Unified Headend Technical Management Of Digital Services

Ran Oz and Amir Bassan-Eskenazi
BigBand Networks, Inc.
David Large
Media Connections Group

Abstract

Historically, cable system designs were based on a one-way broadcast model whereby all signals were sent throughout the network. Even though operators are now introducing digital services and subscriber-specific information streams, they are forced by this broadcast legacy to make fixed allocations of expensive equipment and bandwidth to each new service. This inefficient methodology fails to leverage the inherent advantages of cable's two-way hybrid fiber/coaxial (HFC) networks relative to the networks of competitive broadband service providers.

The authors argue that cable's multiple system operators (MSOs) face an imperative to convert their service delivery architecture from its hard-wired legacy to a truly switched and routed fabric in which resources and bandwidth are dynamically assigned to sessions, of video, data, voice, or any multimedia form, in response to real-time subscriber demand.

The new headend element required to achieve this is a new class of product, the Broadband Multimedia Router, through which all digital services will be controlled.

THE COMPETITIVE EVOLUTION TOWARDS MERGED APPLICATIONS

Homes and businesses access many services from networked connections: telephone, video, Internet access, etc. Traditionally, those services have been delivered by separate suppliers over largely-separate networks.

It is increasingly clear that these networks are converging to carry the full variety of services. For subscribers, this has advantages such as streamlining service provider relationships and the potential availability of new hybrid services that combine various media forms.

Network operators welcome the opportunity to maximize revenues by offering multiple services and to achieve cost efficiencies by combining services into singular, well-managed conduits. In cable there is clear optimism about this future, as indicated by the soaring values of cable systems.

Cable television has been moving aggressively by offering both innovative new video services and non-video services, such as high-speed data and telephony. Video-on-demand (VOD) trials are being widely conducted with a new sense of optimism; millions of people use cable modems for broadband Internet access; and MSOs who offer telephony are achieving very meaningful penetration levels.

The advantages of convergence have not been overlooked by cable's competitors. Direct broadcast satellite (DBS), traditional telcos and fixed wireless operators (multichannel and local multipoint distribution service, or MMDS and LMDS) are all positioning themselves to leverage their network infrastructures to deliver non-traditional services.

The major DBS operators now offer enhanced services including integrated Web access, email and digital video recorder (DVR). While their capabilities do not match cable's, they can provide some VOD-like features, Internet-on-TV display and direct links between television ads and related Web sites.

Major incumbent telephone companies are investing billions of dollars in their plants to provide digital subscriber line (DSL) service at speeds of several megabits per second, capable of carrying high-quality video.[1] Competitive DSL data carriers are now starting to announce video initiatives of their own. These companies want to be full-service suppliers to consumers, and strive to present their offerings as being functionally equivalent or superior to those offered by cable systems.

Finally, MMDS/LMDS companies, which have been only marginally successful in delivery of analog video, are taking advantage of technological developments from some of the most prominent global technology companies in their plans to offer high-speed data.[2] While fixed wireless networks are particularly challenged, from engineering and cost perspectives, in their efforts to achieve broadband speeds, they do hold the enticing promise of enabling mobility in access to converged services.

Clearly, this is a highly competitive and changing marketplace. Each type of network brings its own inherent advantages and disadvantages to this competition. It is incumbent upon MSOs to leverage their HFC advantages in order to compete and prosper.

CABLE'S LEGACY ARCHITECTURE

Historically, cable system architecture was almost exclusively coaxial tree-and-branch, whereby all subscribers were fed from one or a few trunk lines leaving the headend. All services were transmitted in common throughout the network, with various access control methods utilized to limit optional service reception to selected customers.

This network was optimized for the one-way delivery of analog video channels and was constructed as economically as possible for that purpose. For many years the primary innovations in the network were limited to RF bandwidth expansion, reach extension, and access control methods.

While the network was a model of cost-effective engineering and performed its task well as the cable market grew, it was not suited to certain types of new services. In particular:

- There was no effective way to subdivide the subscribers into smaller groups (required to support any significant level of subscriber-specific communications).
- Nearly the entire spectrum was consumed by bandwidth-intensive analog video signals.
- The reliability was limited by the long equipment cascades.
- Two-way communications were limited by both cascade lengths and the large number of network branches.

UPGRADE TO HFC

At its most basic level, a hybrid fiber/coaxial network is one in which optical fibers carry the signals part of the way from a central location to subscribers' terminal equipment. At the local distribution level, it usually means that a portion of the former coaxial trunk network is replaced by a functionally-equivalent linear fiber link, leaving some residual coaxial distribution equipment between each fiber node and subscribers. Depending on the implementation, there may also be higher optical network layers which may carry either analog or digital signals between a headend and intermediate hubs.

The introduction of HFC technology enabled cost-effective bandwidth upgrades and improved reliability. Additionally, it provided the theoretical ability to feed different signals to each of the various segments of the network by re-use of channels, a technique sometimes known as space division multiplexing (SDM).

HFC architectures also make the operation of high-speed, two-way communications practical through reduced sizes, in terms of both cascade depth and number of subscribers, of the individual coaxial distribution sub-networks.

INTRODUCTION OF DIGITAL SERVICES AND STANDARDIZATION OF DIGITAL TRANSPORT

The latest round of downstream spectrum expansion has differed from previous upgrades in that the upper portion of the spectrum is usually reserved exclusively for carriage of digital signals. MSOs are making business commitments to use these assets to tap new revenue streams through introduction of advanced services. This is the first step in the eventual, inevitable, analog-to-digital transition for all cable-delivered signals.

Digital carriage over cable is made possible by several fundamental advances. MPEG compression enables 8-12 video streams to share a 6 MHz channel formerly occupied by a single analog television signal. Cable modem equipment utilizing 256QAM modulation can achieve up to 38 Mb/s gross data rates in a single channel, with the usable data rate divided among customers in each group, where groups may evolve from several combined nodes to parts of a node, as demand increases. Standards such as CableLabs' PacketCable and DOCSIS1.1 facilitate voice-over-IP (VoIP) telephony in which voice traffic can efficiently share bandwidth with data services.

To the credit of the various standards bodies, digital video, data and telephony all use the same MPEG transport protocol, the same RF bandwidth and the same QAM modulation. Nothing precludes these services from being intermixed in the same channels. For example, the DOCSIS interface specification includes a "Downstream Transmission Convergence Sub-layer" which allows multiplexing of video and data MPEG packets. Header information identifies to receiving devices the content of each packet to facilitate multiplexed streams.

Cable operators have developed systems which have a potentially unbeatable combination of characteristics, if properly utilized:

- A theoretical downstream information capacity of several Gb/s
- A two-way wide-band communications path between headends and customers
- The ability to divide customers into small groups
- A common transport protocol for all digital services

REMAINING CHALLENGES

Unfortunately, cable systems currently fail to take full advantage of their potential. They contend with inefficient utilization of precious bandwidth and high capital cost for each service, which is often a barrier to introduction of new services.

The bandwidth and associated equipment, such as modulators and upconverters, for each digital service are typically allocated based on the number of subscribers in a sharing group, the expected service penetration, and the simultaneous usage rates among subscribers. These calculations are typically made for each service, with the assigned bandwidth based on predicted busiest-hour loadings. The bandwidth and resources may be highly utilized, or even insufficient, during periods of high demand, but are grossly under-utilized when demand is low.

Traditional approaches are inefficient on a number of grounds. Peak usage times for various video and non-video services do not occur simultaneously, and thus resource scarcity in one part of the spectrum is likely to coincide with overabundance on another. For example, watching movies is an activity that tends to be more popular on weekends, whereas accessing the Internet tends to be more popular during the work week.

Projecting usage levels for untried services involves, at best, an educated guess. When a popular movie becomes available on VOD, or a highly-anticipated download becomes available via Internet, the demand projection maybe rendered useless. Inexact demand projection leads to excess bandwidth being set aside for penetration increases or peak loading that may only be realized in the future, if ever, compounding the unequal spectrum utilization. More speculative or less popular services may not be offered at all because of the corresponding need to fully dedicate channels and assets.

Finally, these methods ignore synergies possible because video, data and voice all use the same transport protocol.

Figure 1 shows the degree of complexity and functional duplication that results.

Today's methodology means that a new service must somehow present a business case that justifies the use of some integer number, and thus at least one, of RF channels and also justifies the purchase and deployment of all the equipment required for signal processing from signal generation or acquisition through delivery of modulated carriers to combiners. The result is such a high barrier that many promising new revenue streams may never be tapped because of capital and bandwidth constraints and the uncertainty of demand projection. Thus, the "Wonderful new interactive digital service" of Figure 1 may never get launched, even though the cumulative capacity that goes unutilized by other services is more than sufficient to carry the signals associated with the new service.

Assigning dedicated spectrum to each service means that the available spectrum in even a 750 or 860 MHz cable system is quickly being allocated. Furthermore, analog broadcast television will be with us for many years to come while expected must-carry requirements for digital over-air broadcasters will require additional bandwidth, as will new demands coming out of retransmission consent negotiations. Innovations like high definition television (HDTV) may also tax spectrum allocations. Furthermore, many cable operators are constantly under pressure to add new basic service cable networks for both political and competitive reasons.

Operators often find themselves unable to meet external demands for new programming and unable to take advantage of new revenue opportunities, while significant portions of their spectrum are underutilized. Put another way, they are in the ironic position of having the highest per-subscriber bandwidth delivery system of any party in the broadband convergence competition, yet are held back by network sophistication limitations in trying to offer new services.

The authors wish to propose here a more integrated and efficient signal and bandwidth management for switched digital services by elevating the concept of routing in HFC networks.

THE OUTLINES OF A SOLUTION: THE BROADBAND MULTIMEDIA ROUTER

We suggest that dramatic improvements can be made in headend configuration and bandwidth management, in large part through the installation of a single new type of component: a Broadband Multimedia Router (BMR). The BMR is foundation to a service-independent headend management system for all digital services. Its function is open switching and routing of MPEG transport packets of all types, from any advanced digital service, to any node or other grouping of subscribers.

Instead of the headend incorporating a lot of redundant, but under-utilized, equipment, it becomes a switched and routed environment, simplified and streamlined to the configuration in Figure 2. All digital sources feed the multimedia MPEG transport router, while the router feeds a bank of modulators and upconverters which provide the modulated RF signals to the combiner for each node. The BMR would need certain DOCSIS capabilities to support cable modem deployment methodologies.

Any downstream 6 MHz RF channel can then be loaded with whatever mix of MPEG sessions is required to support services requested by customers in a given node, be they VoIP, Internet access, VOD or some other service. The intermixing of different services in shared data streams allows instantaneous spectral sharing and thus the maximum gain from statistical frequency division multiplexing (FDM).

It furthermore means that the mix of communications types may be different in different nodes. Thus, we also gain from statistical space division multiplexing (SDM).

The implications of incorporating the BMR into the headend are extensive, including:

- Elimination of the concept of fixed "channels" and frequency allocations dedicated to particular digital services and sources. The total information capacity that lies within the bounds of each channel can simultaneously carry different services, while packets from any given service can be transported over any channel with sufficient capacity. Dynamic frequency assignment is already provided for in the terminal equipment: video set-top boxes learn about the locations of programs via informational data packets, while cable modems are assigned to frequencies by the CMTS.

- Open and dynamic allocation and re-allocation of spectrum in each node on the basis of real-time demand by customers in that node. Any portion of spectrum can carry any digital service at any time.

- Enhancement of dynamic allocation through phasing out of proprietary equipment and fully embracing such standards as DOCSIS, OpenCable and PacketCable standards that have underlying MPEG transport compliance. This will also facilitate meeting the FCC's retail availability requirements.

Because the BMR is driven by routing tables, it supports new application of intelligence to how the cable plant is utilized. Different communications streams can be assigned appropriate priorities that allow each RF channel to be better used. For instance, high QoS priority sessions (e.g., voice or video) can be distributed over a sufficient number of channels to guarantee a lack of conflict, with the variable remaining space in each channel assigned to traffic whose timing constraints are not as tight.

Using the BMR as the heart of the headend digital management system allows both bandwidth and required equipment to scale with actual customer demand for total utilization of the cable plant, rather than as a function of offered services, each with a projected penetration. Any given service need only provide a standards-compliant transport stream to the BMR input to effect connection with every node in the system. The BMR will sort the packets (including duplication, if required) to feed the appropriate nodes. New transport equipment is only required in response to cumulative plant usage, and not in response to expanding the variety of offered services.

With all digital services now delivered in a switched/routed format, networks can be engineered to provide a total per-subscriber capacity, regardless of what services are offered, which of those services the customer chooses to access, and which services are simultaneously being accessed by others. New services can be added without requiring duplicate equipment and without allocating dedicated bandwidth. It is simply a case of connecting the service to the switching and routing fabric achieved by the BMR.

To the extent that the new services are functional alternates to existing services (for instance, an alternate source of video which is unlikely to be used simultaneously with an existing service) or which have different peak usage times, the total capacity requirements might not change at all and, if they do, it presumably means that customers are more heavily using the network and, therefore, that additional income streams have been tapped. Scaling now becomes a function of total network utilization and the need to guess at the penetration of any given service is greatly reduced.

THE BUSINESS ADVANTAGES OF THE BMR APPROACH

A network configured as described above “looks” unlike a conventional cable television system from a management standpoint. In effect, it changes character constantly: it may become a predominantly a VOD network on Friday nights, a predominantly Internet access network on Monday morning and, perhaps, a predominantly e-commerce shopping network during the Christmas buying period. There can even be node-to-node differences in the “look” of the network at the same time.

Battles over whether some new service is carried no longer happen. Only the signal acquisition or generation equipment is required to add it and, if the service is not popular, it can as easily be discontinued. No significant network engineering is required. Just as easily, a successful service can be scaled up to increase capacity by simply installing more servers or increasing the bandwidth of the connection to the service’s source.

The available spectrum is utilized as efficiently as possible at all times. Therefore, upgrades to increase bandwidth or to split nodes can be delayed, leading to major savings in rebuild capital.

While the BMR can delay node size reduction for some MSOs, others who intend to achieve small sizes, as is being attempted in mini-fiber-node trials, can also benefit from the introduction of switching and routing into the cable plant. Smaller node sizes increase the imperative to distribute service originations with some concentrated upstream and some residing closer to subscribers, according to issues such as service penetration and scalability. Under this scenario, it may be beneficial to distribute the functionality of the BMR between the headend and some lower network level. Increasing control over communications and supporting a variety of protocols, such as SONET, would enable the BMR to facility signal distribution in this heterogeneous environment.

At initial installation, the BMR can simply be added between existing MPEG sources and existing QAM modulators/upconverters. At this level, the significant efficiencies of shared data streams will be realized, while each application can keep its proprietary customer interface. Thus, no significant retraining of headend personnel will be required and the change will be transparent to customers.

Because transmission equipment is now shared across applications, the failure of any individual modulator or upconverter will have limited effect on overall performance. Automatic redundancy switching will simply transfer the load to alternative equipment. Thus, the overall reliability off all services will increase.

In summary, the BMR enables intelligent and dynamic allocation of bandwidth and resources among all digital services. This allows cable to leverage its core strengths to create the greatest effective per-customer information delivery capacity of any multimedia, multi-service provider. Cable, further has the ability to add new services and scale existing services in response to market demand, at the least possible incremental cost. A single new type of device, the BMR, elevates cable to the most sophisticated network for carriage of broadband multimedia convergence.

LONG-TERM BENEFITS AND OPPORTUNITIES

Beyond its immediate implications, the BMR can serve as a foundation for further dramatic innovations to be easily incorporated into cable plant.

Extension of the routing table concept can provide a software management level above automatic routing and resource management. Through this layer, operators can set policies for handling various situations. For instance:

- During times of high capacity demand, the operator may wish to determine a priority by which certain less profitable sessions may be denied, delayed or slowed down.
- Selected video streams (for instance an old movie, but not a live sporting event or high-priced advertisement) may be more highly compressed, with a determined quality reduction.
- Certain customers or customer types may be given access or speed priority.
- During low demand periods, non-real-time sessions (e.g., streaming video to customer premises local caches such as digital video recorders) may be opportunistically initiated.

By monitoring customer response to various network loading situations, these policies can be reviewed and changed as required.

Since not all nodes will require the same total information capacity, nor will the peak demand occur at the same time in every node, a future extension of the proposed scheme could allow the number of modulators and upconverters assigned to a node to vary, based on the total communications needs at any one time. For example, a more commercial part of the system may require a greater total data rate to support small office cable modem service during business hours, but a residential area may more heavily utilize the system for Web browsing and television viewing during the evening hours. By dynamically assigning hardware, as well as spectrum, an operator will need not only less spectrum, but fewer expensive signal processing elements (or looked at in the inverse, provide more services using the same amount of equipment).

New merged services may be enabled by channels carrying a multiplex of several types of communications. For instance:

- Customers may be able to jump between a streamed broadcast video and a related Web site and to optionally initiate a two-way voice conversation with a manufacturer's customer service representative. The "live" program may continue in a window on the screen or even be spooled, so that the customer picks up where he left off when he returns from the Web site. The BMR can smoothly effect the transition between video and data streams and back.

- Ads in broadcast programming may be tailored to the demographics of individual groups of customers, or even to the characteristics of individual subscribers. E-commerce purchasing opportunities may be seamlessly available from any ad. The BMR can simultaneously insert different ads in digital programming being delivered to various nodes.

- Play-along quiz shows that combine interactivity through the Internet with broadcast video may become common. Several precursors to this have emerged, such as the Disney and NFL conducting trials of a parallel Internet game with Monday Night Football in 1999. The BMR can assure coordination in the delivery of streams of different media types to the same subscriber to assure synchronization.

The obituary for analog has already been written by the FCC. While it is uncertain how much longer the patient will live, it is clear that NTSC video will eventually be discontinued. When that happens, operators may well choose to carry all programming on an effectively on-demand digital basis (this can be completely transparent to users, but will result in significantly more effective data capacity, as programs will only be delivered to any given node if somebody served from that node desires to view it). At that time, the full spectrum, capable of delivering multiple Gb/s, can be fully available to deliver services to each grouping of subscribers.

The initial vision for the BMR leverages the ubiquity of MPEG transport protcols. However, with constancy of change a certainty in cable networks, it is important to realize that current standards and protocols may, at some point, be combined with others or displaced. By introducing increased intelligence and communications control, the BMR can extend its benefits to any changes effected in the network. For example, if IP transport is selected by some MSOs for certain services, the BMR would be able to ease this introduction.

If and when desired, services can be offered which are not tied to traditional 6 MHz channelization. This means that the peak data rates available to certain heavily immersive, and as yet unimagined, services or simply higher access speeds for rapid downloads and communications could increase significantly by use of wider RF channels. For instance, a 100 Mb/s data stream would require about 15 MHz of bandwidth.

The BMR would dramatically improve the positioning of MSOs to provide services in the age of broadband multimedia convergence, with both significant near-term implications due to the device's fundamental concept, and far-reaching eventual implications as a foundation for future innovations.

SUMMARY

In summary, cable has moved from its all-analog broadcast roots through the construction of HFC, two-way networks and through the offering of digital services, many of which involve delivering information streams to individual customers.

Allocating a fixed bandwidth and dedicated signal processing equipment to each service based on anticipated peak demand, however, is inefficient of both capital and signal processing equipment.

The authors propose that operators take advantage of the common digital services transport protocol by use of a broadband multi-media MPEG router which will allow dynamic allocation of both spectrum and equipment across many services and nodes. The use of a common headend platform for digital services provides greater spectral efficiency and also allows services to be added and scaled without re-engineering the system. The BMR enables operators to avoid distribution system upgrades that would otherwise be required to increase RF bandwidth or split nodes.

END NOTES

[1] "SBC Orders from High-Fiber Menu," Multichannel News, October 25, 1999.

[1] "Cisco, Broadcom Wireless Gear Boost MCI, Sprint MMDS Plans," Multichannel News, November, 1999.

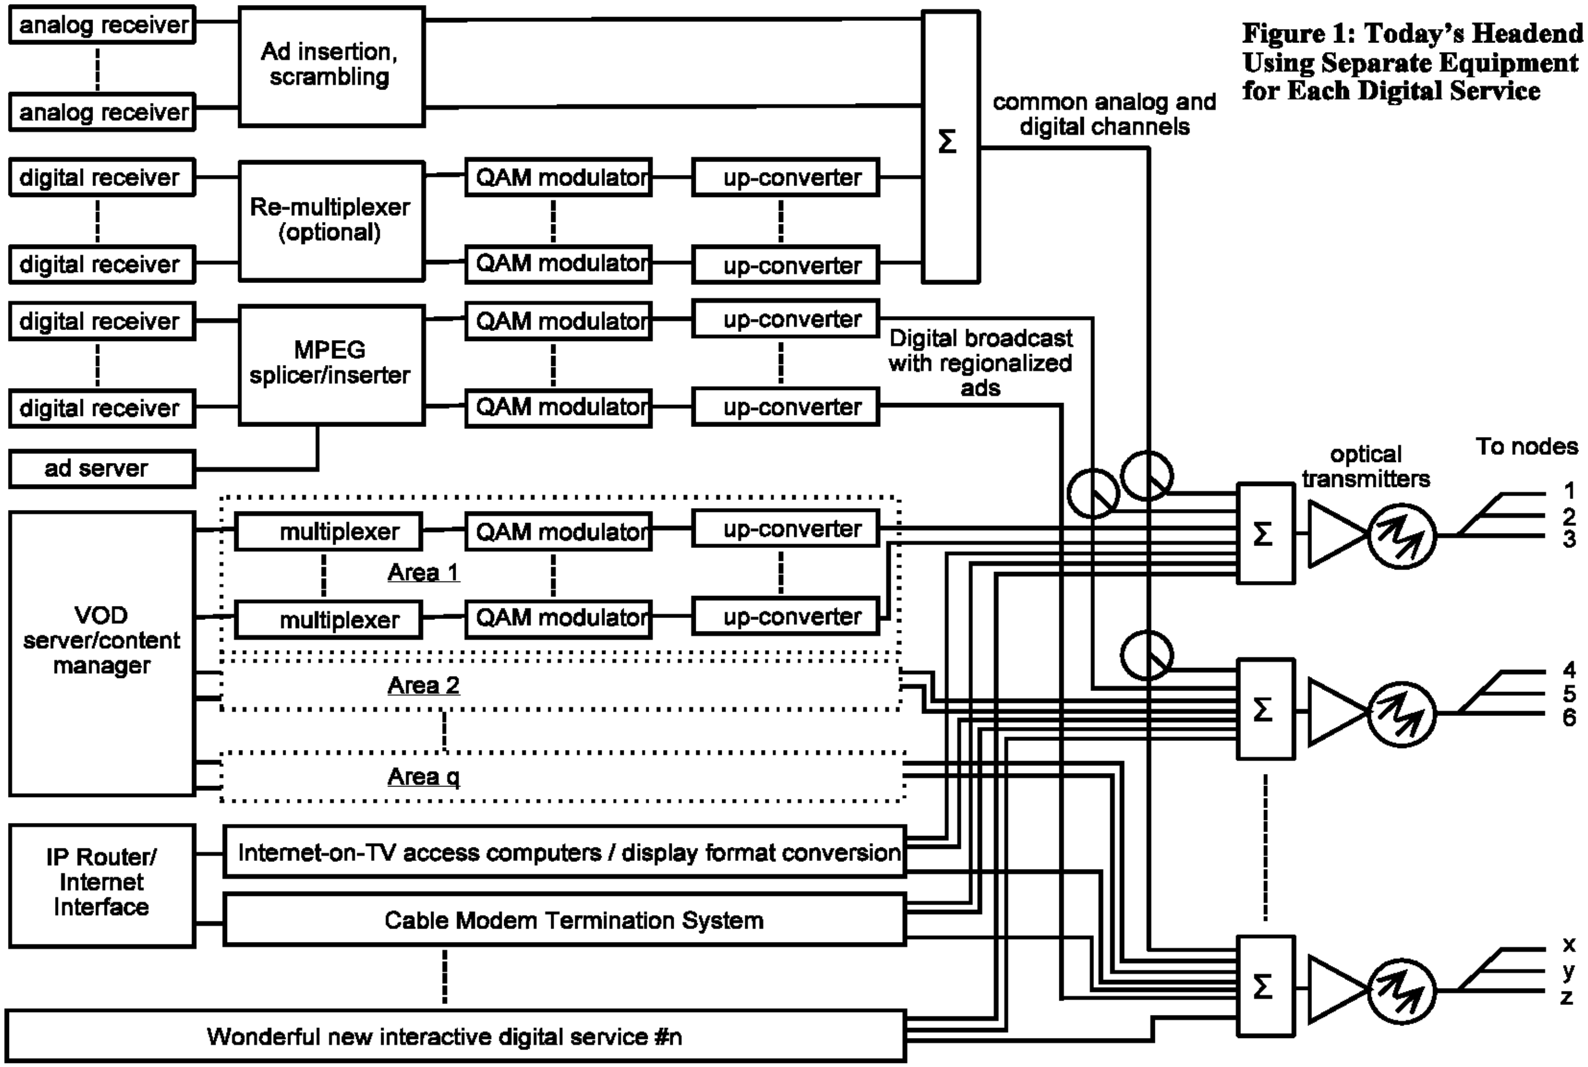

Figure 1: Today's Headend Using Separate Equipment for Each Digital Service

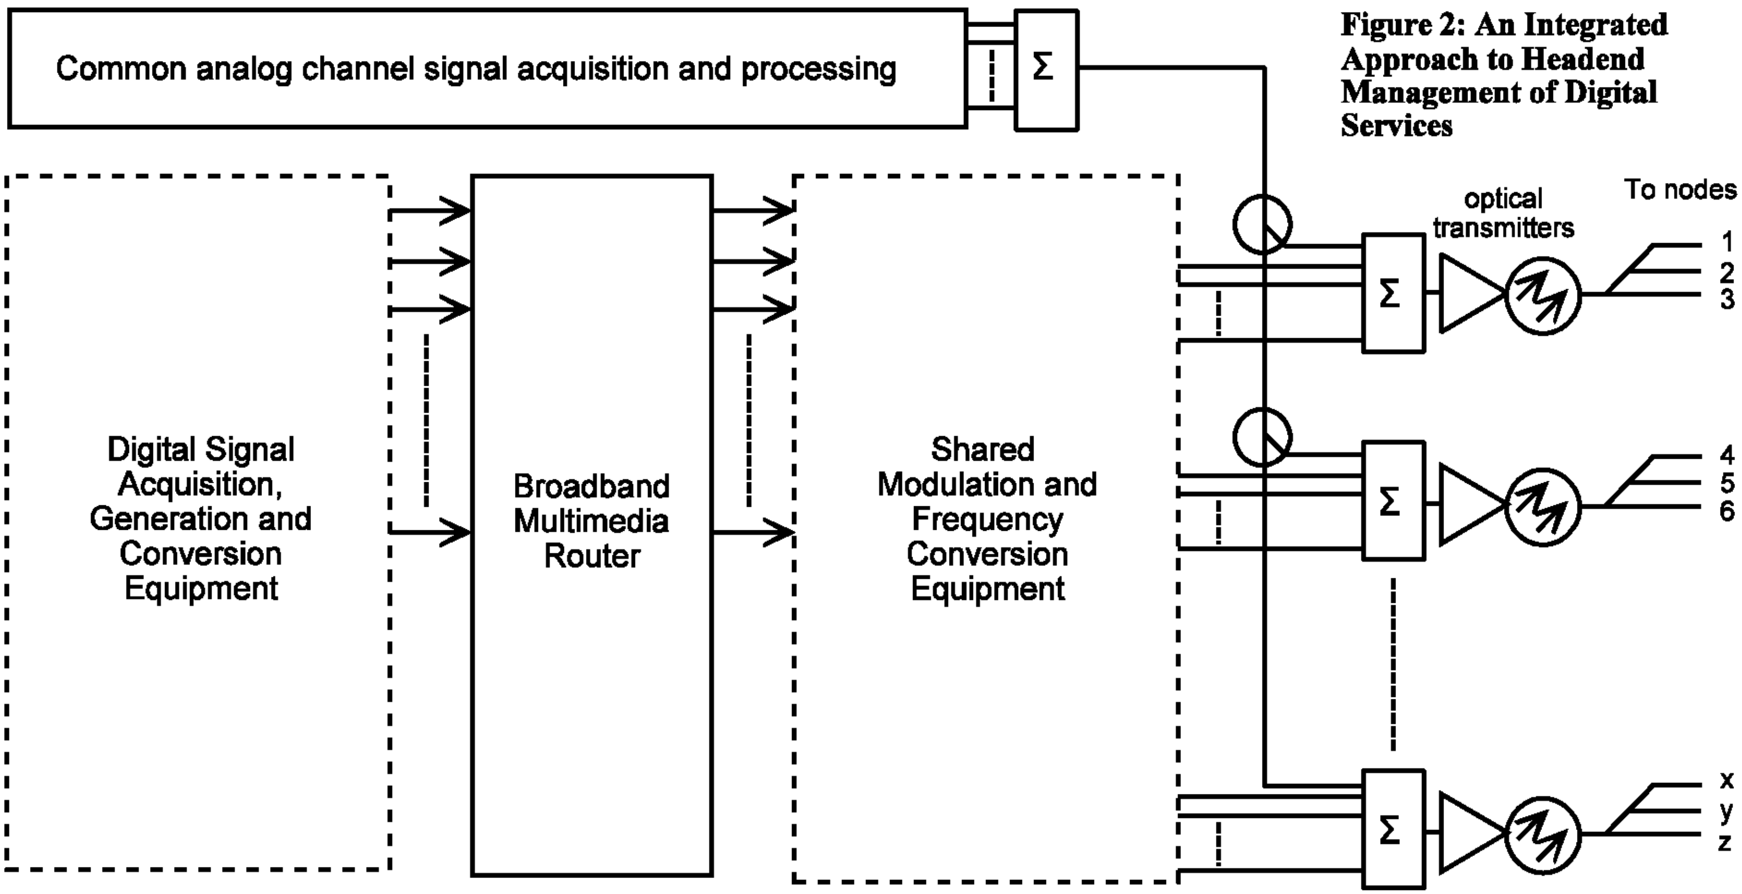

Figure 2: An Integrated Approach to Headend Management of Digital Services