization and system efficiency.

When a subscriber attempts to make a call and no voice circuits are available, no call connection can be made, and this system state is defined as blocking. An effective statistical model (the Erlang B formula for lost-calls-cleared) exists for calculating channel capacity vs. traffic vs. probability of blocking. Any use of contention requires careful analysis to guarantee a minimum grade of telephony service. Contention allows a smaller portion of RF spectrum in the cable system to be used to serve an equivalent number of subscribers as with dedicated channel assignment, but the tradeoff is in potential blocking and customer dissatisfaction during peak traffic periods.

Contention based channel assignment also has another great advantage: it allows dynamic bandwidth allocation on a per subscriber basis. Services such as videophone and high speed data transmission require bandwidth in excess of that provided in the basic voice channel (4 kHz or 64 kb/s). Contention allows assignment of additional capacity to a subscriber on a demand basis, typically in multiples of the basic data rate of 64 kb/s. A subscriber desiring to use one of these advanced services would be given channel capacity for that service only for the duration of the connection, thus saving the subscriber money while making more efficient use of the system's total traffic capacity.

Networks must be designed to meet today's needs and those contemplated in the future. However, it is not always appropriate to fully build out a network today to provide services which may not actually be used until quite some time in the future. It is likely that the number of cable subscribers may be quite high compared to the initial number of telephony subscribers on a cable system. Networks should be designed and built with this in mind, and a network evolution plan should be devised to allow natural growth in traffic capacity as penetration increases.

DISTRIBUTION ARCHITECTURE

Current FSA Network Topologies

Current cable systems utilize a fiber to the serving area architecture (FSA), with each optical receiver typically serving a node of 500-2000 subscribers. This architecture is shown in Figure 3.

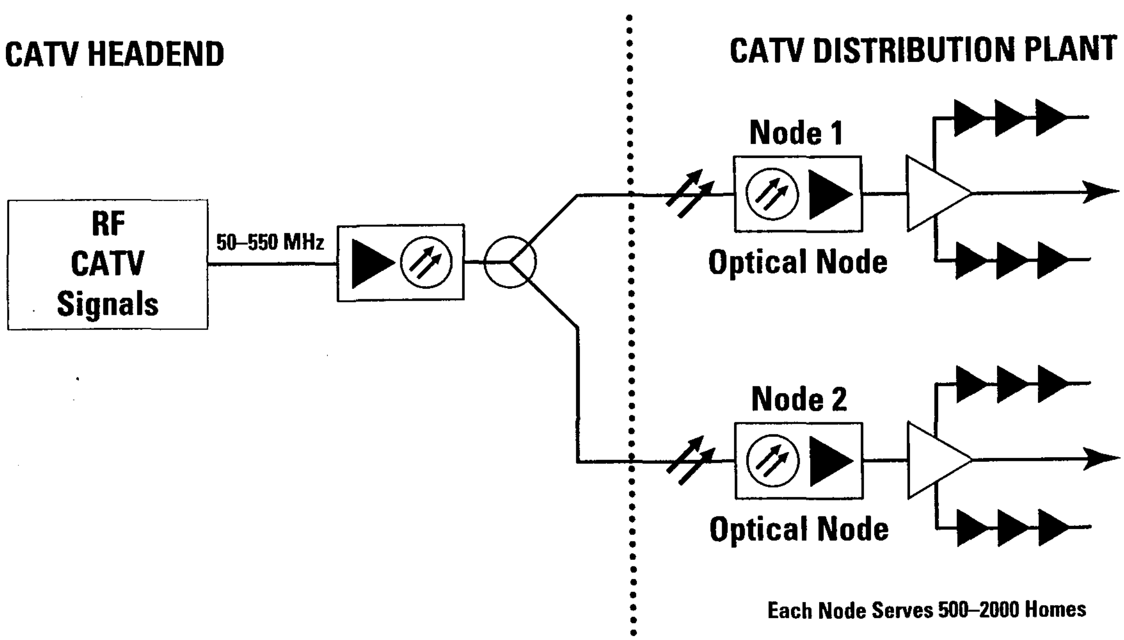

Figure 3. Current FSA Architecture

One laser in the headend may be optically split to serve two or more nodes since the television services provided are the same for each node. This is a cost-effective approach since relatively high output power DFB lasers are available, and only one laser need be used for two or more nodes. This approach is also effective in reducing power consumption, improving distortion and noise performance, increasing reliability, and reducing signal ingress and leakage. The coaxial trunk amplifiers have been eliminated, and the size of the failure group has been greatly reduced. This architecture is very flexible for overlaying additional services and allowing cost-effective system evolution.

FSA Network Topology with Telephony Overlay

Figure 4 shows the same FSA architecture with a fiber telephony overlay. In this case, the two receive nodes are treated as one logical node for telephony services. Note that this approach uses the existing downstream CATV service laser to provide downstream telephony services as well.

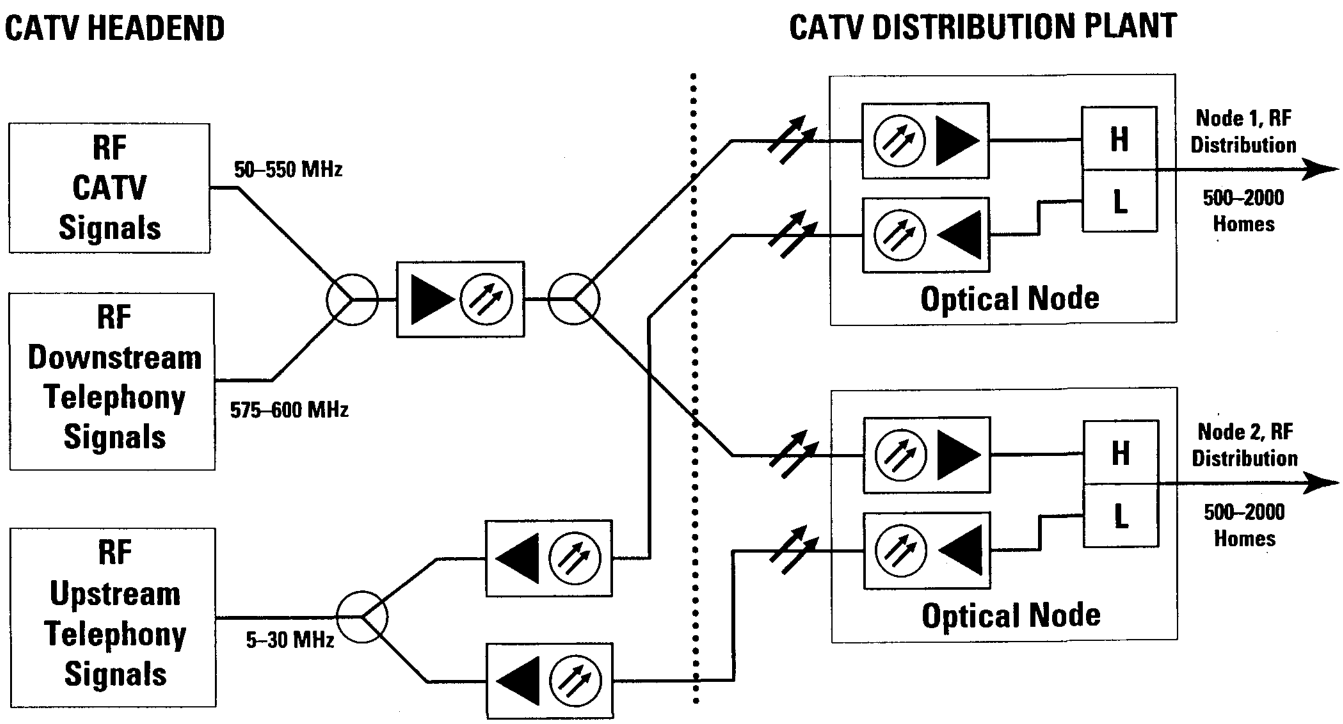

Figure 4. Current FSA Architecture with Telephony Overlay

Providing enough bandwidth is available, the downstream architecture need not be modified at all for this upgrade. The upgrade is facilitated even more since no new fiber need be pulled: dark fibers probably already exist in the cables already feeding the remote receiver nodes. If you are currently in the process of adding fiber to any of your systems, it is wise to plan now for future applications by providing additional fibers in the sheath.

The return path upgrade is readily accomplished in the distribution plant by swapping the original optical node receivers with receivers with return path lasers for the 5-30 MHz band. Many receivers on the market or already installed already provide plug-ins for retro-fitting lasers in receivers already in the field. The only additional equipment required for signal transport consists of two relatively inexpensive optical receivers in the headend for the return path.

Since this topology treats two optical nodes as one telephony node, migration to this configuration is easily and economically accomplished. However, each optical node is now served by only half as many voice channels as are available on from the telephony node at the headend. Using the spectral model shown in Figure 2, 576 voice channels would be available for the two nodes (1,000 to 4,000 homes passed). For initial service offering, telephony penetration may be low, and this architecture can be a cost-effective means of matching capital expenditures with revenue.

When service penetration increases, or if high penetration is anticipated quickly, the two CATV service nodes may be treated as two independent telephony nodes by separating the telephony services for each on different fibers, thus doubling the traffic capacity on each node. Of course, additional telephony hardware must be provided in the headend to support this increased capacity. Using the spectral model shown in Figure 2, 576 voice channels would now be available for each node (500 to 2,000 homes passed).

Figure 5 shows this architecture, which allows yet another cost-effective step in the natural evolution of the network as additional channel capacity is required.

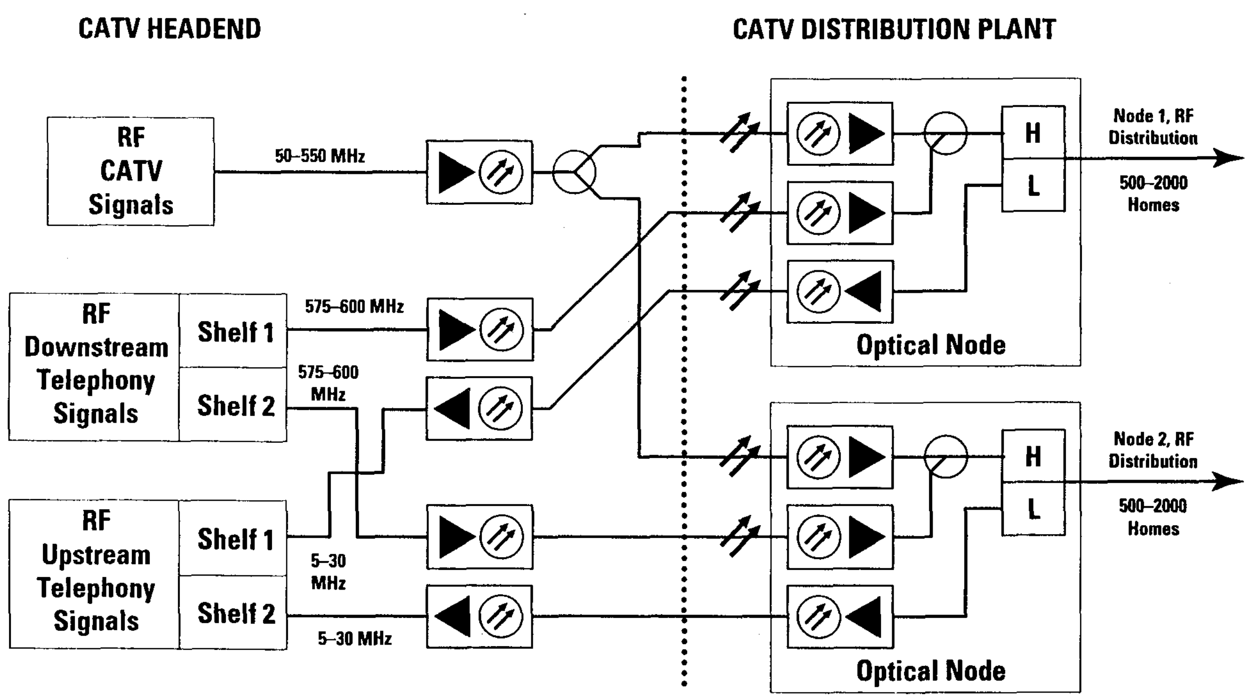

Figure 5. Current FSA Architecture with High Penetration Telephony Overlay

Note that for economic reasons only one DFB laser is used to feed television services to the two nodes. This does necessitate, however, the use of separate downstream and upstream telephony fibers for each node. It should be possible for this application, once again, to use dark fibers already in place. This topology also requires one relatively low-cost laser per node at the headend for downstream telephony and an additional optical receiver at the node for combining the downstream television and telephony services for distribution over the coax. Node receivers such as shown in Figure 5 are already available on the market.

Advanced Services FSA Network Topology

Figure 6 shows an advanced FSA architecture capable of delivering analog television, telephony, and advanced digital services to the node. The bandwidth here need not necessarily be 1 GHz if the advanced digital services are offered over a smaller spectrum or are not included at all: this architecture will work just as well at 600 MHz or 750 MHz (750 MHz amplifier hybrids are currently available, though in short supply; 1 GHz hybrids with adequate performance are still not commercially available).

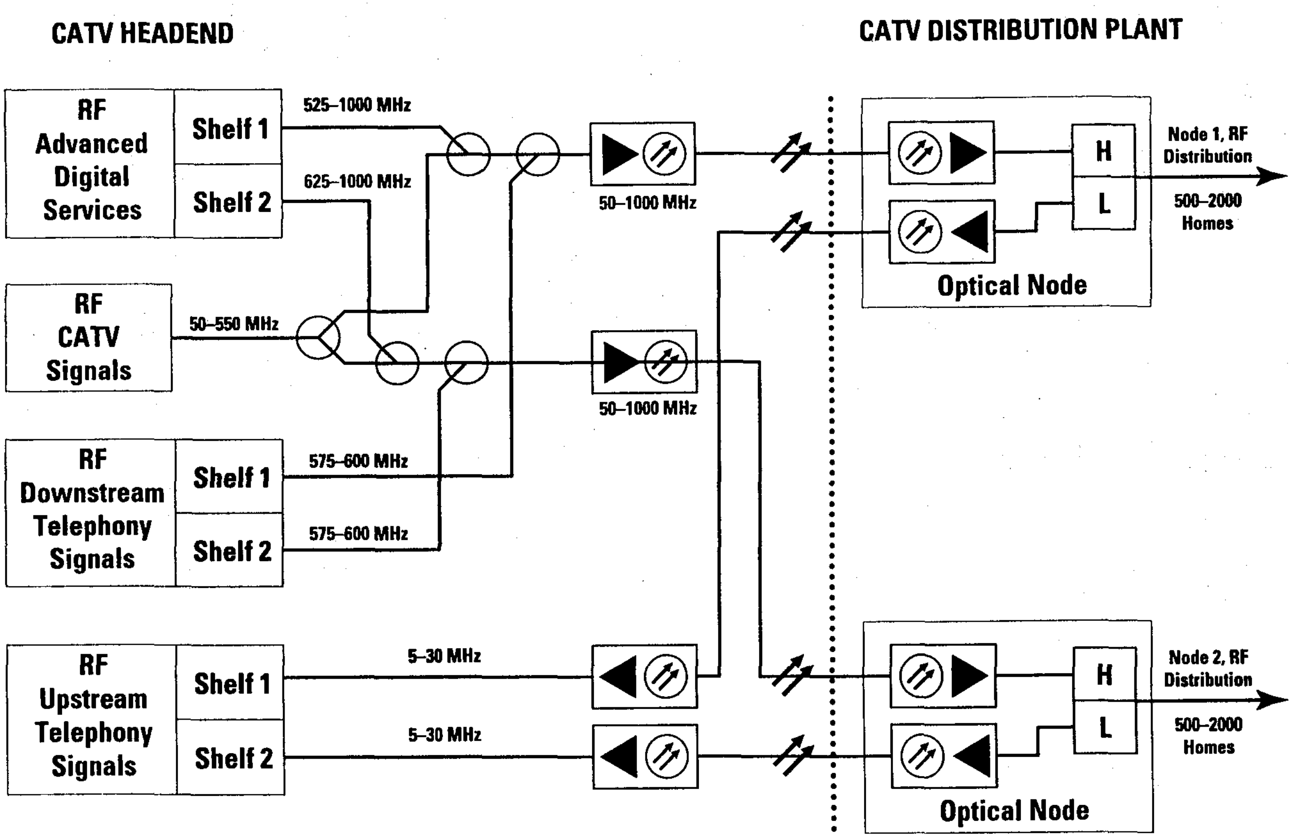

Figure 6. Advanced FSA Architecture with Digital and Telephony Services

Note that two downstream DFB lasers are now used, one for each node. Many envisioned advanced digital services (e.g., VOD) require unique data streams to be provided for each node, and this requires separate downstream lasers for each node (of course, these services could be combined with telephony on a separate laser per node as in Figure 5, providing this laser has adequate capacity).

Although DFB lasers are far more expensive than conventional Fabry-Perot lasers used for transmitting RF data carriers, using one DFB per node allows simplifying the remaining portion of the distribution architecture since these lasers can carry the telephony and advanced digital services as well, and all on one fiber. Of course, this means only one downstream receiver is required per node, as well. Comparing Figure 6 with Figure 5, this latter approach may make more economic sense from the outset (even if advanced digital services are not provided), certainly so if extra downstream fibers are not already present as required in figure 5.

Since the DFB lasers in Figure 6 are not being optically split as in previous examples, lower power lasers will suffice to cover the same distance to the node. Although the trend in the past has been toward higher power lasers to allow greater optical splitting, more systems will begin to require lower power lasers (but many more of them) as advanced architectures are deployed. Hopefully the increased volume and lower power requirements will drive down the cost of these relatively expensive DFB lasers. Long-term, the architecture in Figure 6 should be the more economical approach to delivering advanced services of all types in a hybrid fiber/coax system.

Network Reliability

Since the telephone serves as the fastest, best access a subscriber has to emergency services and information, any system delivering telephony must be highly reliable. In terms of the cable distribution system, many new demands will be placed upon the early detection and correction of potential faults (preferably before an actual failure occurs) and in locating failures and repairing them quickly when they do occur. This will require the deployment of advanced network monitoring and management systems.

However, appropriate monitoring and management systems are not enough to guarantee reliable service. For this reason, all powered devices in the network must have some form of emergency power for operation when utility power is not present. This power can be provided through battery back-up power supplies or gas-powered generators. In any case, switchover to backup power must occur quickly enough to guarantee no interruptions in service take place.

Finally, the distribution system itself must be made more reliable. In part, this can be accomplished through the use of better trade practices and higher reliability products in the construction of the system. Additional reliability can also be gained by providing other backup systems. In the case of fiber optic runs, duplicate fibers using route diversity should be used in case of an accidental fiber cut.

The active devices in the plant must also contain backup circuitry. All amplifiers and optical receivers must utilize circuitry which allows a failed component to be electronically bypassed or replaced by a functional equivalent. This will require all receiver and amplifier modules to be duplicated in each equipment housing, and an intelligent means of switching between these modules must be provided. Furthermore, this intelligent switching device must be tied into the overall network management and monitoring system so that any failure can be reported directly back to the office.

Fully redundant active devices must be deployed in a modular fashion. This allows failed modules to be replaced without bringing down the system, but it also allows optional initial installation of the active devices without the redundancy feature. This allows an operator to plan economically today for easy migration to full backup capability once telephony services are deployed. The backup modules should be operated as cold standbys (i.e., without power applied). This serves three purposes: power consumption is reduced, heat generation is decreased, and standby modules are less likely to be damaged by any surges which may be presented to the power supply. This increases module reliability while keeping operating costs lower.

SUBSCRIBER DROP ARCHITECTURE

Telephony Subscriber Terminal Inside Home

The subscriber drop architecture is critical for economical deployment of telephony services over a cable system. Costs added to maintenance, installation, or hardware at the drop are multiplied by the number of drops in the system. Several options exist for location, powering, and deployment of the Subscriber Terminal Unit (STU), and each of these options has unique advantages and disadvantages. However, regardless of whether the unit is inside the home or on the side of the home, the STU need only be installed when service is initiated. This ties capital expenditures to revenue generation and allows the operator to pay as he goes. A logical choice is to place the STU inside the subscriber's home (Figure 7).

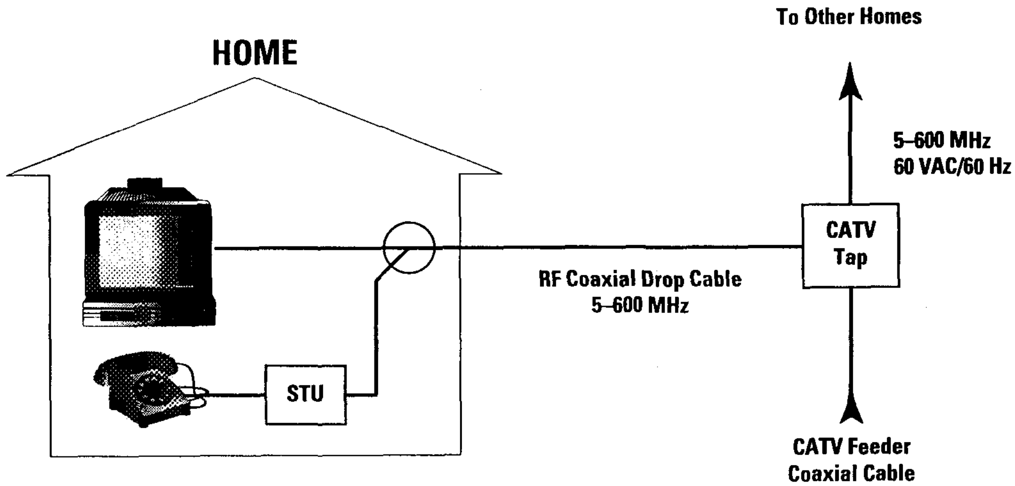

Figure 7. Telephony Subscriber Terminal Inside Home

The cable industry has a long history of placing cable television converters in the home, and many customers are used to their presence. Objections, when raised by customers, tend to focus not on the converter, but on the loss of functionality of consumer products attached to the converter. For an STU, no loss of functions are encountered since the subscriber's telephone equipment works as it always has. And since no interaction is required between the subscriber and the STU, the unit may be placed in a utility closet or any convenient location. Connection to the cable system may readily be accomplished with a directional coupler placed at any convenient point along the internal house wiring or even outside the home.

One strong advantage to placing the STU in the home is increased reliability and lower unit cost. The home provides a well-protected environment, free from temperature extremes and precipitation. An STU designed for in-home installation need not be temperature hardened or environmentally sealed.

Another logical choice in this situation is to power the STU from the home, and once again, the precedent exists for powering set-top converters from the home. The advantage here is that the subscriber absorbs the operational costs for powering the STU. Although an individual STU does not draw very much power by itself , the actual operator costs for powering the STU would not be trivial when one considers that thousands of the units will be operational.

Power failures also present a larger problem when delivering telephony service. When power fails in the home and the cable TV converter loses power, little is lost since the TV set has also lost power. But in the case of telephony, a power failure causing loss of telephone service could be catastrophic since access to emergency services is lost. Therefore, home-powered STU's have to provide some form of battery backup to maintain service during power outages. How long this battery needs to maintain service during a power failure is debatable, but most power failures tend to be relatively short in duration, typically well under four hours. Batteries also need periodic replacement, but modern sealed batteries have lifetimes of several years.

It is also possible to power an in-home unit from the network. However, network powering of the STU does not solve the battery problem, but simply moves the location of the batteries from the STU out into the network since the network itself now must provide battery backup for all the STU's. Fewer, but larger batteries will be required, but maintenance of the batteries should be easier since they will be concentrated at and co-located with the system's standby power supplies, and 24 hour access will be available. During extended power outages, gas-powered generators could be used to extend system operation beyond the batteries' capability.

Network powering, however, does create additional installation problems if power is carried on a separate cable from the drop and adds safety, regulatory, and reliability concerns if powered via the drop. If powered via the drop, power-passing passives must be used in the STU signal path, and care must be taken to block power down any other signal paths in the home. A step-down transformer would also likely be required to provide safer voltages in the home. Under these circumstances, any modifications the subscriber performs to his in-home wiring would likely generate a service call. For these reasons, powering an in-home STU from the network via the drop is not a viable option.

Telephony Subscriber Terminal on Side of Home

Another likely location to mount the STU is on the side of the home. This provides easy access should maintenance be required, but as indicated above, environmental hardening and sealing are now required. Standard twisted pair is used to feed the telephone service from the outside STU to the internal home telephony wiring. Locating the STU on the side of the home also makes it easy to clearly define the network termination point and where home wiring starts.

More powering options exist for a unit mounted on the side of the home. Powering is still possible from the home itself, but in this case a low voltage would be fed out to the STU from a plug-in wall transformer inside the home. This power could be routed to the STU either via a separate small-conductor power cable or by reverse feeding power to the STU up the coaxial cable providing RF to the TV. Of course, powering over any RF coaxial cable requires using power-passing and power-blocking passives where applicable. If powered from the home, separate cabling is the preferred method. But regardless of cabling, a backup battery must be provided with the STU. In this case the battery is always accessible to the operator should maintenance be necessary.

Powering may also be provided from the network via the coaxial drop (Figure 8), but in this case power may be blocked before it enters the subscriber's home, thus avoiding many of the safety and regulatory issues associated with providing 60 VAC down the drop. However, this introduces new problems, primarily at the tap.

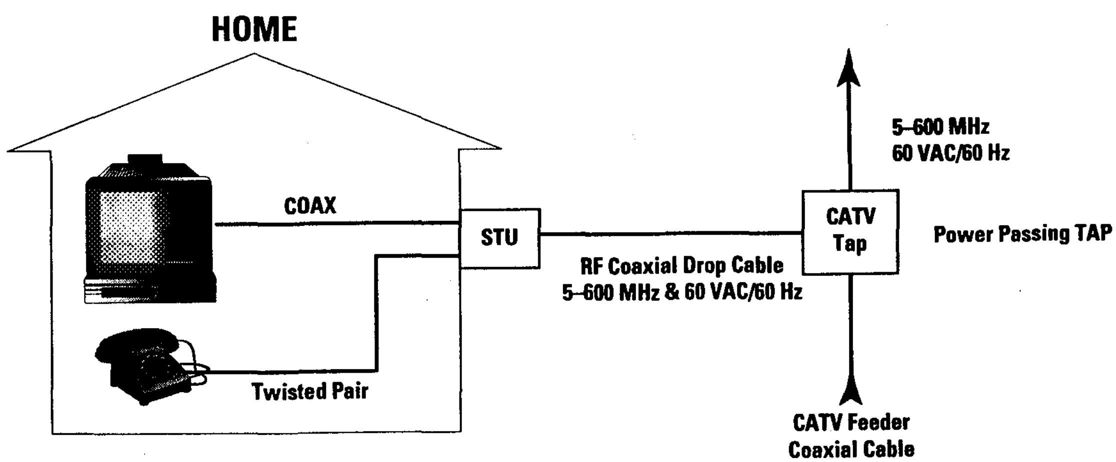

Figure 8. Telephony Subscriber Terminal Outside Home—System Powered on Drop

As new services are added to cable systems, more bandwidth will be required, and this requires higher performance components in the system. Taps are no exception to this, and although 1 GHz taps are now generally the standard when upgrades or rebuilds are performed, designing a tap to pass power down the drop without sacrificing RF performance (primarily in insertion loss and return loss) is next to impossible since power diplex filters must now be added to the main input/output ports as well as to each subscriber port. Decreased insertion loss performance can lead to frequency response problems or may require additional amplifiers to compensate for higher losses (or alternatively, fewer homes could be served per node). Poor return loss performance can potentially lead to intersymbol interference in digital signals.

Additional problems are likely to exist if powering is provided over the cable drop cable. First of all, the drop is required to be connected to the house electrical system ground near the point of entry into the home. The cable system itself is required to be bonded frequently to the power grid ground as well in the feeder system. Because of this, it is not uncommon to see sheath currents on the outer conductor of the coaxial drop cable. These currents can serve to add or subtract, depending on phase, from the power being provided over the drop to the STU. Additionally, several conductor interfaces exist in the drop, none of which have been optimized for power passing. First of all, the F-connector center conductor seizure mechanisms on most devices, while adequate for maintaining RF conductivity, do not provide much surface contact area or surface pressure to maintain good ohmic contact for power applications.

Next, several dissimilar metals are used in these interfaces which may cause corrosion under power-passing conditions. The outer braid and foil of the cable itself are aluminum, and as with aluminum house wiring, are subject to rapid oxidation and increased contact resistance. The F-connector itself is typically brass. The F-port on the tap is also typically brass, but is commonly nickel-plated. The center conductor of the drop cable is typically copper-clad steel, while the F-port seizure is typically tin-plated beryllium-copper.

Finally, moisture ingress at any of these connections will be even more critical to control in power-passing applications. Any decision to provide power via the drop cable will require careful analysis of all these factors if reliability is to be maintained.

Most of these problems can be avoided if network powering is delivered to the home via separate power conductors joined in a Siamese cable with the coaxial cable. This configuration is shown in Figure 9.

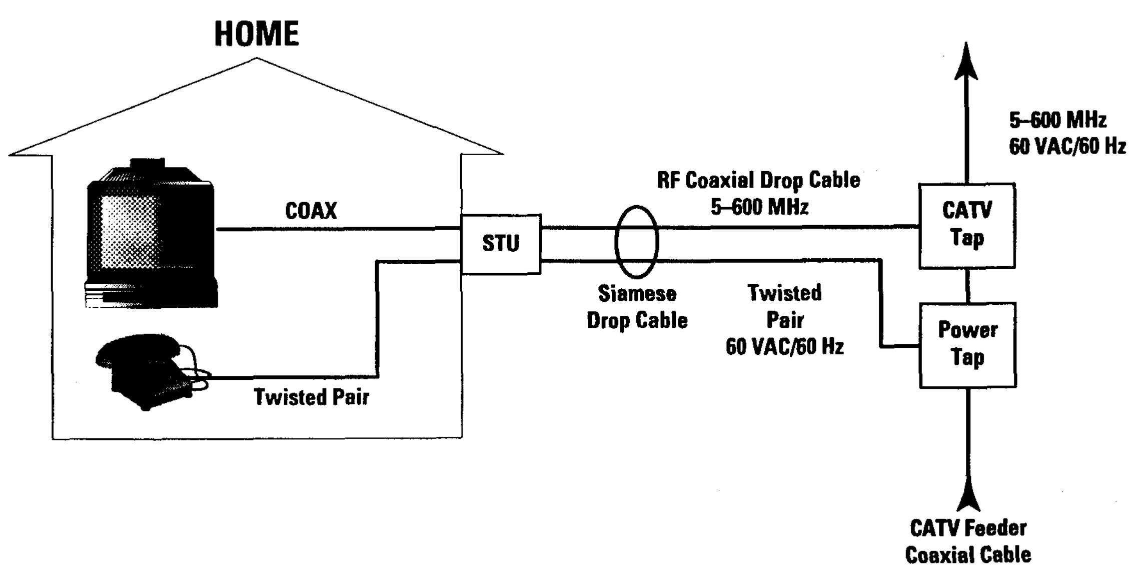

Figure 9. Telephony Subscriber Terminal Outside Home—System Powered on Siamese Cable

In this case power is extracted from the network via a power tap, which functions exactly as a power inserter does in today's systems to inject power into the cable system. This power tap need not be separate from the RF tap, and it is likely that taps will be available shortly which will include this function. Additionally, different regulations may apply regarding the maximum voltages which can be used if powering is not over the drop. If higher voltages can be used, power transmission efficiency can be improved through lower I^2R losses..

Telephony Subscriber Terminal at Curb Side

One other configuration exists for deploying the STU, but in this case the STU is located at the tap and supports multiple subscribers (typically 4, 8, or 16, though any reasonable number could be provided). Th is configuration is shown in Figure 10. In this case, telephony is served to the home over a Siamese coaxial/twisted pair cable.

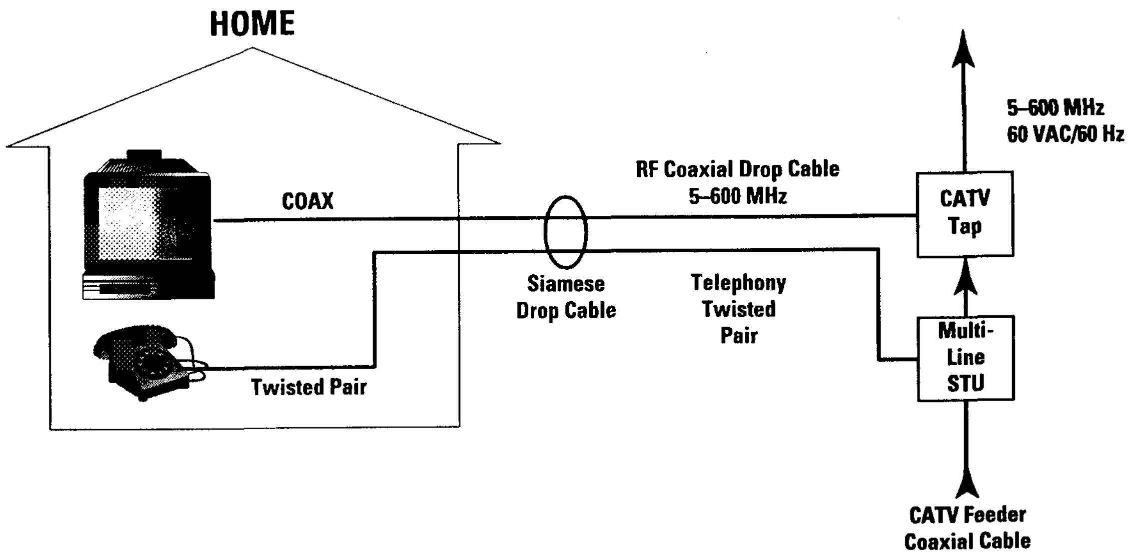

Figure 10. Telephony Subscriber Terminal at Curbside—Telephony on Siamese Cable

This approach has several advantages. First, reliable network powering is easily accomplished since this device is similar to a tap: it has an input and output directly connected to the feeder. Second, from a subscriber perspective this service appears exactly the same as his existing service since no unit need be installed on the subscriber's premises. Third, since multiple subscribers are served from one unit, significant hardware savings can be realized by eliminating the duplication of circuits required at each home when single subscriber units are used. For example, only one power supply is now required to serve several telephone subscribers whereas before one was required in each single-subscriber STU. Fourth, eliminating all these redundant circuits should lead to greatly reduced power consumption, thus lowering a system's powering costs. Last, those circuits which must still be unique to each telephone subscriber can be manufactured as plug-in modules and added to the STU only as new subscribers are signed up .

SUMMARY

Many advanced digital services will soon be deployed on hybrid fiber/coax cable systems. At the moment, this is the only system capable of delivering such services cost-effectively. Telephony will be one of the first services deployed. Many options and architectures exist for providing the signals needed to support telephony. Careful consideration must be made now to make sure both the distribution system and the subscriber drop are designed properly to be ready for this service and to be ready to evolve as more services come on-line.