TELEVISION 3:0 - THE MERITS AND TECHNICAL IMPLICATIONS OF CONTROLLED NETWORK AND CLIENT CACHING

Edmond Shapiro, VP Project Delivery Americas
NDS, Ltd.

Abstract

The cable industry has long debated the merits of using general purpose devices to access cached information in the network (commonly referred to as "cloud storage") as opposed to using cached information stored locally (on a device or within a home network), and in what combinations. In the past these trade-offs have involved the location of video on demand (VOD) and digital video recorder (DVR) storage. Today technical design decisions have become even more complex as engineers grapple with the growing number of caching permutations that will facilitate the deployment of Television 3.0, the next generation of IP based advanced digital cable services.

This paper analyzes network design considerations that cable engineers should consider when architecting Television 3.0, the next generation of IPTV applications using an array of cacheable information that includes: Application logic (cached JavaScript), Presentation logic (Remote and Local User Interface], Content (Adaptive Bit Rate (ABR) and Progressive Download (PDL) files as well as Metadata and Network Interfaces.

By highlighting the importance of an abstraction layer herein referred to as the Television 3.0 Common Service Framework, this paper explores hybrid architectures that permit network operators to dynamically cache information at multiple locations within a network – to dynamically adapt deployed services from one type of device to the next and from one region to the next – constantly evolving as new devices and network resources are made available in a rapidly changing technological environment. Smart software design builds an agile, future-proof foundation to increase deployment velocity of advanced services and enhance the operator's brand through improved system performance and better user experiences. Smart software design also avoids the many pitfalls of the past that have afflicted cable operators – from outdated devices and vendor lock-in, to degrading performance and feature bloat, to network-wide equipment upgrades in support of new services.

Specific applications and services highlighted in this paper include:

- *Content protection solutions – Conditional Access System (CAS), Digital Rights management (DRM)*
- *Content Delivery Networks – global, regional and federated*
- *Content recording and playback – DVR scheduling, resource allocation*
- *Ad insertion –graphic and video*
- *Multi-screen service delivery*

BACKGROUND

Television 1.0

The television application is a communication technology for the sharing of moving images with a group of people: the "mass media". The television transport network is more efficient when it can deliver the same information to more than one person at a time, as was originally the design of the analog terrestrial, cable and satellite networks.

The *Moving Pictures Expert Group* (MPEG) standards enabled the broadcasting of digital television services for the first time, within this paper referred to as the *Television 1.0 service*. As with analog, the Television 1.0 ecosystem was designed to efficiently transport broadcast digital information to a mass of people.

Digital television rapidly increased the amount of content or number of channels that a consumer could view at any moment in time and therefore the Electronic Program Guide (EPG) application was invented to improve content discovery.

With few exceptions (parental controls for one), the EPG user experience remains the same for every viewer. The underlying content changes (Linear and VOD events), but the EPG screens remain constant.

Web 1.0

Though the Internet is defined as the inter-connect of many different private and public IP networks for the purposes of sharing information, in the minds of most consumers, the Internet has become synonymous with the World-Wide Web (WWW) or just Web.

When the Web was invented, the Web user experience was as impersonal as the Television 1.0 experience. The Web was made up of pages written in a HyperText Markup Language (HTML), transferred from one computer to the next using the HyperText Transfer Protocol (HTTP). Each web browser eventually saw the exact same screens (same experience as Television 1.0).

Differences in user experience from one user to the next arose from the distance travelled by the HTML file. If a file crossed too many networks, it might be slowed down or even stopped altogether. The Web doesn't care whether a file is located on a computer in the same city or on the other side of the world. The response time for every Web request is unpredictable because there are no guarantees of reliable transport between the web server and the clients. User experiences on the web are considered "best effort" because of this unpredictability.

The first implementation of these standards came to be known historically as "Web 1.0". As the Hyper*Text* name implies, the first HTML files were simply text files carrying textual information. As long as text was the only content being broadcast on the web, the size of the text file made little difference to the user experience. Best effort was simply good enough, even when accessing files across very slow networks.

However, user experience designers quickly grew frustrated at the impersonal Web 1.0 experience and began to deploy "richer" graphical user interfaces including non-text based files, so called "binary files", such as music, photo and video files, which led to an explosion in the size and number of files managed by web designers (see Figure 1).

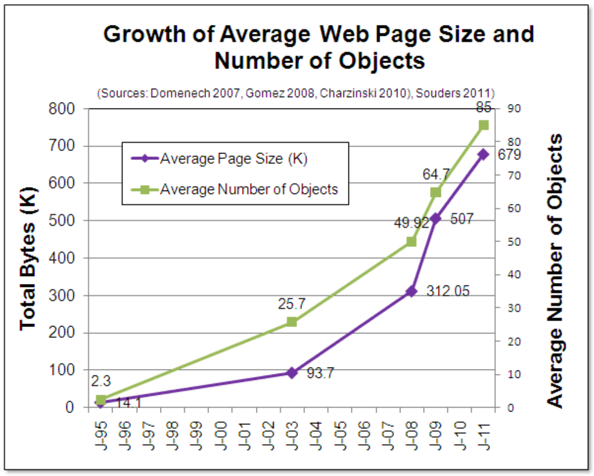

Figure 1: Growth of the Average Web Page

The best effort mode of the Internet was unable to adapt to the demands of the Television 1.0 application developers. Broadcasting (or multicasting) IPTV over the Internet has been notoriously difficult to achieve. Every network between the content distributor and the consumer has to agree to pre-allocate enough bandwidth to carry the fixed bandwidth required by the IPTV service. Only a private network, managed by a single network operator, has ever effectively scaled network capacity to implement this application (e.g. AT&T UVerse).

Web 2.0

User experience designers soon adapted and began to dynamically generate HTML files exposing a richer more personalized user experience for each user and device type. This personalization of web services became known as “Web 2.0”, or the social web, and has been exemplified by the success of web service providers such as Facebook and Twitter.

The successful Web 2.0 service providers learned to work around the “best-effort” design of the Internet by overlaying a virtual network on top of it, called a Content Delivery Network (CDN). The CDN was used to rapidly distribute and store (or cache) the many media files to as many edge networks as possible, as close as possible to the mass of the consumers, thereby avoiding network overload and reducing the number of network hops between the consumer and the media files, and thus reducing the unpredictability of the Web user experience.

CDNs take advantage of specialized algorithms to redirect HTML hyperlink requests to the nearest cache location of the requested media file. These algorithms are constantly and dynamically evaluating network boundaries to avoid bottlenecks and to determine optimal routes.

Over time this CDN virtual network adapted to many different uses as the various types of Internet connected devices exploded. Where once a web application could assume that all web browsers were located on a Personal Computer with a single screen size and a local cache, now mobile devices with smaller screens and no cache had to be accommodated. By delivering different size graphic files and deciding upon remote or local caches, the CDN was able to optimize content delivery to each type of device.

An equally important Web 2.0 development was the widespread adoption of a standardized programming language called JavaScript and the Extensible Markup Language (XML) that enabled web application developers to selectively retrieve parts of the HTML file based on local context or inputs (e.g. user actions, cookies, etc.), as opposed to retrieving the entire HTML file.

Television 2.0

With the growing capacity of CDNs to deliver rich media to Web 2.0 users, demand grew for the delivery of streaming media services such as Radio and Television. As traditional IPTV could not be predictably delivered over the Internet, new CDN friendly techniques were required.

Adaptive Bit Rate (ABR) technology arose from this challenge. Where traditional IPTV content was pushed at a fixed rate, ABR content is pulled at a number of different bitrates that can be influenced by the CDN as well as the local application context.

Instead of broadcasting a common monolithic media file to every consumption device, ABR technology delivers an index file (or manifest file) instead. This abstraction allows for different versions of the media file (size, resolution) to be cached and consumed at any time or place. Where the CDN is able to locate a faster network or when a client connection improves, then the user's viewing experience improves accordingly.

ABR technology has become the foundation for a generation of *Television 2.0* services. These Television 2.0 services have enabled service providers to deliver television to any type of device, and no longer just to televisions. ABR is particularly well suited for unmanaged or minimally managed networks, especially home WiFi networks.

Different ABR standards have competed in the marketplace, being driven by the commercial interests of major CE device manufacturers (e.g. Apple, Microsoft). MPEG recently standardized the *Dynamic Adaptive Streaming over HTTP* or DASH specification which incorporates a number of these approaches.

<u>Cable's Challenge</u>

Cable network operators have watched as popular Over-The-Top (OTT) service providers have taken advantage of ABR technologies to deliver high-quality (even HD quality) services over their fixed and mobile IP networks.

ABR technology is so efficient that it will squeeze nearly every bit of available bandwidth from the access network, limiting alternative services that the cable operator might wish to provide over that same network.

In some countries "net neutrality" regulations restrict cable operators from differentiating or prioritizing any type of broadband service. Cable operators are forced to implement service neutral bandwidth and data caps to control the growth of these OTT services.

The impersonal EPG of legacy Television 1.0 services cannot compare to the personalized and social media experience of many Television 2.0 applications. This is ironic for a cable operator given that the legacy broadcasting Television 1.0 network is generally more efficient at distributing television services.

To compete, cable operators have deployed their own multi-screen Television 2.0 services. However, these same net-neutrality regulations are being tested as subscriber usage limits may be ignored when using the cable operator's own Television 2.0 delivery networks.

From a technical perspective, there is little difference between a private IP sub network and the legacy Quadrature Amplitude Modulation (QAM) as both are based on the same MPEG networks that cable operators used to deliver legacy broadcast and narrowcast television services. All of these services must coexist on the same access network or "last mile".

As ABR encoding (or transcoding) can occur at any control point in the content delivery network, regulators will find it increasingly difficult to make these net neutrality distinctions.

A consumer who chooses to stream ABR content on their own has the option to purchase a CE device, such as Sling Media, to transcode and transmit set-top box content to any other device. A cable operator who delivers the same ABR service from the network will not only save the consumer the purchase price of such a device, but will reduce energy costs for all consumers by centralizing this functionality in the “cloud”.

Television 3.0

OTT services delivered by global CDNs to any network on Earth are ideal for content creators and broadcasters who wish to expand viewing audiences, but only if the access networks are capable of delivering their service.

Broadband service providers who control the access networks will compete on the capabilities of their network infrastructure and will be judged by consumers on the user experience of these OTT services.

Though these various service providers may compete, there is also an incentive for them to cooperate. New *Television 3.0 services* will utilize these same CDN and ABR technologies but in a more network aware fashion.

Where Television 2.0 services utilized CDN overlays and local device context to optimize the user experience, Television 3.0 services will go further, benefiting from a deep and intimate knowledge of the underlying IP networks, and leveraging a *Common Service Framework* to abstract the user experience.

Caching and abstraction techniques led to the advances in Web 2.0 and Television 2.0 content delivery. Extending these techniques with the addition of new communication and collaboration tools will be the foundation for the common service framework of Television 3.0.

For example, publish and subscribe techniques will make it possible for an access network provider to expose to a Television 3.0 service provider the caching resources of the Edge CDN closest to the subscriber. Television 3.0 service providers who design their applications to account for these network variations will inevitably produce a better user experience.

Today most home networks are minimally managed by professionals. This makes the home network the last frontier for a managed IPTV service. Television 3.0 service providers will expand their management of the subscriber’s home networks as well, such that the Edge CDN may very well be within the subscriber’s home.

Television 3.0 services can be extended into the subscriber’s home network only if a home gateway is capable of supporting these services. For instance, a managed gateway supplied by the cable operator could assure the bandwidth needed to supply a 3D service to every 3D capable device within the home network.

Such a technique will make possible other advanced services as well. For instance, managing the home network for the subscriber would enable plug and play video-conferencing, home security and energy management services to co-exist with television on most devices.

A Common Service Framework that is network aware will incorporate all of these advanced services into the service provider’s branded user experience across any device and on any network. Television 3.0 features that are available only in the home network will be disabled on the road. Other features that are only available with a higher subscription tier will be managed in a common fashion across all devices.

The Television 3.0 service benefits from network awareness but should not be dependent upon it. All of the components of the Common Service Framework must be optimized for use within and outside of a managed IP network. For instance, this means that a DRM client application should provide robust enough security to validate a user and their consumption device anywhere in the world, and not just inside of their home or within their cable network.

Finally, as the Internet itself adapts to these new paradigms (e.g. Software Defined Networking), the Television 3.0 service will fully enable consumers everywhere to experience the true "TV Anywhere" service that they desire.

LEVERAGING ADVANCED CACHING TECHNOLOGY

Usage Context

The term "application framework" describes a software structure for developing software applications within a specific operating system or environment. The responsibility of the application framework is to provide "context" to the users of the framework, which are a set of software components called clients or "applications".

In a Television 1.0 service, the application framework that provides a common set of services for accessing the television transport stream is commonly referred to as "Middleware".

Middleware enables a common set of applications, such as the service provider's Electronic Program Guide (EPG), to share set-top box hardware resources, without being tied to a specific set-top box manufacturer. *Digital Video Broadcasting Multimedia Home Platform* (DVB MHP) and CableLabs *OpenCable Application Platform* (OCAP) are two standardized sets of middleware functions or APIs.

The Television 1.0 service has a relatively fixed or static usage context. The users of a middleware framework are expected to reside in a fixed location on a single set-top box attached to a single television. Dynamic events are limited to remote control or front panel user inputs or network control messages that are typically managed by a conditional access function.

The Television 2.0 service has an equally fixed or static usage context as well. OTT services are typically manipulated at the source (in the network) to conform with the requirements of a specific CE (Consumer Electronic) device manufacturer's chosen application framework, typically Microsoft (Windows), Apple (iOS) or Google (Android).

The Television 3.0 service has a much richer usage context, as it is designed to flexibly adapt to a more complex set of environmental variables. A Television 3.0 service may be executing on a fixed device such as a set-top box, or on a mobile device such as a Tablet.

On a fixed device, the Television 3.0 service must adapt to the same usage context as a traditional set-top box, for instance, by supporting a front-panel interface. However, the Television 3.0 service might also support the geospatial feedback that it acquires from the mobile hand-held device.

All Television 3.0 applications whether fixed or otherwise will respond to the same network control messages including subscriber entitlement or service feature changes. The service and content protection function of the Common Service Framework (historically referred to as Digital Rights Management or DRM) must constantly validate the usage context of the media consuming application, ensuring that the content distributor's commercial agreements are respected and that content will not be used for any purposes other than the intended ones.

Within such a complex operating environment, prioritization of usage context becomes critical. For example, in a telephony enabled device, such as a smartphone, the application framework may need to determine for each event whether the television application or the phone application takes precedence. The same Television 3.0 services running on a smart TV, smartphone and tablet will react differently to each type of event, and may in fact be programmed by the user to react differently.

Though service providers have the option of developing a different set of Television 3.0 applications, each optimized for the specific CE manufacturer's application framework and unique usage context, this will inevitably be seen as a costly and infinitely expanding endeavor, as all of these applications will need to be supported and maintained indefinitely.

Alternatively, Television 3.0 service providers will draw upon the experience of Web 2.0 service providers by abstracting their services through the use of the Common Service Framework.

The Television 3.0 service provider will balance the goal of a complete abstraction layer that minimizes device specific development, with the desire to leverage unique device specific capabilities (e.g. larger or higher resolution screens, better memory management, unique man-machine interfaces, hardware security hooks or greater portability).

The HTML5 standard, currently in development, is expected to facilitate the deployment of a common set of applications and services across compatible devices. HTML5 includes richer graphical capabilities and more complex JavaScript application logic. As device specific application frameworks standardize on the JavaScript standard interfaces, Television 3.0 service providers will deploy more features in a common fashion, and thereby reduce their dependency on device specific or downloadable application frameworks.

As an example, HTML5 utilizes JavaScript to abstract the usage context of the local device. Through the use of JavaScript to access the device's local storage, effectively extending the virtual CDN network into the device, Television 3.0 application developers will be able to cache service information that improves the predictability of the user experience, potentially approaching the reliability that consumers experienced in Television 1.0 applications (at least within managed networks).

Built-in DRM functions may already exist in many CE device specific application frameworks. However, in order to achieve a common set of security functions across every device type, the service provider will inevitably require a global set of content and service protection functions to be included within the Common Service Framework. These security functions may leverage device specific security capabilities or usage context, but must never be completely dependent upon them.

For example, to ensure that there exists a trust hierarchy, the Common Service Framework might leverage any hardware based personalization or security functions that are exposed by the device manufacturer (for example Unique Device IDs). Alternatively, DRM clients may be integrated with the device's application framework, such that a DRM application may be downloaded to provide dynamic security hooks that may be leveraged by the trust functions of the Common Service Framework.

Context Control

For a service provider who is designing a Television 3.0 service across fixed and mobile devices using a Common Service Framework, a key decision is whether that service should take advantage of network specific features or whether it should be agnostic to the underlying transport technologies.

If for example a cable operator deploys a network agnostic (Television 2.0) application on a tablet alongside a traditional set-top box application (Television 1.0) over the same access network, then the network resources required to deliver the same quality of service to both devices may end up being twice the resources that would have been required if both devices had implemented a common Television 3.0 application that utilized a network aware Common Service Framework.

Most broadband service providers have already adapted to the demands of Television 2.0 service providers by scaling their access networks to enable ABR to coexist alongside legacy analog and digital cable television services. To avoid the overhead cost of simulcasting ABR video over the same access network that already delivers similar content in a traditional digital video format, the network operator must allow the Television 3.0 Common Service Framework to interface with the legacy network controller systems.

To permit a Television 3.0 service to interoperate with the legacy digital cable plant requires a complex interoperability design that adapts existing Television 1.0 infrastructure to the Common Service Framework. This includes the ability to leverage legacy System Information (SI) and Conditional Access (CA) services already being transported alongside the broadcast digital video service.

The access network operators and service providers must agree on a context control interface to communicate the availability of hybrid or legacy services, and to enable the Television 3.0 service provider to control the access network transcode or transcrypt resources that would be required to convert content to the required consumption format. This includes transferring the service protection metadata of the legacy conditional access system to the Common Service Framework for use by the Television 3.0 applications.

The advantage of deploying a network aware Television 3.0 service is that it can be made to be more scalable by reducing demands on the access network resources. The disadvantage is that interoperability costs may be greater for the network aware Television 3.0 service (see Figure 2).

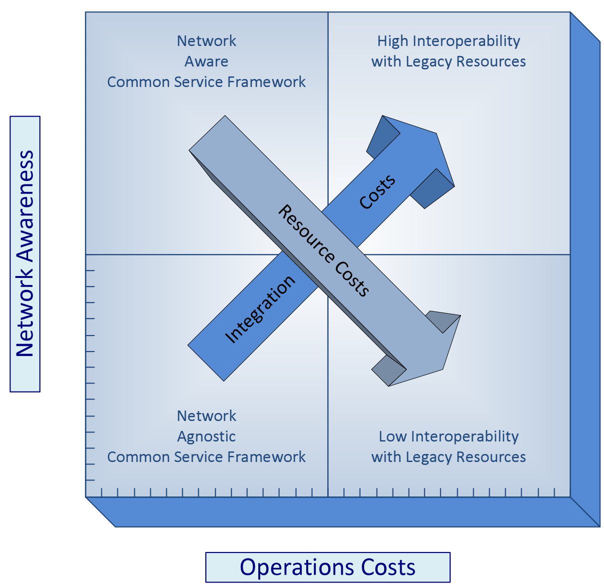

Figure 2: Cost of Network Awareness

Including legacy network awareness into the Television 3.0 Common Service Framework should be built-in from the start as it will become ever more difficult to retrofit fielded applications. Legacy network awareness need not be deployed into the field on day one. These features may exist in the

Common Service Framework, but may be turned off in the first deployments, thereby reducing and/or delaying the legacy interoperability costs.

Inevitably a hybrid gateway device will be required to support Television 3.0 service interoperability with the legacy television services. New Internet TV compatible transport standards such as MPEG DASH may never be feasibly deployed on all existing fielded hardware. Until a service provider forces every subscriber to replace their incompatible legacy equipment, a hybrid device will be useful (and cost-effective) in translating between the legacy transport and the new Internet standard transports.

A hybrid gateway device might be installed within the home or outside of it, but regardless of where it is situated it will serve the same function. The hybrid gateway translates between the legacy service and transport controls, for instance by leveraging MPEG2 transport streams (TS), system information (SI) and conditional access (CA) information to acquire and terminate the legacy broadcast service, and then transcoding, transcrypting and translating these services into the Television 3.0 service exposed to the newer multi-screen applications.

> **Note:** A more detailed discussion of media gateway termination technology is available at: [Architecting the Media Gateway for the Cable Home]

Context Aware Services

As mentioned above, an application may require that the Television 3.0 service will distinguish between conflicting user priorities based on context.

As an example, when viewing television on a mobile phone, the user may choose to pause their viewing in order to answer the phone call. In a pure streaming model, the application would determine availability of pausing by validating whether a “catchup” version of the program is available for bookmarking and later streaming from the paused location.

But if the same mobile application happens to be situated within the user’s home network then a DVR capable application might already be recording a legacy broadcast version of the program to a local storage device, altogether eliminating the need for additional network “catchup” streaming resources. Further, if the Television 3.0 Common Service Framework were capable (e.g. DRM content controls permitted), then the application might be able to stage the content for offline as well as online viewing, enabling the DVR recorded content to be viewed in a park or on a plane.

If the subscriber exposes local storage in the home for the purposes of viewing television content, then the same Common Service Framework (in a managed network) might use Progressive Download techniques (PDL) to persist as many formats of the content in the home as are required by that home’s devices, avoiding the future need for in-home devices to go back to the network for viewing. The same PDL technique may be used to pre-position personalized advertising content.

Another example of a context-aware service framework is the ability to limit or expose service provider resources based on an application’s user privileges. For instance, limiting the quality of the television content might be dependent on a user’s data quota. The service provider might for instance permit the user to limit their household’s access to HD quality content when a certain threshold of usage is met every month, or for a specific household device or user.

Such authorizations may in fact be federated in the Television 3.0 model. For instance, additional personalized metadata about a specific television event might be available to a subscriber only if they subscribe to multiple service providers. As an example, the Common Service Framework may permit the CDN to be managed by one service provider, but the content discovery may be managed by different ones, only sharing a common content identifier. For instance, subscribers to Common Sense media or Rotten Tomatoes might have additional descriptive information about the current movie that the user is accessing from their cable subscription.

To deploy a Television 3.0 Common Service Framework capable of unlimited TV viewing anywhere, as has been described in this paper, the Common Service Framework must be capable of implementing a very robust contextual control interface over the content as well as the content delivery network, whether connected to a network or offline. The context control interface logic itself must be cacheable along with any associated context control metadata including related content and service information required for discovery, protection and transport of the content, so as to permit offline as well as online viewing. The ABR manifest in DASH may be used to implement such a context control interface. The XML manifest file standardized by DASH may be accessed from a stream or from a cache. The DASH manifest file may be adapted and/or extended on the fly by any intermediate control point or it may be kept in its original pristine form, untouched as any associated content is transported through the content distribution network.

For instance a service provider might choose to regionalize, localize and/or personalize the original broadcast manifest file as well as any associated sidecar files. Sidecar files may be used to extend the original manifest or index file, for instance by describing network specific abstractions (and might be required in cases where the original manifest file is write-protected). Examples of content personalization include frame accurate insertion of an overlay graphic, an alternate video replacement, or any other form of advanced advertisement, as well as the inclusion of a user or a group's specific bookmarks.

At each point that the manifest file is transferred from one sub network to another over the entire content distribution network, the manifest or associated sidecar files will be subject to controlled manipulation as required by the context specific needs of the service provider or the end user.

As an example, in the case that a service provider is leveraging a local storage device to permit offline and online viewing of their controlled content, the service provider might pre-stage personalized content or metadata to be used in place of the broadcast content or at other pre-defined interstitial points. Such a Television 3.0 application would not only allow for the display of personalized advertisements, but would indeed allow for any type of personalized content – the same movie might be available in a specific subset of the five parental advisory formats for each user in the household.

Through these ABR synchronization techniques, the Television 3.0 application may be able to access new types of contextual metadata, for instance enabling users to skip through episodes in a series – or articles in a video news journal. The user might even personalize their application to prioritize content based on their location (for instance emphasizing movies shot in Paris over British comedies and then vice-versa as location changes).

Additional Television 3.0 contextual services will be made possible by the development of a robust context control interface. For instance, content discovery and recommendation searches may be persisted and prioritized by users to enable automatic organization of future programs (or versions of programs) in a much more personalized fashion.

The Television 3.0 content will adapt to the use of contextual metadata. Television series will include metadata to allow viewers to automatically catch-up to favorite plots. Movie directors will include metadata to allow viewers to experience their stories differently, depending on personal desires (e.g. family-friendly fare, racy endings, and mood-sensitive plot lines).

Consumers most likely will agree to reductions in privacy in return for more personalized Television 3.0 content – that will be delivered along with more personalized advertisements (which may actually be of interest to consumers). Contextual control of playback will assure advertisers that the consumers have actually viewed their information, and will enable consumers more instant gratification (e.g. immediate purchase of the actress's dress).

The concept of DVR scheduled and recorded content will evaporate over time as all content will be available at any time and in any place. Instead of recorded content, users instead will refer to "My Content Library" in order to distinguish between personally interesting content and everything else. DVR schedulers will evolve into personal recommendation and content discovery tools. All content the user ever viewed will be available to them, but only the recent content most likely to be viewed next will be displayed within personal recommendations list.

Network Aware Services

A few years from now, the content delivery networks of today will be considered as outdated as the Web 1.0 applications of a decade ago.

A service aware contextual gateway application might be deployed at every sub-network interface point on the content delivery network. The content delivery network control application itself might be virtualized and contextualized in the same fashion as the television applications described above. Today network switches and routers are fundamentally constrained by the Open Systems Interconnection (OSI) model to a very limited visibility of application needs within a specific OSI level.

Using the same kind of abstraction model described above for television content distribution, IP packet distribution can be equally freed from the constraints of the existing network control models.

New forms of "network aware services" will be enabled to adapt more easily to the physical constraints of the underlying sub-network. An example of this is a residential gateway that dynamically routes consecutive video packets across both a home wireless and home wire-line network (e.g. MoCA, 80211.AA) depending on temporal noise characteristics and error correction on each physical transport medium.

The poster art representing a movie to be displayed in a television application might be dynamically adjusted not only by the device screen size, but by the capacity of the underlying IP network on which it was transferred to the device.

As each network gateway application is empowered to make service aware optimization decisions, the network controller will coordinate and mediate conflicting needs of applications, service providers and access network operators.

CONCLUSION

In a world where every consumer desires to have their subscription services on any device at any time, service providers must learn to live with an endless variety of devices and an infinite number of services, delivered over any type of network, both offline and online. In the short-term, every service provider must be able to deliver their services in managed, unmanaged and hybrid environments equally well.

The rapid ascent of audio-video services delivered with ABR technologies and the rapid adoption of the MPEG DASH standards exemplify this trend. ABR manifest and associated content and metadata files may be cached and manipulated at every point in the content delivery path to assure consumers access to television services whether in the home or on the go.

Use of similar next generation caching techniques will be extended throughout service delivery platforms to assure that every operator service benefits from similar scalability paradigms, including user interfaces and collaborative communication features.

Just as Television 2.0 took advantage of techniques developed for the social web to optimize delivery of television over unmanaged networks, Television 3.0 applications will adapt those techniques to the needs of network operators who require consistent managed and branded television service to be delivered to any subscriber at any time and any place.