

Watching Video over the Web

Part 2: Applications, Standardization, and Open Issues



Ali C. Begen, Tankut Akgul, and Mark Baugher • Cisco

In this second part of a two-part article, the authors look into applications for streaming, including end-to-end mobile and in-home streaming. They contrast adaptive approaches to other video-delivery paradigms, discuss current standardization efforts, and highlight areas that require further research and investigation.

Many video streaming applications and services exist on the Internet, and we covered several of them in part 1 of this article.¹ The most important applications arguably fall into the mobile and home entertainment categories, which span the range of video presentation formats for portable (PD), standard (SD), and high-definition (HD) video, and these applications include both walled-garden and Internet applications. They also operate on diverse networks, and use both managed and unmanaged end-to-end network services. Here, we identify mobile and home-streaming use cases, and consider the suitability of pull-based versus push-based adaptive streaming for each use case.

Mobile Streaming

People stream movies, YouTube videos, and other media to their mobile devices while on the road, in hotels, on campuses, and otherwise in motion. PD encoding is a video format for mobile devices with lower resolution and bitrate encoding than SD and HD encoding, which are common in homes. PD video delivery scales down to the lower speeds and higher loss found on most mobile networks. Variations in load on network cells, changes in radio-network conditions, and fluctuations in available throughput force video applications to

adapt to various networks – including General Packet Radio Service (GPRS), Enhanced Data Rates for GSM Evolution (EDGE), the Universal Mobile Telecommunications System (UMTS), and Code Division Multiple Access (CDMA) 1x. Thus, it's no surprise that mobile handset vendors and operators have been at the forefront in defining and developing adaptive streaming standards.

When the mobile connection is to a walled-garden server, the network provider can manage network video delivery to a service-compliant, walled-garden client. This is one use case. If the end-to-end connection is to a server on the Internet, the network provider can't provide quality-of-service (QoS) management, and the service isn't managed. This is another use case. A third use case combines mobile and home streaming.

Walled Garden

In the case of a walled garden, the end-to-end path is contained in the mobile service provider's network. The network operator can specialize service for video delivery. The provider can choose to use push-based adaptive streaming services from specialized streamers to get optimal use of the network resources and improved video quality over a range of network conditions. Alternatively, the provider can choose to

reuse its data network services, such as HTTP servers and caches, and employ pull-based adaptive streaming methods such as Apple's HTTP Live Streaming.

Mobile Internet

With mobile Internet, the service provider can't manage the service quality on an end-to-end basis when the video connection traverses the public Internet. In this case, pull-based adaptive streaming has advantages over push-based methods, as we describe in the next section.

Hybrid Mobile–Home Use Case

Here, the mobile service provider can't guarantee a service when the first network hops or the last network hops are on the client's home network. Femtocell networks and home-network services to mobile clients are examples. When an unmanaged home network is a hop on the end-to-end path, and particularly when it's a wireless hop, pull-based adaptive streaming is more desirable than push-based methods.

Home-Network Streaming

Home-network use cases are crucial to video streaming because consumers use video so widely at home. Although in many residences the home network consists of a single Ethernet cable connected to a broadband modem and the home PC, more complex home networks are becoming common in many countries. Increasingly, self-installed home LANs are capable of transporting video.

802.11 (Wi-Fi) LANs are widespread in many regions. One well-known video-delivery problem in Wi-Fi networks is the divergence of network speed and reliability due to network topology, load, and other variables. Metal objects, common home appliances, and other sources of interference mean that a television in one room often can't

sustain the same quality of presentation as a device in another. The reality of diverse network capacity among receivers coupled with the 802.11 medium access control's design also cause a problem for multicast streaming over wireless LANs (WLANs).²

Video-on-demand, not broadcast TV, is likely to be the dominant video application on home networks. Unfortunately, wireless and other types of common home LANs are vulnerable to interference and overload. Although migrating Wi-Fi products to 5 GHz from 2.4 GHz reduces some interference problems, even 802.11n is vulnerable to interference that results in widely differing loss and throughput characteristics on the WLAN.³ Another home LAN technology vulnerable to interference is powerline communications (PLC). In fact, with the exceptions of well-laid Ethernet over Cat5 or cable (as defined by the Multimedia over Coax Alliance [MoCA]), it's hard to imagine a home-networking technology that can offer uniform, high-speed services throughout the home.

Adaptive streaming is crucial for video delivery on home networks, as they are mostly unmanaged and exhibit great differences in network speeds and feeds available to a diverse set of devices. It's unreasonable to expect that a typical home user will manage and select particular resolutions of a movie title based on the quality of network service that a particular device can obtain from a particular spot on the home network.

Although managed video might never be common for in-home streaming, it will continue to provide a superior, trouble-free, and gold-level video service to the home that unmanaged services can only try to replicate. In the use cases that follow, we consider both managed and unmanaged methods and compare

push-based streaming to its pull-based counterpart.

Streaming to a Home Client from a Home Server

This use case is arguably the most interesting but the least common. For years, hobbyists and digital media enthusiasts have connected their TVs to their home networks to play movies from a PC, gaming console, or another server, but this isn't yet a mainstream practice. In the US, a Wi-Fi network can automatically connect an off-the-shelf home network attached server to a Sony PlayStation and a host of other devices implementing Digital Living Network Alliance (DLNA) specifications (www.dlna.org). However, nontechnical consumers don't widely use this capability. Setting up a home media server requires sufficient knowledge about how to load it with commercial movies, which are usually encrypted, match formats with players, manage Digital Rights Management (DRM) when it's exposed or not circumvented, and perform other tasks that people outside the computer industry don't routinely do. Nonetheless, home servers can be quite useful to those who want to play digital movies without experiencing the bottlenecks common to the Internet or ISP networks. Home users should be able to play a file even when their broadband network connection is down, and play their movies on all the TVs or other displays in the house. iTunes has achieved in-home music sharing; Apple and other vendors are currently developing video sharing among home devices as well. On the other hand, standard solutions such as DLNA can open the market to greater product diversity and are likely important to future in-home video use cases. Streaming to a home client from a home server is a use case that's arguably more important to home video's future than it is today.

Whether using proprietary or standard players, home wireless networks are limited by the radio spectrum's available bandwidth and various types of noise and interference; wireless networks, in general, are arguably the weakest link in the video transmission chain. The quality of a video presentation on a wireless home network is subject to all the problems of an unmanaged network plus signal loss, interference, and contention with other traffic on the WLAN.⁴

When home-network owners use push-based adaptive streaming, they must install, configure, and manage a specialized push-based streaming server on the network that isn't managed by a separate business such as an ISP. Moreover, the most popular protocol for video on home networks is HTTP followed by the Real-Time Transport Protocol (RTP) and Real-Time Streaming Protocol (RTSP). HTTP is used on the Web and in home-networking protocols such as DLNA. Many new services, such as UltraViolet, will also use it. Numerous *turnkey* HTTP video applications and commodity servers are available in the market. Thus, pull-based streaming over HTTP has advantages over push-based streaming over RTP.

Streaming to a Home Client from an Internet Server

This use case includes streaming webcasts, YouTube, Hulu, iTunes, and other Internet video content to a home PC or digital TV. As explained previously, a potentially troublesome home-network hop, such as a home WLAN, will often be part of the end-to-end video path whenever the network supports more than a single device connected to a modem. Managed video within the home is unlikely to become common because home networks are usually unmanaged. The public Internet, moreover, supports only unmanaged video delivery. Thus, pull-based adaptive

streaming is the only feasible approach for this use case.

Streaming to a Home Client from a Managed Server

Most digital television services today, including video-on-demand, originate from video service providers. A provider's video service might be a digital TV, IPTV, or HTTP service – in all cases, the provider can optimize its service for managing video over its network to a home device. Managed video generally achieves a level of quality that unmanaged video only attempts to meet – often without success. Naturally, a managed server can use a push-based adaptive streaming method.

HTTP progressive download or adaptive streaming can, of course, run over a managed video network and even use a managed server. Thus, this use case includes both push- and pull-based streaming.

Streaming to a Home Client via Peer-to-Peer Delivery

In certain areas of the world, particularly Asia and some parts of Europe, a popular streaming video method is to use peer-to-peer (P2P) networks. Research has shown that P2P transport could reduce the load on the source servers and provide better scalability for large streaming populations. However, with advances in server and caching infrastructure designs, large streaming capacities are now easily achievable without needing P2P transport. On the business side, uncontrolled and unmanaged P2P systems often fail to provide fair revenue sharing between content providers and ISPs.

Although industry-wide efforts are attempting to address specific issues related to network and provider-friendly P2P transport, the business models that will be profitable for both the content owners/providers and ISPs aren't yet in place. Thus, in this article, we won't go into further detail; interested readers can refer to the December 2007 issue of the *IEEE Journal on Selected Areas in Communications* for detailed coverage of various topics in P2P streaming.

Use Case Benefits

Of the six mobile and home-network streaming use cases, most have unmanaged networks along their end-to-end paths. As Table 1 shows, two cases can benefit from push-based adaptive streaming using specialized servers and clients – that is, the (walled-garden) managed server use cases in which the end-to-end path is over a provider-managed network. Vendors and providers are optimizing managed networks for HD video delivery; managed network devices comply with the technical protection measures major video producers require. Additionally, video service providers have strategic business relationships that give them an early spot in the windowed movie release schedules as well as special pay-per-view video events. This gives service providers a role in premium content delivery of major titles, particularly expensive movies. Managed video is the standard that Internet video technology tries to meet.

How close can unmanaged video come to managed-video quality? And how well can such services

Table 1. End-to-end paths and preferred streaming methods.

	Internet server	Managed server	Home server
Mobile client	Pull-based streaming	Push- and pull-based streaming	Pull-based streaming
Home client	Pull-based streaming	Push- and pull-based streaming	Pull-based streaming

scale? These questions are crucial to the future of Internet video services that use pull-based adaptive streaming. Arguably, the Internet server path shown in Table 1 offers users the greatest access to the most content – including broadcasters' sites, iTunes, Netflix, Hulu, and even social networking sites. If pull-based adaptive streaming can match managed video's quality well enough, then this use case can become the de facto delivery method for successful video streaming.

Of the remaining use cases, streaming from a home server to a home client is much less popular compared to Netflix, Hulu, iTunes, and managed video from service providers. This use case becomes more important, however, as digital files replace digital disks in the home. When a broadband connection becomes unavailable, as broadband connections sometimes do, consumers should at least have access to their titles within the bounds of a functioning home network. Streaming to a home client is also interesting because this use case operates on the home network, and the video-delivery performance characteristics of home networks such as 802.11 are important to the Internet and managed server cases as well. We thus conclude that pull-based adaptive streaming methods have considerable advantages over push-based adaptive methods for most Internet video applications.

Standardization Efforts

Microsoft's Protected Interoperable File Format (PIFF)⁵ and Smooth Streaming Transport (SST)⁶ are complete specifications for an adaptive streaming container and streaming protocol. PIFF and SST use the standard MP4 fragmented file structure. Microsoft has also opened up these specifications for widespread use by offering the Microsoft Community Promise, which provides a free license for its adaptive streaming

protocols. Apple has similarly offered its HTTP Live Streaming, both on the Web and in an IETF draft.⁷ Both the Microsoft and Apple specifications reflect the current adaptive streaming products that the two companies offer.

These and other specification efforts are appearing in various venues. Other efforts include the 3rd Generation Partnership Project (3GPP), MPEG, and industry consortia such as the Digital Entertainment Content Ecosystem (DECE), which announced a new service called UltraViolet in July 2010 (www.uvuu.com). DECE has announced that its UltraViolet service will support a common container format, movie downloads, and streaming with support for multiple DRMs. It seems likely that the final specifications will support at least some state-of-the-art media distribution technologies we've considered in this article.

Although DECE's membership includes dozens of companies that are involved in media distribution, it is a private effort by a consortium. Adaptive streaming technologies have made their way into public standards development organizations including 3GPP and MPEG.

3GPP was the first organization to release a specification related to adaptive streaming over HTTP. It published Release 9⁸ in March 2010, and the SA4 Working Group is currently bug-fixing this specification, taking into account experience from its initial implementations. In addition, 3GPP is preparing an extended version for Release 10, scheduled for publication later in 2011. This release will include several clarifications, offer improvements, and add new features. On the MPEG side, there were 15 submissions to the call for proposals published in April 2010. These submissions were evaluated in July 2010, and, subsequently, a first working draft on Dynamic Adaptive Streaming over HTTP (DASH) was

approved within MPEG. This working draft adopted 3GPP's Release 9 as a baseline specification and started running several evaluation experiments. The DASH Ad Hoc Group has been working on standardizing the manifest file and delivery formats using ISO Base Media File Format and MPEG2 Transport Streams. The DASH Ad Hoc Group has also been coordinating closely with the 3GPP SA4 Working Group to better align their respective specifications. MPEG expects to finalize its specification in July 2011.

At the time of writing, many questions about adaptive streaming had yet to be adequately answered. Although some researchers are conducting early investigations,⁹ most of the questions still remain open, mainly due to the lack of field data required to conduct a rigid analysis. To date, both TCP and HTTP have been studied in great detail for use with conventional applications. Researchers have made proposals, supported by simulations and experiments, for both protocols that would make them a better fit for streaming-like applications. However, making generalizations and drawing conclusions for large-scale streaming deployments based on narrowly scoped studies isn't trivial and could be easily misleading.

Historically, TCP was designed and optimized for FTP-like applications. We must examine the implications of its congestion-control algorithm and currently adopted best practices – such as Nagle's algorithm, delayed acknowledgments, and slow-start restart (see RFC 5681; <http://tools.ietf.org/html/rfc5681>) – in the context of adaptive streaming. Researchers have also carefully studied the effects of running multiple TCP connections simultaneously to overcome the head-of-line blocking problem or using another transport

protocol that inherently provides better support for multiple streams, such as Stream Control Transmission Protocol (SCTP). Most importantly, these studies must represent large-scale deployments with many origin servers, caches, and clients streaming a variety of content over diversely characterized network paths.

A niche area is to develop instrumentation tools to assess adaptive media transport's effectiveness and performance. These tools must be user-friendly and able to provide adequate information for diagnostics and fault isolation. Such information becomes quite handy for service providers in fixing problems and making necessary provisions and enhancements to their networks.

An interesting future research direction is to examine how network elements can help providers achieve better performance – that is, support more clients or deliver a better and more consistent streaming quality. In other words, rather than considering the underlying network as a black box or just a bunch of pipes that carry the bits, having the end points (for example, the origin servers and clients) and middle boxes (such as cache servers) talk to the network, exchange information, and act accordingly could provide many benefits. For example, the network itself would be the first to know about a congestion or failure, helping servers and clients adapt faster and more accurately.

Last but not least, an important open issue is whether the network capacity in a multi-access network can sufficiently support many clients concurrently. In such a scenario, each client competes with others for more bandwidth and is likely to experience a continual bitrate fluctuation, which adversely impacts the quality of experience. Prepositioning the

content or using multicast (especially for live streaming) could prove useful.

In a future article, we hope to provide results from our ongoing studies investigating these open issues. □

Acknowledgments

We thank our colleagues at Cisco, 3GPP, and MPEG for their reviews and feedback for this article.

References


1. A.C. Begen, T. Akgul, and M. Baugher, "Watching Video over the Web, Part 1: Streaming Protocols," *IEEE Internet Computing*, vol. 15, no. 2, 2011, pp. 54–63.
2. D. Dujovne and T. Turletti, "Multicast in 802.11 WLANs: An Experimental Study," *Proc. 9th ACM Int'l Symp. Modeling Analysis and Simulation of Wireless and Mobile Systems (MSWiM 06)*, ACM Press, 2006, pp. 130–138.
3. R. Gummadi et al., "Understanding and Mitigating the Impact of RF Interference on 802.11 Networks," *Proc. SIGCOMM*, ACM Press, 2007.
4. N. Cranley and M. Davis, "Performance Evaluation of Video Streaming with Background Traffic over IEEE 802.11 WLAN Networks," *Proc. 1st ACM Workshop on Wireless Multimedia Networking and Performance Modeling (WMuNeP 05)*, ACM Press, 2005, pp. 131–139.
5. J.A. Bocharov et al., *Protected Interoperable File Format*, Microsoft specification, Mar. 2010; <http://go.microsoft.com/?linkid=9682897>.
6. *Smooth Streaming Transport Protocol*, Microsoft specification, Sept. 2009; <http://go.microsoft.com/?linkid=9682896>.
7. R. Pantos, ed., "HTTP Live Streaming," IETF Internet draft, work in progress, Nov. 2010.
8. *Transparent End-to-End Packet-Switched Streaming Service (PSS): Protocols and Codecs*, 3GPP TS 26.234, Dec. 2010; <http://ftp.3gpp.org/specs/html-info/26234.htm>.
9. S. Akhshabi, A.C. Begen, and C. Dovrolis, "An Experimental Evaluation of Rate-Adaptation Algorithms in Adaptive

Streaming over HTTP," *Proc. ACM Multimedia Systems Conf.*, ACM Press, 2011.

Ali C. Begen is with the Video and Content Platforms Research and Advanced Development Group at Cisco. His interests include networked entertainment, Internet multimedia, transport protocols, and content distribution. Begen has a PhD in electrical and computer engineering from Georgia Tech. He is a member of IEEE and the ACM. Contact him at abegen@cisco.com.

Tankut Akgul is in the Service Provider Video Technology Group at Cisco. His research interests are embedded software design, video compression, and multimedia streaming. Akgul has a PhD in electrical and computer engineering from Georgia Tech. Contact him at akgult@cisco.com.

Mark Baugher is in the Research and Advanced Development Group at Cisco. He has coauthored several widely used international standards in the IETF and ISMA, and is currently cochairing the Internet Gateway Device Working Committee in the UPnP Forum. Baugher has an MA in computer science from the University of Texas at Austin. Contact him at mbaugher@cisco.com.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.

