EFFICIENT HD EXPANSIONS USING SDV, MPEG-4, AND MPEG-2/4 DUAL-CARRY

Robert Kidd
Manager, Applications Engineering
Cisco Systems
5030 Sugarloaf Parkway
Lawrenceville, GA 30044
(770) 236-5266
kiddr@cisco.com

Avinash Ravi
Applications Engineer
Cisco Systems
5030 Sugarloaf Parkway
Lawrenceville, GA 30044
(770) 236-1227
ravia@cisco.com

ABSTRACT

Industry benchmarks for high definition (HD) content are moving from 100 HDs to 150 HDs and beyond, and cable operators are faced with the increasingly challenging task of finding bandwidth on their networks to handle this once seemingly insurmountable load. Fortunately, the advent of MPEG-4 AVC/H.264 video compression adds another, very powerful dimension to operator's bandwidth optimization tool chest. MPEG-4 offers the option of creating new bandwidth efficient MPEG-4 service tiers, either as linear or switched channels. But MPEG-4 also provides another very non-intuitive advantage in terms of bandwidth reclamation. Simulation of bandwidth requirements using channel-change data captured from an actual switched digital video (SDV) system yields the surprising result that in a typical SDV deployment, less bandwidth will be required to dual-carry a service in both MPEG-2 and MPEG-4 formats than in MPEG-2 alone. This dual-carry approach provides an immediate opportunity for bandwidth savings while supporting a measured, gradual migration strategy for the introduction of next generation set-tops.

This paper first provides a review of existing SDV bandwidth management practices followed by scenarios for the extension of these practices using MPEG-4 to efficiently accommodate additional standard definition (SD) and HD content. Next, the results of simulations are presented that estimate the bandwidth required to carry content in both MPEG-2 and MPEG-4 formats as a function of the ratio of MPEG-2 and MPEG-4 set-tops in a system. The goal is to provide a comprehensive MPEG-4 insertion strategy that seeks to optimize both the bandwidth savings provided by MPEG-4 and the operational costs of migrating to an all MPEG-4 environment across an SDV-enabled cable system. Finally, the paper highlights how this approach is readily extended to hybrid or all IP-based systems.

TERMINOLOGY

The paper focuses primarily upon the optimization of bandwidth in the access network. The term bandwidth will generally refer to the downstream Video QAM or DOCSIS bandwidth required to carry RF or IP video content. MPEG-4 Advanced Video Coding (AVC) and H.264 are equivalent video compression standards issued by ISO and ITU, respectively. These standards are referred to simply as MPEG-4 in this paper. Next-generation set-tops refer to cable set-tops capable of decoding both MPEG-2 and MPEG-4 content. Some familiarity with SDV technologies is assumed, but for additional background, the reader is referred to references [1], [2], and [3].

INTRODUCTION

As industry benchmarks for HD content move beyond 100 HD channels, cable operators are aggressively moving forward with bandwidth saving technologies to ensure their offerings continue to meet or exceed those of competing video providers. The efficiency of MPEG-4 video encoding is propelling it rapidly toward critical mass in the cable ecosystem. MPEG-4 encoded content requires typically 50% to 70% less bandwidth than its MPEG-2 counterpart, potentially offering a major windfall to bandwidth-strapped cable service providers. Since HBO's announcement during the 2007 Cable-Tec Expo that it would distribute HBO HD via MPEG-4, the amount of content distributed in this format has increased dramatically. A quick tally from public sources yields around 70 HD services from major media providers including Time Warner, Viacom, Liberty Media, and Comcast, to name a few. Content providers are migrating satellite links to MPEG-4 delivery, headend and transport equipment are increasingly MPEG-4 capable, and next generation (next-gen) set-tops are bringing MPEG-4 capability to the home. All pieces of the MPEG-4 puzzle are quickly falling into place.

Yet for all its promise, there remains one major obstacle to MPEG-4 realizing its true potential – the installed base of set-tops. Over 100 million set-tops that support only MPEG-2 have shipped in North America alone, and a large percentage of these devices will remain in operator's footprints for years to come. How is it possible for cable operators to reap the benefits of MPEG-4 technologies while accommodating the legacy set-top base? Next-gen set-tops are backwards compatible with MPEG-2 and can thus decode both MPEG-4 and MPEG-2 encoded content; therefore, operators have the option of maintaining the status quo and simply continuing to broadcast via the lowest common encoding denominator, MPEG-2. MPEG-2 is a mature technology and maintaining a purely MPEG-2 broadcast infrastructure is readily accommodated by transcoding MPEG-4 distributed content back to MPEG-2 in the headend. However, this status quo approach discards any potential benefits offered by MPEG-4 and does nothing to improve the efficiency of the network. At some point, the requirements of the marketplace to add more content will overwhelm this short-term approach.

A more bandwidth efficient strategy is shown in Figure 1 in which a tiered structure is established such that content with relatively narrow demand is made available only on an MPEG-4 delivery infrastructure. For example, an extended MPEG-4 HD tier could be offered for subscribers that are willing to pay for these premium services, and a next-gen set-top would be provided as part of the premium subscription. In this scenario, Marketing gets their high-value content and Operations is assigned a non-trivial but manageable subscriber base impact. As next-gen set-tops permeate the system, this initial, narrowly defined tiered structure could be incrementally expanded to include more widely viewed content and service tiers.

How might each of these approaches impact a typical system? Let's assume that an operator decides to carry an expanded lineup of HBO HD channels, including *HBO East* or *West, HBO2*, *Signature*, *Family*, *Comedy*, *Zone*, and *Latino*. Assuming an MPEG-4 HD bandwidth allocation of 8 Mbps (HBO's publicly announced bitrate), these 7 channels will require approximately 56 Mbps. The corresponding MPEG-2 bandwidth allocation, assuming re-encoding at 16 Mbps per channel, will be 112 Mbps.

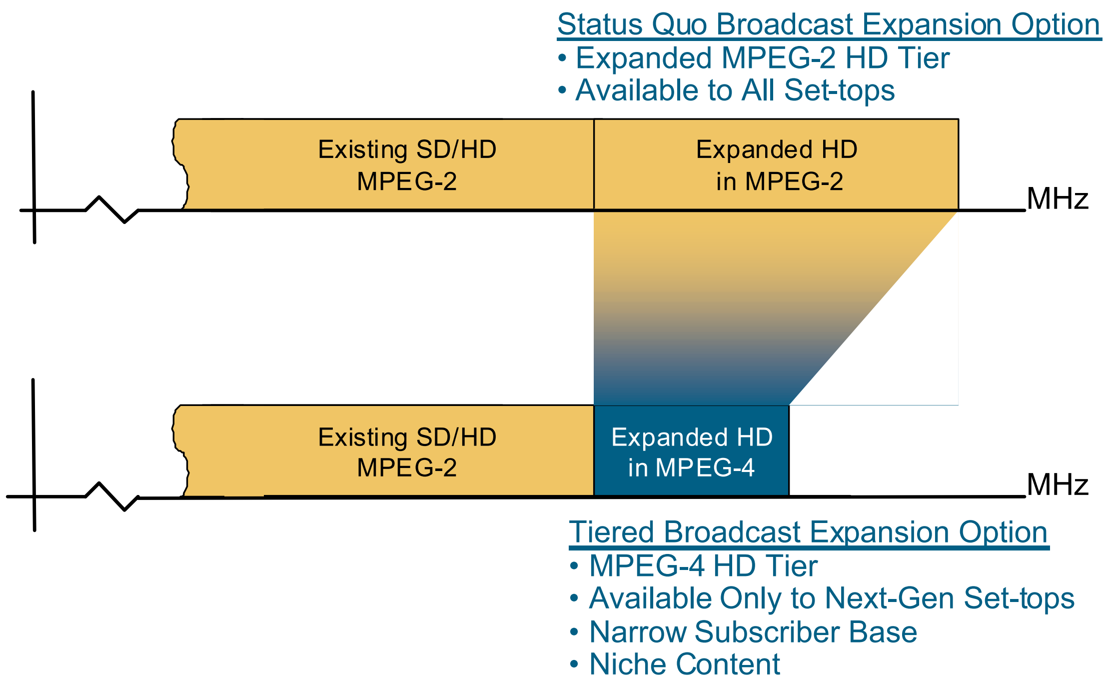

Figure 1. Static Broadcast HD Expansion Options

With the status quo approach, the services are carried via MPEG-2, require 2.9 QAM channels of access bandwidth, and can be decoded by any HD set-top. With the tiered approach, the services are carried via MPEG-4, require only 1.4 QAM channels of bandwidth, but must be received by next-gen set-tops. The tiered MPEG-4 strategy therefore saves 1.5 QAM channels of access bandwidth, but at the expense of the increased subscriber management activities necessary to deploy next-gen set-tops to all subscribers of the premium tier.

The bandwidth savings offered by MPEG-4 in the example above is significant; however, even with this reduced requirement, it's not certain that a typical bandwidth constrained plant can accommodate the additional 1.4 QAMs of bandwidth required. Furthermore, this small handful of programs represents only a fraction of the wide range of new content envisioned. The potential magnitude of desired local and national program expansions can be measured in terms of hundreds and not tens of new programs. The accompanying bandwidth requirement will thus be orders of magnitude greater than that posed in this simple example.

In summary, a dedicated MPEG-4 tier provides a bandwidth efficient alternative to traditional MPEG-2 broadcast for operators willing to absorb the associated operational costs; however, this tiered approach is feasible only for a limited number of programs. It will not scale to accommodate the large numbers of programs that operators are ultimately seeking to deploy. Fortunately, other alternatives exist. In the next section, a migration strategy is proposed that takes advantage of the bandwidth optimization capability of switched digital video (SDV) technology.

MPEG-4 SWITCHED BROADCAST DEPLOYMENT ALTERNATIVES

Anyone acquainted with switched digital video architectures will bristle at the notion of statically broadcasting lightly viewed content, MPEG-4 or otherwise, for this is precisely the type of content that lends itself to a switched solution. SDV leverages the fact that not all programs are being viewed all the time, and specific programs can therefore be switched into a service group (SG) only as necessary to satisfy requests originating from set-tops within that SG. This efficient, demand-driven approach will typically require one-third of the bandwidth that a static, linear broadcast architecture would otherwise require.

An MPEG-4 tier such as the one described in the example above will by definition carry niche or less-viewed programming. After all, the content was selected such that only a select group of viewers and a manageable number of set-tops would be affected. As such, it is likely to be much more efficient to carry this type of content on a switched infrastructure rather than on a dedicated tier. Using a typical 3:1 SDV carriage gain, the access bandwidth required is reduced from 1.4 QAMs to roughly 0.5 QAMs, a much less demanding result.

However, adding MPEG-4 content to the switch further mitigates only the bandwidth side of the equation; it remains necessary to manage the additional segmentation of MPEG-4 set-tops. But what would happen if both MPEG-2 and MPEG-4 versions of the content were carried on the switch, that is, *dual-carried*? We might then have the best of both worlds: reduced bandwidth requirements AND reduced operational requirements. But surely this approach would undo the potential bandwidth savings - after all, the term “dual-carry” itself reeks of inefficiency. It turns out that in a typical switched service group, a large percentage of the switched streams are being viewed by a single user. This characteristic suggests that SDV might provide a happy medium between bandwidth and operational requirements. MPEG-4 set-tops in a service group could use the lower bitrate version, and MPEG-2 set-tops in the service group would continue to access the higher bit-rate version. Both populations would benefit from the efficiencies of switching, and segmentation of the set-tops would not be required. The following analyses and simulations were performed to test this hypothesis.

A BASELINE ESTIMATE OF PEAK SDV STREAM REQUIREMENTS

In order to quantitatively evaluate the impact of dual-carrying both MPEG-2 and MPEG-4 copies of a service on a switched tier, viewership data logged from a production SDV deployment were reduced and analyzed. The fundamental approach was to evaluate individual channel-change start and stop times to build various measures of usage within a service group. In this case, the primary measure is the number of unique (multicast) streams required by a particular group of set-tops. The analysis was structured such that the sensitivity of unique stream count to variables including service group size and specific content groupings over particular time intervals could be captured. One month of set-top channel-change information was collected and filtered to remove content, such as low bitrate music services, that is typically not a candidate for switching. The resulting baseline data set consisted of channel-change data for 364 total video programs including 297 SDs and 67 HDs.

In order to evaluate the sensitivity of stream requirements to service group size, "Virtual service groups" were created containing set-top quantities ranging from 1 to 350. The equivalent tuner counts range from 1 to roughly 500 due to a DVR penetration of 35% within the production system. These service groups were assembled by accumulating set-tops randomly selected from within a node or geographically adjacent group of nodes, with the intent being to retain the behavior of the customer demographic associated with the underlying, physical service group structure.

In order to evaluate the sensitivity of stream requirements to a particular lineup, each video program was ranked according to the total number of viewership minutes that it had accumulated over the one month period. The resulting viewership curve, shown in Figure 2, is then divided into deciles with 36 programs per decile. As expected, the plot roughly follows Zipf's law, with the most popular decile contributing a large portion of total viewership minutes. The remaining deciles form a "long tail" of lower viewership programs that are ideal candidates for switching.

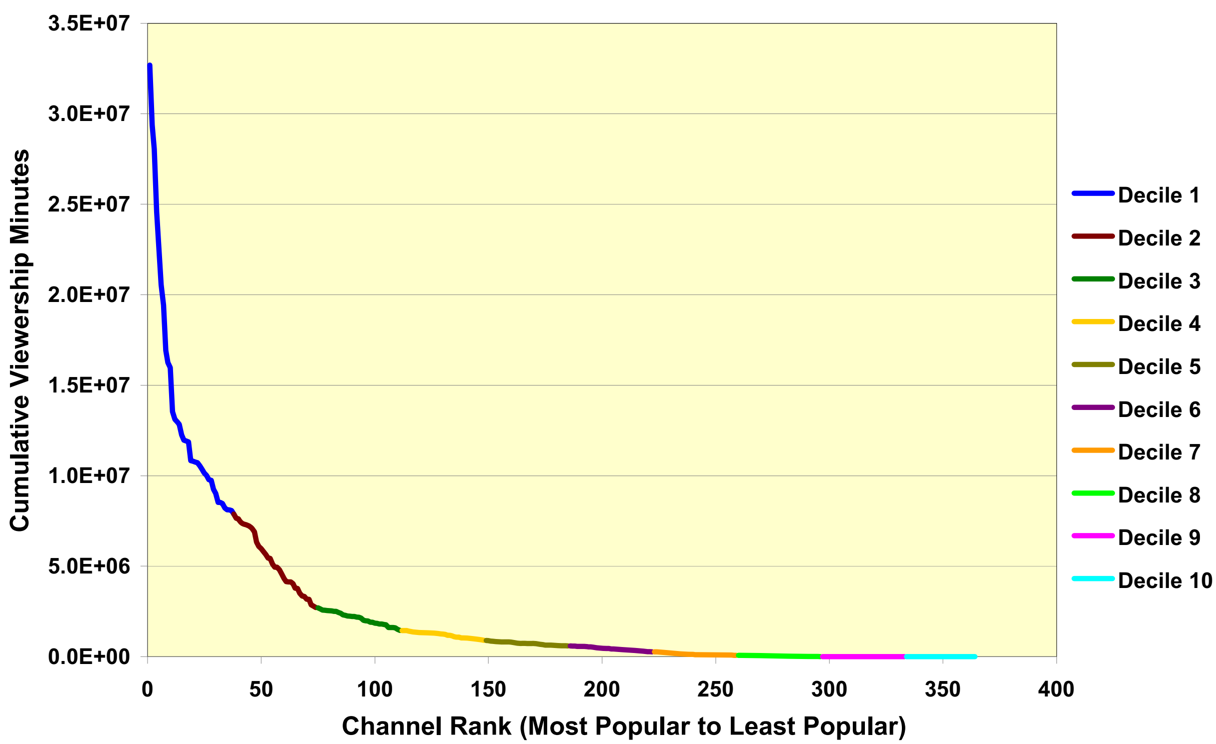

Figure 2. Viewership long tail with popularity deciles

Figure 3 illustrates the current distribution of HD and SD services across the viewership curve. HD services tend to be more highly represented across the center of the viewership curve. Highly viewed local broadcast HD and lightly viewed niche HD bracket the more moderately viewed content. As more HD content is added to a system, and more subscribers buy HD packages, HD services are expected to become more evenly distributed across the curve. For example, a trending analysis performed across several months of data indicates that HD services dominate the list of channels that have increased their relative viewership.

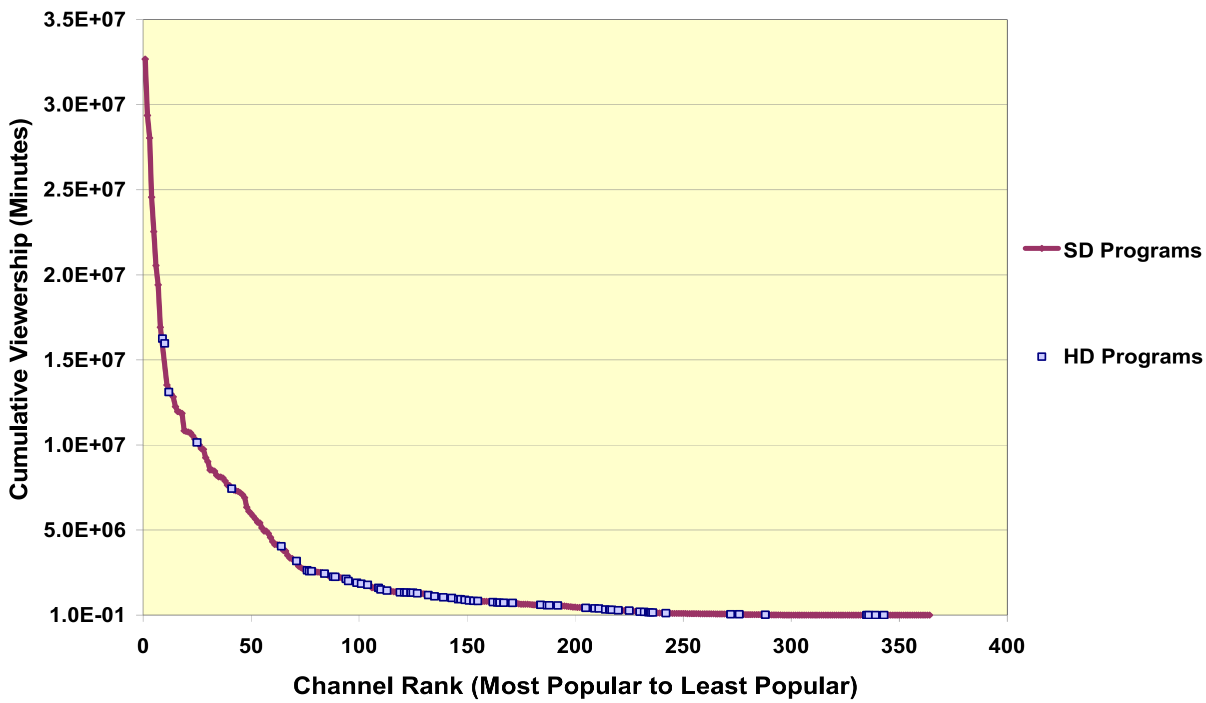

Figure 3. Viewership long tail with SD and HD rankings

In order to understand the impact of dual-carry on peak stream requirements, it is first necessary to capture a baseline estimate for the MPEG-2 case. The peak streams required for a particular virtual service group size and channel lineup were measured by incrementally adding viewership data for a particular set-top and program to a virtual service group and then computing the peak number of unique streams required at each increment. Programs were added in order of least viewed to most viewed, similar to the way they would ideally be added in a production system. The peak values were then averaged across fifty different virtual service groups. The three-dimensional baseline result in Figure 4 displays the number of average peak unique streams required as a function of service group size and cumulative channel popularity. Figure 5 displays the corresponding contour plot. As is typical in SDV deployments, less popular channels generate very little stream demand, even for larger service groups; however, stream requirements and associated bandwidth increase rapidly as the most popular content is added to the switch.

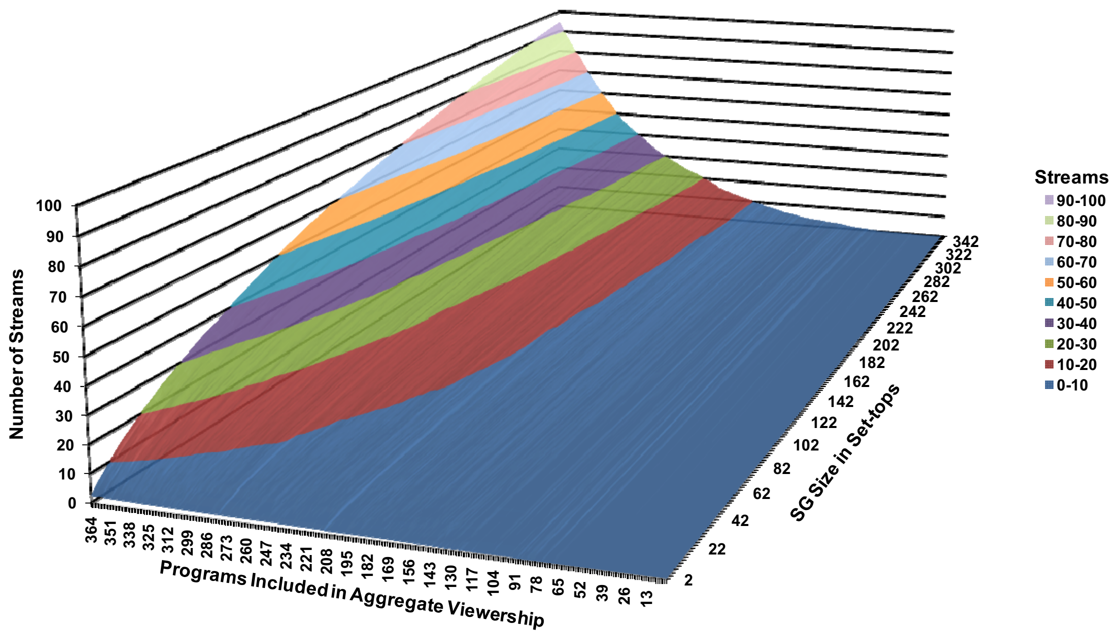

Figure 4. SDV streams required versus SG size and program popularity (3D)

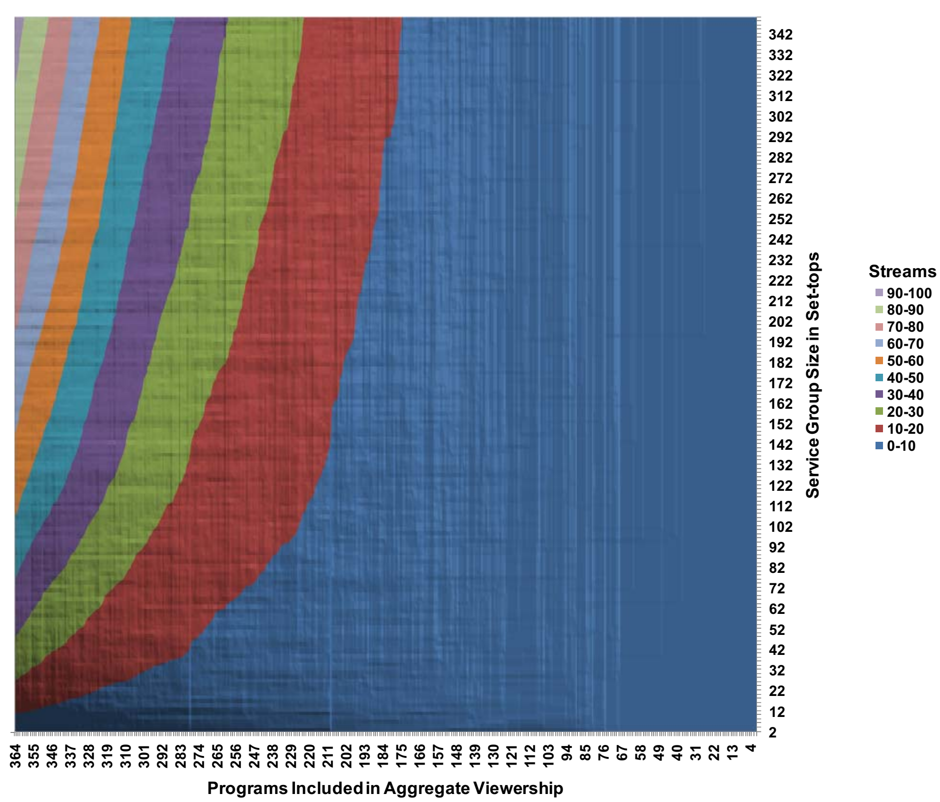

Figure 5. SDV streams required versus SG size and program popularity (contour)

Another alternative view of the peak stream data is shown in Figure 6. This chart illustrates 10 cross sections of Figure 4 with each line representing the number of streams required to carry a given percentile of content. For example the 90th percentile line illustrates the number of streams required to switch the nine least viewed deciles of content shown in the viewership curve of Figure 2. This format is easier to visualize and will be used widely in subsequent analyses.

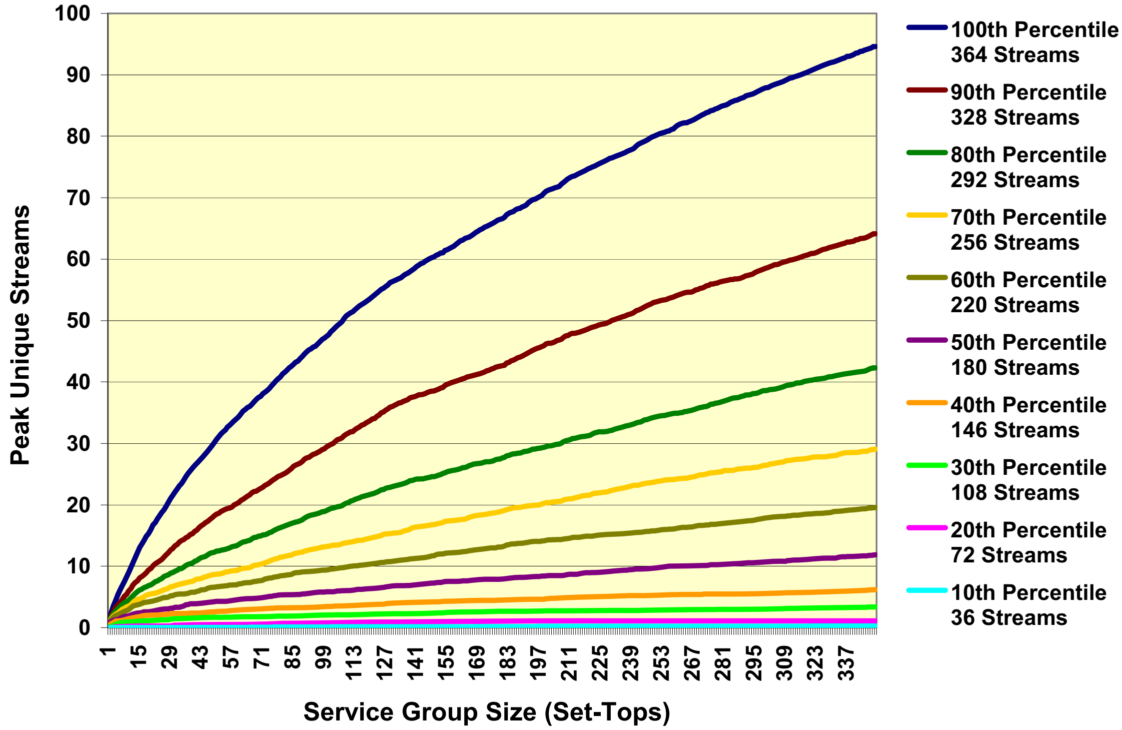

Figure 6. SDV average peak unique stream requirements for percentiles as a function of SG size

The less viewed percentiles of content have little curvature since this content tends to have only a single viewer; i.e. the number of streams required is directly proportional to the service group size, somewhat analogous to the behavior of video-on-demand (VOD); however, as more popular channels are added, the curvature increases, indicating an increase in multicast gain. As each additional user is added to a service group, it is increasingly likely that user will share a popular stream with another member of the service group, thus compressing the overall number of unique streams required.

AN ESTIMATE OF PEAK SDV MPEG-2/4 DUAL-CARRY STREAM REQUIREMENTS

In order to understand the impact on peak QAM bandwidth imposed by the dual-carry approach, a simulation was constructed that builds upon that described above: virtual service groups were created, viewership popularity percentiles were applied, and peak streams were measured. However, for the dual-carry case, an additional variable was added, specifically, the percentage of next-gen set-tops present in the service group. To aid in visualization of the results, the service group size was fixed at 350 set-tops. In the simulation, groups of set-tops from the collected data set were tagged as either MPEG-2 or MPEG-4, and the streams tallied accordingly for the two types of set-tops. Figure 7 below illustrates the number of MPEG-2, MPEG-4, and total streams required to switch 100% of the video lineup.

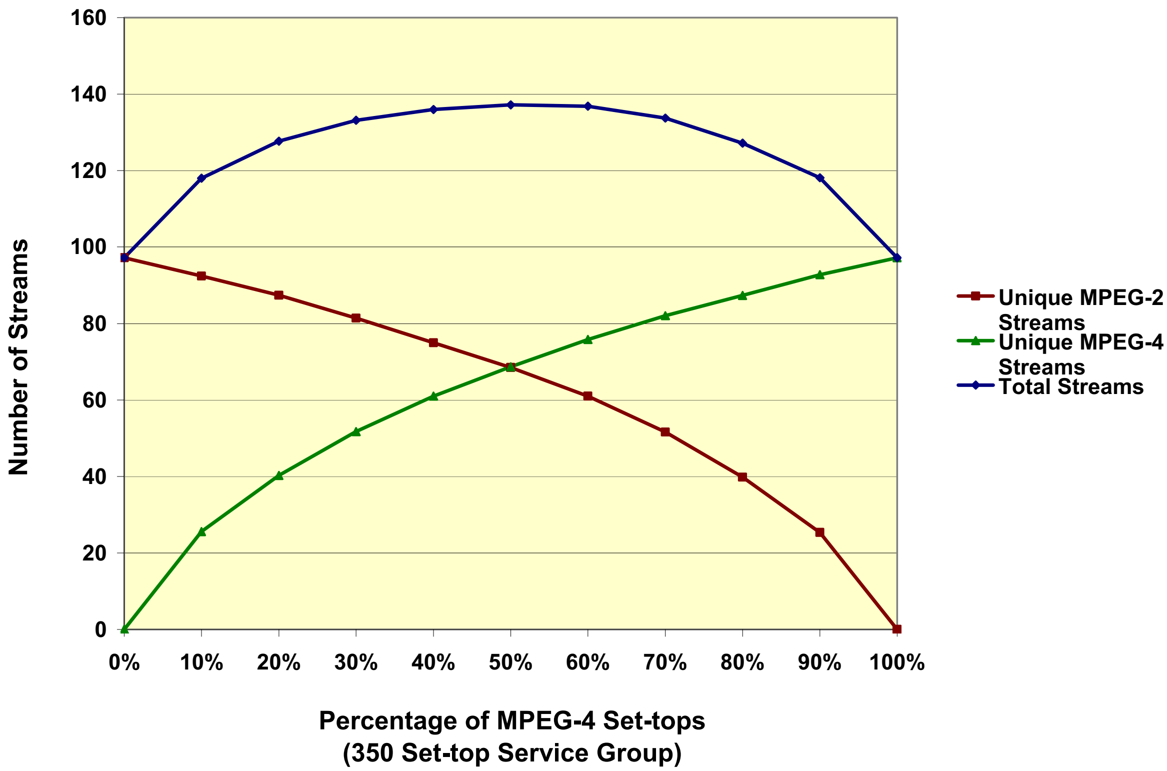

Figure 7. Switched dual-carry average peak stream requirement versus percentage of MPEG-4 boxes

Figure 7 illustrates the worst case scenario in which 100% of the content is included in the dual-carry lineup. The measured number of MPEG-4 streams simply retraces the 100th percentile line of Figure 6, and the MPEG-2 stream requirements follow a similar curve, only reflected about the 50% MPEG-4 point. As the percentage of MPEG-4 set-tops in a service group increases from zero, the number of MPEG-2 streams decreases at a rate less than the rate at which the number of MPEG-4 streams increases. This forces the total number of streams to increase well above the peak stream requirement that would be seen in a homogenous set-top box population. In fact, the total peak stream requirement increases by over 40%. However, since this combined stream population now contains a percentage of bandwidth efficient MPEG-4 streams, the total bandwidth requirement is dampened. The top line of Figure 8 illustrates this effect. Although the total peak stream count increases by over 40%, the peak bandwidth increases by only 7%. The results for the least-viewed curves, which contain the primary candidates for switching, are more interesting. Less-viewed content tends to generate more streams with a single viewer, and any dual-carry stream penalty is virtually eliminated in the lower-popularity curves. The result is that an operator is able to save bandwidth with the first MPEG-4 set-top deployed into a switched dual-carry service group. Only the lower bitrate MPEG-4 version of program must be switched if the program is being viewed only by an MPEG-4 set-top.

This paper is primarily concerned with the *relative* bandwidth required between basic MPEG-2 switching and dual-carry; however, a rough measure of absolute bandwidth required is obtained by making a few simplifying assumptions. All bandwidth results assume SD and HD MPEG-2 bitrates of 3.75 and 15.0 Mbps respectively, with MPEG-4 being half the MPEG-2 bitrates. The results shown in Figure 8 assume a futuristic worst case lineup containing 100% HD. Note that aggregations of services from the least viewed 90% of content - the typical SDV candidates - yield flat or decreasing total dual-carry bandwidth slopes.

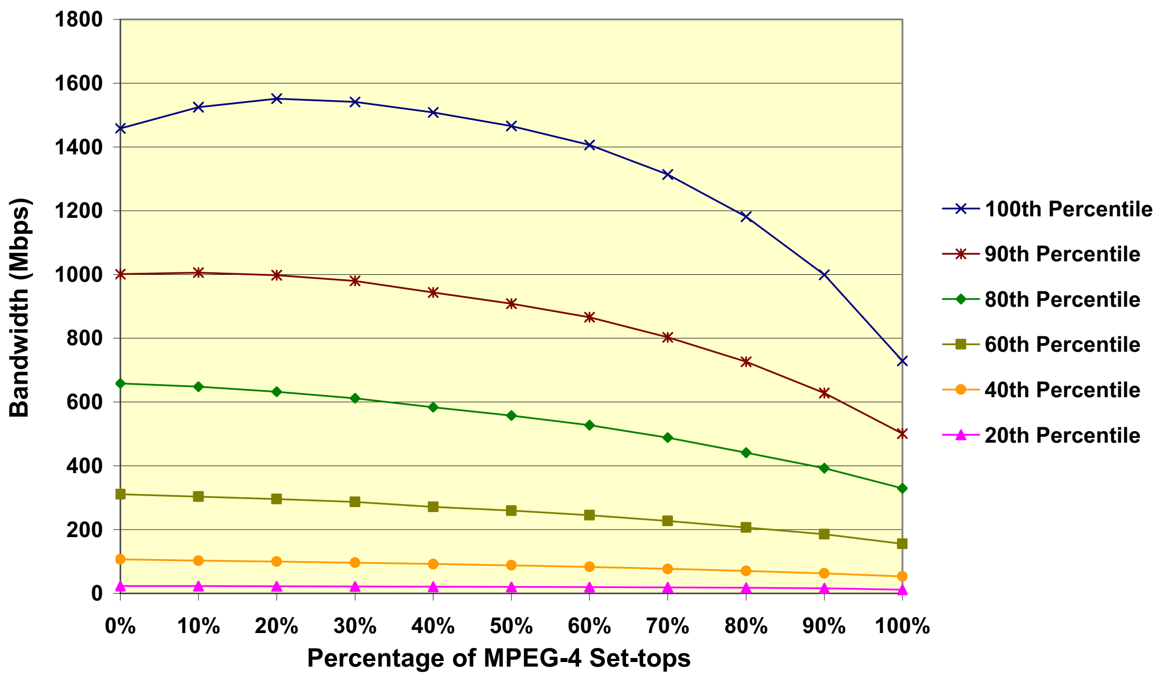

Figure 8. Switched dual-carry average peak bandwidth as a function of MPEG-4 percentage, 100% HD

Figure 9 provides a more realistic near-term scenario in which roughly 33% of programs, 121 programs in this case, are assumed to be HD. The trends for this scenario align with those of the 100% HD case and essentially provide a scaled version of the previous result. The 33% scenario is informative since it provides some measure of requirements in the 100 to 150 HD channel range; however, it relies upon a number of assumptions regarding where these future HD channels might fall with regards to viewership. For example, for the purposes of this analysis, they were assumed to be evenly distributed across the viewership curve. Time will tell the extent to which future HD viewership will validate this assumption, but in the meantime, an additional scenario was evaluated using the production SDV lineup.

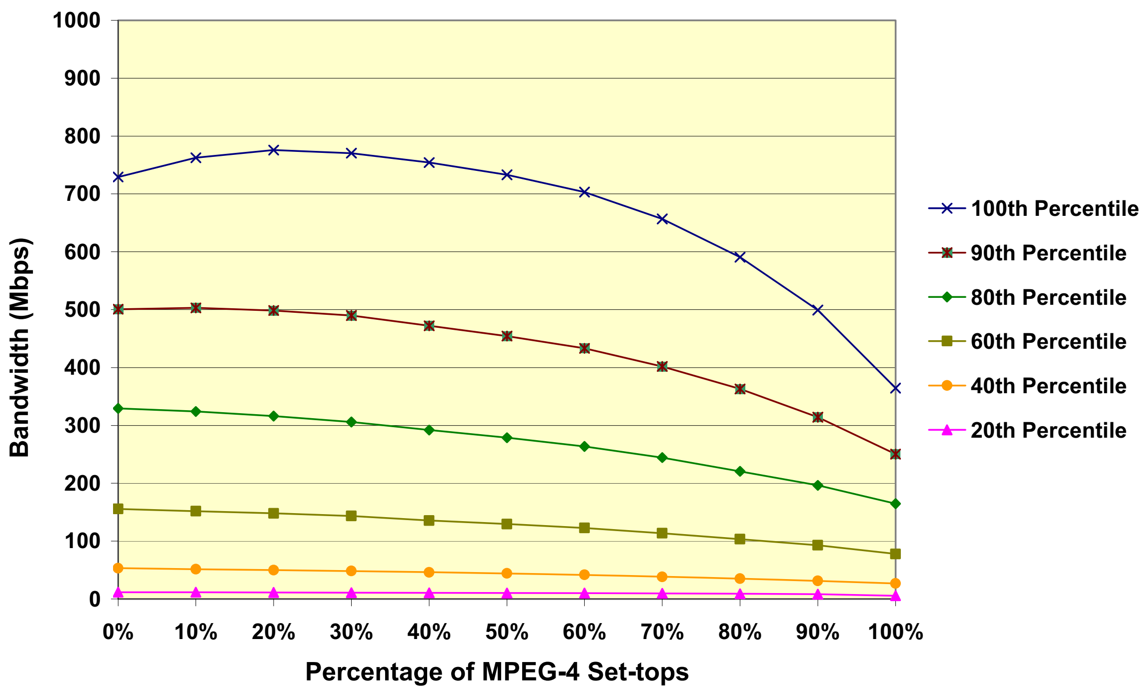

Figure 9. Switched dual-carry peak bandwidth as a function of MPEG-4 percentage, 33% HD

In order to firmly tie the analysis back to current reality, the switched MPEG-2 versus switched dual-carry bandwidth was simulated using the content for the actual SDV channel lineup that was active at the time the data was captured. The results are illustrated in Figure 10 and confirm that rather than imposing a bandwidth penalty, switched dual-carry can offer immediate bandwidth savings. The sooner that MPEG-4 set-tops are deployed, the sooner the additional bandwidth savings can be realized. Furthermore, SDV provides the option for a segmentation-free migration path toward an all MPEG-4 future. An arbitrary collection of set-tops within a service group can pull whatever stream best suits them without the need to establish inefficient static linear tiers or additional operationally intensive subscriber segments.

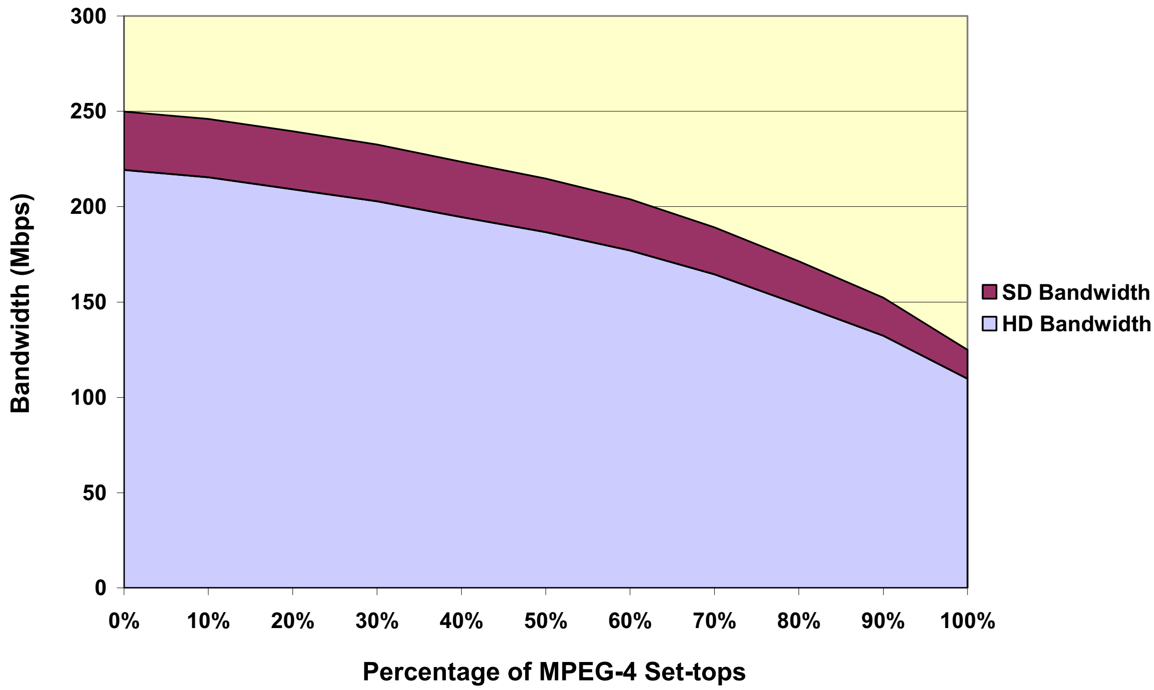

Figure 10. Dual-carry peak bandwidth estimate for a production SDV lineup

AN OPTIMIZED SDV MPEG-2/4 DUAL-CARRY STRATEGY

The analysis of the previous section assumed a simple dual-carry approach: an MPEG-4 capable set-top always receives the MPEG-4 version of the stream, and the MPEG-2 set-top always receives the MPEG-2 version. However, more sophisticated algorithms can be envisioned. For example, if an MPEG-2 program is already present on the switch, next-gen boxes could be directed to tune that program instead of receiving their own MPEG-4 version. This scenario was simulated next, with the results for the 100th percentile case shown in Figure 11.

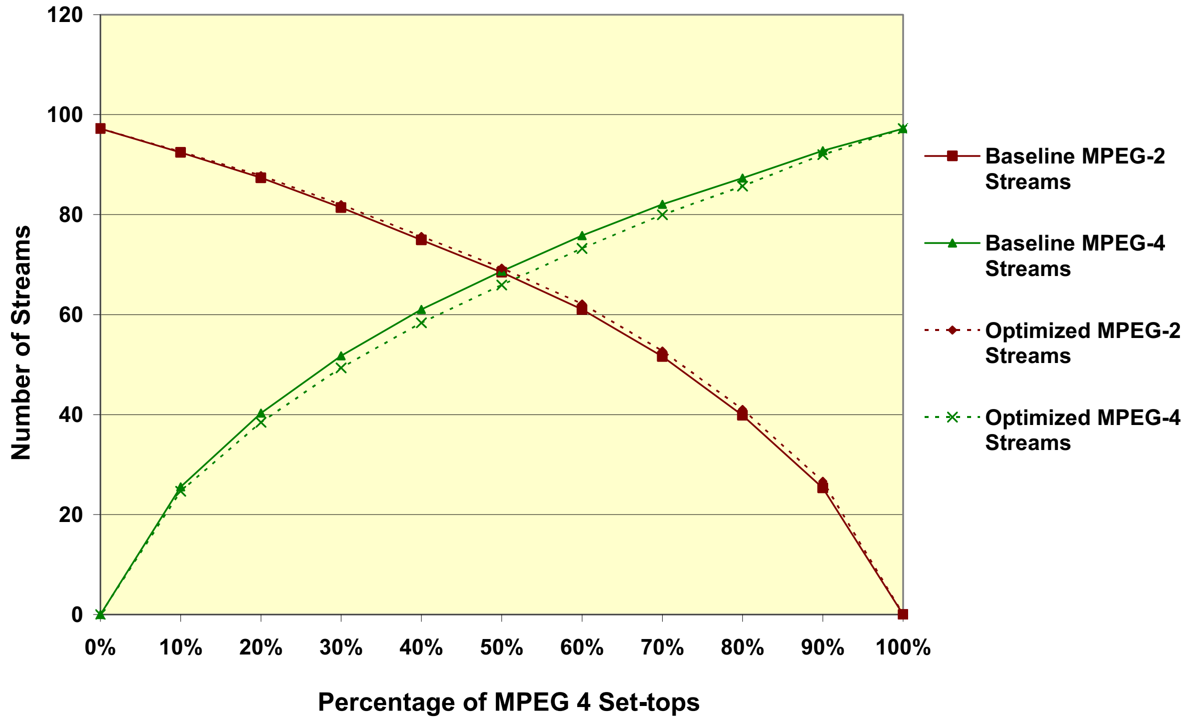

Figure 11. Optimized dual-carry peak stream requirements

The optimized approach requires slightly less dual-carrying of MPEG-2 and MPEG-4 streams; however, it increases the amount of time that higher-bitrate MPEG-2 content exists on the switch; i.e. once built, the MPEG-2 stream tends to be held longer. This result follows from the implementation of a "non-interrupting" approach to stream selection. Specifically, an existing MPEG-2 stream will remain active until all users, either MPEG-2 or MPEG-4, release the stream. MPEG-4 set-tops will not be redirected (and the viewer experience momentarily interrupted) to a newly built MPEG-4 alternative.

Figure 12 illustrates the resulting difference in bandwidth between the basic dual-carry and the optimized dual-carry strategies. The optimized strategy does produce a slightly improved result for lower percentages of MPEG-4 set-tops. However, for higher percentages of MPEG-4 set-tops, the "optimized" strategy actually requires slightly more bandwidth, again, because MPEG-2 streams are held longer than would otherwise be required in the basic scenario. But the general conclusion is that there is little discernable difference between the two methods, and most of the benefit of dual-carry can be achieved with the basic, most easily implemented approach.

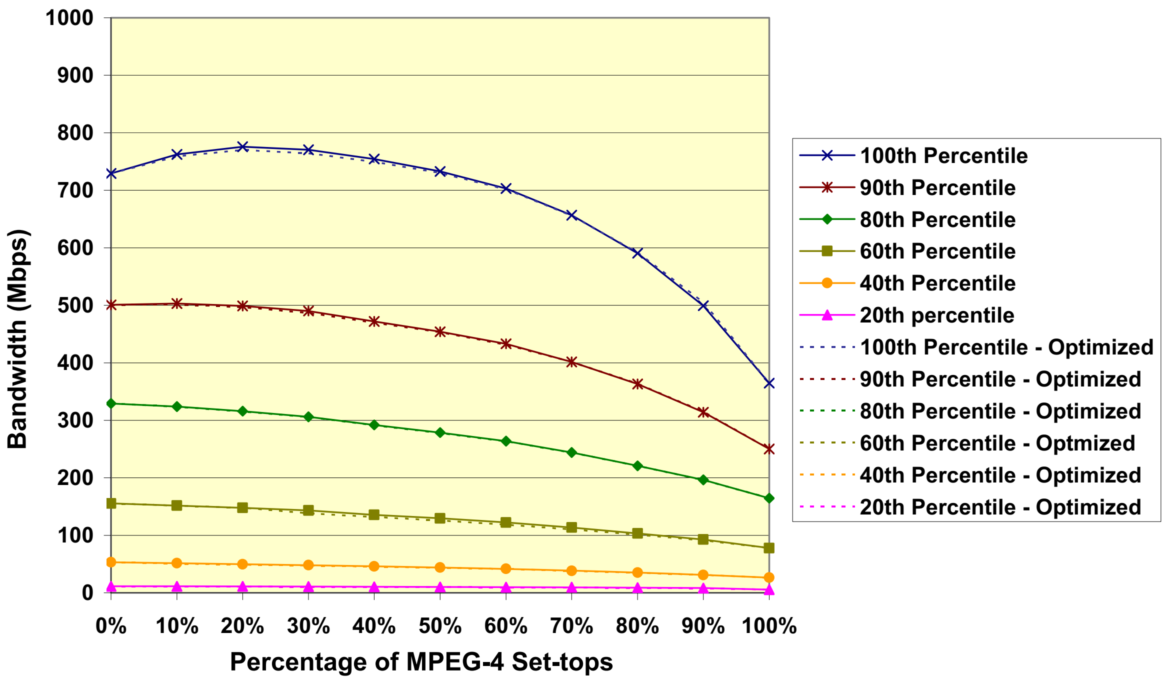

Figure 12. Optimized dual-carry peak bandwidth, HD=33%

SUMMARY AND IMPLEMENTATION CONSIDERATIONS

The results of the above analyses suggest the MPEG-4 deployment approach illustrated in Figure 13. First, the most highly viewed content would continue to be carried on a static, MPEG-2 linear tier. Little if any bandwidth is recovered by switching this content, and retaining MPEG-2 encoding avoids impacting the installed base. The remaining programs would be candidates for switched dual-carry. This approach requires less bandwidth than that required by MPEG-2 switching alone and supports flexible deployment of next-gen set-tops. Finally, the option remains to establish an MPEG-4 switched tier for narrow applications in which an operator wants maximum bandwidth savings and is comfortable with the operational implications of deploying MPEG-4 set-tops to customers with specific niche packages.

Conceptually, implementing dual-carry is as straightforward as converting the received version of a program to its complement and provisioning both on the transport and control systems. But it is never of course this simple in practice. The usual caveats with SDV still apply; for example, legal considerations that may restrict the flexibility with which content can be added to the switch pool. But these typical issues aside, three key considerations standout: (1) the additional intelligence required within the video system control plane to manage the selection of the appropriate version of a service, (2) methods for originating copies of programming in both encoding formats, and (3) support within new or legacy video processing, transport, and ad insertion equipment for MPEG-4.

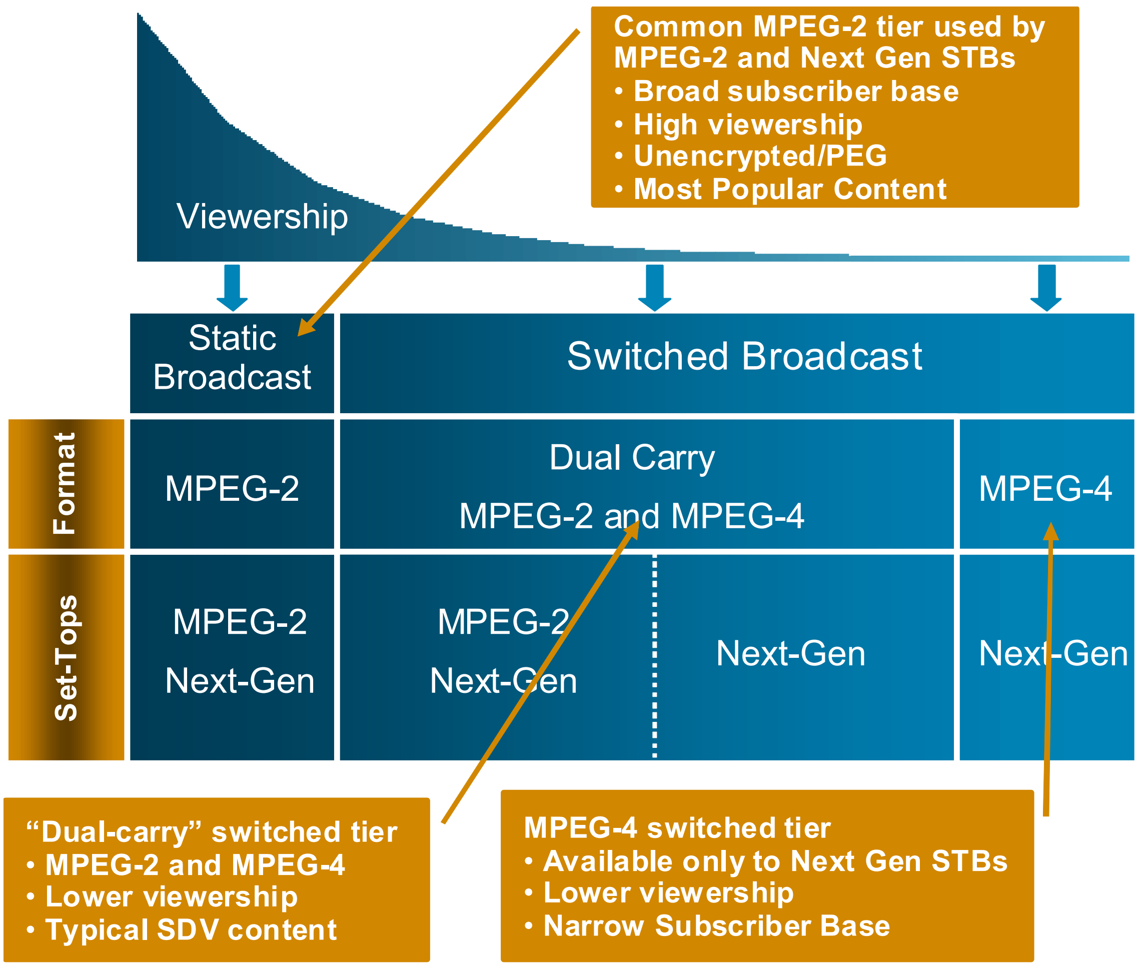

Figure 13. Conceptual allocation of content to linear and switched broadcast tiers

The key SDV dual-carry system elements are presented in Figure 14. From a control system standpoint, the additional logic required in Session and Resource Managers (SRMs) is straightforward and is currently being implemented in standards-based platforms. From a content origination standpoint, a large number of primarily SD services will continue to be distributed in MPEG-2, and those selected for dual-carry must be re-encoded or transcoded to MPEG-4. Fortunately, extremely dense, cost-effective MPEG-2 to MPEG-4 transcoding products are available that are ideally suited to this application. Programming that is already distributed via MPEG-4 imposes a different set of challenges. For many of these programs, dual-carry may impose no additional encoding requirement. Today content received via MPEG-4 is typically re-encoded back to MPEG-2 either within the satellite receiver or with standalone encoders, and both copies may already be available in the headend.

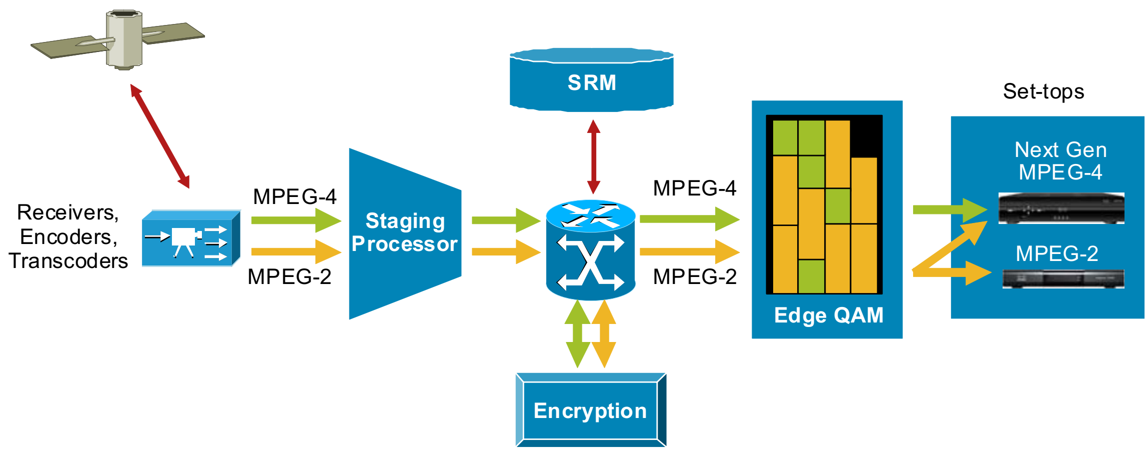

Figure 14. Key SDV dual-carry system elements

An experienced headend engineer can quickly point to a number of potential pitfalls that may be encountered with the above implementation scenarios. For example ad insertion continues to be performed in the MPEG-2 domain. The implication is that the inserted version of the MPEG-4 service must be created from its MPEG-2 counterpart, a requirement that dictates the deployment of relatively more expensive encoding instead of relatively less expensive splicing. But just as the early hurdles with MPEG-2 were overcome, so will the early challenges with MPEG-4 be met. As a critical mass of MPEG-4 capable set-tops builds, so will the importance of infrastructure that enables operators to take full advantage of MPEG-4's capabilities. Necessity drives invention, and the pieces of the MPEG-4 puzzle - set-tops, transcoders, enhanced SRMs, and ad insertion, to name a few - will inevitably fall into place.

EXTENSIONS OF THE MODEL TO IP DELIVERY ARCHITECTURES

Finally, while this paper has focused specifically on SDV applications, the results have obvious applications to not only RF but also potential next-generation IP video delivery architectures. A rigorous analysis of IP requirements is a different topic and will be left to future work, but one simple example will serve to illustrate the applicability of the dual-carry result to IP delivery.

Figure 15 illustrates three typical video delivery architectures: (1) standards-based RF switched digital video, (2) IP delivery via DOCSIS 3.0, and (3) a hybrid of RF and IP. For the sake of simplicity, assume that traditional cable set-tops continue to receive all content via linear or switched MPEG-2, and IP set-tops receive content via MPEG-4. The stream count requirements for a particular size population of these mixed devices can be estimated following an approach similar to that used for SDV dual-carry. In fact, invoking the usual assumptions regarding consistent viewership behavior between IP and RF subscribers, the results of Figure 7 apply directly.

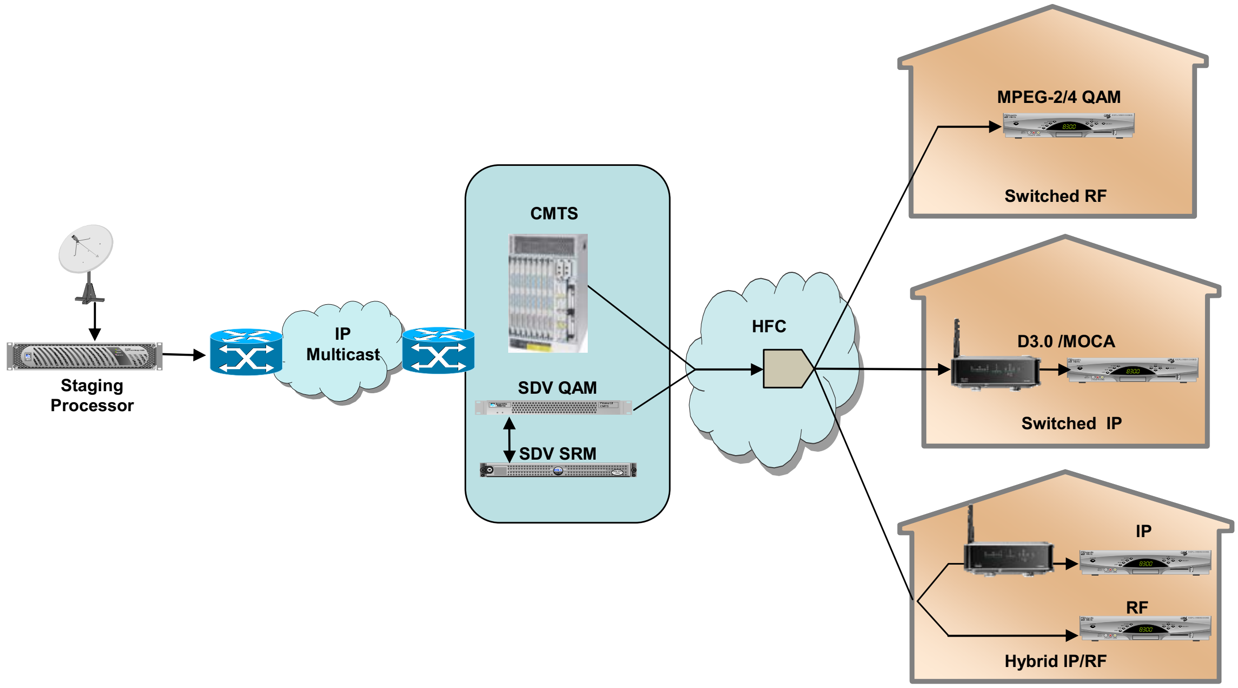

Figure 15. RF and IP switched broadcast delivery scenarios

However, the bandwidth dynamic for this scenario will be even more attractive than that shown in Figure 8. The streams destined for the IP set-tops are assumed to be variable bitrate (VBR) encoded MPEG-4 streams carried in bonded DOCSIS 3.0 channels. This approach requires a small amount of additional bandwidth for DOCSIS overhead, but the overall stream capacity per MHz of allocated downstream bandwidth is significantly increased due to the anticipated efficiency gains of VBR carriage in bonded channels [4]. The MPEG-4 contribution to peak bandwidth in the mixed RF/IP scenario will therefore be less than that required in the RF-only case. The net results are that the already small increase in total bandwidth for the 100% case of Figure 8 will be further dampened in the mixed RF/IP case, and the bandwidth roll-off with IP set-top penetration will be steeper than that of the RF dual-carry scenario.

Of course a number of different scenarios are possible in a mixed RF/IP environment. For example, depending upon what mix of MPEG-2 and MPEG-4 content is offered to each set-top type, triple-carry or even quad-carry scenarios can be envisioned. Each scenario deserves its own due diligence, but in general, for a given service group size, as long as the newer stream types, be they MPEG-4 RF or MPEG-4 over DOCSIS, are significantly more bandwidth efficient than the streams they are displacing, the bandwidth tendencies discussed in this paper hold. As long as the content is switched, carrying multiple formats of content will not impose a significant additional bandwidth requirement, if any, beyond that required for switched MPEG-2 alone.

CONCLUSION

SDV offers an efficient mechanism to allow cable operators to take advantage of the bandwidth efficiencies of MPEG-4 while ensuring a manageable transition of customer premise equipment from MPEG-2 to MPEG-4. Dual-carrying a typical switched lineup in both MPEG-2 and MPEG-4 formats requires less QAM bandwidth than that required for switching MPEG-2 content alone, and potentially far less than that required for simple static, linear broadcast. The savings is achieved with lightweight SRM logic that does not interrupt or otherwise adversely impact the customer viewing experience. Key system elements, including MPEG-4 capable set-tops and dense transcoders, are currently available to implement the strategy. Finally, it is not necessary to introduce operationally intensive measures such as segmentation of the customer base to take advantage of MPEG-4 capability. Dual carrying content ensures not only backwards compatibility with existing set-tops but also support for the latest set-top families. The greater the percentage of MPEG-4 capable set-tops in the system, the greater the bandwidth savings, but these savings accrue naturally as these emerging set-tops accumulate in a switched digital video environment.

REFERENCES

[1] "An Open Architecture for Switched Digital Services in HFC Networks", Luis Rovira, Lorenzo Bombelli, and Paul Brooks, SCTE ET 2006
[2] "Unicast Video Without Breaking The Bank: Economics, Strategies, and Architecture", S.V. Vasudevan, Xiaomei Liu, and Robert Kidd, NCTA 2008
[3] "A Comparison of Edge Bandwidth Requirements For Unicast versus Multicast Delivery of Switched Digital Video Services", Robert Kidd and Michael Shannon, 2007
[4] "Variable Bit Rate Video Services in DOCSIS 3.0 Networks", Xiaomei Liu and Alon Bernstein, NCTA 2008

TABLE OF ACRONYMS AND ABBREVIATIONS

3D	Three Dimensional
DOCSIS	Data Over Cable Service Interface Specification
DVR	Digital Video Recorder
H.264	ITU version of the MPEG-4 standard
HD	High Definition
IP	Internet Protocol
ISO	International Organization for Standardization
ITU	International Telecommunication Union
MHz	Megahertz
MPEG	Moving Picture Experts Group - a family of standards for compressing digital video
MPEG-4 AVC	MPEG-4 Advanced Video Coding - ISO version of the MPEG-4 standard
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
SD	Standard Definition
SDV	Switched Digital Video
SG	Service Group
SRM	Session and Resource Manager
VBR	Variable Bitrate
VOD	Video on Demand