Planning for and Managing the Rollout of Switched Broadcast Services

SCTE 2003 Conference on Emerging Technologies

S.V. Vasudevan
Chief Architect, BigBand Networks
475 Broadway
Redwood City, CA 94063
+1 650 995 5043
vasu@bigbandnet.com
http://www.bigbandnet.com

Paul Brooks
Senior Network Architect, Time Warner Cable
160 Inverness Drive West, Suite 300
Englewood, CO 80112
+1 303 799 5726
paul.brooks@twcable.com
http://www.timewarnercable.com

TABLE OF CONTENTS

Introduction

As new digital services, including those now offered in the bandwidth-hungry HDTV format, clamor for ever scarcer bandwidth on MSO cable TV systems, and with capital-conscious operators still digesting over $60 billion spent in the last sweep of plant capacity upgrades, much attention is being directed towards Switched Broadcast Services. Switched Broadcast is a method of sending selected programming only to nodes where and when subscribers are actively requesting that programming. Such services have the promising potential to promote revenue growth and churn reduction by expanding programming choices and enhancing cable plant bandwidth utilization efficiency. Switched Broadcast is becoming operational reality, with initial trials by major MSOs commencing during the second half of 2002. Results will provide meaningful data for planning broader commercial deployments based on empirical subscriber usage patterns, validating the nature and magnitude of Switched Broadcast's benefits.

In today's cable network, the status quo practice of filling the entire spectrum with live *broadcast* programming at all times has recently been supplemented by *unicast* video in applications like Video On Demand (VOD) and Subscription Video On Demand (SVOD), where stored video programming is directed from a server for consumption only by an individual subscriber. Switched Broadcast represents the further evolution of the cable network by expanding the concept of traditional broadcast programming with a *multicast* programming overlay. This multicast overlay is driven by subscriber requests and can dynamically adjust network transmission capacity resource usage across nodes, service groups, and hubs. This capability efficiently provides an expanded programming experience to television subscribers that is customizable, personalizable and can potentially accommodate a limitless content selection.

While making substantial bandwidth efficiencies possible by itself, Switched Broadcast will be just one of a number of services with statistical usage patterns in a modern cable TV system. Other common statistics-based services include VOD and DOCSIS IP. These services are all fundamentally based on the principle of best effort transmission in a shared capacity medium, although they all offer to varying degrees the capability to provide a guaranteed (non-blocking) grade of service. Additional efficiencies can be tapped by flexibly and dynamically managing the shared bandwidth resource pool for all such services rather than within isolated silos, thus exploiting the statistical nature of viewer distribution fluctuations across all on-demand services over time. This type of advanced video service delivery system will be enabled by introduction of new networking and network management functions on both hardware and software layers of the cable TV network. Thus Switched Broadcast, as a specific implementation of switching resources towards live content according to real time demand patterns, can ultimately be extended to achieve similar efficiencies and benefits in switching dynamically across services and across media.

Switched Broadcast Drivers

A number of factors related to the viewing demands of cable subscribers, the topology of today's modern HFC networks, and innovations in digital video switching technology have

paved the way for Switched Broadcast to become a viable method for delivering expanded content. In a Switched Broadcast paradigm, cable spectrum is no longer looked at as a static collection of channels classified by classical EIA frequency plans, but as a dynamic pool of digital transmission resources that, at any moment, could be serving any broadcast programming from any source. Logically, for even greater content expansion, this could eventually be extended to also encompass narrowcast VOD content, interactive television, or any digital service that could be developed and introduced in the future. Several forces are converging to raise the visibility of a Switched Broadcast delivery model.

Demanding Forces: The Proliferation of Content

While it seems like only yesterday when the major plant upgrades took place, many to 750 and even 860 MHz, operators are quickly finding out that they are again running out of bandwidth for the carriage of traffic. Today's digital cable plant boasts an expanded set of broadcast and interactive video, voice and data services. While these offerings allow cable operators to be well positioned with respect to their satellite TV competitors, the MSOs are finding themselves in a spectrum crunch reminiscent of that which preceded the latest round of upgrades.

SD and HD Digital Broadcast

The rollout of digital television has created a digital tier of programming that subscribers can access via digital set top boxes. Advances in encoding and rate shaping technology have allowed as many as 14 digital programs to occupy the same frequency spectrum as a single analog channel. Yet even this vast expansion of programming capacity has been quickly filled. Many programmers such as HBO, Showtime, Discovery, MTV, and ESPN have expanded their programming packages to offer vertically oriented variants of their standard programming. New programmers have also emerged to offer content that is more narrowly focused. The 500-channel packages of broadcast programming initially articulated by cable visionaries were first realized by DTH satellite operators, and serve as a threat for MSOs trying to sustain a competitive program lineup. The continental footprint and abundant bandwidth resources available to the DTH operators serves as a perennial motivator for cable MSOs to cost-effectively build and sustain a network that is capable of offering an equivalent or greater array of programming.

The recent emergence of HDTV and its rapid deployment by North American operators serves as an additional business opportunity, but also as a further bandwidth burden on the HFC plant. Whereas 10-15 SDTV signals can be carried in 6 MHz of spectrum, the bandwidth consumed by HDTV programming only allows 2 or 3 programs to be carried in the equivalent space. The emergence of HD programming on cable networks can almost be regarded as a *fait accompli*, with commitments offered by the top 10 MSOs to the FCC to provide five simultaneous HDTV signals in the top 100 media markets. With all of the networks and some premium broadcasters offering 8-12 HD programs, operators may find themselves scrambling to find 3 to 6 extra QAM channels to carry HD content. As the cost of HD televisions continue their precipitous drop and the popularity of crystal clear HD

programming explodes, consumers will demand ever-greater portions of their programming in this new format, and vote with their remote controls.

Video On Demand (VOD)

VOD has emerged as both a credible response to the DTH competition and a promising service that can generate incremental subscriber revenues. VOD is widely regarded as the "silver bullet" for the cable industry — its capabilities cannot be practically duplicated by the competing satellite DTH operators. All major MSOs are involved in VOD rollouts in their major systems, and the number of VOD-enabled subscribers in the U.S. is expected to balloon from 7.8 million subscribers in 2002 to 41.7 million subscribers in 2006[1]. However, the rollout of this service also makes a substantial demand for cable spectrum.

VOD is a *unicast* service, which means that a separate digital program stream is transmitted to each individual user, a far cry from the *broadcast* paradigm, where a program is ever-present and available to all users. To size the bandwidth requirements of a VOD service, operators typically make a mathematical assumption about how many digital subscribers will be using the VOD service simultaneously during the peak viewing period. Empirically, this number has typically been 5-10% of the number of digital households. For a headend with 100,000 digital subscribers, this equates to 5000-10000 streams, or 500-1000 QAM channels for VOD alone.

This number may seem daunting compared to the 20 or so QAMs required to carry the digital tier of broadcast programming, but it is mitigated by the fact that the downstream bandwidth needed for VOD can be re-used in each HFC node, or by a clustered group of nodes commonly referred to as a *service group*. The true amount of capacity available for any particular service, like VOD, is actually the product of the QAM count times the number of service groups in the headend system for the particular service.

The re-use of bandwidth by node or service group, a road paved by VOD, provides the perfect vehicle for Switched Broadcast to provide an augmentation of transmission capacity to serve the programming demands of subscribers.

PPV Sports Packages and Niche/Ethnic Programming

The proliferation of programming has led to an interesting conundrum for operators. Event-driven sports programming packages like NBA League Pass™ offer very lucrative incremental revenue opportunities for operators, to the tune of $20 or more per month per subscriber. At the same time, the viewership for these packages by subscribers is limited and sparse. Nevertheless, operators are compelled to broadcast this type of programming, which can consume as many as 4 QAM channels, to all subscribers, even when they know that only a sparse handful of viewers are entitled to receive this package, and a smaller number of subscribers can be expected to be actually watching the programming at any given time,

---

[1] Bilotti, Richard, "The Rise and Fall of Programming and Capital Expenditures", Morgan Stanley, April 5, 2002.

especially competitions between contestants not widely considered compelling and/or from distant locations.

Similarly, many programmers desire to offer international or ethnic programming to a specific demographic of users that comprise their subscriber base. Often, these constituents are largely concentrated in neighborhoods or geographical regions that can be served by a single or small group of hubs. This type of programming by definition will experience low viewership outside its targeted demographic. Switched Broadcast offers the opportunity to re-use bandwidth that would otherwise be allocated system-wide to deliver niche and ethnic programming to all demographic groups under the operators' purview. With Switched Broadcast, it is not necessary to know in advance where the viewership of these programming types is concentrated, nor to define hub or node boundaries accordingly by demographic.

Enabling Forces: HFC Plant Architecture and Bandwidth per Subscriber

The HFC plant architecture used in modern CATV plant is by nature easy to sectionalize into discrete, completely separate communications systems. These small systems, commonly called nodes (with a granularity as fine as 500 homes passed) can and often are aggregated into service groups of 1,000 to 10,000 homes passed by virtue of simple RF splitting and combining. The aggregation ratios (service group sizes) can be different for each service, and are usually optimized based on subscription rates (traffic engineering) and performance (S/N+I) trade-offs. For instance, traffic-based aggregation is commonly done today for the purposes of High Speed Data (HSD) service, where in high-usage cases an individual node may be provided with its own input port on the CMTS (termination) equipment, while in other, lower traffic areas, several nodes may be aggregated.

This architectural capability happened kind of by accident. Because of the limited loss budgets of early linear analog optical links, one laser transmitter could only serve an average of three or four receivers over short links and as few as two or even a single receiver in links of more moderate length. This led most of the industry to build what are essentially "home run" or star topology connections to the optical service area nodes. In conjunction with this, it was also found that the optimum number of RF repeater amplifiers in cascade is typically in the range of two to eight, depending on cost-based coaxial distribution design trade-offs. These short cascades limited the service area sizes to 500 to 2000 homes passed. The combination of home run optics and small node sizes yielded enormous increases in bandwidth per subscriber. This capability has only recently been taken advantage of, first by HSD and more recently by VOD. Switched Broadcast is the logical extension of efficient use of this accidental characteristic of the HFC architecture. If it weren't for this high bandwidth per subscriber, there would not be sufficient spectrum to support the statistical sharing of bandwidth that forms the basis of Switched Broadcast.

Just like modern cellular telephony systems reuse common frequencies in their specific cells, Switched Broadcast implementations can reuse common pools of QAM channels and program bandwidth across service groups to provide a larger logical programming capacity than what is physically available in the traditional broadcast delivery paradigm.

Two-way communication capabilities provided by modern segmentable HFC plants, combined with the present sophistication of real-time interactive set top box designs provide the necessary environment for deployment of Switched Broadcast. It is the belief of these authors that existing set top terminals and existing HFC outside plant architectures DO NOT NEED TO BE PHYSICALLY MODIFIED for the provision of this service.

New Forces: The Changing Nature of Television Viewing

Television has come a long way since January 1953, when 44 million viewers (representing a whopping 72% share of TV households) were simultaneously glued to their sets to experience the birth of "Little Ricky" on the television series *I Love Lucy*. An ever-expanding selection of programming has enabled subscribers to diversify their viewing preferences and habits. Specialized programming is now offered related to cooking, gardening, car racing, education, international/ethnic topics and a host of other subjects. The recent emergence of Personal Video Recorders (PVRs) has further modified traditional patterns of television viewership. Viewing the season finale of *The West Wing* is not necessarily a social viewer experience anymore. While it is true that broadcast networks like CNN and ESPN will always command a certain level of subscriber viewership, the spatial and temporal fragmentation of viewership complicates the task of determining the relative values of less-popular broadcast programming transmitted to subscribers 24 hours a day, perennially consuming scarce spectral resources.

Switched Broadcast Components in the Cable Environment

Three components are essential to the provision of Switched Broadcast. These are: client software resident in the set top, controller functionality at the headend, and, of course, the switch itself. Other issues that must be considered are transport network enhancements to deliver the plethora of content to the switch input, and a conditional access system to allow the operator to control access and monetize the service. This paper focuses on the functionality switch as a distinctive and crucial central element to interoperate with the other components and fundamentally enable Switched Broadcast.

BigBand Networks has developed the functional components to realize a Switched Broadcast implementation for digital cable networks. Central to the implementation is an MPEG-Aware Switch (MAS), BigBand Networks' Broadband Multimedia-Service Router (BMR). The authors will attempt to present this equipment on a generic basis for the remainder of this paper, although certain aspects of the discussion may be specific to BigBand equipment.

Switched Broadcast Client (SBC)

A Switched Broadcast Client (SBC) is a small software application resident on the set top box. This client functionality can just as easily be integrated into the tuning firmware of future set tops. When a Switched Broadcast program is selected, this software component

conveys the channel request via the upstream cable plant, along with information that uniquely identifies the set top box.

Switched Broadcast Manager (SBM)

The Switched Broadcast Manager (SBM) application runs on a machine located in the hub or headend. The SBM uses the channel number received to identify the requested program, and consequently, the MAS input port where the program is being received. Similarly, the SBM uses the set top box ID and associates the service group information to determine the downstream connection (MAS output port) where the subscriber can be reached. In many cases, a subscriber can be reached by more than one downstream QAM channel. This collection of one or more QAM channels, and the dozen or so programs that can be carried in each QAM channel, represent a pool of resources that the SBM has at its disposal to fulfill programming requests within each service group.

When an available downstream QAM and the program resource are identified, the frequency and program information is returned via the downstream out-of-band channel to the set top box, which decodes and displays the program using the normal tuning mechanisms. While the specific frequency and program number for a Switched Broadcast program may vary in time, the channel number as seen by the subscriber will always remain the same.

MPEG-Aware Switch (MAS)

Completing the solution and providing real-time video switching capability is an MPEG-Aware Switch (MAS). In order to perform the switching role effectively, an MAS should exhibit the following characteristics:

- High Input Fan-in: An effective Switched Broadcast solution allows a substantial and theoretically unlimited amount of programming to be switched into a bandwidth-limited set of downstream transmission resources. To achieve this, an MAS should support an input bandwidth that greatly exceeds its per port output bandwidth.
- Fast Switching Time: The amount of time required to switch an input program to an output port should be as small as possible, since this switching time is a component of the overall channel change time. Furthermore, the switching should take place in such a way that the switching process does not syntactically corrupt the output stream.
- Supplemental Media Processing: In addition to the pure switching process that takes place, it is useful for an MAS to have supplemental media processing capabilities as well. Such processing functions include, but are not limited to:
    - Remapping of PIDs and reconstruction of PSI
    - Processing of transport streams to correct PCR jitter
    - Remultiplexing and rate shaping of output streams as necessary
    - Transcoding of elementary streams from other formats as necessary

Switched Broadcast Operation in the Cable Environment

The actions that take place when tuning to a Switched Broadcast program differ from traditional broadcast tuning. In a traditional broadcast environment, *all* programming is sent to the set top box, along with data streams that convey the Program Specific Information (PSI) and System Information (SI). Using these tables, a set top box can determine the specific associated EIA frequency and program number, and command the set top box tuner and MPEG decoder to receive and display the selected program.

In a Switched Broadcast environment, the same tuning methodology is used, with a slight modification. The tables that carry the frequency and program information for the broadcast channels now vary with time. Therefore, the set top box uses the out-of-band channel to acquire the now time-variant tuning information. When this information is retrieved, the tuning process proceeds as normal. To the end user selecting a program via the electronic program guide or the remote control keypad, this is a visually seamless process.

Channel Selection Process Flow

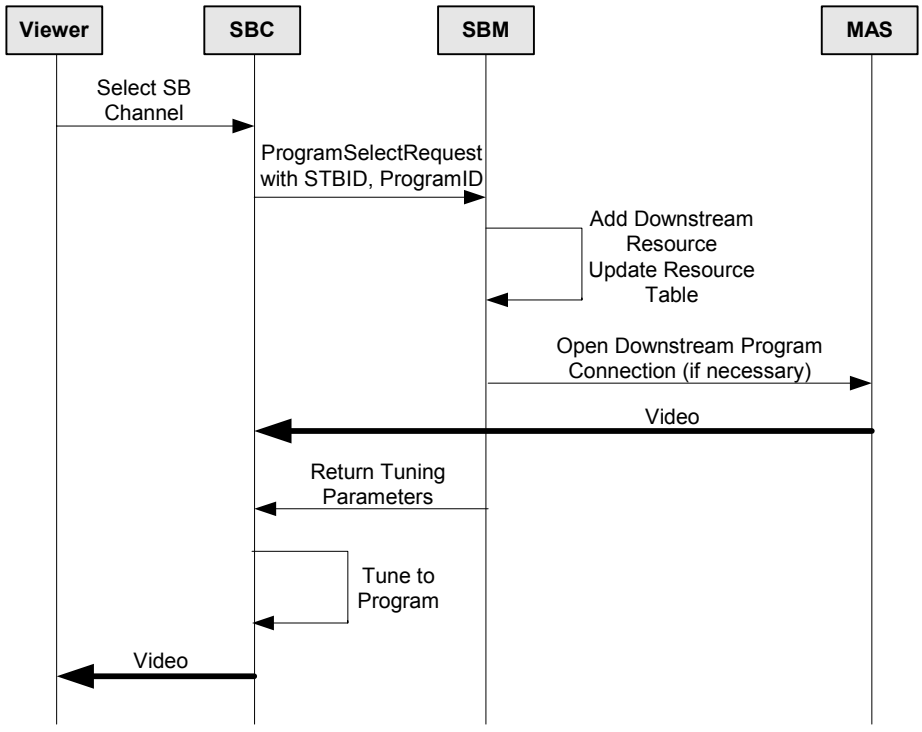

Switched Broadcast – Channel Selection Sequence Diagram

A Switched Broadcast channel change starts with the SBC application on the set top box sending a channel request message upstream to the SBM. The SBM receives this message

and reconciles the request against the available downstream resources. If a downstream resource is available, the SBM instructs the MAS to complete the downstream video connection (if it is not present already) and perform any stream processing that may be necessary. The downstream resource information (the QAM frequency and Program Number) is returned to the set top box, which tunes the program.

Channel selection for subsequent viewers of the same program follows a similar sequence, but the process is simplified. If the requested program is already being transmitted downstream, then the SBM simply needs to return the appropriate tuning information to the set top box.

Channel Leave-off (Stream Dropping and Capacity Recapture)

When a user is selecting a Switched Broadcast program, that user may also be leaving another Switched Broadcast program. Therefore, the channel request message also contains information about which program the set top box is leaving. The SBM uses this information to track the number of users watching a particular program. If the number of users watching a program reaches zero, the SBM can (but does not have to) reallocate this downstream resource for another program request.

In addition to maintaining an accounting of programs currently being viewed, the SBM can also optionally interact with specific set top boxes to confirm viewing information for additional robustness. The ability to recapture QAM program stream capacity is critical to aggressive deployment of Switched Broadcast. In the future, several different means to prioritize the stream dropping process can be implemented, such as program event boundary knowledge, time in stream timers, and set top usage (button push) monitoring and timers. Other mechanisms, such as freeze frame or other alert methods may be used to encourage the subscriber to "push a button" to convey to the SBM the fact that the stream is still being viewed. Such mechanisms would be used prior to dropping of a stream if the viewing status cannot be determined. If such mechanisms are not implemented, the QAM stream inventory can become clogged with streams that no one is using, such as when a subscriber turns off the TV. Because Switched Broadcast is inherently a best effort service, some type of message or barker would be provided to subscribers requesting a stream when none are immediately available.

Viewership Statistics Collection

While there are possible legal implications surrounding the gathering of individual subscriber data, it is advantageous and arguably necessary to the efficient use of transmission resources and subscriber satisfaction to collect viewership patterns. This data would be useful in many ways. Fees in the future could be based on total subscriber viewing hours and/or through a method analogous to web page "hits" where the number of subscribers tuning in is a variable separate from the hours viewed. Cable operators may make programming purchasing decisions based on the actual popularity of the programming. This data also has interesting possibilities in the advertising business.

Busy Signals and Prioritization

The introduction and deployment of Switched Broadcast brings an important new consideration to operators. Given the fact that there will be more input programs available than downstream resources in a service group, this presents the possibility that all downstream resources will be utilized when a Switched Broadcast program is requested.

This risk is mitigated by the fact that in a given broadcast channel lineup, the popularity of programming follows an inverse exponential distribution. In layman's terms, a small handful of programs are watched very often, and a large number of programs are not watched very much at all. It is these programs at the middle-to-far end of the curve that present themselves as ideal candidates for Switched Broadcast.

The risk is further mitigated by the fact that Switched Broadcast is a multicast service; users who watch the same program can share the same stream. This makes a Switched Broadcast implementation very stable in a "flash crowd" environment, where several users suddenly desire to watch the same program.

Furthermore, an SBM implementation can provide a vehicle to prioritize certain programs (or subscribers, for that matter) over others. A set of business rules can be developed, and customized per operator or location, to provide a better assurance that specific programming will be available. Additionally, as necessary, rate shaping techniques can be selectively applied to particular streams.

Switched Broadcast Deployment Topologies

Switched Broadcast can deployed in a variety of configurations, and seamlessly coexists with regular broadcast programs. It can be implemented across an entire headend, or deployed down to the individual service group level. This provides an excellent opportunity to trial its behavior with a small set of programming and grow that programming set over time, and/or subdivide the subscriber group that accesses a common set of downstream resources.

Entry-Level Content Offering Model – Centralized w/SD Rate Shaping

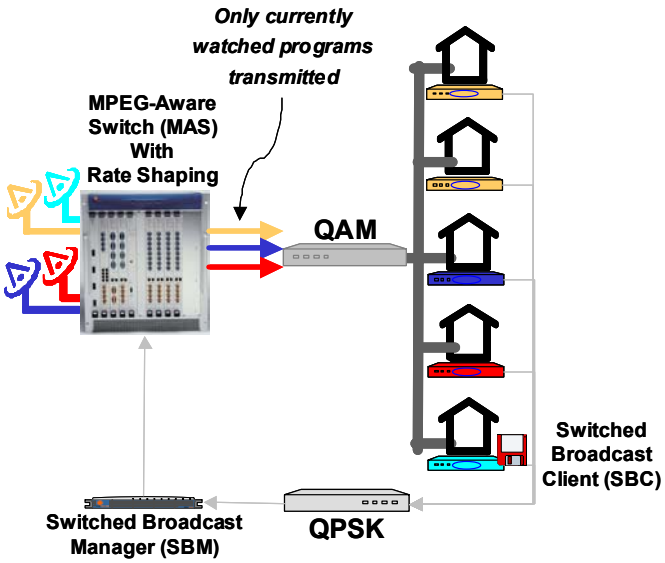

Switched Broadcast – Entry-Level Content Offering Model

In the entry-level Switched Broadcast deployment, the dynamic adding and dropping of programs in a multiplex is combined with a rate shaping engine to produce a solution that allows the seamless replacement of content while maximizing overall video quality.

Large Content Offering Model – Distributed with Acquisition

In larger Switched Broadcast deployments, the number of MASs distributed across hubs presents a more challenging business case to allow rate shaping to be performed on such a wide scale. To optimize the economies of scale, a Broadcast Acquisition system can be located in the headend. The Acquisition system performs a bit rate clamping function, where incoming broadcast programs are rate shaped to a predetermined CBR value. The CBR rate chosen most often is one that mimics the program parameters for VOD streams. Using a Gigabit Ethernet transport structure, the clamped programs are multicast live to all hubs with Switched Broadcast MASs.

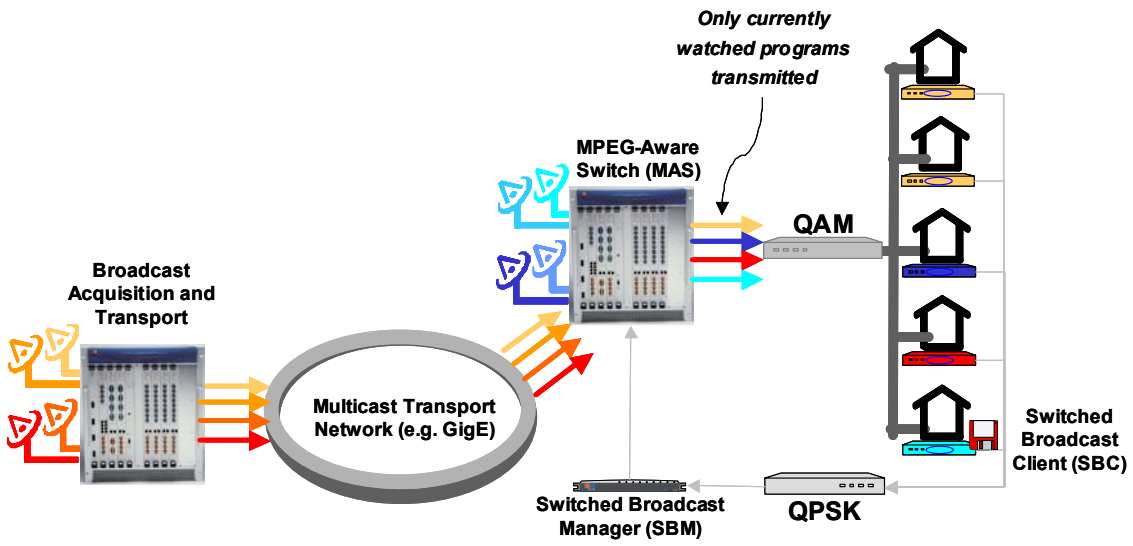

Switched Broadcast – Large Content Offering Model

Employing this model yields a number of advantages. First of all, the rate shaping function is performed once and can be utilized by all hub MASs. Secondly, the new CBR nature of the resulting programs allows for a much simpler downstream resource allocation algorithm. Finally, since the new Switched Broadcast programs "look like" VOD streams, a higher-level bandwidth management function is possible, whereby downstream resources can be dynamically allocated between Switched Broadcast multicast requests, and VOD system narrowcast requests. This allows a larger pool of resources to be created and shared between the two applications, creating greater overall availability, and allowing each service to access a greater pool of resources during their times of peak usage. Tuning latency is also optimized because the addition of a stream to a QAM channel is reduced to a simple TDMA switching operation. It is not necessary to perform further bit rate adaptation for each program stream. As long as the stream is CBR and the rate is known, it is possible to develop allocation rules to optimize the "packing" of streams into each QAM channel.

Bandwidth Management With Other Services

There is a specific benefit that is gained from reformatting broadcast programs into constant-bit-rate streams that "look like" VOD streams. A higher-level resource manager could then manage Switched Broadcast and VOD downstream resources as a common pool, enabling efficiency gains for both services. Again, it is not necessary for each stream to have the same rate cap.

Spectrum Sharing among Switched Broadcast/VOD/xOD

In a VOD deployment, it is common to dedicate a certain number of QAMs to carry VOD streams. The number of QAMs chosen is usually a percentage of the number of digital subscribers served. Whether all, most, some, or none of those VOD QAM resources are

actually carrying VOD traffic, those QAM resources are currently "locked away" and inaccessible by any other service.

Through the use of an intelligent downstream resource manager, the resources required by both VOD and Switched Broadcast can be shared, elevating the paradigm to the next level – switching across services and not just streams within a service. This allows VOD to have access to a larger pool of resources during its peak usage periods, and also allows for Switched Broadcast to access a larger pool of resources during its peak demand periods, which may be more random in their distribution patterns, enabling a more efficient and robust statistical multiplexing effect due to the greater amount of multiplexed content.

Summary

Switched Broadcast represents a shift in the way of considering broadcast service delivery that can cost-effectively enable the next generation of service expansion for operators. By shifting from a content-oriented delivery model, where a cross-section of bandwidth reveals a static flow of information, to subscriber-oriented delivery model where the transmission of programming is a direct function of subscriber demand, operators can scale their delivery platform to track with the number of subscribers served (a pay as you grow model), away from the less-stable model of compounding new programming expenses with additional capital investment to make room the new programming. This growth in transmission efficiency can allow operators to expand their service offering with new revenue-generating and churn-reducing programming, or implement a competitive service offering in systems that (for cost justification purposes) have not enjoyed the benefit of a plant upgrade, or both.

Switched Broadcast can be successfully implemented with niche, low-penetration content, as well as on a mass scale with the entire digital broadcast tier. The dynamic nature of program stream allocation allows Switched Broadcast to be merged with similar services such as VOD to provide an even more efficient provisioning of programming to subscribers. Real-world statistics extracted from ongoing field trials and deployments will provide an even more convincing indication of Switched Broadcast's compelling value proposition.