NETWORK DESIGN FOR A MULTIPLICITY OF SERVICES

Ran Oz
S. V. Vasudevan
BigBand Networks, Inc.

Abstract

Cable plants were initially built as a single-purpose infrastructure for the one-way broadcast of analog television. During the last ten years, in order to take advantage of emerging service opportunities and to head off competitive threats, cable multiple system operators (MSOs) have broadened their offerings to an expanding range of services, including digital broadcast and high-speed data. This progress continues with current trends towards additional growing services such as HDTV and various flavors of cable telephony and video on demand (VOD).

There is an increasingly holistic realization of the cable plant as an integrated multi-service and multimedia infrastructure. This realization holds promise to enhance the efficiency, scalability, functionality and ease of rollout of services, through activities such as the proliferation of open standards, sharing optimized resources, and balancing offsetting characteristics of different services and media.

The key to achieving these benefits is to evolve the cable network from a collection of isolated silos of vertically integrated components supporting particular services, towards more inherent fitness for multiple services and media. Greater inclusion of abstraction layers allows more best-of-breed providers of particular aspects of functionality, including the extension of shared functionality across multiple media and multiple services.

LEGACY CABLE SERVICE METHODS

The architecture of the cable plant traces directly to its broadcast legacy. The very concept of six megahertz channelization in North America exists because that's the amount of spectrum that conventionally carries a single analog NTSC program. Historically, relatively few programs were offered to all subscribers in a cable system. The only service consideration was selective subscription by some subscribers to premium or pay-per-view offerings, for which scrambling techniques were devised.

Not much architectural sophistication was required for a plant built for the singular purpose of pushing out relatively few programs to all subscribers in systems that generally were no larger than a very few hundred thousand homes passed. Because the cable operators initially constructing these plants were generally entrepreneurial and debt-financed, there was an imperative to optimize the economics of plant construction for the task at hand, which subsequently motivated technological innovations in selected areas such as downlinks, amplifiers and taps.

Besides its economic streamlining for the task at hand, a significant historical advantage of cable networks relative to other networks connected broadly to homes (most notably telecommunications networks), is the absolute transmission capacity of the coaxial cable which can carry so many rich media programs. However, telecom-munications networks historically boast other distinctive aspects of functionality such as line powering, two-way transmission, and dedicated switching of content to particular subscribers.

New services beyond analog broadcast television have emerged on the cable platform over the past ten years, and continue to do so today. Each of these services goes through its evolutionary stages of proof-of-concept, economic validation, widespread implementation and scaling over time. Because in part of the fact that mere functional viability must be established first, leading early proponents tend to provide complete vertical integration of the technologies required to provide such services. A byproduct of such integrations is the complete autonomy of that service over a discrete count of six megahertz channels, sometimes with granularity to produce content for the same channel distinctively by node in the hybrid fiber/coaxial (HFC) architecture. The fixed allocation of channels and associated resources from cable's analog broadcast legacy is thus extended to services that exist within their own dedicated silos.

MULTI-SERVICE OPPORTUNITY AND COMPETITIVE IMPERATIVE

From time to time, the telecommunications industry has touted its development of digital subscriber line (DSL) technology as a potential delivery method for video. Also during the last decade, the rapid subscriber growth achieved by direct broadcast satellite (DBS) operators has brought on another very real threat to the traditional cable business franchise. With multiple industries proposing to provide broadband networked connections to subscriber homes carrying video, data and voice content, cable has had to fortify its positioning in the face of emerging competition.

To its credit, and true to its entrepreneurial roots, the cable industry has always looked at expanding multimedia and multi-service offerings as a bona fide business opportunity and not merely competitive reaction. The earliest forays into interactivity and on-demand consumption go back several decades. But with progress made in the satellite and telecommunications industries, as well as the greater application viability of services like VOD, the time has emerged for cable to realize its multimedia, multi-service potential for defensive as well as offensive reasons.

Beyond analog broadcasting, the first services that the cable industry has widely deployed are digital broadcasting and cable Internet access. These services represent directly competitive spaces with DBS and DSL (which has grown as a high speed data service, but has yet to materially provide a video alternative), respectively. Because the deployment of these two digital media services over cable required various enhancements to the cable plant, the last ten years have seen several profound cable industry undertakings, including the development of the HFC architecture (with its node-level multicast granularity), return path capability in the plant, and several widely embraced digital standards such as MPEG-2 transport.

Competing effectively with impressive subscriber growth in both digital broadcasting and cable modem service, cable is currently embarking upon further service expansion, highlighting areas where it can leverage the investments made in the plant, and where it is uniquely or best positioned for distinction. Such services include high definition television (HDTV) and various packagings of VOD including movies on demand (MOD), subscription VOD (SVOD), free VOD (FOD), networked personal video recording (NPVR) and long form advertising (LFA).

Consumer adoption trends of HDTV and VOD offerings are encouraging. At the same time, the cable industry prepares itself for the widespread rollout of telephony services, which have achieved encouraging popularity in their isolated current deployments. Furthermore, the long anticipated interactive television (iTV) space also maintains its promise as technologies are further refined, bolstered by popular indicators such as the success of direct response advertising and audience participation in various forms of reality television.

OPTIMIZING TECHNOLOGIES IN MULTIMEDIA, MULTI-SERVICE ENVIRONMENT

As these trends play themselves out, the cable plant that was initially constructed for one-way broadcast of analog video becomes a much more complicated beast. We increasingly see a single, yet multi-faceted network that carries combinations of video, voice and data content; broadcast and personalized sessions; passive and interactive consumption – among other variations in subscribers' engagement with media and services.

While voice, video, and data may seem like extremely dissimilar services, the composition and delivery of these services share some technical similarities. The general technical elements that are integrated to deliver a service include the source of the service, the switching of the correct content towards the correct destination, the physical transport of content from source to destination, and the media processing of content.

For example, a VOD service requires servers with video storage and transactional and billing applications as service sources, switching capabilities to assure that the right programs are directed to the right service groups corresponding with subscriber sessions, physical transport to get the content to the node (or, in the case of server distribution at hub sites, to propagate the content to storage for later playback), and media processing such as de-jitter, QAM modulation and upconversion.

Some of the earliest VOD offerings integrated all of these elements in a solution offered by a single vendor. This severely limited choices by MSOs in how particular elements of the service could be configured, cost-effective service expansion could be achieved, and how flexible headend architectures over time could be accomplished. In the years since the early VOD deployments, there has been an increasing disaggregation of functionality, to the benefit of specialists in particular elements of functionality, enabling a decoupled system that allows different vendors to provide the server, transport, and modulation functions for VOD installations.

Lack of complete standardization or broad protocol flexibility continues to hem in choices. For example, supported transport protocols are based largely on the output format of a server. But with time, operators are increasingly able to migrate from legacy ASI transport to Gigabit Ethernet transport between facilities, with easier abstraction from which VOD server is selected. And modulation equipment at the network edge will be able to accommodate for the chosen method of transport, be it legacy or contemporary.

RESOURCE SHARING ADVANTAGES

The rich variety of services and media increasingly carried on the cable plant presents opportunities to leverage technologies that can be optimized for their offsetting as well as their common characteristics. Yet, the one-way broadcast legacy of the cable plant by and large maintains its influence, and services are introduced in sub optimal fashions as a result.

One clear example of this legacy is the continuing channelization of spectrum and associated resources on a service-by-service basis. In the same manner that the local affiliate of a major broadcast network is allocated a six megahertz channel, so is such a swath of spectrum is provided on a 24-by-7 basis to approximately ten digital broadcast programs, or VOD capacity for the anticipated peak consumption among approximately 100 digital cable subscribers in a node or service group, or for shared cable modem access by several hundred subscribers to that service.

This means that spectrum, and its associated resources such as QAM modulators and upconverters, are permanently allocated to particular services regardless of those services' use at particular points in time. Because resources and spectrum are generally capable of fungibility across services under common standards, this situation is suboptimal unless demand for services is completely static or correlated.

For example, if a major news story breaks, there may be a lot of demand for high speed data and live broadcasting and relatively little VOD activity. In such a scenario, the operator runs the risk of underallocating high speed data resources, which could deny transactional revenues or frustrate customers (resulting in churn vulnerability). At the same time, capital is being wasted on underutilized resources dedicated to VOD.

Because all of this digital media traffic is based on common standards, such as MPEG-2, it is relatively easy to reap the advantages of dynamic allocation of bandwidth and associated resources in response to real-time demand. Thus a six megahertz channel and its QAM and upconverter can be carrying a heavy load of data services at the moment described above, and more VOD at other times, such as weekends and prime time or particular spikes for such services as when popular titles first become available. Because of the more efficient utilization, fewer overall resources are required. Alternative methods of accommodating scenarios of potential demand peaks through overprovisioning require exorbitant capital expenditures and drive horrendous bandwidth and resource utilization.

Bandwidth and resource sharing across cable services leverages the advantages of statistical multiplexing in which various traffic loads, unlikely to all simultaneously peak, ride together. The result is greater efficiency of shared resources and less likelihood of denial when particular services spike in demand.

Consider a hypothetical situation of services A, B, and C which, respectively, are permanently allocated channels 1-3, 4, and 5-6 in a conventional cable plant as indicated in the image on the left, in Figure 1. The bandwidth available to subscribers of a service is limited by the channels allocated to that application. The left image shows all applications at exactly peak capacity, however, because of the periodic shifts in application consumption patterns (ex: night vs. day, weekday vs. weekend) and short term irregularities that occur (ex: a highly anticipated report is published via Internet, a popular movie becomes available on VOD), in reality a fixed-allocation system cannot be managed to achieve anything close to such high utilization of bandwidth, and fulfillment of demand.

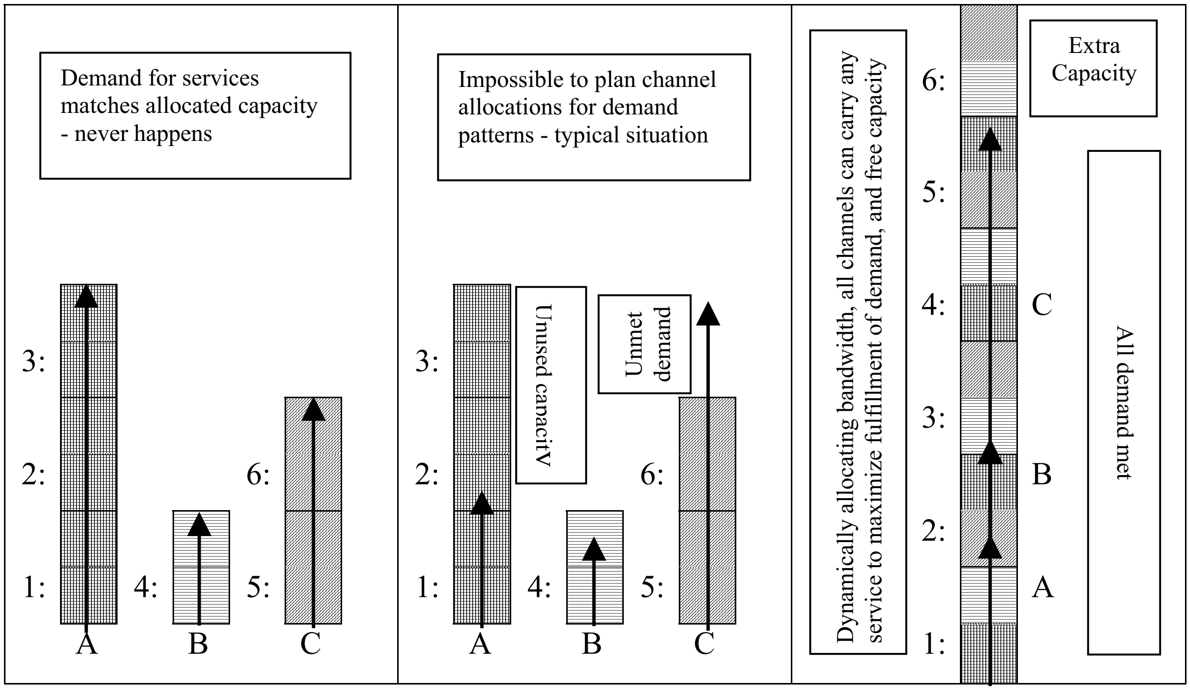

Figure 1: Dynamic allocation of bandwidth across services within channels.

The middle image shows a typical situation in which the demand for service A is less than half its three channels' capacity while demand for service C is well above the spectrum allotted. This results in lost revenues, angry customers, or both due to unfulfilled demand for service C, which occurs at a time when there is free bandwidth in channels 1-3.

By implementing dynamic bandwidth allocation, an MSO can achieve the situation in the right image in which assets are more flexibly managed as all channels, 1-6, can carry any of the services. The demand for A, B and C is combined and mapped onto this one big band. This fulfills all demand with even more capacity available, facilitating introduction of more revenue generating services.

Recent technological advantages have even introduced broadcast television to the pool of services dynamically sharing resources. As hundreds of simultaneous programs are made available to subscribers, it can be assured that not all of them would be watched at any particular point in time in any particular node. Instead of filling spectrum in all nodes with all broadcast programs all of the time, switched broadcast techniques can be utilized to dynamically respond to channel surfing and switch live programming only to the nodes where it is being actively watched, allowing broadcast programming to share spectrum and resources with other, interactive services.

When different services and media share the same bandwidth and resources, offsetting characteristics between them can be leveraged. For example, when video, a real time medium, shares bandwidth with data, which is time-sensitive but not necessarily real-time, peaks in video traffic can be accommodated by shifting data traffic into the troughs of video traffic. Another example is combining variable bit rate HDTV traffic with standard digital video, in which case rate shaping can be applied to the standard video when the HDTV requires more bandwidth, which assures the reliable and robust delivery of both services while maintaining HDTV quality. The combinations of various media within the same channels has the long term potential to enable new generations of services that richly combine media, such as multiplayer games that operate in conjunction with live video programming.

When plant resources in general, and bandwidth in particular, are dynamically and openly allocated across all services in response to real time demand, the processes of service launch and scaling become greatly alleviated. Traditionally, launching programs or services requires discrete determinations of which existing offerings must be stopped or reduced to accommodate spectrum. With open allocation, new services can be layered on, and if capacity contention arises, this can be addressed through implementation of intelligent implementation priorities by service, session or subscriber, or further bandwidth enhancements such as rate shaping to adapt video bit rates.

ARCHITECTURAL OPTIMIZATION

Many advantages are reaped with less vertical integration of services and more resource sharing among them. Specialization and innovation of functional components such as QAM, switching, or transport by particular vendors allows these vendors to hone best-of-breed deliverables in their areas of specialty.

Isolating and aggregating particular elements of functionality allows these elements to be designed for maximum economic scalability. And that scalability should be able to be achieved in an incremental fashion. This improvement, combined with decoupling of components, allows the operator to choose exactly how much of particular elements of functionality are required.

To exemplify the alternative, when systems are vertically integrated, storage may be combined with media processing for some services. Thus the operator's scaling of infrastructure is based on whichever is the bigger driver between total content provided (storage determination) or total session capacity (processing determination), and whichever component doesn't drive the installation is uneconomically overprovisioned as a result.

Decoupling functionality also enables the operator to optimize architectural configurations and not be constrained by vendor designs. In VOD, as an example, some operators express preference for centralized server consolidation, and some prefer to distribute servers to hub locations at the edge. In either case it is generally agreed that media processing such as QAM modulation should be located at the edge, to assure quality of content, with the most economic transport methodologies used to get content to the edge. By having independent QAM components that are separated from the server, this media processing can be maintained at the edge, with servers based in either location.

Servers could even be based both centrally and distributed in a hybrid configuration, which may be desirable either to leverage existing edge servers while moving towards greater centralization, or to allow a form of near-line storage so that more popular titles to be pushed towards the edge to streamline transport utilization, while larger libraries are centralized on economically scalable server resources. Content sourced from both locations share the same QAM resources, optimally placed at the edge, and utilized efficiently whether the temporal demand is relatively higher for the popular titles or the broader library content.

Breaking vertically integrated services towards greater specialization and abstractions between elements of functionality also protects the operator from technological obsolescence. With full vertical integration, modifying one element of functionality may not be viable as it would necessitate a complete forklift upgrade to the system. Alternatively, specialized components can be modified or upgraded with relatively minimal changes, perhaps only requiring modifications in the interface of existing components.

In fact, new and old components can coexist to drive innovation while maintaining existing plant investments. For example, in VOD there is a drive towards Gigabit Ethernet as the best transport method, and Gigabit Ethernet compatible QAMs are emerging to optimally accommodate this trend. Yet, many systems have already invested, in many cases only recently, in ASI QAMs. A solution that avoids any stranded capital is to use the Gigabit Ethernet transport for all VOD traffic, grow capacity with Gigabit Ethernet QAMs, and use Gigabit Ethernet to ASI protocol conversion to prolong the operational life of the already-purchased legacy QAMs.

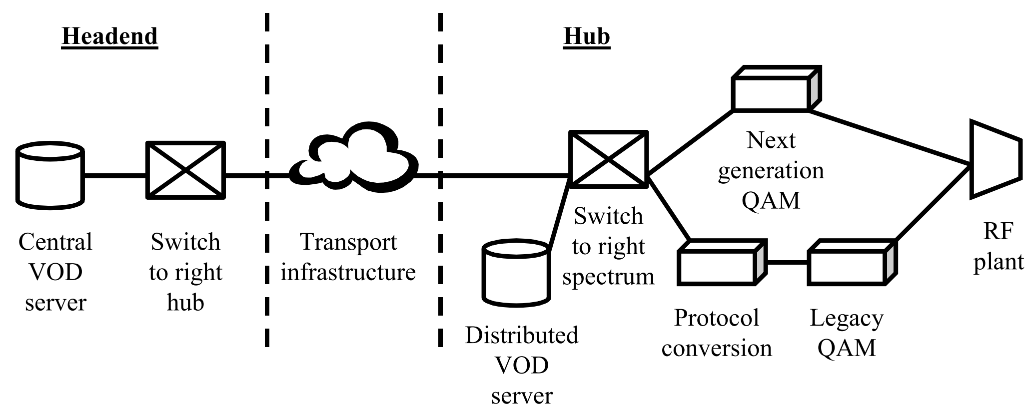

Figure 2: Flexibility for MSO to customize architectural deployment of best-of-breed components openly interfacing to each other (VOD example).

CONCLUSION

There is increasing volume and variety of multimedia services available to be provided over broadband networks. With inherently high bandwidth, increasing sophistication through trends such as migration to packetized digital content and node-level addressability, and established standards such as

MPEG-2, cable has an opportunity to emerge as the leading comprehensive provider of all services and all media.

Situations in cable systems are highly particular due to considerations such as what services are being emphasized, legacy investments already made in the networks, and overall unpredictability over how scenarios will play themselves out going forward. The one entity that can best determine how to architect infrastructure across media and services is the MSO.

The key to cable realizing its multi-service potential is for MSOs to have control over determination of deployments among open, interoperable components. Abstraction should be increasingly implemented among sources of service, switching, transport and media processing. This allows the cable operator to select best-of-breed components, optimize plant efficiency, economically scale resources, and determine precise architectures. So empowered, the cable operator is positioned to compete effectively and seize emerging opportunities.

Ran Oz
Chief Technology Officer
ranoz@bigbandnet.com

S. V. Vasudevan
Chief Architect
vasu@bigbandnet.com

BigBand Networks, Inc.
475 Broadway
Redwood City, CA 94063
650 995 5000
http://www.bigbandnet.com