

Proxy Caching for Media Streaming Over the Internet

Jiangchuan Liu, Simon Fraser University

Jianliang Xu, Hong Kong Baptist University

ABSTRACT

Streaming media has contributed to a significant amount of today's Internet traffic. Like conventional Web objects (e.g., HTML pages and images), media objects can benefit from proxy caching; but their unique features such as huge size and high bandwidth demand imply that conventional proxy caching strategies have to be substantially revised. This article discusses the critical issues and challenges of cache management for proxy-assisted media streaming. We survey, classify, and compare the state-of-the-art solutions. We also investigate advanced issues of combining multicast with caching, cooperating among proxies, and leveraging proxy caching in overlay networks.

INTRODUCTION

With widespread penetration of the broadband Internet, multimedia service is getting increasingly popular among users and contributes a significant amount of today's Internet traffic. Media objects can be accessed similar to conventional text and images using a download-and-play mode; but most users prefer to quickly initiate and then continuously play back a media object while it is being downloaded (i.e., to use a *real-time streaming* mode). We have witnessed the initial and incremental deployment of streaming applications like RealNetworks RealPlayer and Microsoft Windows Media Player in recent years. The performance of such applications, however, is still far from satisfactory, especially during peak hours.


To reduce client-perceived access latencies as well as server/network loads, an effective means is to cache frequently used data at proxies close to clients. This also enhances the availability of objects and mitigates packet losses, as local transmission is generally more reliable than remote transmission. Proxy caching thus has become one of the vital components in virtually all Web systems. Streaming media, particularly those prestored, could also gain significant performance improvement from proxy caching,

given their static nature in content and highly localized access interests. However, existing proxies are generally optimized for delivering conventional Web objects (e.g., HTML pages or GIF images), which may not meet the requirements of steaming applications. In the following, we list some important and unique features of streaming media and discuss their implications to proxy cache design.

Huge size: A conventional Web object is typically on the order of 1–100 kbytes. Hence, a binary decision works well for proxy caching: either caching an object in its entirety or not caching. In contrast, a media object has a high data rate and long playback duration, which combined yield huge data volume. For illustration, a 1-h standard MPEG-1 video has a volume of about 675 Mbytes; caching it entirely at a Web proxy is clearly impractical, as several such large streams would exhaust the capacity of the cache. One solution is to cache only portions of an object, in which case a client's playback needs joint delivery involving both the proxy and the origin server. Which portions of which objects to cache thus has to be carefully managed so that the benefit of caching outweighs the synchronization overhead of the joint delivery.

Intensive bandwidth use: The streaming nature of delivery requires a significant amount of disk and network I/O bandwidth, sustained over a long period. Hence, minimizing bandwidth consumption becomes a primary consideration for proxy cache management, even taking precedence over reducing access latencies in many cases. Moreover, the bandwidth bottleneck limits the number of clients a proxy can simultaneously support; employing multicast delivery and cooperation among proxies thus become particularly attractive for media streaming applications.

High interactivity: The long playback duration of a streaming object also enables various client-server interactions. As an example, recent studies found that nearly 90 percent of media playbacks are terminated prematurely by clients [1]. In addition, during playback a client often expects VCR-like operations, such as fast for-



■ **Figure 1.** A generic system diagram of proxy-assisted media streaming using RTP/RTCP/RTSP.

ward and rewind. This implies the access rates might be different for different portions of a stream, which potentially complicates cache management.

Given these unique features of media objects, novel caching algorithms have been developed in the literature. The objective of this article is to review the state-of-the-art caching techniques dedicated to streaming media caching. We begin with discussions of a generic proxy caching architecture and some protocol considerations. Next, the caching strategies for streaming media are classified, examined, and compared. We also investigate some advanced issues. Finally, we conclude the article.

ARCHITECTURE AND PROTOCOLS FOR STREAMING CACHING

Streaming applications generally support diverse client-server interactions and have stringent demands on packet delay and jitter to ensure discontinuity-free playback. To meet these requirements, the Internet Engineering Task Force (IETF) has developed the RTP/RTCP/RTSP protocol suite. A generic system diagram of proxy-assisted media streaming using this suite is depicted in Fig. 1.

In this system, the basic functionalities for data transferring are provided by the Real-Time Transport Protocol (RTP), including payload identification, sequence numbering for loss detection, and timestamping for playback control. Running on top of UDP, RTP itself does not guarantee quality of service (QoS), but relies on its companion, the Real-Time Control Protocol (RTCP), to monitor network status and provide feedback for application-layer adaptation. The Real-Time Streaming Protocol (RTSP) coordinates delivery of media objects and enables a rich set of controls for interactive playback. For proxy-assisted streaming, the proxy has to relay these control messages between the client and the server. The problem is particularly involved if only part of a media object is cached at a proxy. In this case, the proxy must reply to the client PLAY request and initiate transmission of RTP and RTCP messages to the client for the cached portion, while requesting the uncached portion(s) from the server. Such fetch-

ing can be achieved through an RTSP Range request specifying the playback points. The Range request also enables clients to retrieve different segments of a media object from multiple servers or proxies if needed.

Besides this classical client/server paradigm, peer-to-peer streaming and other overlay streaming paradigms have also attracted much attention recently; these will be discussed later.

CACHING STRATEGIES FOR STREAMING MEDIA

Due to the aforementioned features of streaming media objects, media caching has many distinct focuses from conventional Web caching. On one hand, since the content of a media object is rarely updated, management issues like cache consistency and coherence are less critical in media caching. On the other hand, given high resource requirements of media objects, effective management of proxy cache resources (i.e., space, disk I/O, and network I/O) becomes more challenging. In this section we survey the state-of-the-art media caching strategies for both homogenous clients and heterogeneous clients, with an emphasis on how the strategies minimize resource demands.

STREAM CACHING FOR HOMOGENEOUS CLIENTS

Most existing caching algorithms focus on homogeneous clients that have identical or similar configurations and capabilities behind a proxy. As such, a single version of an object would match the bandwidth and format demands of all requests to the object. Nevertheless, what to cache (which portions of which objects) and how to manage cache (e.g., cache placement and replacement) at the proxy remain challenges. According to the selection of the portions to cache, we classify existing algorithms into four categories: *sliding-interval caching*, *prefix caching*, *segment caching*, and *rate-split caching*.

Sliding-Interval Caching [2] — This algorithm caches a sliding interval of a media object to exploit sequential access of streaming media. For illustration, given two consecutive requests

Streaming applications generally support diverse client-server interactions and have stringent demands on packet delay and jitter to ensure discontinuity-free playback. To meet these requirements, the IETF has developed the RTP/RTCP/RTSP protocol suite.

Segment caching generalizes the prefix caching paradigm by partitioning a media object into a series of segments, differentiating their respective utilities, and making caching decision accordingly.




Figure 2. An illustration of sliding-interval caching. The object consists of nine frames, each requiring one unit time to deliver from the proxy to a client. Requests 1 and 2 arrive at times 0 and 2, respectively. To serve request 2, only two frames need to be cached at any time instance: a) time 0: request 1 arrives; b) time 1–2: frames 1 and 2 accessed by request 1 and cached; request 2 arrives; c) time 2–3: frame 3 accessed by request 1 and cached; frame 1 read by request 2 and released.

for the same object, the first request may access the object from the server and incrementally store it into the proxy cache; the second request can then access the cached portion and release it after the access. If the two requests arrive close in time, only a small portion of the media object needs to be cached at any time instance, and yet the second request can be completely satisfied from the proxy (Fig. 2). In general, if multiple requests for an object arrive in a short period, a set of adjacent intervals can be grouped to form a *run*, of which the cached portion will be released only after the last request has been satisfied.

Sliding-interval caching can significantly reduce network bandwidth consumption and start-up delay for subsequent accesses. However, as the cached portion is dynamically updated with playback, sliding-interval caching involves high disk bandwidth demands; in the worst case, it would double the disk I/O due to the concurrent read/write operations. In addition, its effectiveness diminishes with increased access intervals. If the access interval of the same object is longer than the duration of the playback, the algorithm is degenerated to the unaffordable full-object caching. To address these issues, it is preferable to retain the cached content over a relatively long time period. Most of the caching algorithms to be discussed in the rest of this section fall into this category.




Figure 3. A snapshot of prefix caching with workahead smoothing.

Prefix Caching [3] — This algorithm caches the initial portion of a media object, called the *prefix*, at a proxy. Upon receiving a client request, the proxy immediately delivers the prefix to the client, meanwhile fetching the remaining portion, the *suffix*, from the server and relaying to the client (Fig. 3). As the proxy is generally closer to the clients than the origin server, the start-up delay for a playback can be remarkably reduced.

To ensure discontinuity-free playback with a start-up delay of s , the proxy has to store a prefix of length $\max\{d_{\max} - s, 0\}$, where d_{\max} is the maximum delay from the server to the proxy. If cache space is abundant, the proxy can also devote some space to assist in performing *work-ahead smoothing* for variable bit rate (VBR) media [3]. With this smoothing cache, the proxy can prefetch large frames in advance of each burst to absorb delay jitter and bandwidth fluctuations of the server-to-proxy path. The delay of prefetching can be hidden by prefix caching. Similar to sliding-interval caching, the content of the smoothing cache is dynamically updated with playback. However, the purposes are different: the former is to improve cache hit for subsequent requests, while the latter is to facilitate workahead smoothing.

Segment Caching [1, 4–6] — Segment caching generalizes the prefix caching paradigm by partitioning a media object into a series of segments, differentiating their respective utilities, and making a caching decision accordingly (Fig. 4). Vari-





Figure 4. An illustration of segment caching.

ous segment caching algorithms have been proposed in the literature by employing different segmentations and utility calculations. Wu *et al.* [4] suggested grouping the frames of a media object into variable-size segments, with the length increasing exponentially with the distance from the start of the media; that is, the size of segment i is 2^{i-1} , which consists of frames 2^{i-1} , $2^{i-1} + 1, \dots, 2^i - 1$. The motivation is that a proxy can quickly adapt to the changing access patterns of cached objects by discarding big chunks as needed. The utility of a segment is calculated as the ratio of the segment reference frequency over its distance from the beginning segment, which favors caching the initial segments as well as those with higher access frequencies. Chen *et al.* [1], however, argued that neither the use of a predefined segment length nor the favorable caching of the initial segments is the best strategy for reducing network traffic. They suggested postponing segmentation as late as possible (called *lazy segmentation*), thus allowing the proxy to collect a sufficient amount of access statistics to improve the effectiveness.

A salient feature of segment-based caching is its support of VCR-like operations, such as random access, fast forward, and rewind. As an example, Fahmi *et al.* [6] proposed to cache some key segments of a media object, called *hotspots*, which are identified by content providers. When a client requests the object, the proxy first delivers the hotspots to provide an overview of the stream; the client can then decide whether to play the entire stream or quickly jump to some specific portion introduced by a hotspot. Furthermore, in fast forwarding and rewinding operations, only the corresponding hotspots are delivered and displayed, while other portions are skipped. As such, the load of the server and backbone network can be greatly reduced, but the client will not miss any important segments in the media object.

Rate-Split Caching [7] — While all the aforementioned caching algorithms partition a media object horizontally along the time axis, the rate-split caching partitions it vertically along the rate axis: the upper part will be cached at the proxy, whereas the lower part will remain stored at the origin server (Fig. 5). This type of partitioning is particularly attractive for VBR streaming, as only the lower part of a nearly constant rate has to be delivered through the backbone network. For a QoS network with resource reservation, the bandwidth reserved should be equal to the peak rate of a stream; caching the upper part at the proxy clearly reduces the rate variability and improves the backbone bandwidth utilization. A critical issue here is how to select the cutoff rate or, equivalently, the size of the upper part for caching. Zhang *et al.* [7] studied the impact of the cutoff rate for a single stream through empirical evaluation, and found that significant bandwidth reduction can be achieved with a reasonably small cache space. They also formulated the multiple-stream case as a knapsack problem with two constraints, disk bandwidth and cache space, and developed several heuristics (e.g., caching popular objects only, or caching those with high bandwidth reduction).



■ Figure 5. An illustration of rate-split caching.

Summary and Comparison — Table 1 summarizes the caching algorithms reviewed for homogeneous clients. While these features and metrics provide a general guideline for algorithm selection, the choice of a specific streaming system also largely depends on a number of practical issues, especially the complexity of the implementation. In fact, only a few simple algorithms have been employed in commercial systems, although recently built prototypes have practically demonstrated the viability and superiority of the intelligent algorithms, such as lazy segmentation [1].

In addition, we emphasize that these algorithms are not necessarily exclusive with each other, and a combination of them may yield a better performance. For example, segment caching combined with prefix caching of each segment can reduce start-up latency for VCR-like random playback from any key segment. Combination with conventional data caching algorithms has also been examined.

STREAM CACHING FOR HETEROGENEOUS CLIENTS

Owing to diverse network models and device configurations, clients behind the same proxy often have quite different requirements on the same media object, in terms of streaming rates or encoding formats. To accommodate such heterogeneity, a straightforward solution is to produce replicated streams of different rates or formats, each targeting a subset of clients. Although widely used in commercial streaming systems, the storage and bandwidth demands of this approach can be prohibitively high [8]. An alternative is to transcode a media object from one form to another of a lower rate or a different encoding format on demand [9]. The intensive computation overhead of transcoding, however, prevents a proxy from supporting a large diverse client population.

A more efficient approach to this problem is the use of layered encoding and transmission. A layered coder compresses a raw media object into several layers: the most significant, the *base layer*, contains the data representing the most important features of the object, while additional *enhancement layers* contain data that can pro-

We envision the FGS-based streaming and caching as a very promising solution to media steaming over the Internet comprising highly heterogeneous end-systems. Several caching algorithms have been proposed to minimize the bandwidth consumption and/or improve the client utility.

		Sliding-interval caching [2]	Prefix caching [3]	Segment caching [1, 4–6]	Rate-split caching [7]
Cached portion		Sliding intervals	Prefix	Segments	Portion of higher rate
VCR-like support		No	No	Yes	No
Resource Demand	Disk I/O	High	Moderate	Moderate	Moderate
	Disk space	Low	Moderate	High	High
	Sync overhead	Low	Moderate	High	High
Performance improvement	Bandwidth reduction	High ¹	Moderate	Moderate	Moderate
	Start-up latency reduction	High ¹	High	High ²	Moderate

¹ There is no reduction for the first request in a run.
² Assume the initial segment is cached.

■ **Table 1.** Comparisons of the caching algorithms for homogeneous clients.

gressively refine the quality. Thus, a client can subscribe to a subset of cumulative layers to reconstruct a stream commensurate with its capability. For layered caching, Kangasharju *et al.* [10] assumed that the cached portions are semi-static and only completed layers are cached. To maximize the total revenue, they developed effective heuristics based on an analytical stochastic knapsack model to determine the cache content. In their model the client population and the distribution of their capacities are known a priori. For layered streaming under dynamic conditions, Rejaie *et al.* [11] studied segment-based cache replacement and prefetching policies to achieve efficient utilization of cache space and available bandwidth (Fig. 6a). The main objective is to deal with the congestion problem for individual clients. To this end, the proxy keeps track of popularities of each object on a per layer basis. When the quality of the cached layers is lower than the maximum deliverable quality to an interested client, the proxy sends requests to the server for missing segments within a sliding prefetching window. On cache replacement, a victim layer is identified based on popularities, and its cached segments are flushed from the tail until sufficient space is obtained.

A critical drawback of the existing layered streaming systems is that the number of layers is pretty small, typically only two or three; hence, their adaptation granularity remains coarse. Fortunately, recent development in the coding area has demonstrated the possibility of fine-grained post-encoding rate control. An example is the MPEG-4 Fine-Grained Scalable (FGS) coder with bitplane coding. FGS generates embedded streams containing several bitplanes, each of which can be partitioned at any specific rate. As such, for narrowband clients the proxy can reduce the streaming rate using a bitplane filter; for wideband clients the proxy can fetch some uncached portion (i.e., higher-order bitplanes) from the server and assemble it with the cached portion to generate a high-rate stream. As illustrated in Fig. 6b, the available bandwidth of a client can be almost fully utilized, and, more importantly, both the filtering and assembling

operations in FGS can be done with fast response. Hence, we envision FGS-based streaming and caching as a very promising solution to media steaming over the Internet comprising highly heterogeneous end systems. Several caching algorithms have been proposed to minimize bandwidth consumption and/or improve client utility [8].


ADVANCED ISSUES

So far we have considered a standalone proxy with only unicast delivery. While it can noticeably reduce access latencies and backbone bandwidth demands, the scalability and robustness of this simple architecture are still restricted. We now discuss two effective enhancements, multicast and proxy cooperation; we also address the role of proxy in the recently popularized overlay communication paradigms.

COMBINING PROXY CACHING WITH MULTICASTING

Like caching, multicasting also explores the temporal locality of client requests. Specifically, it allows a media server to accommodate concurrent client requests with shared channels through batching, patching, or periodic broadcast. However, multicast delivery suffers from two important deficiencies. First, to save more bandwidth, it is better to accommodate more requests in one multicast channel by using a large batching/patching window; however, this leads to long start-up latencies. Second, while IP multicast is enabled in virtually all local area networks, its deployment over the global Internet remains limited in scope and reach. Hence, it is unlikely that a multicast streaming protocol can be used for geographically dispersed servers and clients.

Interestingly, both deficiencies can be alleviated through the use of proxies. Specifically, a request can be instantaneously served by a cached prefix while waiting for the data from a multicast channel [12, 13], and proxies can bridge unicast networks with multicast networks (i.e., employing unicast for server to proxy delivery while batch-



■ **Figure 6.** Caching for layered streaming: a) coarse-grained layering; b) fine-grained layering.

ing and/or patching local accesses). Wang *et al.* [13] derived the optimal length of the prefix to be cached for most typical multicast protocols, and showed that a careful coupling of caching and multicasting can produce significant cost savings over using the unicast service, even if IP multicast is supported only at local networks.

COOPERATIVE PROXY CACHING

In general, proxies grouped together can achieve better performance than independent standalone proxies. Specifically, the group of proxies can cooperate with each other to increase the aggregate cache space, balance loads, and improve system scalability [14, 15]. A typical cooperative media caching architecture is MiddleMan [14], which operates a collection of proxies as a scalable cache cluster. Media objects are segmented into equal-sized segments and stored across multiple proxies, where they can be replaced at a granularity of a segment. There are also several local proxies responsible to answer client requests by locating and relaying the segments. Note that in cooperative Web caching, a critical issue is how to efficiently locate Web pages with minimum communication costs among the proxies. This, however, is not a major concern for cooperative media caching, as the bandwidth consumption for streaming objects is orders of magnitude higher than that for object indexing and discovering. Consequently, in MiddleMan, a centralized coordinator works well in keeping track of cache states. On the other hand, while segment-based caching across different proxies facilitates the distribution and balance of proxy loads, it incurs a significant amount of overhead for switching among proxies to reconstruct a media object. To reduce such effects as well as to achieve better load balance and fault tolerance, Chae *et al.* [15] suggested a Silo data layout, which partitions a media object into segments of increasing sizes, stores more copies for popular segments, but still guarantees at least one copy stored for each segment.

STREAMING CACHING IN OVERLAY NETWORKS

So far, we have focused on the client/server paradigm for media streaming, with proxies acting as intermediaries between them. Generalizing the proxy functionalities into every end host will shift the system to the recently popularized overlay communication paradigms, such as peer-

to-peer communication or application-layer multicast. There have been many pioneering efforts on overlay streaming, which have demonstrated the superior scalability and deployability of these overlay systems; the enormous buffer capacities distributed in end hosts also enable efficient client-side caching and sharing to improve content availability as well as support asynchronous streaming [16–19].

Nevertheless, we are aware that, in contrast to reliable and dedicated servers or proxies, loosely coupled autonomous end hosts can easily crash, leave without notice, or even refuse to share their own data. Given that a media playback lasts a long time and consumes huge resources, we believe dedicated proxies will still play an important role in building high-quality media streaming systems; in particular, strategically placed proxies may effectively assist the construction and maintenance of large-scale overlays. On the other hand, we may also leverage the overlay paradigm in proxy design. As an example, Guo *et al.* [20] suggested a proxy and its clients be structured into a peer-to-peer system to collaboratively serve local streaming requests. Their work focused on local area collaboration. We believe that it can be extended to a two-level streaming overlay: a cluster of proxies forms an overlay in wide area networks while each proxy collaborates with local client caches to form a local overlay. These two overlays complement each other well. The proxy-level overlay provides dedicated storage and reliable service, while the local overlay provides scalable storage for caching and significantly reduces the load of the proxy.

CONCLUDING REMARKS

Proxy caching is an effective means to reduce access latencies as well as resource consumption for networked applications. Due to the unique features of media objects like huge size and high bandwidth demand, a number of novel streaming caching solutions have been reported in the literature. This article serves as a pioneer survey in this field, although it by no means covers all aspects. Plenty of research issues have yet to be addressed, such as caching over the wireless mobile Internet, for large-scale dynamic overlays, and with advanced video coding schemes such as multiple description coding [18], as well as security and privacy for cached media objects,

Generalizing the proxy functionalities into every end-host will shift the system to the recently popularized overlay communication paradigms, such as peer-to-peer communication or application-layer multicast.

We envision that streaming media caching remains a fertile area, and both theoretical and practical solutions to the listed problems are urged with rising demands on ubiquitous multimedia services throughout the world.

to name but a few. We envision streaming media caching as a fertile area, and both theoretical and practical solutions to the listed problems are urged with rising demands on ubiquitous multimedia services throughout the world.

REFERENCES

- [1] S. Chen *et al.*, "Designs of High Quality Streaming Proxy Systems," *Proc. IEEE INFOCOM '04*, Hong Kong, China, Mar. 2004.
- [2] R. Tewari *et al.*, "Resource-Based Caching for Web Servers," *Proc. MMCN '98*, San Jose, CA, Jan. 1998.
- [3] S. Sen, J. Rexford, and D. Towsley, "Proxy Prefix Caching for Multimedia Streams," *Proc. IEEE INFOCOM '99*, New York, NY, Mar. 1999.
- [4] K. L. Wu, P. S. Yu, and J. L. Wolf, "Segment-based Proxy Caching of Multimedia Streams," *Proc. World Wide Web Conf.*, Hong Kong, China, May 2001.
- [5] Z. Miao and A. Ortega, "Scalable Proxy Caching of Video Under Storage Constraints," *IEEE JSAC*, vol. 20, no. 7, Sept. 2002, pp. 1315–27.
- [6] H. Fahmi *et al.*, "Proxy Servers for Scalable Interactive Video Support," *IEEE Comp.*, vol. 43, no. 9, Sep. 2001, pp. 54–60.
- [7] Z.-L. Zhang *et al.*, "Video Staging: A Proxy-Server-Based Approach to End-to-End Video Delivery over Wide-Area Networks," *IEEE/ACM Trans. Net.*, vol. 8, no. 4, 2000, pp. 429–42.
- [8] J. Liu, X. Chu, and J. Xu, "Proxy Cache Management for Fine-Grained Scalable Video Streaming," *Proc. IEEE INFOCOM '04*, Hong Kong, China, Mar. 2004.
- [9] X. Tang, F. Zhang, and S. T. Chanson, "Streaming Media Caching Algorithms for Transcoding Proxies," *Proc. CPP '02*, Aug. 2002.
- [10] J. Kangasharju *et al.*, "Distributing Layered Encoded Video Through Caches," *IEEE Trans. Comp.*, vol. 51, no. 6, June 2002, pp. 622–36.
- [11] R. Rejaie *et al.*, "Multimedia Proxy Caching Mechanism for Quality Adaptive Streaming Applications in the Internet," *Proc. IEEE Infocom 2000*, Tel Aviv, Israel, Mar. 2000.
- [12] S. Ramesh, I. Rhee, and K. Guo, "Multicast with Cache (Mcache): An Adaptive Zero-delay Video-on-demand Service," *Proc. IEEE INFOCOM '01*, Anchorage, AK, Apr. 2001.
- [13] B. Wang, S. Sen, M. Adler, and D. Towsley, "Optimal Proxy Cache Allocation for Efficient Streaming Media Distribution," *Proc. IEEE INFOCOM '02*, New York, NY, June 2002.
- [14] S. Acharya and B. C. Smith, "Middleman: A Video Caching Proxy Server," *Proc. NOSSDAV 2000*, June 2000.

- [15] Y. Chae *et al.*, "Silo, Rainbow, and Caching Token: Schemes for Scalable, Fault Tolerant Stream Caching," *IEEE JSAC*, vol. 20, no. 7, Sept. 2002, pp. 1328–44.
- [16] D. Xu *et al.*, "On peer-to-peer Media Streaming," *Proc. IEEE Int'l. Conf. Distributed Comp. Sys.*, Wien, Austria, July 2002.
- [17] S. Jin and A. Bestavros, "Cache-and-relay Streaming Media Delivery for Asynchronous Clients," *Proc. 4th Int'l. Wksp. Networked Group Commun.*, Boston, MA, Oct. 2002.
- [18] V. N. Padmanabhan *et al.*, "Distributing Streaming Media Content Using Cooperative Networking," *Proc. NOSSDAV '02*, Miami, FL, May 2002.
- [19] Y. Cui, B. Li, and K. Nahrstedt, "oStream: Asynchronous Streaming Multicast in Application-Layer Overlay Networks," *IEEE JSAC*, vol. 22, no. 1, Jan. 2004, pp. 91–106.
- [20] L. Guo *et al.*, "PROP: a Scalable and Reliable P2P Assisted Proxy Streaming System," *Proc. IEEE Int'l. Conf. Distrib. Comp. Sys.*, Tokyo, Japan, Mar. 2004.

BIOGRAPHIES

JIANGCHUAN LIU [S'01, M'03] (csljc@ieee.org) is an assistant professor in the School of Computing Science, Simon Fraser University, BC, Canada. He was an assistant professor at the Chinese University of Hong Kong from 2003 to 2004. He received a B.E. degree (cum laude) from Tsinghua University, Beijing, China, in 1999, and a Ph.D. degree from Hong Kong University of Science and Technology in 2003, both in computer science. He is a recipient of a Microsoft research fellowship (2000) and a Hong Kong Young Scientist Award (2003), and is a co-inventor of three European/U.S. patents. His research interests include media streaming, service overlay networks, and ad hoc and sensor networks. He serves as EC or TPC member for various international conferences, including IEEE INFOCOM '04 and '05, and is a guest editor for the ACM/Kluwer *MONET Journal* special issue on wireless sensor networks.

JIANLIANG XU [S'02, M'03] (xujl@comp.hkbu.edu.hk) is an assistant professor in the Department of Computer Science at Hong Kong Baptist University. He received a B.Eng. degree in computer science and engineering from Zhejiang University, Hangzhou, China, in 1998, and a Ph.D. degree in computer science from Hong Kong University of Science and Technology in 2002. His research interests include mobile and pervasive computing, Web content delivery, and wireless networks. He has served as a session chair and program committee member for a number of international conferences including IEEE INFOCOM. He is co-editor of a book, *Web Content Delivery* (Kluwer, to be published).