UNICAST VIDEO WITHOUT BREAKING THE BANK: ECONOMICS, STRATEGIES, AND ARCHITECTURE

S.V. Vasudevan, Xiaomei Liu, and Robert Kidd
Cisco

Abstract

Driven by competition and consumer demands, linear video delivery is following a trajectory from broadcast to multicast and ultimately to unicast. Traditionally, video delivery has been broadcast only. Today, cable operators are deploying switched digital video (SDV), which uses multicast technology to improve the bandwidth efficiency of HFC networks. The next logical progression to unicast delivery is on the horizon and is positioned to become tomorrow's video delivery mechanism.

Unicast delivery of linear content is an incremental extension of the multicasting approach used in SDV implementations. The incremental investment in bandwidth resources to support unicast delivery can be offset by the contribution of preferentially valued advertising opportunities, reduced subscriber churn, and the ability to attract new subscribers through differentiated service offerings.

This paper analyzes the unicast value proposition, including cost, revenue potential and return on investment. SDV field trial viewership statistics will be reviewed, and used to shed light on the cost sensitivities related to channel popularity and HD penetration. Best case and worst case scenarios for HFC bandwidth consumption will be explored and analyzed, along with the cost structures associated with each of them. Cost mitigation and revenue improvement strategies will be explored, demonstrating how cable operators can optimally combine unicast and multicast approaches in order to maximize overall return on investment.

Based on the results of this analysis, a switched arch#itecture will be presented for cable operators to smoothly migrate their networks to support unicast delivery mechanisms for linear video services. The proposed architecture accomplishes the strategies for cost-effective unicast delivery and supports:

- *A flexible combination of multicast and unicast delivery mechanisms*
- *Traditional ad insertion based on geographic ad zones or a new generation of targeted ad insertion based on demographic profiles*
- *Fast channel change and personalization of unicast content*

INTRODUCTION

Switched Digital Video (SDV) is now a mainstream technology that is delivering on its promise to dramatically improve upon the bandwidth efficiency of the traditional linear broadcast model. Aggressive SDV oversubscription ratios (the ratio between the number of SDV programs offered and the number of stream resources provisioned) have been observed and will continue to increase as more niche and HD content is added to cable MSO service tiers. Switching technology has proven to be a powerful addition the cable operator's bandwidth management capability.

This material was originally presented at NCTA, The Cable Show, 2008. Reprinted with permission.

Yet there remains unlocked potential within the SDV infrastructure. Current generation systems support open standards, allowing the insertion of new technologies and applications. The session and resource managers (SRMs) that respond to subscriber channel-change requests provide a level of intelligence and network awareness that was previously unavailable. SDV systems maintain a real-time accounting of programs being viewed as well as the number of subscribers viewing those programs. This awareness of program usage comes concomitant with knowledge of bandwidth allocation. This knowledge is powerful, for even with existing switched multicast, there remains significant room for improvement in bandwidth utilization. Furthermore, *Switched Unicast*, an advanced form of SDV, is drawing increasing interest. This emerging SDV architecture offers exciting opportunities to introduce new revenue-generating services, but it also has the potential to overwhelm available access bandwidth. Only with the knowledge of user demand provided by advanced SRMs is this latest architectural challenge tractable.

It is first helpful to clarify the requisite terminology. *Switched multicast* refers to a video delivery architecture where an MPEG program, typically in the form of a single program transport stream (SPTS), is IP-encapsulated and transported on a distribution network via IP multicast. A system session and resource manager (SRM), acting upon channel change requests from subscribers, may then instruct an IP-attached edge QAM to join the multicast. The edge QAM rebuilds a multi-program transport stream (MPTS), containing content requested by multiple viewers, and modulates the content onto the HFC network. The key differentiator of this approach is that multiple set top boxes within a service group can share a stream that is active within that service group; if N viewers within a service group watch the program *MTV*, only one instance of *MTV* is switched into that service group. Thus, the content is delivered over the IP network using IP multicast and over the HFC plant with a stream sharing, RF multicast mechanism resembling that on the IP network.
Switched unicast refers to a delivery mechanism in which, regardless of the IP transport and routing mechanisms deployed, the stream on the HFC side of the plant is destined for a single tuner within a single set top; i.e. set tops within a service group no longer share MPEG streams; if N viewers within a service group request *MTV*, N instances of *MTV* are switched into that service group. Typically the transport mechanisms on the IP network will also be via IP unicast, but hybrid solutions can be envisioned in which, for example, ad servers at the edge join multicast IP content, insert a targeted ad, and then send the new stream via IP unicast to the edge QAM.

The drivers for deploying switched unicast are compelling. Switched unicast offers the opportunity to personalize video. Since each tuner now receives an individual video stream, media processing techniques can be used to modify the stream to suit the preferences of an individual subscriber. These modifications may include graphics overlays on the screen that are tailored to the subscriber's preferences. For example, a ticker that displays preferred stock quotes, sports scores, localized weather information, etc.

But perhaps the most compelling driver for switched unicast is in the ability to personalize advertising. The North American cable industry has a long history in spot advertising as a revenue source, and many systems perform a limited level of localization by dividing a cable system into zones and offering spot insertion on a zone by zone basis. However, the ability to transcend beyond zones and offer advertising on a *personalized* basis offers MSOs the opportunity to charge a premium for these slots – indeed the success of Google in its ability to personalize advertising in the online space has made all participants in the advertising delivery chain sit up and take notice. Furthermore, with personalized advertising, ads need not be restricted to 30-second spots – they can additionally take the form of graphics logo overlays, with additional possibilities enabled from the incremental ability to launch interactive ads.

This raises an interesting question which is the subject of this paper – switched unicast has the ability to increase MSO revenue through the incremental benefits of personalized advertising, yet to reap this benefit requires an investment in delivery infrastructure to enable these personalized capabilities. Is the investment worth it?

This turns out to be a challenging question to answer, as the answer depends on a number of variables, including subscriber viewing patterns, service group sizing, equipment costs, and the anticipated premium to be expected from placing a personalized ad as opposed to a zoned ad. To better understand this topic, we propose to study the subject from 3 perspectives: analyzing data from a real SDV deployment to understand multicast and unicast resource sensitivity based on program popularity, service group size, and other factors; developing a financial model of switched unicast, which allows ROI to be analyzed based on a number of factors; and exploring alternate models of SDV that enable unicasting on an opportunistic basis.

SUBSCRIBER VIEWING PATTERNS AND SWITCHED UNICAST

It is difficult to make purely analytical predictions regarding SDV efficiency, since efficiencies ultimately depend on subscriber viewing patterns, which in turn is driven by the behavior of human beings. However, by analyzing the pattern of channel change messages from SDV, it is not difficult to infer what the expected system behavior would be if the system were unicast instead of multicast.

To better understand this relationship, viewership information was extracted and analyzed as part of a SDV trial with a major North American MSO. By aggregating and post-processing channel change information from server log files, it is analytically straightforward to determine not only the relative difference in resource requirements between multicast and unicast, but it is also possible to analytically modify the "virtual" size of the service group to better appreciate the sensitivity between service group size and resource utilization for both the multicast and unicast scenarios.

The viewership study was conducted over several weeks and included channel-change data for 247 broadcast video services comprising 228 standard definition programs and 19 high definition programs delivered to 680 tuners in a service group. The study included four major steps: (1) computation of the viewership long-tail; (2) segmentation of the long-tail into popularity quintiles; (3) segmentation of the settops into "virtual service groups", and (4) computation of unicast (total) and multicast (unique) streams required to deliver cumulative quintiles of programming to a range of service group sizes.

The concept of a "virtual service group" warrants some discussion. One important result desired from the analysis was the expected variation in video stream resource requirements with service group size. In most production systems, a target service group size is established and the inside and outside plants designed to that target size. If an analysis is performed only with the existing service group structure, bandwidth requirements can be predicted for only a very narrow range of service group sizes. Therefore, instead of using the existing system's service group structure, the nodes were regrouped into sets of virtual service groups containing tuner counts that span the range of interest.

Figure 1 provides viewership results for one week of data and illustrates the classic long-tail viewership phenomenon. The graph is generated by summing the number of seconds viewed for each broadcast video program and then ranking the channels in order of decreasing total viewership [1]. Once the programs are ranked, the viewership curve is segmented into five sections, each of which includes 20% of the offered programs. These quintiles are used to evaluate stream requirements as a function of program popularity.

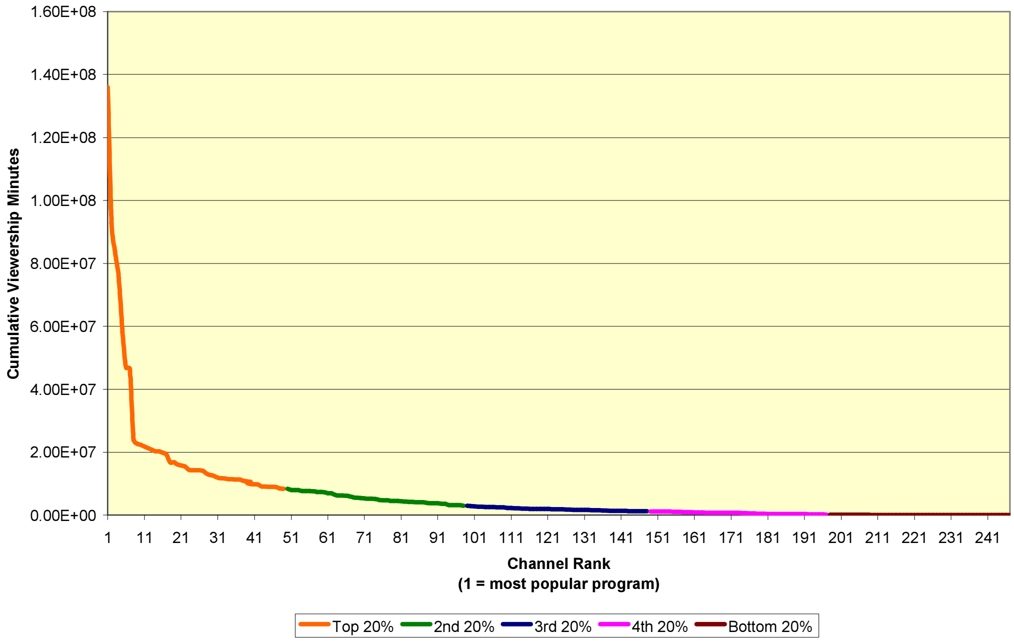

Figure 1. Program Viewership Ranking

Figure 2 illustrates switched multicast stream usage and displays the number of peak unique simultaneous streams required as a function of service group size and popularity of content. The horizontal axis displays the number of tuners per service group, and the vertical axis displays the number of peak unique simultaneous streams. Five curves are included on the plot, each of which illustrates peak stream requirements for a particular grouping of content: the bottom curve illustrates the peak unique streams required for the least viewed 20% of content; the next-to-bottom curve illustrates peak unique streams required for the least viewed 40% of content; and so on up to the top curve that illustrates the peak unique streams required for the entire broadcast video lineup.

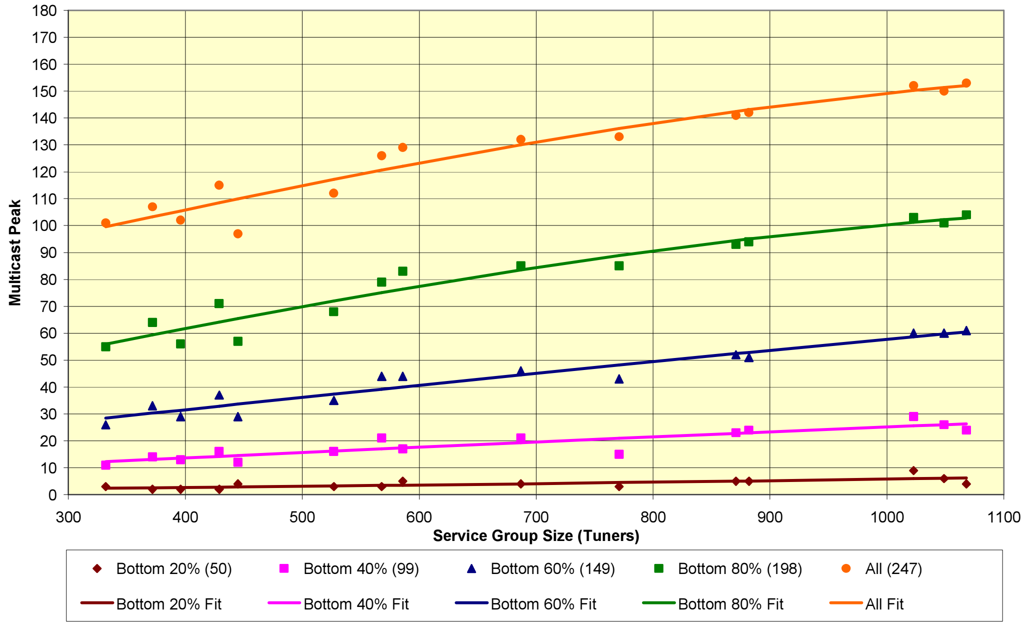

Figure 2. Switched Multicast Peak Stream Usage

Each curve is generated from a regression analysis of the raw data points, also included on the plots, that result from processing each virtual service group for a specific grouping of content. Two-sided 95% prediction intervals (not shown in the figure) were also computed. Future individual peak stream count results are expected to fall below the upper bound of this prediction interval 97.5% of the time.

Figure 3 is the unicast equivalent of Figure 2. In this case, the vertical axis displays the number of peak total simultaneous streams as opposed to the number of peak unique simulcast streams; otherwise, the data analysis is equivalent to that described above for the multicast case.

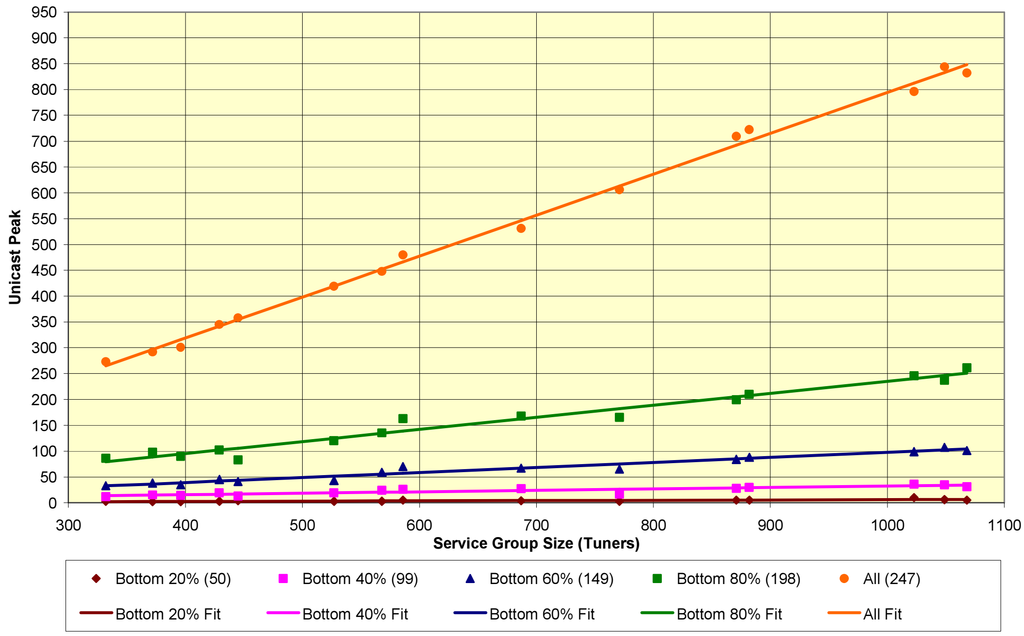

Figure 3. Switched Unicast Peak Stream Usage

Figure 4 summarizes the stream count dynamics between unicast, multicast, and basic broadcast for the entire 247-program lineup.

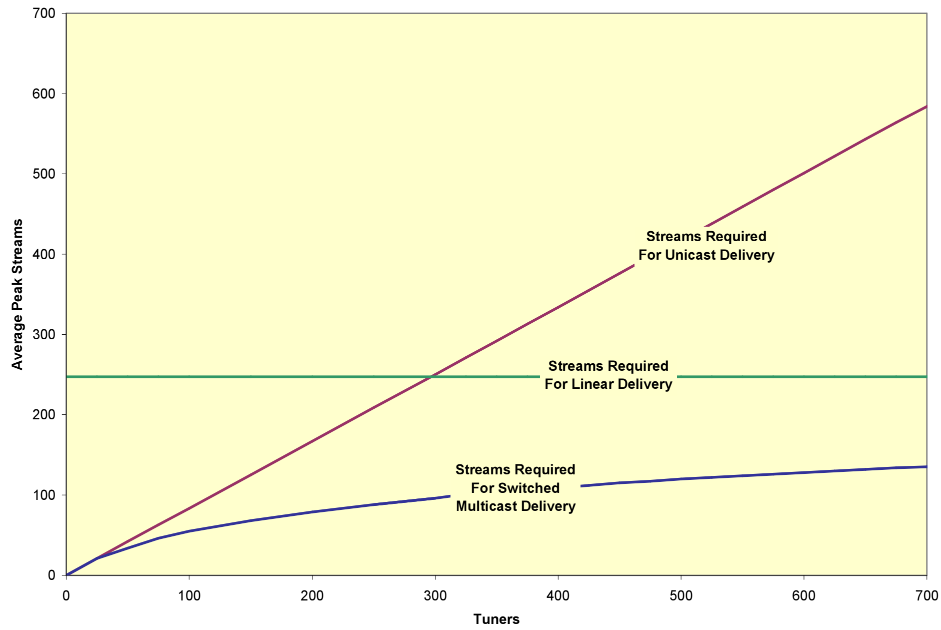

Figure 4. Stream Requirement Overview

As expected, the different delivery mechanisms require dramatically different HFC bandwidth allocations. The number of streams required to deliver the lineup via traditional linear broadcast is of course constant and therefore independent of service group size. The number of streams required to deliver the same lineup via switched multicast is significantly less since settops may share streams within an HFC service group. The number of streams required to deliver the lineup using switched unicast is linearly proportional to service group size. This latter viewership curve is similar to that for VOD except that the slope of the unicast demand curve for broadcast content is much steeper than that experienced with VOD.

For delivery of the most popular content to larger service groups, unicast requires significantly more bandwidth than not only multicast SDV but also simple linear broadcast. For example, in order to offer the entire broadcast video lineup to today's typical 500 tuner service group, the unicast model of Figure 4 requires approximately 417 peak streams, the multicast model requires approximately 120 peak streams, and the broadcast model requires 247 streams. Making the simplifying assumption of a 50/50 SD/HD split for the unicast streams, approximately 104 QAM carriers would be required to carry the unicast content.

Clearly the capacity to unicast the entire lineup is not available on a typical service group in today's hybrid digital/analog HFC systems; however, there are a number of revenue-enhancing opportunities that can be supported today by the surgical insertion of unicast technologies. These surgical deployments can be more fully developed as the industry continues the inexorable push toward smaller service groups and increased digital delivery. Service group sizes are trending towards a future where 250 tuner service groups will be the norm, and only 209 streams will be required to unicast the entire the 247 channel broadcast lineup, fewer than the number required for simple linear broadcast.

One key caveat should be raised at this point regarding unicast: viewership statistics are nonstationary from a statistical standpoint, that is, they change with time, and this fact has significant practical implications. As an example, consider the case of a weather or news channel, a service with average viewership sufficiently far down the long-tail to support its inclusion into a switched tier. As time goes by, a hurricane or other major news event will inevitably emerge, and the popularity of this previously moderately-viewed programming skyrockets. If the channel is offered on a multicast tier, viewers that flock to the channel share a single stream; however, if the channel is offered on a unicast tier, viewers receive their own stream, and the required edge bandwidth mushrooms beyond that which may have been predicted based upon prior viewership studies. A unicast tier is much more sensitive to the choice of selected content than a multicast tier and is therefore less stable from a bandwidth planning perspective. In order to mitigate this risk, a key potential feature of a unicast system would be the ability of the system to automatically promote and demote between the unicast and multicast tiers.

Finally, in the above discussion stream counts are used as a proxy for bandwidth requirements and the two are often used interchangeably. However, if different streams have different bandwidths, aggregate stream counts and their associated total bandwidths may not be directly proportional. For example, if viewership of HD services is consistently higher than that for SD services, the most viewed 20% of content may require substantially more bandwidth than the stream estimates alone would predict. In the system under consideration, 19 of the 247 broadcast video programs offered are HD. The percentages of these HD services in the 20%, 40%, 60%, 80%, and 100% quintiles of content were 0%, 21%, 42%, 11%, and 26% respectively. Thus the spread of HD content is slightly skewed towards the most viewed groupings; however, the effect of this skew on bandwidth requirements is muted given the relatively small number of HD programs offered at the time of the study.

SWITCHED UNICAST ARCHITECTURE AND ROI MODELING

It was previously mentioned that perhaps the incremental revenue generated from switched unicast could fund the investment in the necessary delivery infrastructure. In reality, this is a complicated problem with numerous factors contributing to the analysis. Before going into details of ROI analysis, we will start with some basic assumptions of a switched unicast architecture.

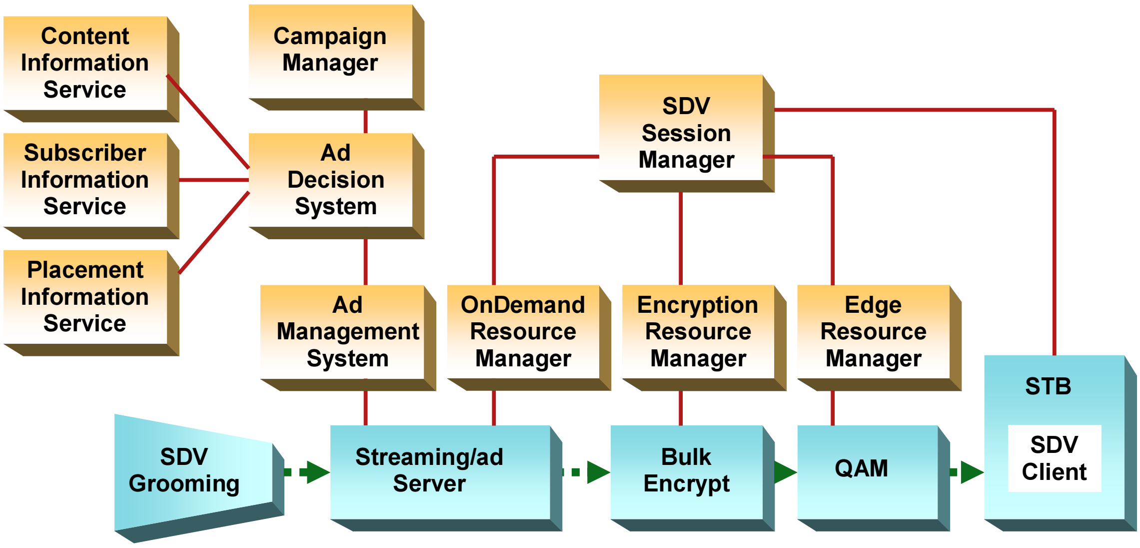

Figure 5. Switched Unicast Architecture

A switched unicast system architecture builds upon the existing, widely deployed switched multicast architecture and represents an evolutionary path. Existing components may be retained and augmented, minimizing the incremental investment. In the example switched unicast architecture shown in Figure 5, switched video content flows from a SDV groomer (performing VBR-CBR rate clamping) to STBs passing through streaming/ad servers, bulk encryptors and QAMs. The key difference from traditional switched digital video is the introduction of a streaming/ad server in the data path. The streaming/ad server is defined as a component that constantly ingests live linear content and streams it out as requested. The streaming/ad server also detects ad placement opportunities and splices ads. In the control plane, the SDV session manager manages SDV channel changes and orchestrates SDV resource allocations through various resource managers. The targeted advertisement control plane components are shown and are defined by the SCTE-130 specification. An ad decision system makes an intelligent decision on which ads to insert for an ad placement opportunity based on information from the functional components labeled content information service, subscriber information service and placement opportunity information service. The streaming/ad server detects placement opportunities, notifies the ad decision system about the placement opportunities via the ad management system and inserts ads based on decisions made by the ad decision system.

An analytical model representing switched unicast sizing and pricing was developed to evaluate optimal dimensioning parameters for future deployments, and quantitatively examine targeted advertising development opportunities that could provide an attractive return for the infrastructure investment.

The analytical model for switched unicast considered the following parameters on the "expense side", shown in Table 1. The graphs that appear later are computed from the default values indicated in the tables.

Cost Modeling Parameters	Default Value
Service group size	1000 HHP
Subscriber penetration	60%
Digital penetration	60%
Tuners per household	1.8
Total channel offered	249
HD channel percentage	20%
Channel popularity	Extracted from trial log data
SD channel bandwidth	3.75 Mbps
HD channel bandwidth	15 Mbps
Spectrum available for SDV	30 RF channels
QAM channel bandwidth	38.8 Mbps
AVC STB penetration	20%
Node split cost	$15,000
Transport cost	$9/Mbps
QAM cost	$13/Mbps
Streaming server cost	$15/Mbps

Table 1. Cost Modeling Parameters and Default Values

The following parameters were considered on the "revenue side", shown in Table 2.

Revenue Modeling Parameters	Default Value
Ad revenue per subscriber	$300
MSO per sub ad revenue	$60
Possible ad revenue share with programmers and networks	30%
Targeted ad percentage	30%
CPM improvement of targeted ads	2
Targeted advertising operating margin	80%

Table 2. Revenue Modeling Parameters and Default Values

The strategy for the analysis was as follows: using the number of channels, the extracted program popularity, and the bandwidth of the programming, the total aggregate bandwidth of the expected unicast streams was calculated. Based on this information, the optimal service group size was calculated. If this optimal value was less than the existing service group size, the model would factor in the price of node splits to compute the infrastructure costs to achieve the proper service group size. Once this value was known, industry-current figures for QAM, streaming server, and transport costs (normalized to a $/Mbps factor) were used to calculate expected investment costs.

To migrate to switched unicast, operators incur both a data plane cost and control plane investment cost. The data plane costs largely consist of capital equipment (QAMs, streaming servers, etc.) The control plane costs largely consist of software enhancements to existing SDV server platforms to enable unicast signaling. Taking advantage of the fact that control plane components such as the SDV session managers and resource managers are already essential ingredients of switched multicast video services, the additional cost for providing switched unicast control plane infrastructure is not significant. For our analysis, the cost of SCTE-130 targeted advertisement control plane components was factored in with an advertising operating margin of 80%. Major data plane expenditures

came from spectrum, QAMs, streaming servers and other transport costs.
Comprehensive analysis has been done evaluating the cost and benefits of various technologies that can squeeze more bandwidth out of the cable plant [2]. Among 1 GHz node upgrades, all digital conversions, node splits, advanced video coding, 1024QAM etc, node splits stand out as the attractive approach with an appealing cost to benefit ratio. In our cost analysis, we assume the plant starts with a 1000 HHP service group size, which is the equivalent of 500 tuners per service group with 60% subscriber penetration and 60% digital penetration. Assuming a maximum 30 RF channels are available for digital linear video service in a service group, node splitting to smaller service groups will be used until there is enough spectrum for providing the switched unicast service. The estimated cost of node splitting is $15,000 per split.

Although advanced video coding saves 50% bandwidth when compared with MPEG-2 video, it is not currently widely used in the linear video service because of the broadcast nature of the service today. A lowest common denominator is chosen in terms of STB capabilities. In other words, as long as legacy MPEG-2-only STBs are in the field, the video must be offered in the MPEG-2 format. This situation is changed with the introduction of switched unicast. With switched unicast, AVC coding can be used when delivering unicast video to newer AVC-capable STBs, while the MPEG-2-version of the program can be used when delivering to legacy MPEG-2-only STBs. In the cost analysis, we assumed 20% of STBs as AVC capable.

The result of per subscriber cost analysis is shown in Figure 6.

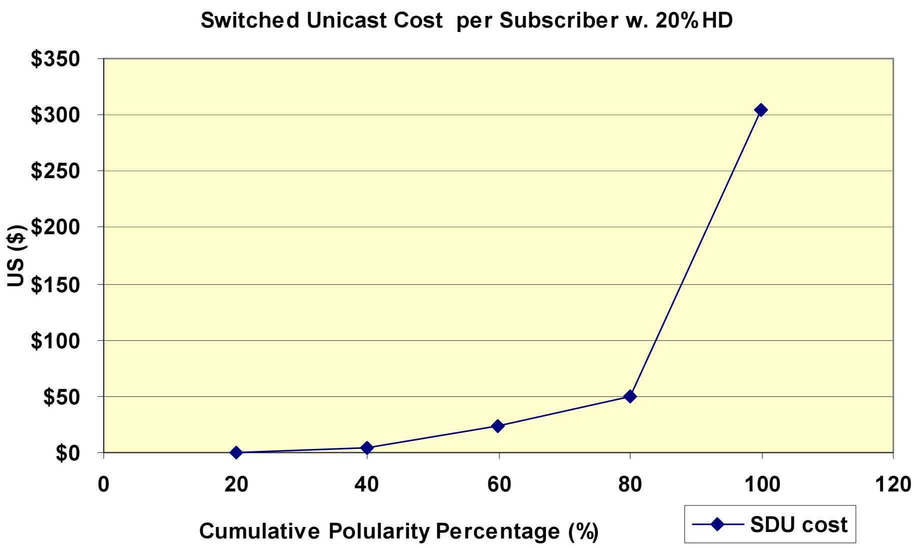

Figure 6. Switched Unicast costs per popularity quintile

As one would expect, switched unicast per subscriber costs increase as more popular content is offered as unicast. The biggest jump comes when the top 20% most popular channels are offered as unicast.

Targeted advertising is often touted as the next revenue growth engine for cable MSOs. Switched unicast is an enabling platform for targeted advertising. In 2006, worldwide cable TV advertisement spending totaled $24 billion, with nearly $5 billion contributing to MSO revenues. With only a 20% share of the total spend, targeted advertising could be an effective vehicle to improve this number by providing a higher-value product. [3]

When estimating the ad revenue potential of switched unicast, we apply the advertisement CPM (cost per thousand impressions) improvement of 2x for targeted advertising. Currently, in a typical cable network [4], per subscriber advertising revenue is $60. Since the MSO's share is just 20% of per subscriber ad revenue, the total per subscriber ad revenue is calculated as $300. The ad revenue modeling assumes that 30% of placement opportunities in switched unicast are targeted and assumes that MSOs have a 30% placement opportunity split with broadcasters and cable programmers for non-local ads. Then the potential MSO per subscriber ad revenue increase can be derived. However, the uneven distribution of unicast viewers with regard to the channel popularity complicates the calculation. The additional ad revenue of a node is calculated next as the potential per subscriber ad revenue times the unicast viewers in the node. Lastly, this ad revenue of the node is averaged to compute the switched unicast ad revenue per subscriber. The projected ad revenue is shown in Figure 7.

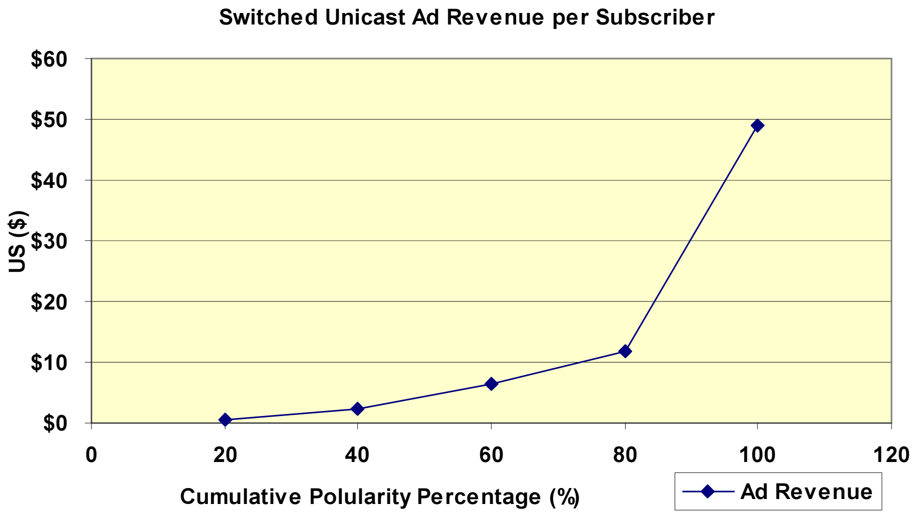

Figure 7. Switched Unicast potential revenue per popularity quintile

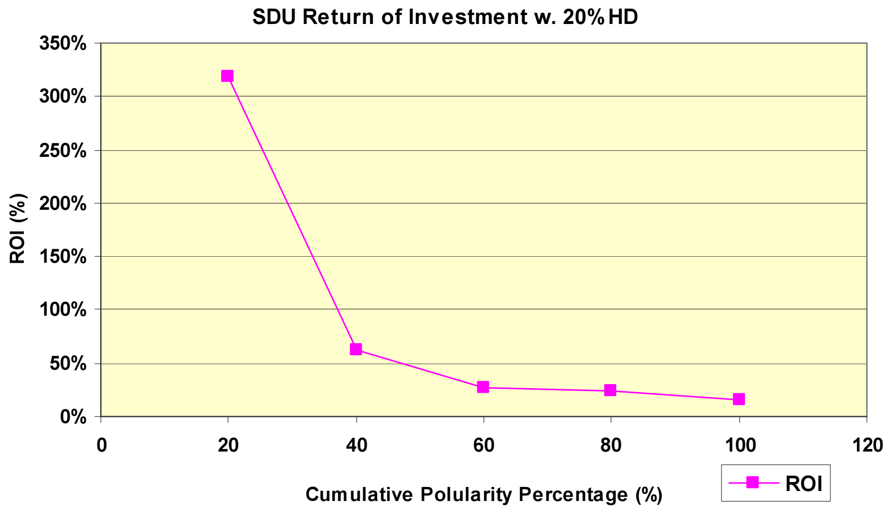

Figure 8. Switched Unicast ROI per popularity quintile

The return on investment results are illustrated in Figure 8. The diagram reveals that the ROI decreases as more popular channels are offered as unicast. In spite of the decrease, even when all channels are offered as switched unicast assuming 20% HD channels, with the submitted parameters the ROI can be demonstrated to be as good as 16%.

One interesting fact to notice is the sensitivity to HD channel percentage in the channel lineup. Figure 9 clearly demonstrates that as the percentage of HD channels in the offering increases, the cost for switched unicast increases dramatically. In fact, if more than 30% of the 247 channels are HD channels, bandwidth and spectrum requirement will push the service group size to below 125 tuners if the majority of the STBs are legacy MPEG-2 only STBs. This can be intuitively understood by the fact that an HD program can consume 4 times the bandwidth of an SD program, but may not necessarily command a 4x premium for spot ad insertion.

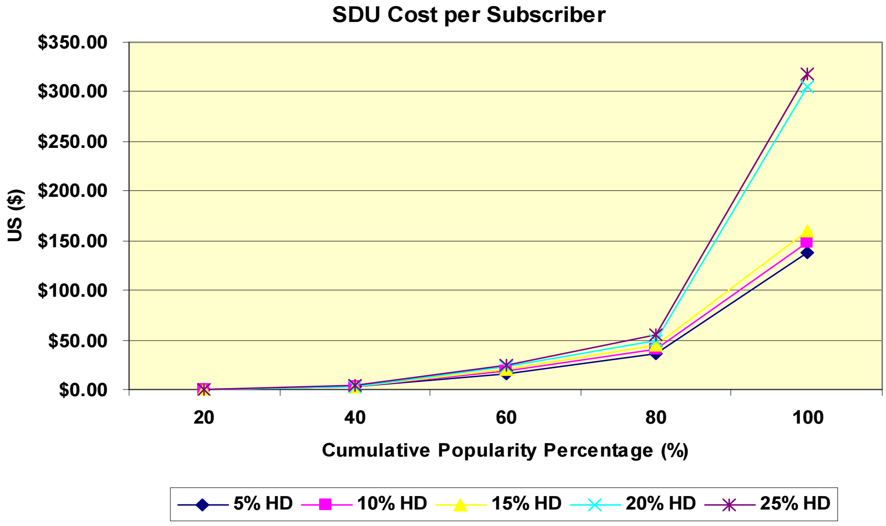

Figure 9. Switched Unicast costs with varying percentage of HD programming

What is the best strategy for MSOs?

The previous analysis provides insight that supports a recommended switched unicast strategy. First, although there is little doubt that the ad revenue potential of switched unicast is well worth the investment, switched unicast can be offered gradually with the least popular channels offered in unicast first. This approach is even more appealing if we consider that the least popular channels also have more local ad inventory accessible to MSOs. In a switched video architecture, a switched digital video session manager could implement a policy control to prioritize the less popular channels in unicast first.

Second, the number of HD channels offered in the switched unicast must be carefully evaluated to maximize the ROI. On one hand, it is tempting to offer more HD channels in switched unicast. On the other hand, HD channels consume much more bandwidth than SD channels. For MPEG-2 video, the HD bandwidth is roughly four times the SD bandwidth. Unfortunately, in today's advertisement arrangement, there is no revenue premium to insert ads in HD channels as opposed to SD channels. The proliferation of AVC STBs does alleviate the problem by reducing the HD bandwidth need by half. Additionally, MSOs may find creative ways of getting more revenue from HD ad insertion.

Third, even though the spectrum capacity can be increased by node splitting to get an HFC plant unicast ready, a better way might be to reclaim analog channels first to free up some spectrum. In the cost analysis, we assume that there are 30 RF channels available for linear digital video services since most of the RF spectrum is used by the analog tier. With analog reclamation, more spectrum will be available for linear digital video services and fewer node splits would be needed.

HYBRID UNICAST-MULTICAST DELIVERY SYSTEMS FOR MAXIMUM BANDWIDTH OPTIMIZATION

It has been observed that an appropriate premium incentive for a targeted advertisement service can justify the infrastructure investment. What guidance does such a statement provide to an MSO? Should an operator wait until a certain CPM threshold is crossed before deploying a targeted advertising solution?

It is possible to consider another variation of switched unicast. This variation can be easily understood by examining Figure 10.

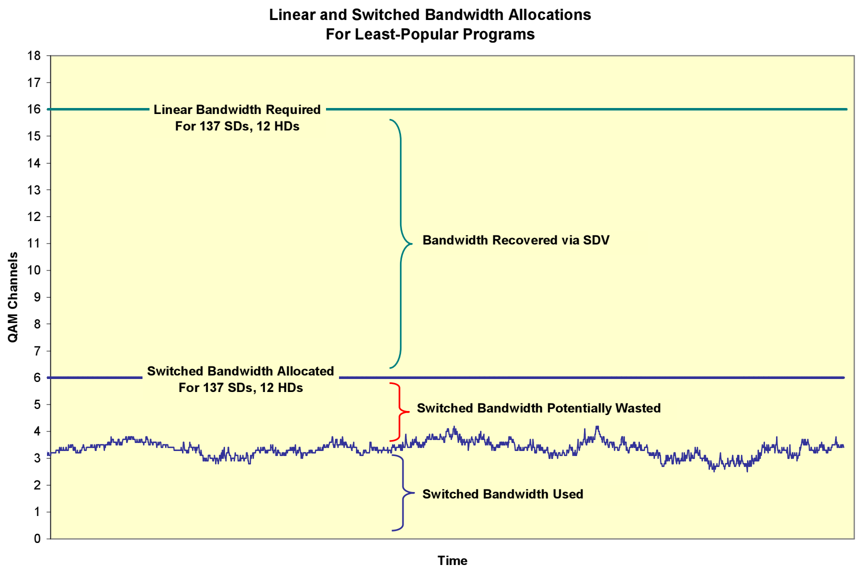

Figure 10. Switched Multicast Bandwidth Utilization

Indeed, Figure 10 shows that SDV saves bandwidth over its linear video equivalent, but what of the spectrum capacity that is allocated but not used? This represents a lost opportunity for better bandwidth utilization. To be more demonstrative, one could argue that this too is *wasted* bandwidth. The counter-perspective is that the unused bandwidth provides a "buffer zone" to protect against possible resource overflows. This is true, but the design of SDV systems demands that enough QAM resource be provisioned to handle the peak SDV consumption periods (which is typically during the evening/prime time hours). During non-peak hours, the bandwidth utilization can be quite low.

A hybrid multicast-unicast switched digital delivery system can provide the optimum blend of content personalization and bandwidth utilization in every situation. In such a system, a SDV server monitors resource utilization, and uses this indication to determine whether to respond to a channel request with tuning parameters for a unicast or multicast stream. When bandwidth utilization is low, for example in non-peak hours, the system can respond to channel requests by creating a personalized unicast channel. If bandwidth resource utilization reaches a (configurable) threshold, the SDV server can respond to subsequent channel requests with tuning information for a non-personalized multicast or "shared" stream. Since the multicast stream is shared by all of the subsequent users, this provides an effective "safety valve" to cap the stream usage when necessary, while offering maximum opportunities for personalization. As the subscriber population churns through channel changes, the aggregate number of active unicast channel streams will reduce by attrition, and the system can vary the unicast/multicast stream mix through a natural feedback process to manage the bandwidth utilization and the personalization opportunities to an optimal level.

What would a personalized vs. non-personalized channel look like? An example (one possible embodiment) is shown in Figure 11 [5].

Figure 11. Screenshot of linear channel vs. personalized channel

Multicast/unicast stream mix is an example of policy-driven resource management. In this example, allocation of edge resources is driven by a policy that seeks to optimize the insertion of unicast streams while observing rules on maximum bandwidth limits. But the rules do not need to be this simple. Future evolutions of this stream selection function could use more deeply sophisticated decision-making algorithms, weighing factors such as subscriber profile, program type, advertising value, time of day, and other factors to hyperoptimize the allocation of shared resources.

CONCLUSION

Switched Unicast is an extension of Switched Digital Video that enables content personalization and targeted ad insertion. While the main goal of most SDV deployments is centered around cost-effective programming expansion, switched unicast offers a revenue generation opportunity through targeted advertising and interactive services. Analysis of actual subscriber channel-change log data can provide valuable insight to the viewing patterns that might be expected in a switched unicast environment. Hybrid multicast/unicast implementations offer an opportunity to incrementally explore the value proposition of interactive programming by enabling fractional levels of personalization with a modest incremental investment to current SDV systems.

REFERENCES

[1] “A Comparison of Edge Bandwidth Requirements For Unicast versus Multicast Delivery of Switched Digital Video Services”, Robert Kidd and Michael Shannon.

[2] ”An Evaluation of Alternative Technologies for Increasing Network Information Capacity, Ron Shani and David Large”, NCTA 2005

[3] ”Irreconcilable Differences or a Match Made in Heaven? The Future of Advanced TV Advertising”, Ben Hollin and John Morrow, SCTE ET 2008

[4] Comcast 2007 Financial Results

[5] Screenshot courtesy ICTV and Turner. Use of this screenshot does not necessarily imply endorsement by ICTV or Turner of Cisco and/or the concepts presented in this paper.