MODELING SWITCHED BROADCAST VIDEO SERVICES

Joseph Weber, Jiong Gong
Cable Television Laboratories

Abstract

Switched broadcast video services can be used to offer many more broadcast programs, using less bandwidth, than traditional broadcast services. A typical 750 MHz cable plant could theoretically offer over a thousand broadcast digital programs to subscribers, compared to a few hundred programs using traditional broadcast.

We have developed a model to predict average and peak usage of bandwidth for a switched broadcast system given service area size and expected popularity of various programs. This model can be used to plan for the equipment and capacity needed to offer a switched broadcast service.

INTRODUCTION

Cable operators continue to increase the number of broadcast channels they offer their subscribers. In a broadcast network, the number of channels that can be offered is limited, however, by the bandwidth capacity of the last mile of the hybrid fiber coax (HFC) plant to the home. Operators will encounter this limitation as they begin broadcasting more channels, and more high-definition content. While more efficient codecs can increase the number of channels offered to subscribers, with traditional broadcast services the number is ultimately limited by available network bandwidth.

In reality, as the number of programs being offered grows, fewer subscribers are actually watching some of the more obscure programs. A switched broadcast network gives the operator the ability to offer an almost unlimited number of new programs to subscribers by taking advantage of the finite number of simultaneous programs actually being watched by subscribers on an individual network.

In this paper we develop a mathematical model for predicting the network capacity necessary to deliver a given number of video programs to a network of a given size. The model could be used for predicting capacity requirements, or determining when it is economically advantageous to use a switched broadcast service over a traditional broadcast service.

We also extend the model to identify which channels are eligible for adding to a switched broadcast tier, and which should remain on traditional broadcast.

OVERVIEW OF SWITCHED BROADCAST

Traditional broadcast video services transmit all available programs to all subscribers on the network all of the time. Given the variability in content across channels, it is highly likely that during some parts of the day many programs or channels are not being watched by any subscriber. With digital services, the number of possible programs carried by the network increases dramatically, and therefore, so does the probability that programs are being broadcast but no subscriber is actually watching.

In a switched broadcast model, programs are broadcast on the local network only when requested by a subscriber. When a subscriber selects a program for viewing via their interactive program guide, the application determines if that program is currently being broadcast on the local network. If so, the consumer premises equipment (CPE) simply tunes into that broadcast. If the program is currently not on the network, the CPE makes a request to the server application to begin broadcasting that program on the network. When the subscriber switches to another program or is no longer watching the selected program, the server application is again notified. When the server application determines that there are no more subscribers on the node viewing a particular program it can remove it from the network and, thus, free the bandwidth for other uses.

Switched broadcast, therefore, has the potential to reduce the amount of bandwidth required to support large numbers of broadcast programs. With normal broadcast service, the amount of bandwidth required grows linearly with the number of programs. For analog services this requires 6 MHz of spectrum for every program, while digital services require about 0.6 MHz of spectrum per program. With switched broadcast the bandwidth grows much slower with the number of channels.

Like VOD, a switched broadcast video service requires a two-way cable plant. Using their remote controls, subscribers select a program from an interactive program guide (IPG). In switched broadcast services, only those programs that other subscribers on the local network are currently watching are present on the local network. Therefore, a request for a new program not currently on the network must make its way back to the cable headend so it can be added to the local network. At the headend, the new program is loaded onto the local network, and the location of the new program is added to the local channel map.

The Potential of Switched Broadcast Video

The model developed in this paper can be used to estimate the bandwidth required to offer a given number of switched broadcast programs, given the network size. With traditional broadcast service, a 750 MHz plant typically can broadcast about 250 programs (a mix of analog and digital services). Even if all video services are broadcast in digital, the maximum number of programs the operator could offer is about 700. With switched digital services, the model presented here predicts that an operator could offer over a thousand broadcast programs on that same network (See Table 1).

Tier	Present Broadcast Allocation		All-Digital Broadcast Allocation		Switched Broadcast Allocation	
	Programs	6 MHz Channels	Programs	6 MHz Channels	Programs	6 MHz Channels
Basic (always analog)	26	26	26	26	26	26
Standard	50	50	130	13	140	14
Premium	8	8	200	20	500	15
Digital	120	12	340	34	1000	25
HD	8	4	14	7	40	20
Total	212	100	710	100	1706	100
Key:	**Digital**		**HD**		**Switched**	

Table 1. Example Spectrum Allocations for Present, All-Digital and Switched Broadcast Systems

Video-on-Demand (VOD) and network Digital Video Recording (network DVR) are a form of switched broadcast sometimes called unicast or narrowcast. For these services, each subscriber can request a unique program received only by the subscriber. But both services require significant network bandwidth because a stream would be dedicated to only a single user. Switched broadcast, on the other hand, uses a multicast approach. If more than one subscriber has requested a particular program, they all share a single broadcast stream, therefore using the same bandwidth for one viewer as for one hundred. Only in the unlikely case that a large number of subscribers on the network each request a unique program would capacity limitations come into play.

Events such as Thursday night primetime and the Super Bowl, which could cause concerns for a network DVR service because of the large number of users on the system, actually place a lower demand on switched broadcast since most viewers can be served with only a few multicast streams (although users do not have the same pause, rewind, and other recording features available with VOD and network DVR services).

Network Requirements

Implementing a switched broadcast system can require a significant addition of hardware into the network, as well as software in both the CPE and the network. In particular, with traditional broadcast service a single QAM modulator can be used for the entire network. The same QAM-modulated signal is sent to every service group on the network. With switched broadcast, each service group would need its own QAM modulator since the services available on each service group's local network is different depending on usage.

A MODEL FOR SWITCHED BROADCAST

Switched broadcast takes advantage of the fact that the popularity of various programs is non-uniform: some shows are more popular than others. Therefore, in a network with a finite number of active viewers at any particular time, some of the more popular programs will be watched by multiple viewers, while some of the less popular programs will not be watched by any viewers. Thus, if the operator offered m different programs to viewers, most likely less than m of those programs are actually being viewed at any one time. With traditional broadcast, as the operator increases the number of programs offered for viewing, the bandwidth requirement grows linearly with the number of programs regardless of the number of subscribers on the network. With switched broadcast, the bandwidth requirement grows slowly with both the number of viewers on the network and the number of programs being offered. As a trivial example, if the network has only one viewer, then only one program needs to be broadcast. If there are two viewers, a maximum of two, but sometimes only one program needs to be broadcast at any one time. As the number of viewers grows, many will be watching the same popular programs, so the number of programs required will be less than the number of viewers.

We developed a model to estimate the number of simultaneous programs that are being viewed by at least one subscriber, and the amount of bandwidth needed for switched broadcast delivery on a network as a function of the number of programs offered and the number of subscribers on the network. The model uses rating information to identify the relative popularity of each channel on the network. This is combined with the variation in the number of active viewers during the day to produce estimates of the number of unique programs being watched as a function of the number of channels being offered, the number of subscribers on the network, and the time of day.

Program Ratings

The popularity of various cable and broadcast channels clearly changes with time. Weekday primetime might show a large number of viewers tuned to network broadcast channels, while Saturday evenings might have a preponderance of viewers tuned to cable movie channels.

Nielsen uses the following terms for identifying the relative popularity of various channels:

- HUT = number of households watching TV / total number of TV households

The HUT measure varies throughout the day, with a peak around primetime and a low early in the morning. HUT is always a number between 0 and 100%. We designate the variable HUT by $h(t)$ where t is the time of day.

- Rating = number of households tuned to a particular channel / total number of TV households

Rating numbers also range from 0 to 100%. Ratings give the relative popularity of individual channels. They can also be used to estimate the number of households watching a particular channel by multiplying by the total number of TV households. We designate the ratings for a series of m program channels by r_i for $i = 1,2,...,m$. Note that at any particular time $\sum_{i=1}^{m} r_i = h(t) \le 1$.

- Share = number of households tuned to a particular channel / total number of households watching TV

Share numbers also range from 0 to 100%. They indicate what fraction of households actively watching TV are tuned to that particular channel. We designate the shares for a series of m program channels by s_i for $i = 1,2,...,m$. Note that at any particular time $\sum_{i=1}^{m} s_i = 1$ and $s_i = \frac{r_i}{h(t)}$.

Effective Channel Capacity

Given the ratings curve, r_i, and HUT values as a function of time, $h(t)$, we can estimate the number of unique programs being viewed on a particular local network (service group). This model is independent of the method used to obtain the probabilities. They can be either empirical or estimated using a Geometric distribution. Given a service group in a network with n' homes passed, the number of customers within households actively watching television is given by the product of n' and the HUT, or $n = n' * h(t)$. We extend the definition of HUT beyond that used by Nielsen, since a household might have multiple televisions and VCRs, resulting in a HUT above 1.0. Typically, HUT remains well below 1.0 during the early morning hours.

The n customers watching television on the network can choose any of the m switched programs being offered. The key to switched broadcast video economics is the fact that at any one time, many programs are not being viewed by any subscriber. The number of those customers watching a particular channel depends on the popularity, or share, of that channel at that time. In switched broadcast systems, the exact number of customers watching a channel is not important, only if there is at least one customer watching. Although a channel has to be provisioned in real-time even when only one customer is viewing a channel, the probability of such an event occurring is still affected by the popularity of a channel program.

For our purpose, the state of a particular program channel *i*, which we denote as x_i, can be characterized as either on or off, with "on" meaning at least one person is viewing that program channel (requiring a corresponding channel to be provisioned), and "off" meaning otherwise. x_i is then naturally a Bernoulli random variable that can be defined as taking one value [*1*, *0*] with the associated probabilities $[1-(1-s_i)^n, (1-s_i)^n]$. The first probability characterizes the event that at least one customer is watching channel *i*, and the second probability characterizes the event that no one is watching channel *i* at all. As a result, the efficiency of turning the program channel *i* into a switched program channel is entirely driven by the probability of its state being off, which is in turn determined by the channel's popularity, or share, s_i. The total channel usage over all the *m* channels, which we denote as *l,* is then the sum of *m* Bernoulli random variables

$$l = \sum_{i=1}^{m} x_i \tag{E2.2}$$

For sufficiently large *m*, the Central Limit Theorem says that this summation series can be approximated by a Normal distribution.[1]

[1] For a comprehensive treatment of Central Limit Theorem, see for example "*Probability Theory*", by M. Loeve, p 268-383, D. Van Nostrand Company, Princeton, New Jersey, 1963. Central Limit Theorem applies to the mean of a series of random variables. It can also apply to the sum of a series of random variables, provided that the summation series is bounded. Since *m*, the number of program channels in our case, cannot go to infinity, the summation series is thus bounded.

As a result, the actual number of channels needed, which we call *effective channels* and denote it by *q,* is then a selected integer number where the probability of *l* exceeding *q* is pre-specified for this normal distribution. That probability is commonly known as the blocking probability *p*. Under these conditions, the effective channel bandwidth required is as follows:

$$q = E(l) + a\sqrt{Var(l)} \tag{E2.3}$$

where

(1) a is such that $p = 1 - \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{a} e^{-y^2/2} dy$ (E2.4)

(2) $E(l) = m - \sum_{i=1}^{m} (1-s_i)^n$ (E2.5)

(3) $Var(l) = \sum_{i=1}^{m} (1-s_i)^n - \sum_{i=1}^{m} (1-s_i)^{2n}$ (E2.6)

The first part of the equation E2.3 above gives the expected number of channels being watched. Because channel watching is a stochastic process, not only the expected value but also the variation in the number of channels being watched is important. Knowing that total usage follows a Normal distribution allows us to calculate the upper bound of that variation given a pre-specified blocking probability as defined in the second part of equation E2.3.

In practice, (E2.4) can be calculated by using a standard Normal distribution table.[2] For example, setting $a = 2$ yields a blocking probability $p = 2.28\%$. Setting $a = 3$ yields a blocking probability $p = 0.13\%$. Conversely, given a *p* value, *a* can also be easily found from the table.

[2] One commonly used table is from P.G. Hoel, "*Introduction to Mathematical Statistics*, " 4th ed., New York, Wiley, 1971.

The capacity requirement equation can be further simplified if we make the share uniformity assumption that all the program channel shares are the same. For example, suppose $s_1 = s_2 = \cdots s_i = s_m = \frac{\mu}{m}$, where μ represents the total share of these m programs. Then the above capacity equation can be simplified to:

$$q = m[1-(1-\frac{\mu}{m})^n] + a\sqrt{m(1-\frac{\mu}{m})^n[1-(1-\frac{\mu}{m})^n]} \quad \text{(E2.7)}$$

The share uniformity assumption should be a good approximation for several reasons, aside from the benefit of simplified calculation. First, the tail of the program share distribution curve, which most likely represents the channels which are candidates for switching, is indeed close to forming a uniform distribution in reality. Second, E2.7 gives the upper boundary of the expected usage of channels regardless of the share distribution of the viewing profile. Consequently, this approximation also gives the least amount of bandwidth savings in a switched broadcasting environment. Or stated differently, this approximation provides the most conservative estimate to channel capacity provisioning, especially useful in the absence of accurate data on program share numbers.

Expected Channel Usage

Expected channel usage plays a dominant role in the effective channel capacity requirement equation. Generally, as we increase the number of switched channels, the bandwidth efficiency gain from switching tends to increase. Figure 1 below shows the expected number of unique channels being viewed with 250 and 500 subscriber groups as a function of the number of channels offered. As expected, the number of actively viewed channels grows slowly as the number of channels increases beyond the number of subscribers. However, the concave nature of these curves means that capacity requirements do not increase linearly with program channels, but at a much slower rate, especially when the service group size is small. As a result, the bandwidth efficiency gain increases as the number of switched program channels increases. Traditional broadcast bandwidth requirements are linear in the number of channels offered and are shown for comparison.

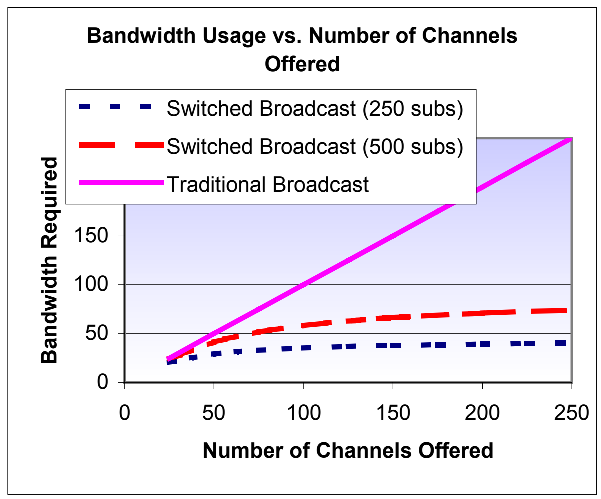

Figure 1 Number of Channels Broadcast on Switched and Traditional Broadcast Services as a Function of the Number of Channels Offered

We would expect that the number of channels being actively watched would grow along with the number of subscribers on the network. This is because the probability of the state of a program channel being "on" increases exponentially as the size of the service group increases. The number of independent channels being watched will also depend on the shape of the program ratings curve. In Figure 2 we graph the expected number of channels as a function of the number of subscribers on the network for the average 'Primetime' ratings curve, and a typical Monday night football curve. The number of channels actively being watched approaches the total number of 200 channels as the service group size increases.

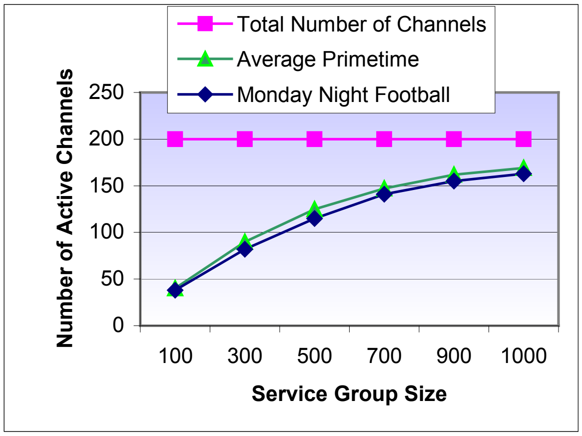

Figure 2 Number of Channels Actively Being Watched as a Function of Service Group Size during Primetime and Monday Night Football Nights

Bandwidth Savings

The primary benefit of switched broadcasting video is saving bandwidth so that more program channels can be provisioned. Figure 3 shows the extent of bandwidth savings based on our effective channel capacity requirement, as a function of service group size and the number of channels switched, respectively. The percentage of bandwidth savings is defined as one minus the ratio between the provisioned effective channel capacity and the number of actual available program channels. For example in Figure 3 when 100 program channels are switched for a service group size of 500 customers, percentage of bandwidth savings reaches 36%, meaning that only 64 channels are needed to offer 100 programs. If the service group size is further reduced to 250 customers, percentage of bandwidth savings increases to 58%, meaning that only 42 channels are needed for offering 100 programs. Here we assume 1% blocking probability, 25% total viewership share for the switched channels, and HUT=0.6, which is a typical weekday primetime HUT value.

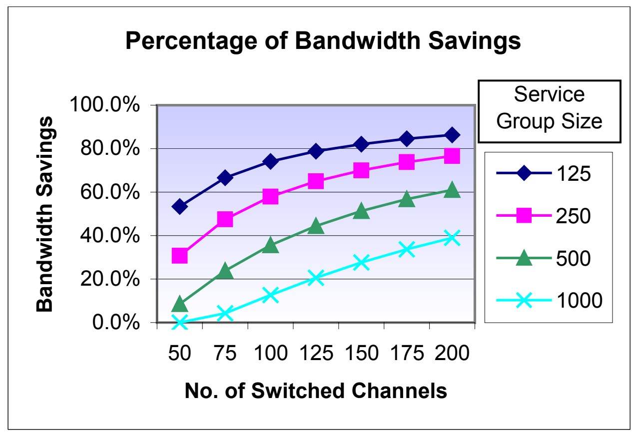

Figure 3 Bandwidth Savings Versus Service Group Size

Figure 4 shows the extent of bandwidth savings as a function of variations in the probability of blocking, while assuming service group size is 500. As can be seen, there is not a large effective channel capacity variation between 1% and 2% blocking probability. But achieving 0.1% blocking requires considerable more capacity.

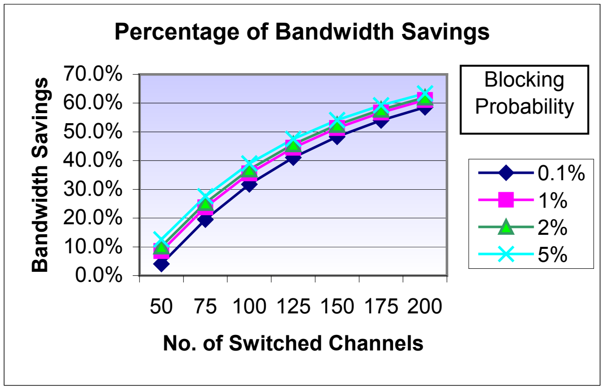

Figure 4 Bandwidth Savings Versus Blocking Probability

Adding a Time Element

The HUT value varies throughout the day, peaking usually during primetime. We can model the variation in HUT as a sinusoid with a period of 24 hours, and peaking around 10 pm. Adding the time-varying HUT to our model, we can simulate a node throughout the day.

It is, therefore, possible to dynamically predict the peak usage as a function of time during the day, and dynamically allocate bandwidth resources appropriately. This would allow a switched-broadcast tier to use even less bandwidth during off peak times freeing bandwidth for other uses.

OPTIMAL DEPLOYMENT MODEL

Not all channels are candidates for inclusion in a switched broadcast service. Deploying a switched broadcast service is more expensive than traditional broadcast because it requires additional QAM modulators and two-way CPE. With traditional broadcast, a single QAM device can serve an entire cable system, while with switched broadcast each service group needs a dedicated QAM device. In this section, we create a model for determining which channels are candidates for being included in a switched broadcast service.

Selection of Switched Channels

We assume adding switched broadcast capability to a cable network incurs a fixed cost in terms of additional QAM modulators, new software on existing digital set top boxes (STBs), and switching controllers within the network. The total additional system cost for a service group is divided by the number of customers within that service group to obtain a per-subscriber cost. We assign the variable C as the cost per subscriber of adding a channel to the switched broadcast service.

Moving a group of channels to a switched broadcast service frees up bandwidth that can be used for other services, including additional channels. We assign the variable R to the additional per-subscriber revenue that can be obtained by the additional free spectrum. This can also be seen as an opportunity cost for not freeing up the spectrum by remaining on a traditional broadcast service.

The model states that the number of channels freed by using a switched broadcast service is:

$$m-q=m-E(l)-a\sqrt{Var(l)}=\sum_{i=1}^{m}(1-s_i)^n$$
$$-a\sqrt{\sum_{i=1}^{m}(1-s_i)^n-\sum_{i=1}^{m}(1-s_i)^{2n}}$$

(E3.1)

From this equation the incremental savings in bandwidth of adding program channel k to a program portfolio consisting of $k-1$ switched program channels is:

$$(1-s_k)^n - a\{\sqrt{\sum_{i=1}^{k}(1-s_i)^n - \sum_{i=1}^{k}(1-s_i)^{2n}} - \sqrt{\sum_{i=1}^{k-1}(1-s_i)^n - \sum_{i=1}^{k-1}(1-s_i)^{2n}}\}$$

(E3.2)

The first term above captures the first order difference in usage of moving from a portfolio of k program channels to that consisting of $k-1$ program channels. That is essentially the expected usage for program k. The second term in parenthesis above captures the second order difference. That is the difference in standard deviation of usage between k program channels and $k-1$ program channels. In reality, the second order difference is actually very small and can be ignored. Under that assumption, the incremental savings in bandwidth of adding a program channel k can then be simplified as $(1-s_k)^n$.

When the incremental revenue, in terms of saved bandwidth, exceeds the incremental costs of adding a channel to switched broadcast service, then that channel should be added to the switched broadcast service. This is equivalent to

$$(1-s_k)^n R > C \text{ or } (1-s_k)^n > C/R \qquad \text{(E3.3)}$$

If the s_k are rank ordered from largest to smallest, $(1-s_k)^n$ is an increasing function of k. Therefore, at some channel K, the value of $(1-s_K)^n$ will become larger than C/R, meaning that all channels from K to m should be placed on the switched broadcast service, while channels 1 to $K-1$ should remain on traditional broadcast tier. Channel K is the first channel where $s_k < 1-(C/R)^{1/n}$.

Note that since every s_k must be between zero and 1, adding switched broadcast service is only viable if $C/R<1$, that is if the incremental revenue of the freed bandwidth exceeds the costs of creating the service.

<u>Partitioning Channels Example</u>

As an example of how channels can be partitioned, we take the model approximation of prime-time share values using a Geometric distribution and assume a 250-subscriber service group. QAM device costs are set at \$600 for the service group, and hardware and software costs are set at \$3 per subscriber. Consequently, the value of C is \$600/250 + \$3 = \$5.40 per subscriber. These numbers are for illustrative purposes only and do not represent actual pricing.

The present value of all future cash flows for a single channel worth of bandwidth is taken be to \$8, which is equivalent to \$0.80 per year revenue stream discounted at 10%. The value of R is therefore \$8. Again this number if for illustrative purposes and does not represent any typical expected revenue for bandwidth.

Given these values of C and R, equation E3.3 predicts that channels 1 through 37 should remain on traditional broadcast, while channels 38 and higher should be put on the switched broadcast service. Figure 5 illustrates the result.

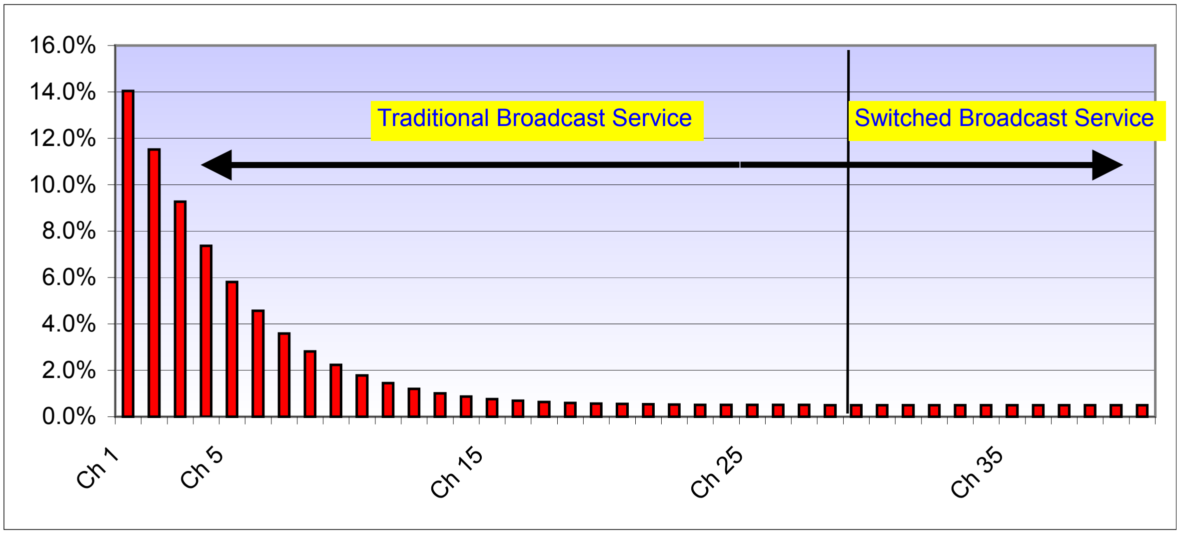

Figure 5 Identification of Which Channels Should be Migrated to Switched Broadcast Service Based on Economic Model

CONCLUSIONS

In this paper we developed a model for estimating channel capacity requirements for a switched broadcasting video service. This model uses the expected household television viewership and program share data from Nielsen to approximate the likelihood of cable customers tuning to a particular channel. Given a desired level of blocking probability, the total channel capacity requirement is then derived as a function of the service group size and the number of channels on the switched service. This model can be further used to derive the optimal number of program channels to be switched, given a valuation of the bandwidth savings as a result of moving some channels from traditional broadcast to switched broadcast tiers. We further validate our model against simulated empirical data from a real cable system. The data collected from the live system indicates that the model appeared to predict an upper bound on the service usage, and the bandwidth required to ensure minimal blocking of the service.

Switched broadcast video can be a valid element of the next generation cable network architecture.

Our model appears to provide an accurate estimate of the amount of bandwidth that can be saved by adding programs with low viewership shares to switched broadcast service instead of traditional broadcast. Because typical channel lineups have 75 to 80% of viewership share concentrated in the top 30 channels, the remaining channels can have very low viewership shares. As a result, many of the channels with low share are not being watched for extended periods of time. Switched broadcast can be used to deliver many of those remaining channels using significantly less bandwidth than traditional broadcast methods. If the saved bandwidth can be used for other revenue generating services, switched broadcast can economically offer many more channels to subscribers.

Bandwidth savings can be at least 36%, and go over 60% for a service group size of 500 digital subscribers

The channel capacity model we developed provides a very conservative upper bound of channel usage given a blocking probability. Even under conservative assumptions and blocking probabilities at 1%, bandwidth savings from a switched broadcast service can be substantial. In a realistic example of a service group size with 500 digital subscribers, bandwidth savings during peak hours can be 36% when 100 channels are placed on the switched tier, and go over 60% when 200 channels are switched. By reducing the service group size, bandwidth savings can be even higher.

Bandwidth savings is an increasing function of the number of switched channels, and a decreasing function of the service group size

Bandwidth savings in a switched broadcast environment is driven by the service group size, and the number of switched channels. In one example where 100 channels are switched, bandwidth savings jumps from 36% to 58% as the service group size is reduced from 500 subscribers to 250 subscribers. However, reducing service group size incurs a cost in terms of more QAM equipment and possibly the need for node splitting.

The efficiency gain from switched broadcast video also increases as the number of switched program channels increases. The more channels in the service, the greater the probability that some of those channels are not being watched. However, efficiency gain tends to diminish marginally as we move up the rating curve to include more popular channels into the switched channel lineup.

Switched broadcast video provides a strategic competitive advantage over DBS providers in terms of adding program channels and efficiently using bandwidth

Switched broadcast video is an ideal architecture for broadcasting hundreds of new program channels if each of those channels has a small rating share. The switched broadcast architecture may be viewed as an interim step towards full on-demand television.

Switched broadcast is a competitive advantage for cable, due to the interactive nature of the cable network architecture. It is unlikely to be matched by DBS providers.