



Enterprise Streaming: Different Challenges From Internet Streaming

John Apostolopoulos, Mitchell Trott, Ton Kalker, Wai-Tian Tan
Mobile and Media Systems Laboratory
HP Laboratories Palo Alto
HPL-2005-98
May 18, 2005*

video streaming,
enterprise
streaming,
streaming media
content delivery
networks

Media streaming over the best-effort public Internet has been a focus of research for over a decade. Enterprise or corporate streaming is another area of media streaming that is practically very important and has a different set of challenges and feasible solutions. For example, the quality and reliability requirements for enterprise streaming are much stricter than for typical Internet streaming. Furthermore, in typical enterprise streaming scenarios, a single entity has control over most elements of the system, including the endpoints and the infrastructure. This entity has the powerful ability to monitor, adapt, and deploy new infrastructure as necessary. The goal of this paper is to describe enterprise streaming and identify the basic differences between typical Internet streaming and enterprise streaming, and how these differences alter the challenges that must be overcome for enterprise streaming to be successful. Specifically, we examine enterprise streaming media content delivery network design and operation, video conferencing, peer-to-peer networking (P2P), voice over IP (VoIP), and briefly touch upon wireless and security issues.

* Internal Accession Date Only

To be published in IEEE International Conference on Multimedia and Expo, 6-8 July 2005, Amsterdam, The Netherlands

Approved for External Publication

© Copyright 2005 IEEE

ENTERPRISE STREAMING: DIFFERENT CHALLENGES FROM INTERNET STREAMING

John Apostolopoulos, Mitchell Trott, Ton Kalker, and Wai-Tian Tan

Streaming Media Systems Group, Hewlett-Packard Labs, Palo Alto, CA, USA

ABSTRACT

Media streaming over the best-effort public Internet has been a focus of research for over a decade. Enterprise or corporate streaming is another area of media streaming that is practically very important and has a different set of challenges and feasible solutions. For example, the quality and reliability requirements for enterprise streaming are much stricter than for typical Internet streaming. Furthermore, in typical enterprise streaming scenarios, a single entity has control over most elements of the system, including the endpoints and the infrastructure. This entity has the powerful ability to monitor, adapt, and deploy new infrastructure as necessary. The goal of this paper is to describe enterprise streaming and identify the basic differences between typical Internet streaming and enterprise streaming, and how these differences alter the challenges that must be overcome for enterprise streaming to be successful. Specifically, we examine enterprise streaming media content delivery network design and operation, video conferencing, peer-to-peer networking (P2P), voice over IP (VoIP), and briefly touch upon wireless and security issues.

1. INTRODUCTION

Media streaming over the Internet is both an important practical application and a challenging research area [1]. Media streaming within a corporate infrastructure, referred to as enterprise streaming [2], is also an important practical application, but has a different set of challenges and tradeoffs.

The delivery of video and audio for corporate applications dates back to the video conferencing systems of the 1980s. In the 90s there was increased excitement about delivering audio, video and other forms of rich media over enterprise IP networks to employees' desktops. However, the costs of building the necessary network infrastructure and the content creation, delivery, and consumption chain initially limited the deployment and use of enterprise streaming. Since then, technology improvements and cost reductions have accelerated enterprise streaming. For example, improved video and audio coders, such as H.264/MPEG-4 AVC, have significantly reduced the storage and transmission cost of video. Maturity of signaling protocols such as Session Initiation Protocol (SIP) have enabled Voice over IP (VoIP) and other networked applications. The continuous

decline in cost and improvement in performance of wire-line and wireless networks, personal and server computers, cameras, displays and content creation tools, have made enterprise streaming both technically and financially feasible.

The goal of this paper is to describe enterprise streaming and identify the basic differences between typical Internet streaming and enterprise streaming and how these differences alter the challenges that must be overcome for enterprise streaming to be successful. This paper continues by briefly identifying some of the similarities and differences between Internet and enterprise streaming. The network infrastructure for streaming is then discussed, as well as enterprise streaming media content delivery network (SM-CDN) design and operation. This is followed by examining the important topics of video conferencing, peer-to-peer networking, voice over IP, and briefly touching upon wireless and security issues.

2. INTERNET & ENTERPRISE STREAMING

The goals of enterprise and Internet streaming are similar yet there are many factors that cause streaming systems to have sometimes drastically different designs in the enterprise than in the Internet. Broadly speaking, Internet streaming involves adapting streaming for a given infrastructure with no control, whereas enterprise streaming allows co-adaptation of streaming and infrastructure. We identify in more detail some of the key attributes that underlie design differences between typical Internet and enterprise streaming systems in Table 1.

First and foremost is the level of control in the design of the streaming system. For the Internet, media content is typically created by one party, hosted by a second party, and then delivered by a third party before finally consumed by a fourth party. Even more parties may be involved: Internet-wide routing is achieved by collaboration of many autonomous systems (AS) that are often owned by competing entities. The existence of multiple parties with different interests limits the ability to realize end-to-end improvements by any single party. In contrast, enterprise streaming typically involves a single entity that controls everything. Coordinated upgrades to effect end-to-end streaming improvement is therefore possible and can often be mandated over an entire enterprise. In enterprise streaming, the content largely stays within a single AS, and there is ability to monitor, con-

Internet Streaming	Enterprise Streaming
Limited control	Control of end-points & infrastructure
Limited financial resources	Flexible financial allocations
Reliability important, but depends on cost	Reliability is critical
Quality is important, but depends on cost	Quality is very important
Limited network resources, must be shared	Overprovisioned
Firewalls	Generally within firewall
Focus on individual	Focus on the “public good” for all users
Supporting fast mobility is difficult	Facilitates mobility
High diversity of content	Limited diversity of content

Table 1. Key differences in Internet and enterprise streaming.

trol, and change the infrastructure to improve streaming.

Detailed knowledge about network conditions available in enterprise streaming can be used, for example, to achieve lower latency and to reduce bandwidth. If a link fails or becomes congested the network monitoring system can react by rerouting over a clean link. In the Internet, protection against congestion-induced packet loss requires some form of redundancy to “average over” random packet erasures. Such coding typically increases latency and consumes extra bandwidth. The higher the required reliability the greater the bandwidth or latency increase. Typically, for Internet streaming data link impairments must be handled end-to-end while for enterprise streaming impairments can also be mitigated through improved signaling, routing, and control.

Additional differences between enterprise and Internet streaming include the following. Generally, there are limited financial resources available for Internet streaming, and reliability and quality are achieved in a best-effort basis given a limited budget. In contrast, since enterprise streaming is used for business purposes, requirements on reliability and quality are inelastic, and financial resources are allocated as necessary. Internet streaming is often afflicted by limited network resources which must be shared; on the other hand, enterprise networks are often overprovisioned. While Internet streaming often requires traversing firewalls, enterprise streaming does not require firewall traversal since the stream stays inside the enterprise. A great degree of mobility (ranging from less than a second to imperceptible) can be supported in enterprise settings, while it is difficult to support in Internet settings. Finally, an enterprise typically offers a limited diversity of content, as opposed to an Internet setting where a high diversity of content is supported, and their typical workloads can also differ significantly [3].

Enterprise streaming is typically designed and operated to achieve a business aim for the good of the enterprise as a whole; there is no need to finely distinguish which user pays for which equipment. In the Internet, however, equipment and services that generate “public good” are not generally offered unless their cost can be recovered, e.g., by charging individual users. This relates to the notion of perceived fair-

ness. In the Internet, perceived fairness is important, and an important issue in P2P systems such as BitTorrent [4] is enforcing fairness constraints such as equality of upload and download data volume. P2P systems that don’t have fairness constraints generally decline, e.g., Kazaa [5]. For the community of users in the same enterprise, the insistence upon such a notion of fairness is generally weaker, allowing more degrees of freedom in system design.

Some important consequences of the differences mentioned above are as follows. First, enterprise streaming systems generally have fewer constraints than Internet streaming, but the supported services have more constraints on reliability and quality. Thus, in the Internet, services are commonly downgraded until they match the available network performance, while in the enterprise the infrastructure can be upgraded to meet the requirements of services. Second, the knowledge and control of the network afforded by enterprise streaming allows better and global optimization of infrastructure and end-points, as opposed to Internet streaming’s limited knowledge and end-point only optimization. This leads to a much larger space of possibilities and improved performance for enterprise streaming. In subsequent sections, we study several important aspects of streaming, and discuss how the differences between Internet and enterprise streaming manifest themselves.

3. NETWORK INFRASTRUCTURE AND ENTERPRISE STREAMING MEDIA CDNS

In contrast to an Internet setting, in an enterprise setting a single entity often owns and/or controls the source-end-host or origin-server, first mile to corporate core network, corporate core or backbone, last mile to end hosts, and the end hosts themselves¹. This dramatically effects how the infrastructure can be monitored, adapted, and changed. For example the monitoring and management capability provides the ability to identify what content is viewed, for how long, by which end hosts, how is it traversing the network, is it caus-

¹For enterprise streaming we use “first” and “last” miles to denote the LANs connecting the end-nodes and the enterprise high-speed networks.

ing problems to the conventional data traffic, should rerouting be performed, etc. Since the end hosts are controlled by the enterprise, they can often be upgraded relatively easily, often with minimal end-user involvement (e.g., for software upgrades). In the enterprise setting changing the infrastructure is involved but doable, however this is often not an option in Internet streaming.

One of the simplest ways to enhance the enterprise infrastructure to improve media streaming is through the use of a content delivery overlay network, where nodes are placed on top of (overlaid on) the existing network in order to provide new or improved services while leveraging the existing network infrastructure. A streaming media CDN corresponds to an overlay network designed for streaming media delivery. It is composed of a number of overlay nodes placed at strategic locations in the network, generally close to the end users. Each overlay node consists of one or more overlay servers, and a manager which may be part of the server or may be a separate entity. Some of the key functions provided by the streaming media CDN include [6]: content distribution to the servers and caching at the servers, server selection or re-direction, streaming from the selected server to the end host, application-level multicast, media transcoding to support diverse client capabilities and heterogeneous networks, mid-session handoff between streaming servers or transcoding servers, resource monitoring and management, and other media services.

There are many advantages in having a streaming media CDN. These include improved reliability and scalability to a larger number of users, improved performance in terms of both faster startup time and higher delivered bandwidths and associated media quality. SM-CDN can also reduce demand on the network infrastructure since the streaming occurs over shorter network distances, and through support of application-layer multicast that is suitable for popular live-content such as quarterly reports or CEO announcements to the entire company.

An enterprise streaming media CDN provides significant freedom in placing caching/streaming overlay nodes, as compared to Internet streaming. This is because in the enterprise setting a single entity often controls the entire infrastructure and therefore can place overlay nodes at any desired location. This is not possible in Internet streaming where different parties own different portion of the network. This has an important impact on many key SM-CDN design and operation problems, including: how many servers to use, where to place them (server placement problem), and how to use them (server selection problem). By controlling the infrastructure, the enterprise setting allows a much larger range of possible solutions which can result in significant improvements in resource usage, e.g., increased freedom in server placement may enable fewer servers to support the same number of requesting clients.

Both the overlay server nodes and underlying network nodes (e.g., routers) can monitor what is happening in the

network. Their logs provide valuable insights on the operation at many layers, and they are collected, analyzed, and used to adapt or redesign the infrastructure to improve performance. In addition, real-time monitoring and analysis is used to assess infrastructure operation, identify problems, and dynamically adapt the infrastructure operation (routing, server selection, etc.) to overcome the problems.

The control of both the overlay network and the underlying network provide significant monitoring and control possibilities, which may lead to improved performance, faster convergence (e.g., of routing updates) and prevent oscillations between overlay network action and underlying network re-action (which can hamper performance and produce system instability in certain cases) [7]. This is not possible in Internet streaming when the overlay network and the underlay network(s) are independently controlled by different entities with potentially different policies and goals.

4. VIDEO CONFERENCING

Video conferencing has long been an important application for the enterprise. However, as of today it has achieved limited success. Despite significant time and financial overheads associated with travel, a trip followed by meeting with another party in the same room is often preferred to remote video conferencing. Current Internet conferencing systems typically offer small video picture size with high delay and occasional picture breakup or freezes. Such video quality cannot support a feeling of working together to solve a problem, which is essential to video conferencing. Several aspects of video conferencing need to be improved to support a sense of working together, including higher media quality, and better “in-contact” with the other participants.

For media quality, excellent video quality corresponding to DVD quality and above depending on display size and viewing distance, and excellent audio quality without echo and other highly perceptible artifacts, are necessary. In addition, the end-to-end delay (including delay jitter) for the media streams should be less than roughly 100 msec to achieve interactivity. These requirements mean that the network infrastructure must support several Mb/s of compressed media, with no loss and low delay, and provide reliability corresponding to less than, e.g., one glitch per hour. Clearly such level of QoS is not attainable in the wide-area Internet. For a typical enterprise network, where bandwidth is generally more plentiful, the task of provisioning such QoS alongside other corporate data traffic may still prove impossible. One possible approach for achieving such stringent QoS, as afforded by the higher budget for enterprise networks, is to strategically dedicate leased lines or virtual private networks (VPNs) for video conferencing between different sites at different geographic locations. There are two advantages of this approach. The use of leased lines or VPNs provides superior QoS in terms of bandwidth, delay and loss by shielding the effects of traffic outside the en-

terprise. The use of a dedicated video network eliminates uncontrollable or unpredictable cross traffic such as corporate email inside the enterprise.

Improving a sense of being “in-contact” is another aspect of providing an experience of working together in the same room. There are several examples of such efforts mainly with the enterprise environment as the target. For instance, the use of multiple cameras and multiple large displays at each end enable the distant viewers to be displayed at approximately the same size as the participants in each room. The real-size image greatly facilitates natural interaction. Similarly, research in image-based rendering (IBR) has attempted to solve a variety of important additional problems such as eye-gaze correction [8, 9]. In addition, for video conferencing to replace conventional meetings, individuals at different sites need to collaborate by exchanging, viewing, and editing documents and other artifacts; effectively providing such a collaboration environment is an important area of current research.

In summary, technologies used in video conferencing over the Internet attempts to minimize cost, where users are complacent with the limitations of the existing network and equipment. In contrast, video conferencing in the enterprise is a business need and has to effectively support an experience of working together, or it will not be used. The larger budget for enterprise video conferencing also enables higher quality and more sophisticated video and audio compression, high-bandwidth and low-delay delivery of the streams, removal of audio artifacts and eye-gaze problems, large displays to make the participants feel like everyone is in the same room, and collaboration capabilities on documents and other artifacts.

5. PEER-TO-PEER (P2P) NETWORKING

P2P systems have been used for many years (e.g., Internet routing and Usenet), however Internet P2P systems have recently received considerable attention. P2P systems are architectures and applications where distributed resources (computers and associated storage, computation, and bandwidth) are used to perform a task in a decentralized manner. Popular Internet P2P systems include the file sharing systems Napster, Gnutella, and Kazaa (which primarily differ in how they locate which peer has the desired file), as well as Voice over IP P2P system Skype (discussed in Sec. 6) and GRID computing networks such as SETI [10].

Internet P2P systems typically have the following attributes: (1) peers are self-organizing (they discover one another), (2) peers are generally autonomous and the system operates in a decentralized manner, (3) peers may come or go (churn) with no guarantees on availability or services provided, and (4) peers generally have equal standing (peer equality), though some P2P systems attempt to provide improved performance by classifying some nodes as supernodes based on attributes such as reliability and network band-

width. These attributes provide a number of benefits that are important in the context of the Internet, including improved scalability and reliability (avoiding dependence on centralized resources since peers directly interact with each other), and of course reduced cost. However, many of these attributes are not necessary and are even undesirable in an enterprise setting. For example, in an enterprise setting it is desirable to have centralized authentication and access control, centralized databases that provide increased reliability and timely updates, etc. In addition, it is unlikely that enterprise peers will come and go, and it may be possible to specify the availability and services provided by each node (although their available resources may vary unpredictably over time). Also, peer equality is of lesser importance, and the server/client model can simplify various applications. Furthermore, many of the issues that motivate P2P systems in the Internet, such as limited budget, lack of mutual trust between nodes, high churn rate, and only limited importance of the service, are almost opposite in typical enterprise settings where the required reliability and importance of the provided service(s) often outweighs the cost, nodes are generally trustworthy, and churn can be limited. An excellent discussion on when a P2P system is appropriate is given in [11].

P2P systems for media delivery generally have two steps: (1) discovery of the location(s) of the requested content and (2) delivery or processing of the content. While the discovery step may be centralized (e.g., centralized directory such as used by Napster) the delivery or processing may be decentralized. For example, computational resources may be aggregated across distributed computers to solve large computational problem, e.g., graphics rendering for Hollywood films, or distributed storage across multiple nodes to provide increased reliability for important content or to increase storage capacity. These decentralized processing steps can be controlled in an enterprise setting to achieve a desired performance objective, however it is difficult to achieve similar objectives in an Internet setting. Finally, as mentioned in Sec. 2 perceived fairness is very important for Internet P2P systems, however it is not really a concern for enterprise P2P systems which generally are designed to benefit all.

6. VOICE OVER IP (VOIP)

Over the last few years there have been major economic motivations to move to VoIP for speech communications. By bypassing the established telecommunication infrastructure, VoIP provides the consumer with dramatically reduced costs for international and domestic calling. In the enterprise VoIP reduces costs, allows the convergence of the data and phone networks to a single data network, and makes it practical to unify voice, instant messaging, and email communication. Both Internet and enterprise VoIP systems use gateways to connect VoIP end hosts to conventional circuit-

switched telephones.

The required reliability for Internet VoIP and enterprise VoIP are very different. In exchange for lower cost the home consumer may accept sporadic quality degradations, the need to retry a connection, or even rare service outages. Enterprise users (and their customers and business partners) have no such patience. The overarching need for high availability drives architectural choices in the enterprise that are not available on the Internet. For example, within an enterprise routers can be configured to prioritize voice traffic ahead of data traffic, and to partition voice and data onto separate virtual LANs (VLANs).

Enterprise users demand not only more quality but more features than typically required for Internet VoIP, including for example sophisticated call transfer options, full interoperability with legacy analog phone systems, and fine-grained routing control to direct a call to the gateway that minimizes long distance calling charges. Many of these challenges are addressed at the signaling layer rather than the data link layer.

Firewalls and network address translators (NATs) are a problem for Internet VoIP that does not typically limit enterprise VoIP. Overcoming NATs and firewalls is critical for VoIP over the Internet, where different hosts are connected in different ways. For example, the popular Skype VoIP P2P application allows users to work behind NATs and firewalls [12, 13] by first probing the network, then, if necessary, employing an additional Skype peer as a relay. In contrast, within the enterprise, hosts can usually communicate directly. When direct connection is not possible, usually only a small number of NATs or firewalls—all owned by the enterprise—need to be reconfigured for VoIP.

Security and privacy are a significant concern for both enterprise and Internet VoIP. The data path and signaling traffic can both be protected using encryption but it can be difficult to eliminate all vulnerabilities in the signaling path, particularly when interoperating with legacy equipment. A basic security measure is physical access to the LAN. Paradoxically, however, the enterprise can be more vulnerable than the Internet: a typical enterprise has unprotected ethernet ports in every conference room, while the LAN of a home Internet user is exposed only to houseguests. (A poorly configured wireless home router reverses this situation.) Both enterprise and Internet VoIP are vulnerable to viruses, worms, and spyware, though to the authors' knowledge no such VoIP-specific attacks have yet appeared. The network management tools in an enterprise make it much easier to detect misbehaving machines though not necessarily to physically locate them.

7. WIRELESS

The wireless landscape is rapidly improving for both Internet and enterprise streaming. Wireless Internet streaming is most commonly provided over either a cellular network

or over a wireless LAN based on IEEE 802.11. Enterprise wireless streaming is done almost exclusively using 802.11. Wireless can be configured as a “last mile” access solution or, alternatively, the entire network can be wireless, as for example in emerging mesh solutions.

The IP data services offered by current cellular networks have very high and variable latency, making them marginally suitable for conversational communication. As a result, most IP streaming targeted at cellular is one-way and has a large playout buffer. Next generation networks, based for example on proprietary formats, on 802.16 (WiMax), or on UMTS-HSDPA, have much better latency characteristics.

Until quite recently the deployment of 802.11 wireless was hindered by limited bandwidth, random installation of multiple access points (AP) to form a network, up to 1/2 second congestion delays at the AP (preventing conversational applications), and broken security. Fortunately these problems are being dramatically reduced or overcome via the introduction of 802.11a/g and upcoming 802.11n for improved bandwidth, 802.11e for QoS, 802.11i for security, and a number of standardization efforts and tools for 802.11 network monitoring, management, and control. These improvements may enable conversational video and audio over wireless enterprise networks.

Enterprise wireless has problems not present in last-mile Internet 802.11 streaming. Scalability is foremost among them: enterprises must support a large number of simultaneous users without service degradation. The spectrum available to 802.11 is large but not unlimited, and a dense enterprise deployment can exhaust it. To efficiently use available spectrum is an exceptionally challenging problem, and typically requires joint optimization of the physical layer, the multiple access layer, channel assignment, power control, and packet scheduling. In Internet streaming this optimization is performed within just one access point; in the enterprise it is done campus-wide, across dozens of access points with overlapping regions of coverage. Although “switched WLAN” systems that implement some of these centralized mechanisms are now available commercially, many of the basic questions of how to optimize for latency-sensitive data remain open.

New media-enabled and wireless-enabled end hosts are changing how content is created, shared, accessed and searched. For example, both emerging cell phones and handheld computers, e.g., HP iPAQ, include integrated high-resolution image and video cameras, video and audio real-time encoding, and wireless 802.11 for local connectivity and 2.5G (and soon 3G) cellular for wide-area coverage. These capabilities enable any user to create and upload content to the infrastructure to be shared with others. This capability may change the operation of many businesses, including medical, insurance, support and repair, emergency response, etc., and it has significant effects on how the large amount of end-user created media content is identified, stored, indexed, and searched for. Automated and accurate techniques to solve

these problems are necessary.

The centralized infrastructure of enterprise streaming offers the possibility of seamless mobility for wireless clients. With suitably-designed wireless infrastructure, all jointly controlled, the physical layer handover delay can be reduced to milliseconds, and the IP address of a mobile client can be preserved throughout the campus. Thus a streaming application need not even be aware of mobility. Internet streaming generally lacks such supporting infrastructure. Mobility generally requires a streaming session to be torn down and restarted after a user moves to a new area. Moreover, for P2P networks, directory services will take several seconds or minutes to “catch up” with the new location of the user.

Finally, an interesting assertion made by security experts is that because of the recent focus on the security vulnerability of wireless networks, the wireless portion of today’s enterprise networks may be more secure than the wired portion.

8. SECURITY

Important concerns for Internet streaming are piracy protection, confidentiality of end-user communication, secure but easy consumption (rendering, copying, editing, etc.) and security against attacks. In the enterprise setting the primary security goals are typically confidentiality of important information (or piracy protection from non-employees), secure logging, access control and protection from attacks.

The first line of enterprise security typically involves distinguishing between employees and non-employees (people inside or outside the enterprise). The first step toward this goal involves separating the enterprise network from the outside networks, with proxies and firewalls carefully designed to limit or completely prevent access in, and possibly limit access out. For nodes inside an enterprise network there is often associated a sense of trust, and therefore the average media within the enterprise (including email) is generally not encrypted during transit but potentially associated with access rules. Controlled access of media is evolving into an enterprise-specific digital rights management (DRM) system. Whereas Internet DRM systems (or better, consumer oriented DRM systems) are oriented towards a hostile environment with very few allowed user interactions (mostly rendering and storing), enterprise DRM systems typically live in a less hostile environment (i.e., less emphasis is needed on secure implementations), but the number of allowed interactions is typically much larger (e.g., editing, splicing, sampling, etc). As a matter of fact we observe that enterprise DRM systems and consumer DRM systems constitute two separate worlds: an interesting area of research is the creation of DRM systems that span both worlds, allowing a gradual transition from enterprise to consumer. In particular, in the area of professional content intended for consumer use it would be useful to have a security system available that would span the whole value chain,

from creation and aggregation to delivery and consumption.

9. SUMMARY

This paper provided an overview of enterprise streaming with the goal of describing the basic differences between enterprise streaming and Internet streaming and how these differences alter the challenges that must be overcome for media streaming in the enterprise. While it might at first seem that enterprise streaming is easier than Internet streaming, due for example to a system-wide control available in the enterprise, the demands on quality, reliability, and scalability can make it harder. Moreover, the availability of control does not answer the question of what should be controlled or how the control should be carried out. Questions about global optimization and control of streaming services, networks, and end-points within the enterprise provide a rich and active area of research.

10. REFERENCES

- [1] J.G. Apostolopoulos, W. Tan, and S.J. Wee, “Video Streaming: Concepts, Algorithms, and Systems,” *HP Labs Tech Report 2002-260*, Sept 2002.
- [2] D. Turaga, M. Schaar, and K. Ratakonda, “Enterprise multimedia streaming: Issues, background and new developments,” *IEEE ICME*, July 2005.
- [3] L. Cherkasova and M. Gupta, “Analysis of enterprise media server workloads: Access patterns, locality, content evolution, and rates of change,” *ACM/IEEE Transactions on Networking*, October 2004.
- [4] <http://www.bittorrent.com>.
- [5] <http://www.kazaa.com>.
- [6] S. Wee, J. Apostolopoulos, W. Tan, and S. Roy, “Research and design challenges for mobile streaming media content delivery networks (MSM-CDNs),” *IEEE ICME*, July 2003.
- [7] Y. Liu, H. Zhang, W. Gong, and D. Towsley, “On the interaction between overlay routing and traffic engineering,” *IEEE Infocom*, March 2005.
- [8] C. Zhang and T. Chen, “A survey of image-based rendering—representation, sampling, compression,” *Signal Processing: Image Communications*, January 2004.
- [9] H. Baker, N. Bhatti, D. Tanguay, I. Sobel, D. Gelb, M. Goss, J. MacCormick, B. Culbertson, and T. Malzbender, “Computation and performance issues in coliseum, an immersive videoconferencing system,” *ACM Multimedia*, November 2003.
- [10] <http://setiathome.ssl.berkeley.edu>.
- [11] M. Roussopoulos, M. Baker, D. Rosenthal, T. Giuli, P. Maniatis, and J. Mogul, “2 P2P or Not 2 P2P,” *3rd International Workshop on Peer-to-Peer Systems (IPTPS)*, February 2004.
- [12] www.skype.com.
- [13] S. Baset and H. Schulzrinne, “An analysis of the skype peer-to-peer internet telephony protocol,” *Columbia University Technical Report CUCS-039-04*, September 15, 2004.