IMPLEMENTING ADDRESSABLE ADVERTISING IN LINEAR NETWORKS

Steve Riedl, Principal Architect, Time Warner Cable
Doug Jones, Chief Architect, BigBand Networks

Abstract

Although the potential for addressable advertising to increase revenues is widely accepted, the infrastructural changes needed to support it are still being examined. Existing network functionality needs to be enriched to allow the following abilities to be implemented:

- *Replication of program streams such that more than one copy can be simultaneously supported on the network – since bandwidth-intensive HDTV is increasingly being offered to subscribers the ability to place its replication closer to the edge of the HFC network could conserve IP distribution network capacity;*
- *Managing the multiple copies within the current systems without causing issues with the large increase in sources;*
- *Implementing IGMPv3 to insure only viewed networks are distributed on the IP sections of the operator network;*
- *Association of a particular program stream with viewers of like demographics to enable relevant ads to be delivered to those viewers – grouping these viewers can be accomplished either when theviewer initiates a channel change or by force-tuning that viewer before the ads are to be played;*
- *Addressing of program streams to a single user in order to deliver a particularly relevant ad to them;*
- *Selective insertion of advertising into a program stream to match viewers' demographics – ad insertion can occur either through seamless splicing or a playlist;*
- *Interoperability between the HFC resource management system and the advertising decision manager to ensure that the available capacity of the HFC network is taken into consideration;*
- *Instantiation of an overall control mechanism to coordinate these activities – the SCTE DPI committee has begun the process of developing standardized interfaces to meet this need.*

The delivery of advertisements should include extensions of switched digital video known as microcast and unicast, and set-top based ad insertion mechanisms. These methods are complimentary because content can be delivered either over the network or from a local hard disk drive in the set-top box. This paper unifies these topics, describes an architecture designed to satisfy the requirements listed above, and explains its operation. The intent of the authors is to prove a technology framework for converged resource management and addressable advertising for emerging marketing services.

The paper also provides relevant context on the addressable advertising service including an architecture for a next-generation advertising system. The authors show how ads can be stored and how the developing standard interfaces will enable next-generation ad delivery systems. Additionally, it identifies a variety of parameters that can be used to size the advertising system by leveraging projections of initial deployments.

INTRODUCTION OF ADDRESSABLE ADVERTISING

Addressable advertising is the selective insertion of advertising into a program to match viewers' interests, thereby making that advertising more relevant to the viewer. Ad relevancy is important, because if presented with a choice many viewers would likely opt-in to

minimize or eliminate certain types of advertising that does not match their interest; for example, parents without kids might prefer not to see another baby diaper ad.

With addressable advertising, sets of viewers can be targeted at home watching a program like ESPN's SportsCenter and delivered different ads more relevant to there interest at precisely the same time. For example, one set of viewers who prefer full-size high-performance vehicles can be shown a commercial for the new Chevy Blazer and another set of viewers who prefer mid-size fuel-efficient vehicles can be shown a commercial for the all-new Chevy Equinox. This type of viewer ad targeting was pre-planned by the advertiser and tools are becoming available in the network to make it a reality based on aggregate geographic, demographic, psychographic or other characteristics of the consumers residing within specific areas.

Making advertising more relevant to viewers can have several benefits. First, viewers will come to see advertising as part of the programming, and will most likely be happier and more interested in viewing it. Additionally the relevancy will become more valuable to the advertisers making cable programming a better place to spend their advertising dollars. Finally, intelligently choosing advertising benefits the advertiser by minimizing the placement of irrelevant advertising in front of viewers.

While there are many methods to address advertising to viewers, this section will introduce two widely available methods; ZIP+4 and Prizm Codes. Both methods are based on the premise that "birds of a feather flock together," and that people with similar demographic traits tend to behave in the same way in the marketplace.

ZIP+4 codes are specific from a single to a few dozen households, which clearly is more specific than the average cable service group which can pass several hundred to a thousand homes. Commercial databases are available which will provide specific demographics for the households within all ZIP+4 areas and advertisers already use these for mass mailings. This advertising principle can easily be carried over to cable if the proper advertising infrastructure were in place to provide advertising to just a few dozen homes, or the viewers associated with those homes.

Prizm Codes are associated with several demographic parameters. There are 66 defined Prizm Codes and one or more can be assigned to each cable customer. In fact, one of the 66 Prizm Cluster Codes is already assigned to every address in the U.S. Prizm Codes are based on one of four urbanicity categories which is determined by the population density of an area and it's neighboring areas. Within each urbanicity category, segments are further sorted into groups based on affluence, another powerful demographic predictor of consumer behavior.

CABLE PRIVACY ACT

A discussion of addressable advertising requires an up front disclosure on the implications on customer privacy. Addressable advertising implies a knowledge about a customers demographics and with all the information available about consumers, a major issue in addressable advertising is that cable operators must comply with some very strict privacy guidelines, including the Federal Cable Privacy Act of 1984 and other specific laws passed by individual states. Nothing in this paper should be taken as legal advice. Always check with your company privacy attorneys before implementing any system.

TODAY'S CABLE ADVERTISING LANDSCAPE

This paper primarily focuses on Linear programming but many of the concepts will also apply to On Demand programming. Linear programming is traditionally a broadcast medium. On a cable plant, broadcast bandwidth is a very precious resource. Narrowcast bandwidth can be created by reducing node size and/or adding additional QAM bandwidth to a node by removing analog networks. Once the node size reaches about 500 tuners, then it is usually more economical to expand the number of QAMs per node.

Since unicast services such as VOD and multicast services such as SDV share this narrowcast bandwidth, an Edge Resource Manager needs to be put in place to allocate these resources. This paper show how this service plays an important role in utilizing this bandwidth for advertising purposes.

Not all channels are insertable; it generally depends on the programming contract. There are about 60 ad-insertable networks today. Most operators insert on 40 channels, four 30 second spots per hour is normal – and only 12 hours of the day matter. So while there are about 4,000 insertion opportunities each day, only about 2,000 are useable for any significant revenue. The top ten networks provide about 80% of the ad revenue and the top 20 networks provide over 90% of the revenue.

Given that cable local / spot advertising is estimated to be a $5 billion business in 2008 (NCTA), the last 10% is not small change at $500 million using current methods. By making that smaller number of eyeballs addressable, measurable and interactive, this 50% of the avail space could bring in well over $5 billion.

Some programs are addressed by their very nature. The typical ESPN viewer is different than the typical Lifetime Movie Network viewer. Others networks like CNN attract a wider demographic of viewers. Even a network that attracts a niche demographic will have some variation and if third-party data can be used, like if the car lease is expiring, then targeting on those networks has real value.

Today the most sophistication associated with advertising is geographic zoning, where customers in one part of a city will see different ads than customers in another part of the city. The degree of geographic zone subdividing within the cable system would generally be measured in tens of square miles associated with different regions of a city. Figure 1 shows an example of carving a city into five geographic zones.

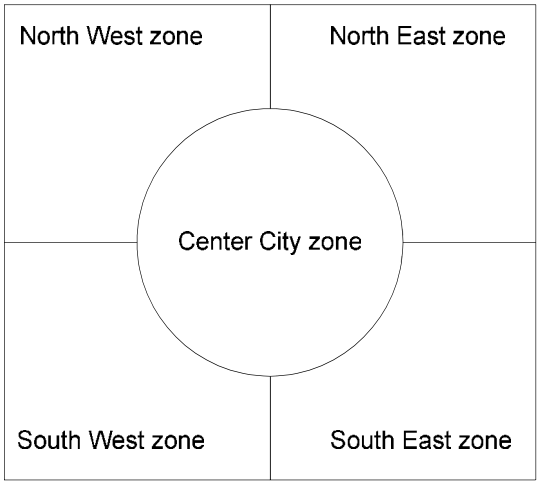

Figure 1 – Geographic Zoning

To deliver this type of zoning, the ad-insertable programming is replicated, generally in the headend, and each copy of the programming is sent through equipment which inserts ads for a particular geographic region. Then that programming is distributed to the appropriate geographic zones and is delivered to the viewers.

Migrating to more personalized addressability requires pushing the functions of program

replication and ad insertion further out to the edge of the network, and coupling these functions to a system which can make decisions about which streams to which viewers to personalize and how.

ARCHITECTURES FOR NETWORK-BASED ADDRESSABLE ADVERTISING

This section introduces both a method to do addressable advertising with current network infrastructure as well as an emerging architecture to deliver addressable advertising.

The architecture being developed with today's network components for linear addressable advertising is based on modifications to the On Demand system. Based on how VOD is deployed today, most of the resources are centralized at one or more super headends while other processing can be moved out to the hub to manage the metro network bandwidth utilization. If the transport capabilities are inexpensive enough and can handle the required streams, then centralizing the resources makes more sense. If the transport is more expensive or just not possible, then functions can be pushed to the edge, which is a key component of the emerging system to be described following the rest of this example.

Figure 2 shows one logical flow through this architecture. The network signals are received off the satellite which then goes to a stat-mux/splicer to be rate-shaped either as a feed for SDV, network PVR or as part of a multi-program multiplex. The splicer portion is used to insert advertisements at every available opportunity.

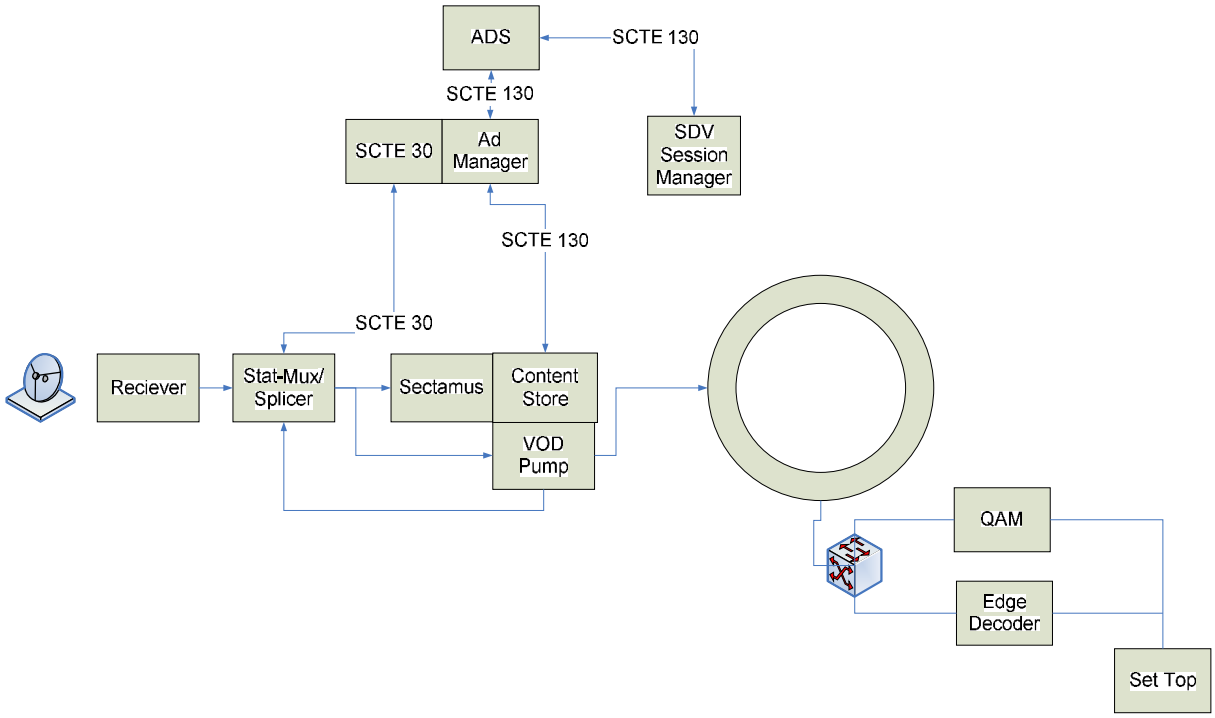

Figure 2 – Linear Addressable Advertising Based on VOD Modifications

These will typically be the advertisements for the largest geographic ad zone. By performing this initial splice, the network stream is prepared for easier content segmentation. This feed is then distributed onto the metro ring to the hubs to provide a constant reliable network feed source. This processed feed also then goes into the On Demand server complex for real-time acquisition. If additional ad zones are required, real-time feeds can be set up out of the On Demand server and new ads can be inserted by using playlist techniques from the On Demand server.

Figure 3 shows the emerging method to implement an addressable advertising system which is not based on the On Demand system. With this new method, the equipment is located closer to the viewers such as in a Distribution Hub. In this example, a new category of network component, the media services platform, is capable of the traditional stream replication and ad insertion, but also more advanced forms of personalization including managing bound applications (Enhanced Television, or ETV, and OCAP bound applications) as well as the possibility for creation of personal mosaics.

The media services platform allows the operator to grow the new addressable advertising service on purpose-built equipment without impacting existing services such as VOD. In this example, the media services platform ingests program streams, national or local feeds which come with default ads, and has the capability to personalize these streams by replicating them and inserting advertising specific for the viewers.

The media services platform interfaces with a personalization engine which is the decision maker that instructs the media services platform when and how to personalize streams. The personalization can include inserting addressed advertising, or a particular enhanced programming or to create a mosaic. The personalization engine can make decisions on what streams to personalize based on the household viewing of that content, their geographic location, assigned Prizm Codes or any number of other factors known about that household including specific opt-in opportunities.

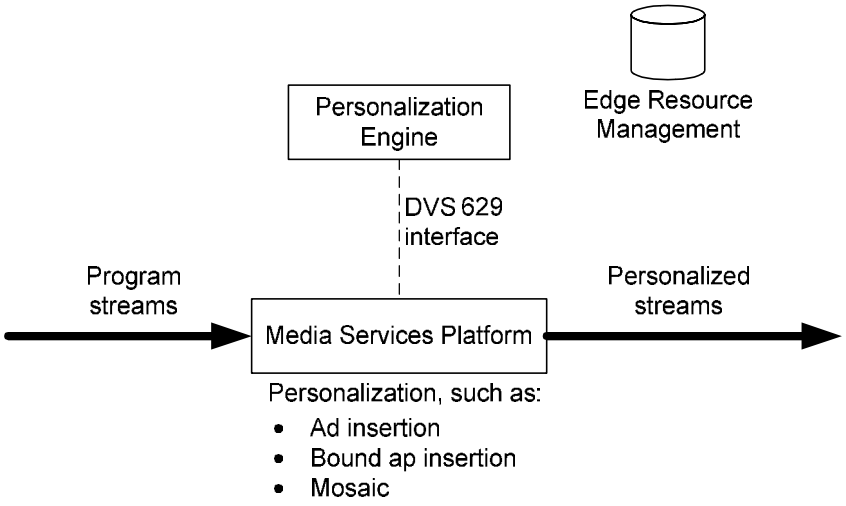

Figure 3 – Personalization System

When either a customer changes channels or an advertising avail comes up in a program stream, the personalization engine needs to make a real-time decision to decide if:

a) The program is ad-insertable;
b) There is an available campaign for this subscriber on this network;
c) Is there enough QAM capacity on the service group to do something like support replicating the stream for just that one viewer.

Separating the decision making from the actual delivery system allows for open systems. The SCTE (Society of Cable Telecommunications Engineers) DVS (Digital Video Subcommittee) is developing an interface standard, SCTE 130, which supports this separation of functions and will allow innovation to occur within both the personalization engine and media services platform while allowing operators to choose the components independent of each other.

Finally as shown in Figure 3, the addressable advertising system has to interface with an ERM (Edge Resource Management) system which is used to manage the digital QAM bandwidth to subscribers. Because the personalization system can cause program streams to be replicated on a service group, it uses edge QAM capacity more so than a geographic zoned ad system. As such, the personalization system needs feedback from the ERM system as to how much QAM capacity is available for personalization. During most parts of the day, the QAM service group is underutilized and there is excess QAM capacity available for personalization. It is only at peak viewing times that the QAM service group can become full and if the operator wants to continue high levels of personalization during those times, additional QAM capacity should be considered.

With a system as show in Figure 3, the options for delivering addressable advertising in linear programs are expanded from just geographic zoning to include three basic types of addressability which should cover the broad spectrum of advertising models to be developed in the foreseeable future; broadcast, groupcast and unicast linear ad insertion models.

Network-based addressable advertising uses SDV (Switched Digital Video) as an enabling technology for several reasons. First, SDV has every channel change request go through a server which allows subscribers to receive the appropriate personalized programming each time they change channels. Additionally, tuning information from set-top boxes, including both switched and broadcast networks, can be collected and used to understand how to apply demographics to viewers.

Broadcast is the current ad environment and is addressed through the specific network, the show, time of day, etc. As shown in Figure 4, a single version of the program stream is delivered to all viewers who all see the same ads. This solution is the basis for today's advertising business and cable is continually adding better audience measurement capabilities that will allow us to better utilize this segment.

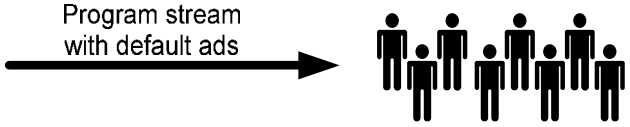

Figure 4 – Broadcast Advertising

Groupcast is a solution where viewers of like demographic are "grouped" onto a common copy of a program which is then enhanced for them. Take a program such as CNN which can have a wide variety of viewers, the personalization system would create several copies of CNN, one for each demographic which an advertiser is trying to reach and each copy of CNN would then be personalized with ads for that particular demographic.

An example of groupcast is shown in Figure 5 where there are three copies of the program stream created in the media services platform and viewers of particular demographic are grouped onto those copies which each carry ads specific for that demographic. In this case there are copies of the program generated for demographics A and B and a third copy of the program is available for viewers which do not fit these demographics.

Groupcast is very effective at demographic targeting-based techniques such as ZIP code or Prizm code. A household is either in a ZIP or Prizm cluster and doesn't move (at least during the duration of the program). Groupcast is not as effective as unicast at hyper-targeting households for certain ads based on a specific household (such as a specific credit card holder or not).

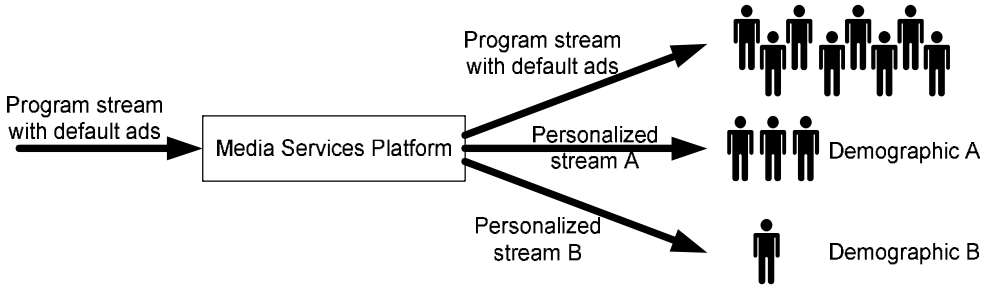

Figure 5 – Groupcast Advertising

Before a new groupcast feed is created for a particular demographic, bandwidth availability can be determined by polling the ERM. If bandwidth is not available then the additional program feed is not created which can result in a slightly lower yield, but requires no additional edge bandwidth spending. Reporting can indicate if this is a frequent occurrence and the operator might consider adding QAM capacity to those particular service groups.

Unicast is a technology where a single program stream is customized for a single viewer, as show in Figure 6.

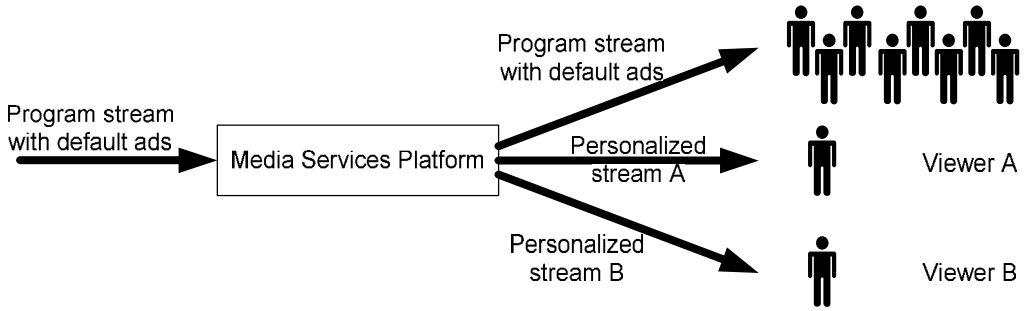

Figure 6 – Unicast Advertising

With unicast, whenever a viewer changes channels, the personalization system has to determine if a new unique copy of that program should be created just for the viewer. This solution is narrowcast, where there is a stream for each viewer and can be used to fill available capacity on a QAM service group.

While unicast means that every subscriber could have a dedicated stream, it does not necessarily mean that they will have a dedicated stream, just that if their demographics were such that they matched a specific set of available ads to play and bandwidth is available on the service group, a custom stream could be created for that single viewer.

Groupcast is a good example where the system can dynamically deliver advertising using both broadcast and unicast, as well as the specific case of groupcast. When used fully as a bandwidth saving multicast service, groupcast resembles broadcast. When QAM capacity is available, the demographic parameters used to replicate programs can approach unicast mode where it may be possible to create completely customized programs for each viewer.

The migration to unicast needs to be managed by taking advantage of available service group capacity, as more QAM capacity is added over time. Unicast does not have to imply that a separate stream is available for every viewer. Unicast may mean that separate streams are available for only some viewers. For example the QAM capacity of a given city may support 80,000 narrowcast SDTV slots, but the peak TV viewing population of that city may surpass that at certain times of the day. When this happens, it is simply not possible to deliver all unicast streams and some blend of broadcast, groupcast and unicast will be most efficient. It is the personalization engine which makes these decisions to best optimize the delivery of advertising and personalization while ensuring customers can watch as much programming as desired.

The goal is to add addressable streams based on bandwidth, demographics and ad availability. At startup, the system sets up an initial stream for every insertable network and assigns a base demographic to that stream. The system could also create a "spare" stream of every ad insertable network with no assigned demographic. This spare stream could be switched onto the service group if the personalization system determines it is appropriate to customize it for a new viewer (at which time a new "spare" stream would be created). If on the other hand there is no bandwidth available (or above a limit set for ad purposes) the new viewer gets tuned to an existing groupcast which best fits their demographics.

SDV ADVERTISING PHASED IMPLEMENTATION APPROACH

The following is a proposed phased approach to addressable advertising based on groupcast using switched digital video. The first three phases can be done with existing equipment and small modifications to the SDV session manager. Note that in this example transition the edge QAM resources will be managed through the SDV system to ensure the additional copies of programs do not overflow the QAM resources available on the service group. A more detailed discussion on QAM resource management is presented following the phased implementation approach.

Step 1 – Study the current ad insertable networks to determine which demographics are needed. Also learn which demographic groups watch the networks and during what times. For example if the goal is to personalize programming for Senior Citizens, don't advertise on Nickelodeon.

Step 2 – Use the data from step 1 to pick a small number of networks that typically have four demographics watching them. For a deployment of this scope, the SDV server can perform the demographic selection based on a look-up table with set-top box IDs, networks and the demographic groups. Ad selection in this case can be done with current traffic and billing systems by generating four schedules or, the system could choose to use an SCTE 130-based system. Alternatively in Step 2 one could use fixed demographics such as ZIP+4 or political party affiliation (democrat, republican, libertarian, unknown) as a stable demographic selector.

Step 3 – Allow the demographics on any given network to change during the day to better represent both the viewing patterns and the available ad inventory. This additional degree of flexibility will allow better addressability, albeit the combinations can become more specific and migrating to a true personalization engine may be prudent at this time.

Step 4 – Extend this groupcast example to all ad-insertable networks. This will maximize the revenue by bringing addressability to all the networks where the rights are available to insert ads.

Step 5 – Extend groupcast to unicast when there is an available ad campaign. Unicast advertising most closely resembles the Internet advertising model and promises the greatest revenue per ad.

With respect to QAM service group bandwidth, the following group of figures show the utilization of the QAMs in a typical SDV service group over a week, note these figures do not yet include addressable advertising but they can still be instructive to understand how viewers are using the linear lineup. While only about 50 networks are ad inserted, typically 150 or more networks are carried and only a percentage of viewers are watching ad inserted networks that would use the additional bandwidth if available. The non ad inserted networks would always use multicast and only one version per service group.

Figure 7 shows how the amount of traffic on the switched tier is allocated between unicast traffic (a single unique viewer on a program) and multicast traffic (2 or more viewers on a program) and indicates that there is a significant percentage of time when there is but a single viewer on a program which could provide an opportunity for unicast advertising without having to replicate an additional version of the program on the service group.

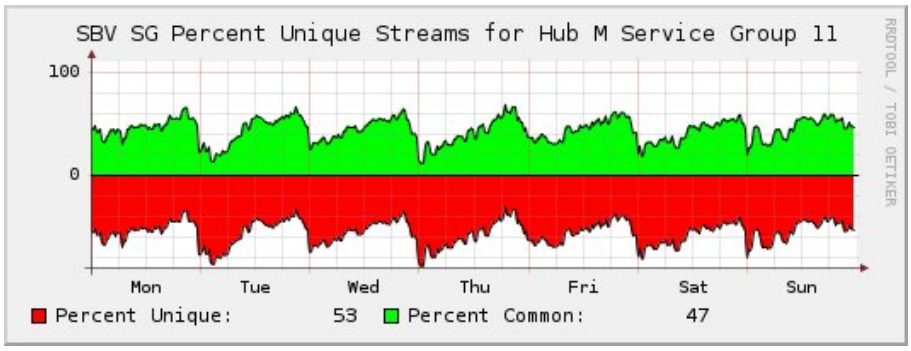

Figure 7 – Percent Unique Streams per Service Group

Figure 8 shows the number of active viewers on a service group. Note the periodicity of the graph with peaks building up through during prime-time each day.

Viewers can be reached at any time during the day with addressable advertising; however, the largest numbers of viewers are present during prime-time.

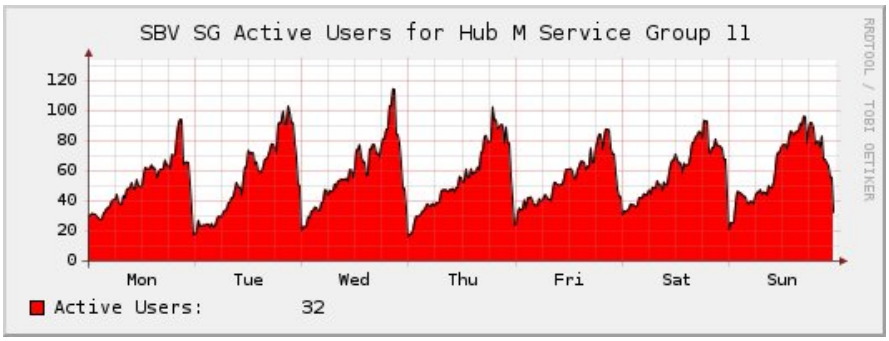

Figure 8 – Active Viewers per Service Group

Figure 9 shows the percentage of bit-rate used on a service group during the day. Again note the periodicity of the graph with peaks building up through during prime-time each day. As noted earlier in the paper, the stream personalization needed for addressable advertising uses more capacity in the service group then otherwise switched digital video would. Hence, overlaying addressable advertising onto this service group would require more bandwidth throughout the day

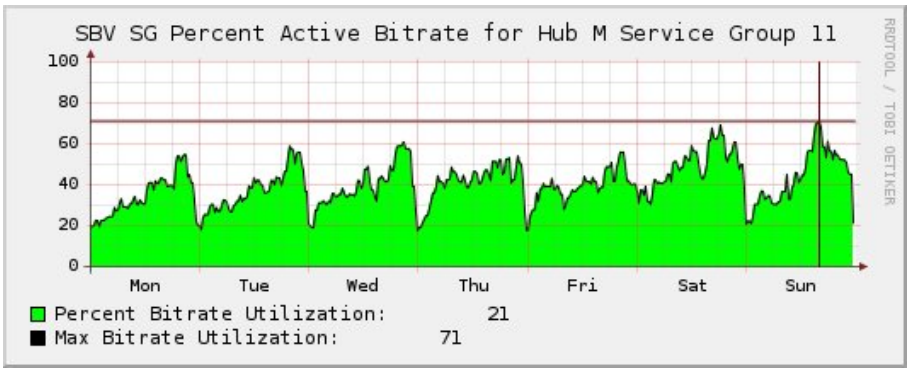

Figure 9 – Active Bit Rate per Service Group

And when coupled with the information from Figure 8, that more viewers are on the system during prime-time, it should be clear that to reach the most viewers with addressable advertising that additional service group capacity will be needed. However, because there is extra service group capacity available during off-peak times, it is possible to get started with addressable advertising with a switched digital system as engineered today. The addressable advertising system can make more efficient usage of service group capacity by filling up an otherwise empty service group with revenue generating addressable advertising during the times of the day when the QAMs might otherwise sit empty and un-used.

To more fully reap the benefits of addressable advertising as the system begins to pay for itself with additional ad revenues, the operator can grow additional QAM capacity into the service group to deliver additional addressable advertising during the times of peak service group usage. And once those QAMs are available, they can be used for other types of personalization as well including personalized guides and mosaics.

In general, a working philosophy for edge QAM utilization is to pick a value that if the QAM group is above the threshold, additional streams will not be used for advertising (unless the required network is not already in the Service Group). The current number is believed to be between 80 and 90% utilization. The following examples illustrate how the operator can get started with addressable advertising with the currently available QAM service group capacity.

Examples

Scenario 1: A viewer tunes to a program and there is already a version of that program on that service group assigned to demographic 3, but the addressable advertising system decides that this viewer should get demographic 7 in order to play different ads for additional revenue. The service group is currently only 20% full, so the SDV manager allows the creation of a new feed of the program assigned to demographic 7, places that program onto the service group and feeds the tuning info to that viewer.

Scenario 2: Same as scenario 1, except that the service group is now 93% full meaning there is no additional capacity to create a new feed of the program. In this case, the SDV server will join the viewer to a program that already exists and most closely matches their demographic.

Scenario 3: The viewer selects a program which is not currently switched onto the service group; however, the service group is 93% full. This viewer should get the demographic 7 version of the program. Since this is a tuning request for a program which is not already on the service group, the SDV session manager will add that program with a demographic 7 version and then tune the viewer to that program.

Because the addressable advertising system interfaces with the edge resource management system, the net effect of addressable advertising on edge bandwidth therefore should be effectively zero. By monitoring service group usage, if over-time there is not enough edge QAM bandwidth to run the scheduled ad campaigns, then the operator should consider a bandwidth expansion on specific service groups where capacity is an issue. But at this point in time, there should be definite revenue numbers associated with the addressable advertising which will offset the bandwidth expansion.

SUMMARY

Cable local and spot advertising is currently a $5 billion business, and it is widely accepted that adding addressability will grow that number. This paper presents an architecture for linear addressable advertising which builds upon existing linear splicing by adding Switched Digital Video as a means of personalization. Using SDV allows the operator to offer addressable advertising in a number of contexts including both switched groupcast and switched unicast.

The paper identifies several areas where planning may be necessary for addressable advertising including reviewing transport network capacity and switched service group usage. Planning can help ensure that the service meets the needs to better target advertising while giving the customer a better viewing experience, more HDTV and a more user friendly guide experience.