DYNAMIC INSERTION FOR SHORT-FORM VIDEO ON DEMAND ADVERTISING

John Chandler-Pepelnjak and Brent Roraback
Advertiser and Publisher Solutions, Microsoft

Abstract

Historically Video On Demand (VOD) advertising has required that programming content and advertising content be encoded and delivered together. In 2006, we deployed two field trials of "dynamic insertion," the run-time assembly of advertising and programming content.

This paper details the execution of these trials from an ad-server perspective. We explore the requirements of an ad-serving solution, on both the execution and reporting aspects of a campaign. Sections 1 and 2 provide introduction and background material. Section 3 details the technical requirements for campaign management and execution. Section 4 details the data that is reported on and best practices for its analysis.

INTRODUCTION

What Atlas does

Video On Demand (VOD) advertising has been labeled an "emerging media" channel for far longer than one would think it takes to "emerge." The delay has been due, in part, to the impressively complicated systems that underlie the delivery of cable television. Another part has been the reluctance of cable operators to innovate and thereby risk disturbing service. Finally, VOD advertising has been challenging for the agencies and advertisers that would be likely to use it—markets are small, deployment is tricky and creative requires long lead times.

In this paper we detail a solution that solves many of these agency and advertiser issues through the Atlas On Demand Media Console. Atlas, founded in 2001 and acquired by Microsoft in 2007, makes software for agencies and advertisers to plan, manage, deploy, track, report on, and optimize online marketing campaigns. For the dynamic insertion VOD trials outlined here, we used this technology to facilitate campaigns in two cable markets.

A brief history of VOD advertising

Historically, VOD advertising has been very similar to linear television advertising, at least from the perspective of the advertiser. Typically ad creative is produced, sent to a network where it is encoded with programming and distributed to operator VOD systems. Views of the programming content run via a request from the set-top box (STB) and the collective viewing of the ad asset and programming content is recorded. At the end of the campaign advertisers and agencies receive some data detailing number of views and the reach (the unique count of subscribers or STBs viewing the content). Critically, views of ad content are not separated from programming content views, obscuring the critical information on whether or not the advertisements were fast-forwarded or even viewed at all. Moreover, modifying the scheduled creative mid-campaign is impractical, if not impossible, due to the lead times required to re-encode and distribute the updated content.

For example, movie studios often take their creative material from the finished version of the film. These shots are sometimes ready only one or two weeks before the film is released. With typical VOD lead times of six weeks, VOD advertising becomes untenable for these advertisers. Additionally, it is often advantageous for these advertisers to change creative after the opening weekend, touting reviews from critics or other achievements (e.g., "Number one movie of 2008"). Given that the ad content and programming content are joined, simply swapping creative is impossible. There are many

other industries for whom the ability to change creative depending on external circumstances is valuable.

The final drawback of the current VOD system is perhaps the most damaging in the long-run. The internet's quick ascension in the marketing mix is partially due to the ability to target. Meaning, when an ad impression is called for, the content provider or network can use information about the viewer to determine the most effective creative messaging. In the current incarnation of VOD, this is not possible as all viewers of a piece of programming content will receive the same ads. The ability to make television addressable through dynamic insertion is a critical feature both for advertisers and for the medium itself.

The requirements for a new solution are clear. Advertisers must have the flexibility to insert creative into placements on short notice. This allows creative swaps, different ads targeted to different viewers, and evaluation of creative performance independent of programming performance.

CAMPAIGN MANAGEMENT AND EXECUTION

Background

"Campaign Management" refers to the set of activities concerned with the definition and management of advertising campaigns. Sellers of inventory describe the format and characteristics of media they have for sale. Agencies and advertisers record the inventory they have purchased from various sellers and the terms under which it was purchased. These business terms include information such as the cost method used for describing the unit of media (e.g., CPM or "cost-per-thousand" impressions, time-based costing), the cost per unit (or "cost basis"), target or guaranteed quantity of impressions to be delivered, date ranges, acceptable types of inventory (e.g., length of spots), etc.

Along with information describing the inventory purchased, the second important aspect of campaign management relates to the ads associated with the campaign. Information about the ads (e.g., identification codes such as ISCI or AdID, the names of the assets, the asset durations, etc.) is specified along with the inventory with which they are associated.

Campaign Management Solution

In dynamic advertising scenarios, campaign management becomes an active exercise: users receive frequently updated information regarding the status of their campaigns and have the ability to make changes to the campaign while the campaign is being delivered. In a non-dynamic scenario, such as broadcast television, the same degree of active management and detailed reporting does not exist.

Atlas has created a solution for agencies and advertisers to present instructions to ad execution and management systems directing these systems to display specific ads when a particular piece of content or inventory is delivered. Atlas also collects information about the viewing of advertising content, calculating metrics based on this viewership along with the business terms and goals under which the inventory was purchased.

Our solution, Atlas On Demand, interfaces with the ad execution and management systems over secure connections via APIs defined with our technology partners in this space (e.g., Aaris, SeaChange; Atlas is also a provider of inventory and ad management solutions for sellers of inventory, integrated with other partner systems). The set of services and message structure for managing this communication leverages

emerging industry standards such as SCTE 130 (DVS 629).

Within Atlas, campaigns are created and ads assigned to purchased inventory. Instructions are then published to the ad execution and management systems. When ad-supported content is requested by subscribers the content is assembled based, in part, on these instructions. Ad and program content are "seamlessly spliced" together and streamed down to the subscriber's set-top box (STB).

Campaign Execution

For a variety reasons – disparate, closed network systems; manual or semi-automated processes; emerging standards; pre-existing workflows; etc. – executing dynamic VOD advertising campaigns is still a very complex process. Tight coordination across a range of partners at multiple levels, from senior sales executives and content owners to network engineers and ad operations personnel, is required.

In order for Atlas to be able to communicate with the ad execution and management system deployed at the operator a secure connection must be established, such as through a VPN concentrator or other secure web service connection. The operator's endpoint, transport and protocol are managed through configuration settings in Atlas.

As noted above, program content and ad content have traditionally been encoded together as a single on demand asset. With dynamic VOD advertising, programs and ads are treated as separate assets, "seamlessly" spliced together at runtime and streamed to the viewer. In such a model, content providers and distributors must account for "ad free" versions of the content suitable for dynamic VOD ad campaigns, as well as the ad assets themselves, insuring that the complete package of program and ad assets is distributed appropriately.

To insure the viewing experience in dynamic VOD is of the same quality as other on demand viewing, content encoding standards must be rigorously followed. In our trial campaigns, CableLabs OD encoding standards formed the basis for theses specifications but extra care was taken to insure that bit and frame rates, resolutions, and audio were identical for all assets. Assets were also required to start and end with silent frames of black to ease the transition between assets. The "seamless splicing" of assets, mitigating any remaining discontinuities between MPEG files, was either accomplished in software by the VOD system or in the edge device (i.e., the QAM) level, depending on the VOD system provider's approach.

Aside from system configurations and content preparation, the standard campaign workflow generally begins with media negotiation: sellers of media (content networks, operators) offer packages of inventory to buyers of media (agencies, advertisers). Rates, schedule and other terms are negotiated and agreed to through terms and conditions, insertion orders, etc.

After the contract is finalized, ad campaign information is configured in the respective inventory and campaign management systems of the buyers and sellers (information may already exist in the seller's system, enabling the seller to forecast and book inventory). Ad assets are distributed, generally to the inventory seller but in some cases to the operator directly. Ad and program assets are encoded per the on demand specifications and distributed to the operator and headend systems, usually through existing "pitch – catch" mechanisms but potentially through IP-based distribution.

Once campaigns have been configured and ad and program content distributed, ad

instructions are published by the agency/advertiser to the inventory/ad management system at the operator. Instructions are validated by the operator system and acknowledgements returned to Atlas. If the instructions are valid they will be referenced when ad-supported content is assembled.

Viewers request ad-supported VOD content, initiating sessions with the VOD system. The Atlas instructions are referenced in the assembly of the on demand content, ad and program content is seamlessly spliced together and streamed down to the viewers. Viewer interactions with the content ("trick mode" activity, such as fast-forwards, pauses, rewinds) are recorded by the VOD system. We collect detailed information regarding viewing and playback of ad assets. Data is imported in Atlas' reporting system and metrics are calculated.

Atlas users view statistics related to their campaign's performance through online reports. Statistics are updated several times per day. This granular viewing data allows Atlas to calculate and display multiple metrics describing campaign performance. Impressions (i.e., views irrespective of playback speed), "Brand Exposure Duration" (BXD) (i.e., viewing duration at normal playback speed), completed plays, reach, trick mode counts, and more may used to compare the performance of ads and/or their associated inventory. Campaign performance is assessed by analyzing the metrics corresponding most closely to the advertiser's campaign goals. Our users may then apply this information by "optimizing" their campaign: changing the ads assigned to their campaigns and/or the business rules governing ad rotation to maximize performance. New instructions are then published to the VOD system and enforced during subsequent viewing sessions.

REPORTING AND DATA ANALYSIS

Data description

As mentioned previously, one of the principal benefits of dynamic insertion is the ability to measure ad performance separately from programming performance. In other words, as users fast-forward (FF), rewind (RW) or pause their ad programming, we collect data on each trick-play and can use that to measure viewership.

The data that Atlas collects come in two different styles. The simplest conceptually is what we call "event-level records". These records detail every trick play event and capture the following pieces of data (or meta-data) associated with the event:

- Date
- Time
- Operator
- Headend
- Masked MAC address
- Ad Asset Name
- Ad Asset ID
- Event (Setup, Play, FF, RW, Pause, Teardown)
- Event Speed (1 for play, positive for play or FF, negative for RW)
- Programming Content Name
- Programming Content ID

In data of this format, one row of data represents one event of ad viewing. A viewing session is defined as beginning with a Setup event followed by a Play and ending with a Teardown. There can be any number of interstitial trick play events (FF, RW, Pause). The most complicated data field is masked MAC address. Typically the unique identifier of a STB must be masked for privacy reasons. It is important that the masking algorithm be 1-1 so that no two MAC addresses can be mapped to one masked MAC address so that reach and frequency can be accurately calculated.

Occasionally the detailed data format is not available. In that case, there is an alternative data format ("aggregate-level") where one row of data represents one (potentially partial) view of an ad. In order to use this data, additional fields must be added. These include the following:

- Start Date
- Start Time
- End Date
- End Time
- View time (amount of time asset was viewed in normal playback speed)
- Fast Forward Count
- Rewind Count
- Pause Count

The key metrics that can be derived from these reporting fields are worth pointing out explicitly. All traditional campaign measures—GRPs, impressions, reach, and frequency—are available. Additionally, we can look at viewership patterns by time of day and day of week. Finally, a variety of user-level reporting metrics are available and these will be detailed later in the section.

Measures of Performance

The VOD landscape is crowded with many metrics used to evaluate performance across many dimensions. Ads, placements, asset lengths, and content providers are judged by a variety of yardsticks. Though our dynamic insertion trial, two metrics emerged as the most critical for evaluating these campaign attributes. Because VOD is an accountable media, we take the time to highlight these two different performance metrics and detail their implications.

The first measure has already been mentioned: Brand Exposure Duration (BXD). This is simply the amount of time an asset is viewed in normal playback speed. If you sum up all the BXD values for every impression on a campaign you have "Total BXD". While simple at first glance, Total BXD is a powerful omnibus metric, combining length of asset, asset average view time, and total number of views. Increasing asset length has the effect of decreasing the average view time (jumping from a 15 second to a 30 second spot diminishes viewership) but our research indicates that net viewership typically increases with increasing asset length. We can also look at BXD in several different ways. Average asset BXD can be used to optimize creative—if this method is followed total viewing time will be maximized. Alternatively, if all creative are in rotation in a placement (say, the first commercial break of an ad-supported VOD program) then average BXD will determine the value of the placement and the suitability of that program's audience to the creative message. All else being equal, BXD will tend to choose longer commercial assets. From the data we have seen across VOD on television and video on the web, doubling an asset's length rarely cuts the average percentage of the asset viewed by half. Finally, BXD is a useful cross-platform measurement. BXD can be calculated for video across any screen and is a useful measure of engagement across platforms. Note that the efficacy of *any* video metric diminishes if the ability of users to FF is disabled.

The alternative to BXD that emerged during our trials was "completed plays". This metric can be defined in multiple ways but the simplest is a view of an asset with no FF activity. (An alternative is to look for every second of an asset being viewed at least once in normal play mode, though this is more complicated.) Assets are compared with each other based on the percentage of views resulting in a completed play. Whereas BXD tended to reward longer assets, completed plays unequivocally skew results towards short assets. Completed plays are, however, the only metric that makes sense for certain classes of creative assets. Commercials where the call-to-action or brand message is delivered in the closing seconds require

optimization based either on completed plays or on a "weighted BXD" where certain portions of the asset are worth more than others.

Results from Data Analysis

The two trials detailed in this paper had different compositions. The first had two pre-roll creative assets for one advertiser running over two different time periods across one content provider. This trial allowed us to prove the concept in a simple environment. The second trial was much richer from a data analysis perspective: two advertisers and two content providers running multiple assets of varying lengths in both pre-roll and post-roll positions. In the results that follow we focus on this second trial. It is easier to follow the results when speaking about a specific campaign and the results themselves are much deeper with these data.

As mentioned above, two advertisers took part in the second VOD trial. The first advertiser had several assets in rotation, all 30-second spots, running in both pre-roll and post-roll. The second advertiser had pre-rolls of both 15 seconds and 30 seconds, followed by post-rolls of 60 seconds and 120 seconds. These assets were all in rotation on both content providers.

When reporting on short-form, dynamically-inserted VOD campaigns, there are a number of standard metrics that barely need mentioning in this forum. Although fundamental to campaign evaluation, metrics like reach, frequency, and impressions (divided into various time ranges and publisher and placement groups) are straightforwardly defined elsewhere. Instead we will focus on a series of analyses we performed during the trials that were unique to the dynamic VOD environment.

Initially there were two pieces of "conventional wisdom" that we wished to analyze regarding pre- and post-roll advertising. Are longer pre-rolls fast forwarded more often than shorter pre-rolls? Do viewers stick around to watch post-roll advertising?

It goes without saying that in our data the pre-roll ads received a higher impression count and higher completed plays. This is nearly tautological. Due to the greater length of the post-roll, however, the post-roll commercials resulted in longer BXD. Again, BXD is the average or aggregate duration in minutes that an ad or a brand (if multiple ads) is watched. This mirrors research we typically see with digital video: longer assets perform better from a BXD perspective, shorter assets perform better from a completed play perspective.

Viewership on post-rolls was surprisingly high, including the number of completed plays. This would indicate an undervaluation of post-roll ads given the common assumption of little to no viewership. One factor contributing to the longer BXD was the longer durations of the creative used in the post-roll positions compared to the creative used in the pre-roll positions. On average post-rolls were viewed approximately 36% of the way through (versus 42% for pre-rolls).

Unsurprisingly, our analysis revealed that the 30 second pre-rolls were fast forwarded more than the 15 second pre-rolls. Since one advertiser had only 30 second spots, it is possible there was some burn-out, although this behavior (more FF activity on longer spots) is not atypical.

There was evidence to suggest that viewership of post-roll ads can be augmented by pre-roll ads from the same advertiser. In other words, "bookended" placements (with a single advertiser in both positions) are more valuable than pre- or post-roll placements alone or from distinct advertisers.

From a viewer perspective, post-roll viewership of bookended placements might be perceived in much the same way as mini-long-form advertisements. Indeed, our research indicates that, if maximizing BXD is the goal, advertisers would do well to treat post-roll viewership as a "conversion" and focus on pre-rolls that are most likely to achieve viewership of the post-roll. Given the viewers discretion to watch or FF in this context post-roll viewership starts to look like a click or performance-type media on the web. Our research also indicates that longer pre-rolls (:30s over :15s) make viewers more tolerant of longer post-rolls.

Another question we asked of the data was, Is there an interaction effect between advertiser and content provider? A display of the data appears in Figure 1.

This figure shows the interaction between the advertiser's commercial content and the programmer's content adjacent to which the commercials run. On the y-axis we see BXD expressed as a percentage of asset length—higher numbers indicate more of an asset was watched. As we can see, Programmer 2 performed better overall, but there is an interaction between the advertiser's content and the programmer's content (indicated by the crossing of the lines). This is important: not all programming is optimal for all advertisers and advertisers may reasonably value different pieces of VOD inventory differently. Currently many

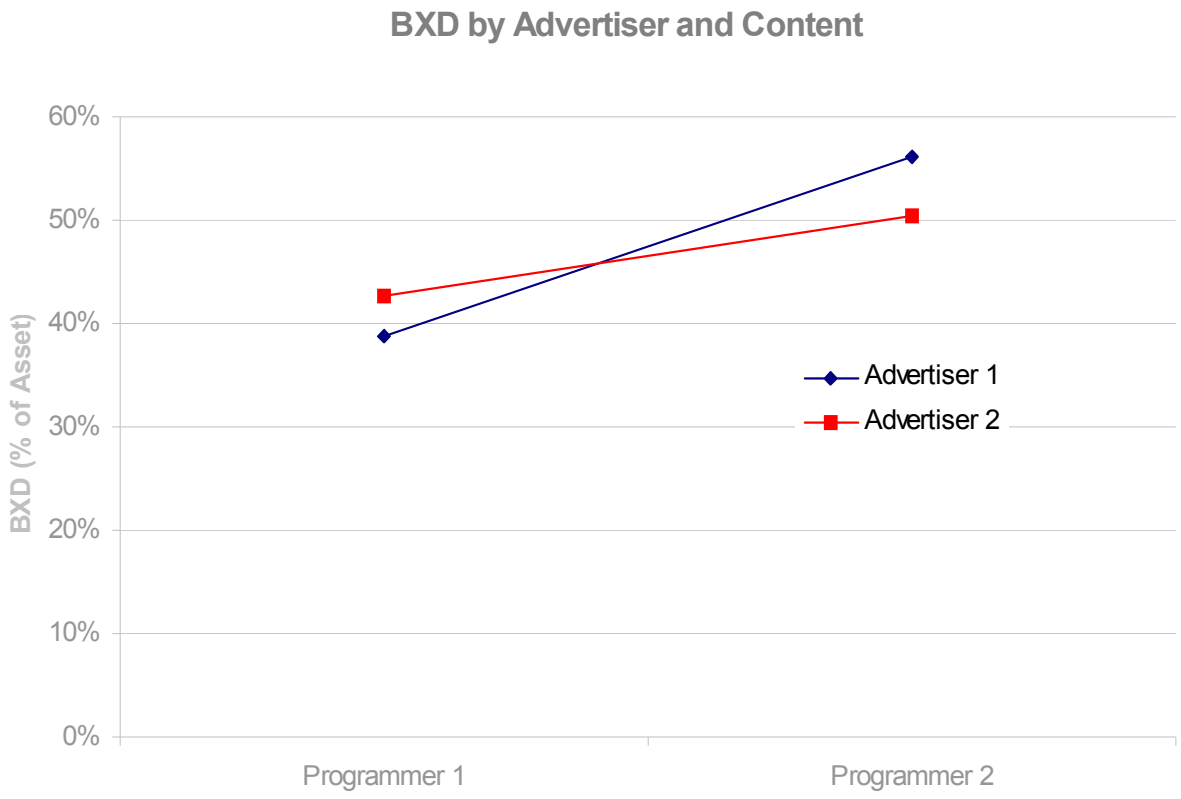

Figure 1: In this figure we see the interaction between the advertiser's commercial content and the programmer's content adjacent to which the commercials run. On the y-axis we see Brand Exposure Duration (BXD) expressed as a percentage of asset length—higher numbers indicate more of an asset was watched. As we can see, Programmer 2 performed better overall, but there is an interaction between the advertiser's content and the programmer's content (indicated by the crossing of the lines).

marketers believe that pre-roll video viewing is dependent entirely upon the content following the spot. This research contradicts that view.

There was, however, a tradeoff to consider between reach and impressions on the one hand and increased duration on the other.

As for asset lengths, longer advertising assets generally resulted in longer viewing durations and more total minutes viewed, shorter commercials enjoyed more completed plays and were watched in their entirely a higher percentage of time.

Although you might assume that the average BXD percentage might decrease as the length of the commercial increases, it was found that the percentage of :120 spot viewership was higher than that of the :60s.

Post-rolls were watched approximately the same percentage of the time regardless of asset length (roughly 35%). This confirms previous research conducted by Atlas on long-form VOD advertising. One interesting effect of post-roll viewership was noted. For one advertiser, viewership of :120s in the post-roll position increased by 20% if the pre-roll was a :30 instead of a :15.

The :30 commercials of one participating advertiser performed better in the pre-roll than in the post-roll in terms of both BXD average minutes and as a percent. Some of this effect could be due to the repetition of the same ads given that if the same spot appeared in the pre-roll position as in the post-roll position a viewer would be predisposed to fast forwarding through the second appearance.

CONCLUSION

There are two critical components in the deployment of a short-form VOD campaign. The first is the ability to package ad content separately from programming content. This is fundamental to many needs of advertisers including creative management, creative swapping, decreasing creative lead times, accurate reporting. The second is detailed measurement of individual ad viewing duration. The first requirement is a near-prerequisite for the second, but it is only with this measurement that the power of the addressable television medium is achieved.