

Exploring Large-Scale Peer-to-Peer Live Streaming Topologies

Chuan Wu, Baochun Li
Department of Electrical and Computer Engineering
University of Toronto
{chuanwu, bli}@eecg.toronto.edu
and
Shuqiao Zhao
Multimedia Development Group
UUSee, Inc.
zhaosq@uusee.com

Real-world live peer-to-peer (P2P) streaming applications have been successfully deployed in the Internet, delivering live multimedia content to millions of users at any given time. With relative simplicity in design with respect to peer selection and topology construction protocols and without much algorithmic sophistication, current-generation live P2P streaming applications are able to provide users with adequately satisfying viewing experiences. That said, little existing research has provided sufficient insights on the time-varying internal characteristics of peer-to-peer topologies in live streaming. This paper presents *Magellan*, our collaborative work with UUSee Inc., Beijing, China, for exploring and charting graph theoretical properties of practical P2P streaming topologies, gaining important insights in their topological dynamics over a long period of time.

With more than 120 GB worth of traces starting September 2006 from a commercially deployed P2P live streaming system that represents UUSee's core product, we have completed a thorough and in-depth investigation of the topological properties in large-scale live P2P streaming, as well as their evolutionary behavior over time, *e.g.*, at different times of the day and in flash crowd scenarios. We seek to explore real-world P2P streaming topologies with respect to their graph theoretical metrics, such as the degree, clustering coefficient, and reciprocity. In addition, we compare our findings with results from existing studies on topological properties of P2P file sharing applications, and present new and unique observations specific to streaming. We have observed that live P2P streaming sessions demonstrate excellent scalability, a high level of reciprocity, a clustering phenomenon in each ISP, and a degree distribution that does *not* follow the power-law distribution.

Categories and Subject Descriptors: C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed Applications*

General Terms: Measurement

Additional Key Words and Phrases: Peer-to-peer streaming, topology characterization

1. INTRODUCTION

Based on the peer-to-peer (P2P) communication paradigm, live P2P multimedia streaming applications have been successfully deployed in the Internet, with millions of users at any given time. Prominent examples that are better known to

the research community include CoolStreaming [Zhang et al. 2005], PPLive [Hei et al. 2006; 2007] and TVAnts [Silverston and Fourmaux 2006; 2007]. The design philosophy of P2P streaming applications focuses on shifting the bandwidth burden to end systems of the Internet, away from dedicated streaming servers. They have effectively achieved such a shift by allowing peers to exchange data blocks among themselves. The successful commercial deployment of P2P streaming applications has made it possible to stream volumes of legal content to the end users, with hundreds of live media channels.

As a commonly adopted design in most recently emerged successful P2P live streaming applications, a live media stream is divided into blocks, and delivered over a mesh overlay topology in each session, with reciprocal exchanges of useful blocks among multiple peers. The practical success of these applications has validated important advantages of such a mesh-based design over tree-based streaming topologies: mesh topologies may achieve better resilience to peer dynamics, better scalability in flash crowd scenarios, more efficient use of bandwidth, as well as simplicity with respect to topology maintenance.

It is interesting to observe that most current-generation P2P streaming applications employ relatively *simple* protocol designs, including heuristic peer selection strategies. They typically use central tracking servers to obtain initial knowledge of existing peers in the channels, and periodically exchange peer lists among peers themselves. As mesh-based streaming topologies play such an important role towards the commercial success of P2P live streaming, it is critical to acquire a thorough and in-depth understanding of the topological characteristics of these P2P meshes. It would be an intriguing research challenge to investigate and understand how the constructed topologies actually *behave* in practice, dynamically *evolve* over time, and *react* to extreme scenarios such as huge flash crowds.

Unfortunately, although Internet topologies have been characterized extensively at the IP layer, there has been little literature on the discovery and analysis of P2P topologies in live P2P applications. As a well-known work on P2P topology characterization, Stutzbach *et al.* [Stutzbach et al. 2005] explored topological properties of the file-sharing Gnutella network. However, short-lived queries in Gnutella are fundamentally different from long-lived streaming sessions based on block exchanges among peers, which may well lead to different topological properties. In addition, like most of the existing measurement studies, their measurements rely on the methodology of crawling the Gnutella network, by which the speed of crawling plays a significant role in deciding the accuracy of the constructed topologies.

In this paper, we seek to explore and chart large-scale P2P live streaming sessions, and to gain in-depth insights and a complete understanding of their topological characteristics. We present *Magellan*, a comprehensive study of graph theoretical properties of real-world P2P streaming topologies in UUSee Inc. [UUSee Inc.], one of the main P2P live streaming solution providers in mainland China. Our observations are primarily based on over 120 GB of traces and more than 10 million unique IP addresses that we have collected over a two-month period, from September to October, 2006. In addition, we have further validated our discoveries using a more recent set of traces in 2007. Different from the crawling methodology, our traces are collected as periodical reports from the peers in the network, as we collaborate with

UUSee to add measurement and reporting facilities into their client software. Such peer reports construct *instantaneous* streaming topologies of the entire network throughout the trace period.

With emphasis on their evolutionary nature over a long period of time, we have utilized and extended classical graph measurement metrics — such as the degree, clustering coefficient, reciprocity and likelihood — to investigate various aspects of the streaming topologies at different times of the day, in different days in a week, and in flash crowd scenarios. We also compare our discoveries with existing results related to file sharing applications, with further insights unique to P2P streaming.

The original insights that we have brought forward in this paper are the following. First, we show that the current-generation P2P streaming platform scales very well, even in large flash crowd scenarios. Second, we observe that the degree distribution towards active neighbors in a P2P mesh does *not* follow the power-law distribution. Third, we argue that ISP-based clusters are formed from the dynamic peer selection process carried out during live streaming sessions. Fourth, we believe that the high-level reciprocity among peers in exchanging useful blocks plays a key role towards the success of such P2P applications. Finally, we believe that there do not exist super-node hubs in the streaming topology.

The remainder of this paper is organized as follows. In Sec. 2, we review existing work in P2P measurements and topology characterization. In Sec. 3, we outline key elements of the peer selection protocol in UUSee, and discuss our trace collection methodologies. In Sec. 4, we analyze the topological properties from a number of important perspectives, and discuss the implications of our discoveries. Finally, we summarize our findings and conclude the paper in Sec. 5.

2. RELATED WORK

There have been a number of measurement-based studies on various P2P applications in recent years. With respect to KaZaA, Gummadi *et al.* [Gummadi et al. 2003] characterized its workload as compared to web traffic, and Liang *et al.* [Liang et al. 2006] studied the connectivity and dynamics of its two-tier overlay structure. With respect to the previous-generation Gnutella network, which is constructed as a flat unstructured overlay and uses TTL-scoped flooding for search, earlier work [Adamic et al. 2001; Jovanovic et al. 2001; Ripeanu et al. 2002] reported the discovery of power-law degree distributions and strong “small-world” properties, *i.e.*, small network diameters and peer clustering. They also analyzed the impact of such graph properties on the performance of Gnutella’s routing and searching mechanisms.

With respect to the topology characterization of modern Gnutella network, which features a two-level architecture, Stutzbach *et al.* [Stutzbach et al. 2005] utilized a fast Gnutella crawler to capture back-to-back snapshots of the Gnutella network, in order to study its graph properties and dynamics. They argued that the discovery of power-law peer degree distribution in earlier work was a result of distorted snapshots captured by slow crawlers, while the degree distribution is more accurately described as a two-piece power-law distribution with spikes in between the two segments. They also reported that modern Gnutella topologies exhibit the “small-world” phenomenon, but with a comparably smaller degree of clustering than the

previous-generation Gnutella topologies.

The focus of our investigation in *Magellan* is on modern P2P live streaming topologies, and is significantly different from existing work on P2P topology characterization. First, current-generation P2P streaming applications are based on a BitTorrent-like block distribution mechanism, which includes block exchanges, centralized peer selection by tracking servers, and the construction of a loosely-connected mesh overlay topology featuring dynamic reconfiguration. No previous work exists for topology characterization of such BitTorrent-like P2P applications. Second, mesh-based P2P live streaming applications have fundamentally different requirements as compared to BitTorrent-like file-sharing applications, especially with respect to peer selection and topology construction, which usually features a biased protocol design towards timely delivery of media content. Third, we base our topological study on the instantaneous snapshots of the entire P2P streaming network, which are constructed with instantaneous peer reports and represent a much lower level of distortion over the time domain, as compared to the snapshots collected by a crawling method.

Topology characterization aside, towards measurements of various modern P2P applications, Pouwelse *et al.* [Pouwelse et al. 2005], Izal *et al.* [Izal et al. 2004] and Guo *et al.* [Guo et al. 2005] have investigated the performance of BitTorrent based on different scales of traces with various metrics. Peer churns and session characteristics in *KAD*, a Distributed Hash Table, were explored by Stutzbach *et al.* [Stutzbach and Rejaie 2006] and Steiner *et al.* [Steiner et al. 2007]. Various aspects of *Skype*, a P2P VoIP application, have been explored with respect to its protocol [Baset and Schulzrinne 2006], traffic characteristics [Guha et al. 2006], and user satisfaction [Chen et al. 2006].

Targeting commercial P2P streaming applications, Hei *et al.* [Hei et al. 2006; 2007] carried out a measurement study of *PPLive*, from both global views of peer distribution characteristics and local views of media traffic characteristics. In [Hei et al. 2007], they investigated the time variance of peer numbers and the geographic distribution of peers (which we also summarize as the first step of analyzing UUSEE's traces), but none of the other important topological properties. Ali *et al.* [Ali et al. 2006] presented an analysis of the control traffic, resource usage, locality, and stability of two P2P streaming protocols: PPLive and SOPCast. Silverston *et al.* [Silverston and Fourmaux 2006] studied the upload/download traffic generated by another P2P IPTV system, TVants, while users were watching the last FIFA world cup. In their recent work [Silverston and Fourmaux 2007], they further compared the traffic patterns among four popular P2P IPTV applications, PPLive, PPStream, SOPCast, and TVants, with the traces collected during the FIFA world cup. In addition, Cheng *et al.* [Cheng et al. 2007] have studied the end user experience in a P2P video-on-demand system, GridCast.

A preliminary report of this work appeared in ICDCS 2007 [Wu et al. 2007]. This paper represents a substantial revision, with complete details of the UUSEE peer selection mechanism, and more extensive topological characterization of the UUSEE network. Prominent examples of new studies include the topological characterization of peers grouped based on their geographical areas, in addition to the ISP-based topological studies, and the exploration of the existence and connectiv-

ity of supernodes in the UUSee streaming network. In addition, we have further validated our discoveries made from the 2006 trace set with a more recent collection of traces in 2007.

After our conference paper [Wu et al. 2007] was published, we are aware of a more recent measurement study on PPLive [Vu et al. 2007], that has also investigated two graphical properties of the P2P streaming application, *i.e.*, degree and clustering coefficient. Nevertheless, their work is significantly different from ours in a number of aspects: (1) Their measurements are collected based on the crawling of the PPLive network, while we believe our approach of trace collection — using periodical peer reports by collaborating with the P2P streaming solution provider — obtains more accurate topology snapshots in a much larger scale, both spatially and temporally, without possible distortion in the time domain. (2) Their degree investigation is limited to the average outdegrees of nodes in one media channel, while we have studied more cases extensively for both peer indegree and outdegree. (3) They study the clustering coefficient within a small portion of the crawled peers, while our clustering coefficient study is based on the entire P2P streaming network. (4) Based on the crawling methodology, they can only derive a list of known partners from each crawled peer, but not the knowledge whether a partner is active (with instantaneous block exchanges) or not. In contrast, we are able to collect the number of blocks delivered and the throughput achieved over each P2P link in each measurement interval, and thus derive the *active instantaneous* topologies of UUSee, consisting only of active links. For a P2P live streaming application, we believe it is more meaningful to study such “live” P2P streaming topologies.

Therefore, to the best of our knowledge, this paper represents the first complete and in-depth measurement study to characterize large-scale P2P live streaming topologies, with an emphasis on their graph theoretical properties. Nevertheless, we also seek to compare our new insights with discovered topological properties from traditional file sharing P2P applications, to identify properties that are unique to mesh-based P2P streaming, and also make comparisons with related results discussed in other measurement studies on P2P streaming applications.

3. BACKGROUND

As this paper focuses on the graph-theoretical analysis of large-scale traces from real-world P2P live streaming sessions offered by UUSee Inc. [UUSee Inc.], we believe that it is important to first introduce the P2P streaming protocol being used, as well as our methodology for collecting traces.

3.1 P2P Streaming Solutions from UUSee

UUSee Inc. [UUSee Inc.] is one of the leading P2P streaming solution providers in mainland China, featuring exclusive contractual rights to most of the channels of CCTV, the official Chinese television network. When compared to PPLive (which is better known in the research community due to existing measurement studies), it features a rich collection of *legal* content, encoded and broadcast live over the Internet. The UUSee P2P streaming framework consists of a media encoding server, which encodes the media channels to high quality constant-rate streams around 400 Kbps using its own proprietary codec, and a large collection of dedicated streaming servers, which receive encoded media channels from the media encoding server and

serve the P2P network composed of all the users of UUSEE. With its large collection of dedicated streaming servers around the world, UUSEE simultaneously sustains over 800 media channels. With a growing user base, it routinely serves millions of users in any given day. Its Windows-based P2P streaming client represents one of the most popular downloads in this category.

Similar to most current-generation P2P streaming protocols, UUSEE's streaming protocol (henceforth referred to as *UUSEE*) is designed with the principle of allowing peers to serve each other by exchanging blocks of data, that are received and cached in their local playback buffers. The buffer at each peer represents a *sliding window* of the media channel, containing blocks to be played in the immediate future. After a new peer joins a channel in UUSEE, the initial set of a small number of *partners* (up to 50) is supplied by one of the tracking servers. The peer establishes TCP connections with these partners, and *buffer availability bitmaps* are exchanged periodically. During this process, it detects the *round-trip delay* with each partner by using handshaking protocol messages, measures the TCP throughput of the connection, and also executes an estimation algorithm based on such measurements to predict the partner's availability to serve itself. It then ranks all known partners, and selects the best 30 peers from which it actually requests media blocks. A synchronous playback mechanism is employed in UUSEE, by which each newly joined peer always starts to buffer the media blocks that are to be played 20 seconds later than the current playback time of the media channel at the media encoding server, and as thus, all peers in the channel are following a similar playback progress.

The buffer size at each peer in UUSEE is 500 media blocks, and each block represents 1/3 second of media playback. The new peer starts the media playback from the first buffered block after 20 second, if a satisfactory buffering level has been reached during this time period. Otherwise, it will restart its initial buffering process for another 20 seconds, and then start the playback from the first block that has been buffered during this new 20 second period. Therefore, the initial startup delay at the peers in UUSEE is usually 20 seconds, and may be up to 40 seconds or 1 minute depending on the availability of media blocks in the network. During the playback at each peer, blocks to be played in the immediate future are continuously requested and cached in the playback buffer, and those that are not retrieved in time for playback are skipped. There is further a policy employed in the buffering process, that a peer will stop filling up its buffer when the buffer has reached around 75% of its total size.

Compared to other P2P streaming protocols such as PPLive, UUSEE incorporates a number of unique algorithms in its peer selection strategy, in order to guarantee smooth media playback at the peers. The main objective in P2P streaming is to shift the bandwidth burden away from dedicated streaming servers, by maximally utilizing peer upload bandwidth. During the initial startup phase, each peer in UUSEE employs an algorithm to estimate its maximum download and upload capacities. During the streaming process, each peer continuously estimates its aggregate instantaneous receiving and sending throughput from and to all its partners. If its estimated sending throughput is lower than its upload capacity for 30 seconds, it will inform one of the tracking servers that it is able to receive new connections from other peers. The tracking servers keep a list of peers that are able

to accept new connections, and bootstrap new peers with existing peers that are randomly selected from this set.

The number of available blocks in the current playback buffer is used in UUSee to represent the current streaming quality of the peer, which is referred to as the *buffer count*. During the streaming process, in addition to exchanging new media blocks with each other, neighboring peers also recommend known peers to each other. The buffer count is used as an important criterion for such recommendations. When a peer i finds that another peer j has a low buffer count, *i.e.*, an insufficient number of blocks in its buffer, peer i will recommend its known peers with larger buffer counts. As a last resort, when a peer has a low buffer count for a sustained period of time, it will contact the tracking server again to obtain additional peers with better streaming qualities.

3.2 *Magellan*: Collection of Traces

To discover and chart live P2P streaming topologies, we instrument and measure a wide range of performance metrics and protocol parameters in the UUSee streaming protocol. Besides the measurements that have been mentioned earlier, for each active partner with which it has a live TCP connection, each peer also actively measures the number of sent/received blocks, as well as the sending or receiving throughput of the TCP connection.

A standalone trace server is responsible for the collection of measurement reports from existing peers. Each report includes an extensive collection of vital statistics within the peer. It includes basic information such as its IP address, the channel it is watching, its buffer availability bitmap, as well its total download and upload bandwidth capacities. In addition, the report also includes a list of all its partners, with their corresponding IP addresses, TCP/UDP ports, number of blocks sent to or received from each partner, and current sending/receiving throughput on each connection. The reports are transmitted to the trace server periodically using UDP datagrams.

A new peer sends its initial report after 20 minutes, and sends subsequent reports once every 10 minutes. This ensures that the reports are sent by relatively long-lived peers in the channels. Even though the reporting peers represent only a subset of all existing peers, they constitute a stable “backbone” of the streaming topologies, and are more representative in our topological studies. In addition, since each peer reports a large number of its active partners (up to hundreds), there is a high probability that transient peers may appear in the partner lists of at least one reporting peer.

Since September 2006, we have commenced collecting these measurements to UUSee’s trace server, by upgrading all existing UUSee clients to the new release that produces periodic reports. To visualize our traces, in Fig. 1, we illustrate an example topology snapshot constructed from the reports of three representative peers, taken at 10:08:45 a.m., September 5, 2006. In a two-month span, we have collected more than 120 GB of traces on the trace server with more than 10 million unique IP addresses, constituting continuous-time snapshots of P2P streaming topologies throughout this period. In this paper, we take advantage of the traces to focus our studies on the inference and discovery of global topological properties and their evolutionary dynamics.

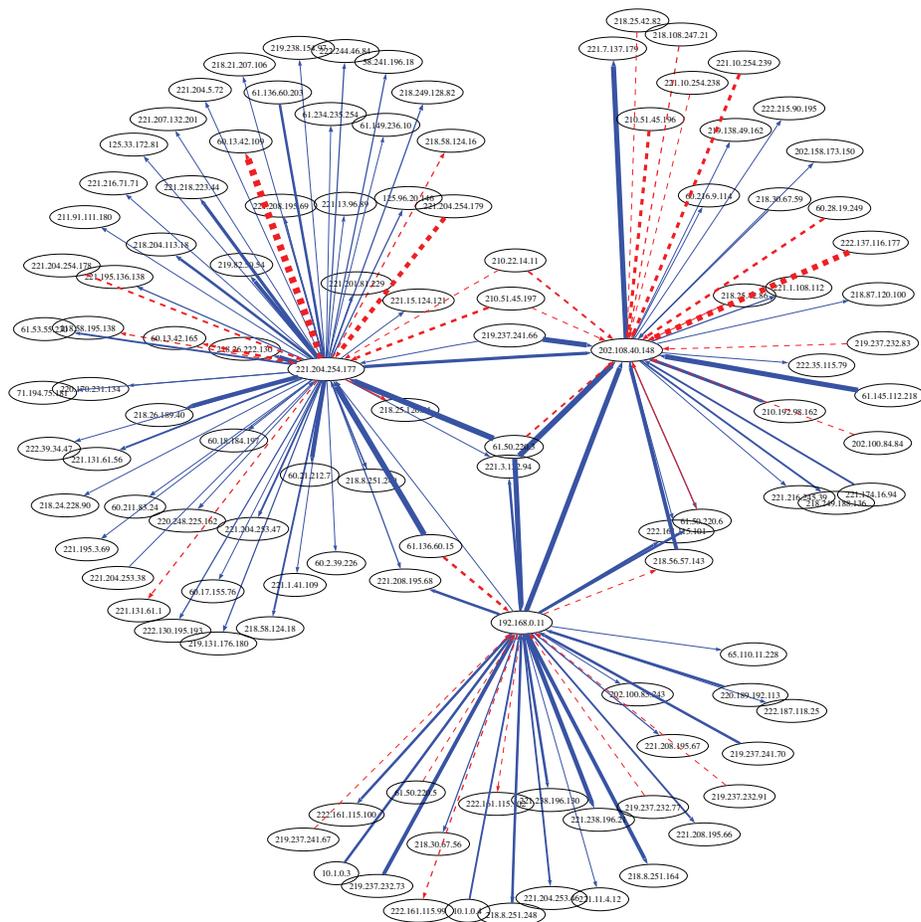


Fig. 1. A topology snapshot involving three reporting peers and their active partners, visualized from traces collected at 10:08:45 a.m., September 5, 2006. While widths of both types of lines represent bandwidth, the dashed links have 10 times higher bandwidth per unit width than the solid ones.

4. CHARTING LARGE-SCALE P2P STREAMING TOPOLOGIES

We now present an in-depth investigation of the topological properties of UUSEE P2P streaming overlays, as well as their dynamics over time. While we have analyzed the entire two-month trace period, we choose to show results from two representative weeks, from 12:00 a.m. October 1st, 2006 (GMT+8) to 11:50 p.m. October 14th, 2006 (GMT+8). We have observed that the selected periods have included all typical scenarios that we wish to present.¹

¹When we are preparing this journal paper, we have collected UUSEE traces over 9 months. To validate the general applicability of our topological property discoveries from the earlier trace set, we will also present representative topological characteristics from a more recent trace period, from 12:00 a.m. February 13th, 2007 (GMT+8) to 11:50 p.m. February 19th, 2007 (GMT+8).

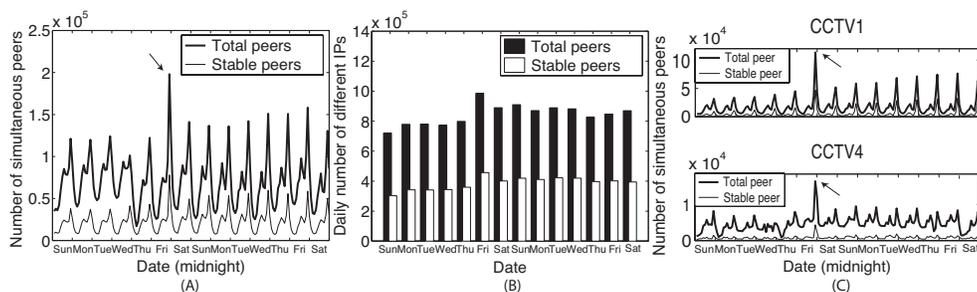


Fig. 2. Daily peer number statistics.

4.1 Scale of UUSee topologies

As a natural first step, we investigate the scale and general streaming performance of the UUSee application, with respect to the number of peers in different categories and their streaming rates.

4.1.1 Overall number of simultaneous peers. First of all, we are interested in addressing the following two questions. (1) How many concurrent users are usually online in the UUSee streaming overlay? (2) What percentage do the stable peers (whose reports are received by the trace server) take in the peer population, as compared to transient peers (whose reports are not received)? To answer these questions, we summarize the IP addresses from which reports were received and recorded in the traces, and all the IP addresses that appeared in the traces, including IP addresses of peers that have reported and peers in their partner lists. The peer number statistics are shown in Fig. 2(A).

The statistics indicate that there are around 100,000 concurrent peers at any time in the UUSee streaming overlay. There is a daily peak around 9 p.m., and a second daily peak around 1 p.m., which identify the same daily peer number pattern as given in the study of PPLive [Hei et al. 2007]. Different from [Hei et al. 2007], we observe only a slight peer number increase over the weekend, considering the weekly variance trend. In addition, we have clearly observed a flash crowd scenario around 9 p.m. on October 6, 2006, which was the mid-autumn festival in China, and the flash crowd was caused by the broadcast of a celebration TV show on a number of CCTV channels.² Comparing the number of stable peers to the total number of all peers, we discover that the former is asymptotically 1/3 of the later.

To obtain a better idea of the scale of daily users of UUSee, we further summarize the number of distinct IP addresses that appeared in the traces on a daily basis in Fig. 2(B). The statistics show that UUSee serves up to 1 million different users each day.

4.1.2 Number of simultaneous peers in two representative channels. As mentioned earlier, UUSee provides over 800 channels to the users. We have investigated and discovered that the topologies of different channels are largely disjoint. In this

²Note that in all figures throughout the paper that present the temporal evolution of a metric, a small arrow is drawn to indicate the occurrence time of this flash crowd.

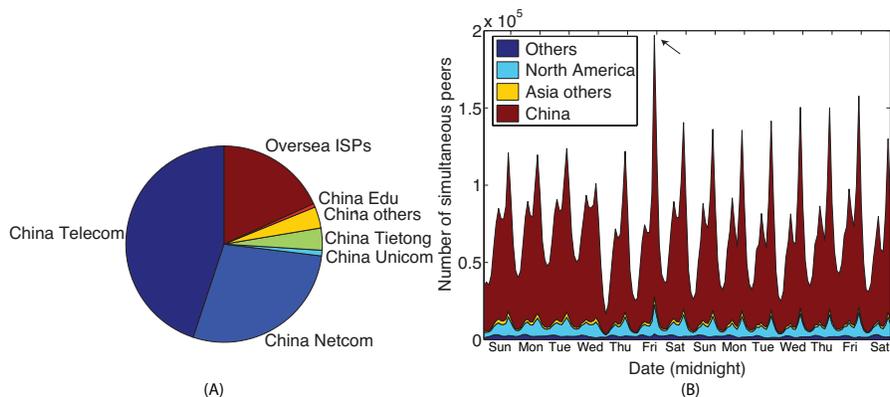


Fig. 3. Peer number statistics for different ISPs and geographic areas.

case, we wonder whether the observations we have made earlier regarding the global topology also apply to individual channels. To this end, we select two representative channels broadcast by UUSee, CCTV1 and CCTV4 (both from the official Chinese television network), where CCTV1 is among the most popular channels sustained in UUSee and CCTV4 has less popularity. Their peer number statistics are shown in Fig. 2(C). The different scales of the concurrent peer numbers clearly demonstrate the popularity difference of the two channels. Nevertheless, the evolutionary pattern of peer number in both channels is very similar to that in the global topology. In addition, both curves reflect a flash crowd scenario on October 6, as both channels broadcast the celebration TV show. The CCTV1 curve further exhibits a more distinctive daily peak on evenings, as is the prime time for China news broadcasting in the channel.

4.1.3 Number of simultaneous peers in different ISPs and areas. Throughout the paper, we also emphasize on the mapping of the abstract streaming topology to the real world scenario, with respect to the ISP and geographic area each peer is located at. For this purpose, we have obtained a mapping database from UUSee Inc. that translates ranges of IP addresses to their ISPs and geographic areas. For each IP address in China, the database provides the China ISP it belongs to and the city/province the user is located at; for each IP address out of China, it provides a general ISP code indicating foreign ISPs, and coarse geographic information of the continent the address lies in.

Using this mapping database, we have determined the ISP membership and geographic distribution of simultaneous peers at any time. With respect to ISP distribution of peers, we have discovered that it does not vary significantly over the two-month period that we collected the traces. Therefore, we only depict the averaged shares of peers in major ISPs over the trace period in Fig. 3(A). For peer geographic distribution, Fig. 3(B) depicts its evolution. We observe that while users in China dominate the UUSee streaming network, there still exist a substantial number of peers from overseas, especially from the North America. Another interesting discovery is that the evolutionary pattern of the number of concurrent North American users is similar to that of users in China. We identify the reason

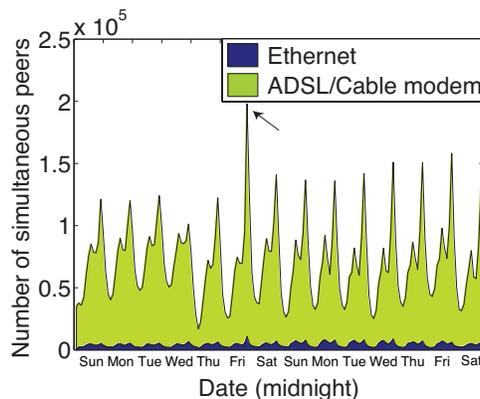


Fig. 4. Peer populations in two categories: Ethernet and ADSL/cable modem.

to be that the majority of UUSEE users are watching CCTV channels, and their popular programs are broadcast live according to the same Beijing time (GMT+8).

In our subsequent studies, when we explore characteristics of global topologies, we include all peers from all ISPs; when we investigate ISP-based topological properties, we will mainly focus on the ISPs inside China.

4.1.4 Population of Ethernet and ADSL/cable modem peers. Making use of peer download capacity estimations included in the traces, we investigate the distribution of peers in two categories, classified based on the fact that the download bandwidth of the fastest ADSL connections in China is at most 3 Mbps: (1) Ethernet peers, for those with download capacity higher than 384 KBps; (2) ADSL/cable modem peers, for the remainder. The statistics in Fig. 4 exhibit that ADSL/cable modem peers constitute the majority of UUSEE users. Especially in the flash crowd scenario, the number of ADSL/cable modem peers peaked, but there was little increase of Ethernet peers. We infer that the reason is that users tend to watch the broadcasts at home, and most of the residential high speed services in China are based on ADSL or cable modems.

4.1.5 Streaming quality. To explore the streaming quality of UUSEE, we make use of two measurements collected at each peer in different channels — the number of available blocks in the current playback buffer (buffer count) and the aggregate instantaneous receiving throughput. With the example of CCTV1 and CCTV4, Fig. 5(A) depicts the percentage of peers in both channels whose buffer counts are no lower than 75% of the total buffer size, and Fig. 5(B) shows the percentage of peers in both channels whose receiving throughput is higher than 90% of the channel streaming rate.

With respect to the buffer count metric, we observe that the buffering level for the popular channel CCTV1 is very encouraging, as for most of the time, around 95% of its participating peers have a satisfactory playback buffer count. Considering the UUSEE buffering policy we have mentioned in the previous section, that a peer will stop buffering when its buffer has been roughly 75% full, this shows most of the peers in CCTV1 are having a “full” buffer for continuous playback. For

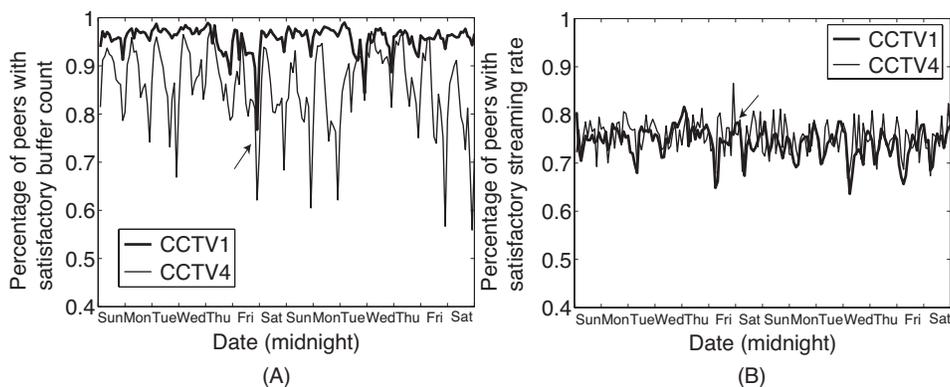


Fig. 5. Percentage of peers with satisfactory streaming quality.

the less popular channel CCTV4, the peer buffering level represents much larger fluctuations, but still over 80% of its peers achieve satisfactory buffering for the majority of the times. Such a difference between popular and less popular channels exhibits that peers watching a popular channel can usually retrieve media blocks more easily as there are more supplying peers in the channel. A closer look at the plots further reveals that the buffering level is generally lower at the peak hours of a day and during the flash crowd on October 6, 2006, especially for the less popular channel.

In terms of the aggregate receiving throughput metric, we can see that around 3/4 of all viewers in both channels achieve satisfactory streaming rates, and no evident daily pattern is discovered for the evolution of the percentages. Comparing the two quality metrics, peer buffer count and aggregate receiving throughput, we would suggest that the buffer count metric better represents the actual streaming quality experienced at the peers, as it captures the block receiving quality in a recent period of time. On the other hand, the aggregate receiving throughput is measured instantaneously at each peer, and thus may be less accurate. Nevertheless, results from both metrics are positive, exhibiting that the UUSEE peer selection protocol scales quite well to a large number of peers.

4.2 Degree distribution of stable peers

In our traces, each stable peer reports the IP addresses in its partner list, and also the number of blocks sent (received) to (from) each of the partners. With this information, we are able to categorize partners of each peer into three classes: (1) active supplying partner, from which the number of received blocks is larger than a certain threshold (*e.g.*, 10 blocks); (2) active receiving partner, to which the number of sent blocks is larger than the threshold; (3) non-active partner, in all the other cases.

With reports from stable peers in the streaming overlay, we investigate their degree distributions with respect to the number of active supplying partners (indegree), the number of active receiving partners (outdegree), and the total number of partners in the partner list including both active and non-active partners. Note

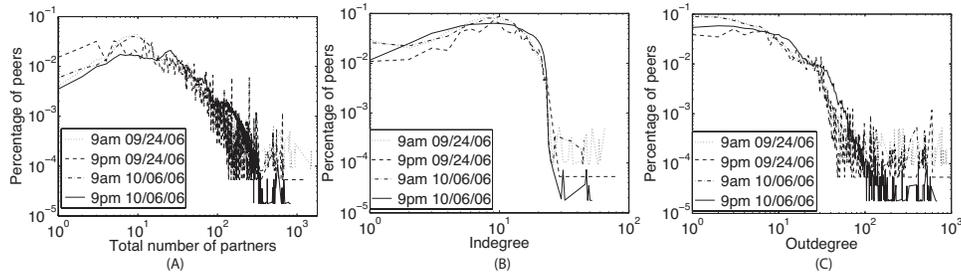


Fig. 6. Degree distributions of stable peers in the global topology.

that in a mesh network, it is common for a partner to be both a supplying partner and a receiving partner of a peer at the same time. In this case, it is counted into both peer active indegree and active outdegree.

4.2.1 Degree distribution in the global topology. Most existing research on peer-to-peer topologies reported power-law degree distributions. In their studies on modern Gnutella topologies, Stutzbach *et al.* [Stutzbach et al. 2005] pointed out that its degree distribution does not follow a power-law or two-segment power-law distribution, but has a spike around 30, as the Gnutella client software tries to maintain 30 neighbors for each peer. From Fig. 6(A), we observe that the distributions of total number of partners at the stable peers in the UUSee network do not follow power-law distributions either, with spikes whose corresponding degrees vary at different times. For the distributions observed in the morning, the spikes lie around a partner number of 10; for those observed in the daily peak hour, 9 p.m. at night, the spikes are located at larger values. During the flash crowd scenario around 9 p.m., October 6, 2006, the distribution peaks around 25. These reveal that at peak times, each peer is engaged with more partners. In addition, the results also exhibit that although each peer has an initial set of around 50 partners upon joining, the number of partners at different peers varies a lot during the streaming process, according to neighbor dynamics, the block availability at known neighbors and the available upload capacity at each peer.

For the peer indegree distribution shown in Fig. 6(B), we observe spikes around 10, and the spike is at a slightly higher degree in the flash crowd scenario. For indegree distributions at all times, they drop abruptly when the indegree reaches about 23. According to the UUSee peer selection protocol, a peer only accepts new upload connections when it still has spare upload capacity, and thus the upload bandwidth on each upload link is guaranteed. Besides, during streaming, the aggregated download rate at each peer is limited by the streaming rate in UUSee. All these explain the observation that the number of active supplying peers at each peer, which guarantees satisfactory streaming rates, is relatively small in the UUSee overlay, as compared to other file sharing applications.

The peer outdegree distributions in Fig. 6(C) are closer to two-segment power-law distributions, with a joint point around degree 10. The curves for peak times exhibit a flatter first segment, which implies that peers have higher outdegrees when there are more requesting peers in the network.

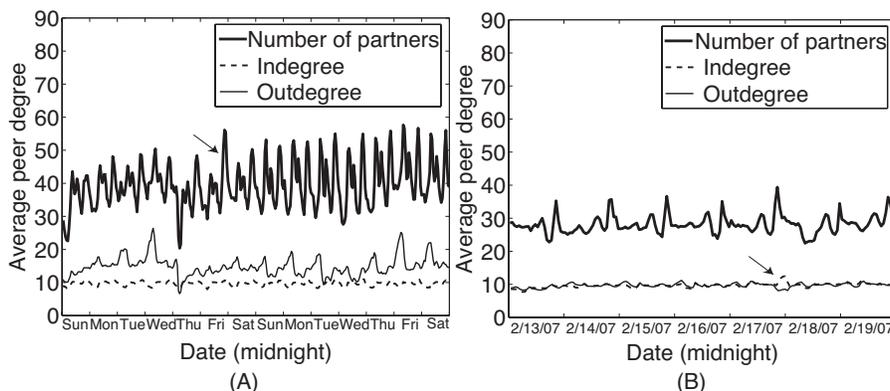


Fig. 7. Evolution of average degrees for stable peers in the global topology. (A) From October 1, 2006 to October 14, 2006. (B) From February 13, 2007 to February 19, 2007.

4.2.2 Degree evolution in the global topology. We next show the evolution of average degrees of stable peers during the two-week period in Fig. 7(A). We observe that the averaged total number of partners peaks at the peak times, but the average peer indegree is consistently around 10. Given that UUSee does not explicitly impose such an upper bound for active incoming connections at each peer, we explain this phenomenon as follows. At peak times when a large number of peers coexist in the network, there are abundant supplies of upload bandwidth, considering the streaming rate of around 400 Kbps is lower than the upload capacity of most current ADSL/cable modem peers in China, which constitute the majority of UUSee users. Combining with our previous results that large portions of peers can achieve satisfactory streaming rates at the peak hours, we conjecture that many peers will be able to offer help to others, and thus “volunteer” themselves at the tracking server, or get known by other peers when peers exchange their useful partner lists. This leads to the result that each peer knows a large number of other peers. Nevertheless, each peer does not actually need to stream from more peers to sustain a satisfactory streaming rate.

To validate our discoveries with the more recent UUSee traces, we investigate the peer degree evolution during the week of February 13, 2007 to February 19, 2007, which also includes a flash crowd scenario due to the broadcast of Chinese New Year celebration show on the evening of February 17, 2007. From the results in Fig. 7(B), we can see the average number of partners and average outdegree per peer are smaller than those in Fig. 7(A), while the average indegree is at a similar level. The reduction of partner number may be attributed to the ISP upgrade of the access link bandwidth, which occurred during these months, so that peers can achieve satisfactory streaming rates without knowing many other peers. Nevertheless, the degree evolution patterns remain similar. In addition, we have also investigated the degree distributions at each specific time during the new trace period, which we have found represent similar shapes to those in Fig. 6, and are thus omitted for presentation.

In Vu *et al.*'s PPLive measurement study [Vu et al. 2007], they have derived an ACM Transactions on Multimedia Computing, Communications and Applications, Vol. V, No. N, October 2007.

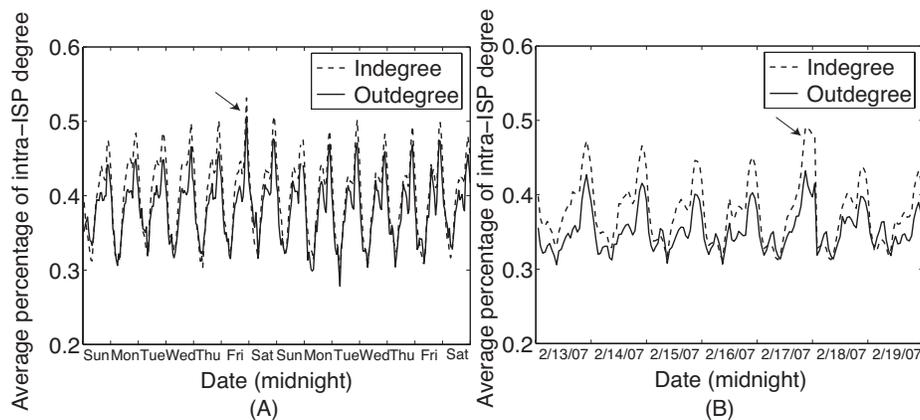


Fig. 8. Evolution of average intra-ISP degrees for stable peers in the network. (A) From October 1, 2006 to October 14, 2006. (B) From February 13, 2007 to February 19, 2007.

average node outdegree within the range of 30 to 43. In their measurements, the outdegree at each peer includes all the partners that may retrieve from the peer, not necessarily only the ones that are actively streaming from it at each specified time, as how our outdegree is measured. Therefore, a fairer comparison would be to compare their results with our total number of partners, which are at a similar magnitude for both P2P streaming applications.

4.2.3 Intra-ISP degree evolution. To better understand the connectivity among peers in the same ISP and across different ISPs, we further investigate the active indegrees and outdegrees at each peer that are from and to peers in the same ISP. At each stable peer, we calculate the proportion of indegrees from partners in the same ISP to the total indegree of the peer, and the proportion of outdegrees toward partners in the same ISP to its total outdegree, respectively.

Fig. 8 plots the evolution of the intra-ISP degree percentage, averaged over all stable peers in the network at each time. From Fig. 8(A), we observe that the percentages for both indegrees and outdegrees are around 0.4. Considering that many ISPs coexist, this exhibits that the majority of active supplying/receiving partners of each peer are within the same ISP. Although UUSEE does not take ISP membership into consideration when the tracking server assigns new partners to a peer and when neighboring peers exchange partners, this exhibits the “natural clustering” effects in the P2P streaming overlay over each ISP. The reason behind such clustering is that, as connections between peers in the same ISPs have generally higher throughput and smaller delay than those across ISPs, they are more inclined to be chosen as active connections.

In addition, Fig. 8(A) shows that the percentages for both indegree and outdegree peak at the daily peak hours and during the flash crowd scenario. This implies that each peer has more partner choices when the network is large, and it is always able to choose high throughput connections that are largely intra-ISP.

All the above observations are further validated by the investigation results using the more recent traces in February, 2007, as shown in Fig. 8(B). This exhibits the

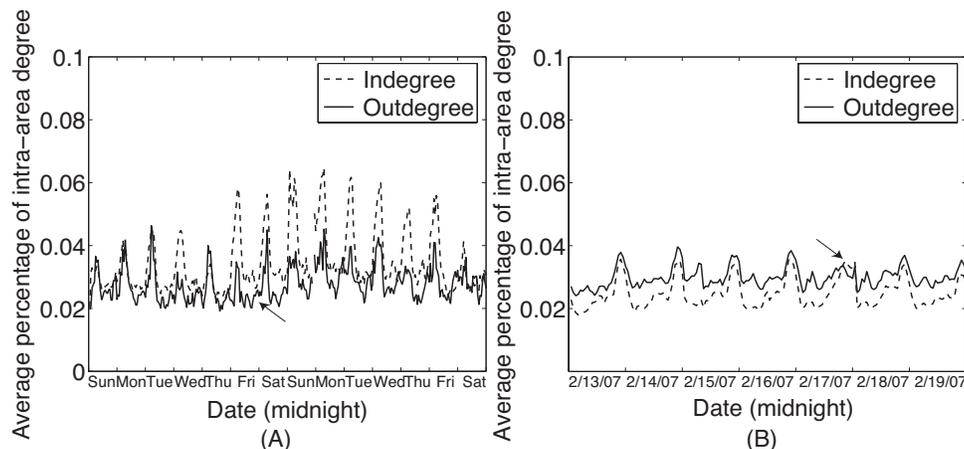


Fig. 9. Evolution of average intra-area degrees for stable peers in the network. (A) From October 1, 2006 to October 14, 2006. (B) From February 13, 2007 to February 19, 2007.

general applicability of our conclusions over a long period of time.

4.2.4 Intra-area degree evolution. Similarly, we further investigate the connectivity among peers in the same geographic location and across different areas. For IP addresses in China, they are in the same geographic area if they are located in the same province; for IP addresses out of China, they are grouped based on the continent they belong to. We compute the percentage of indegrees and outdegrees at each stable peer that are from and to peers in the same area at each time, and the evolution of averaged intra-area degree percentages (over all stable peers in the network at each time) is shown in Fig. 9.

From the results from both trace periods, we notice that the intra-area degree percentage is very low for both indegree and outdegree (less than 0.062 at all times), implying no area-based clustering in the streaming topology. As link TCP throughput is one major criterion for peer selection in UUSee, for connections inside China, this may also reveal that there does not exist a significant throughput difference between connections in the same province and across different provinces. For peers outside China, they may not have been able to efficiently select peers in nearby regions, which may require further improvements of the UUSee peer selection protocol.

4.2.5 Peer sending throughput vs. outdegree. As the throughput along each P2P connection varies, a peer with a large degree may not necessarily indicate that it is a supernode, *i.e.*, a peer with high throughput. To further explore the relation between peer throughput and degree, we make use of the throughput data along each P2P link as contained in the traces, and compute weighted peer degree distributions, with link throughput as the link weight. The weighted indegree of a peer at each time is indeed its instantaneous receiving throughput, which is consistently around the streaming rate in such a P2P streaming application. The weighted outdegree of a peer is its instantaneous sending throughput, which is a better indication

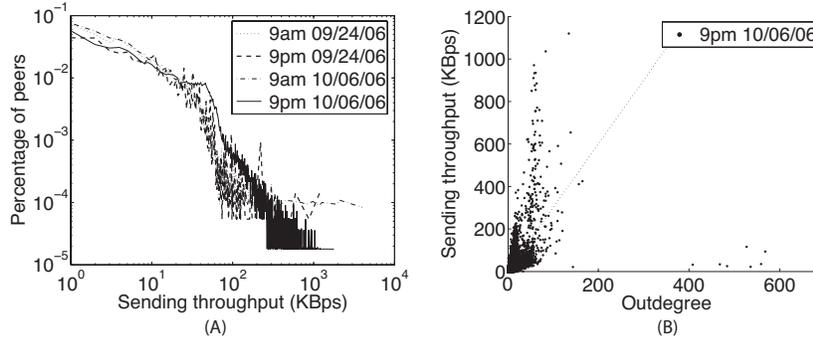


Fig. 10. (A) Sending throughput distribution of stable peers in the global topology; (B) Correlation between outdegree and sending throughput.

of the peer’s resource availability as a supernode. Therefore, we first investigate the sending throughput distribution across the UUSee overlay, and then examine its correlation with the peer outdegree. The results are given in Fig. 10.

Comparing Fig. 10(A) and Fig. 6(C), we notice that the sending throughput distributions represent similar shapes to corresponding outdegree distributions. Similar to outdegree, sending throughputs in the streaming overlay tend to be larger at peak hours and in the flash crowd scenario. In addition, although occupying a small fraction, there do always exist peers with large upload capacities in the overlay. For example, in the overlay of 9 p.m., October, 6, 2006, there were about 730 peers with higher than 384 KBps aggregate sending throughput, out of 87340 stable peers in the network.³

While the similarities between Fig. 10(A) and Fig. 6(C) may imply that a larger outdegree indeed indicates a larger sending throughput in UUSee, we further validate this point by plotting the correlation between the outdegree and sending throughput in Fig. 10(B). The plot reveals a positive linear correlation, which is also proven by the calculated Pearson product-moment correlation coefficient between outdegree and sending throughput, at the value of 0.4871.

Finally, we note that such a correlation does not exist between the peer indegree and aggregate receiving throughput, as the receiving throughput is consistently around the streaming rate, while the peer indegree varies significantly, as shown in Fig. 6(B).

4.3 Clustering

Studies on the Gnutella network have pointed out that both previous and current generation Gnutella networks exhibit “*small-world*” properties, *i.e.*, peers are highly clustered with small pairwise shortest path lengths, as compared to a random graph of similar peer numbers and link densities. To investigate whether an undirected graph g is a small-world graph, a clustering coefficient is calculated as

³Note that we do not include the set of dedicated streaming servers in our topological studies throughout the paper, and thus the possibility that any of the discovered supernodes might be a dedicated streaming server is excluded.

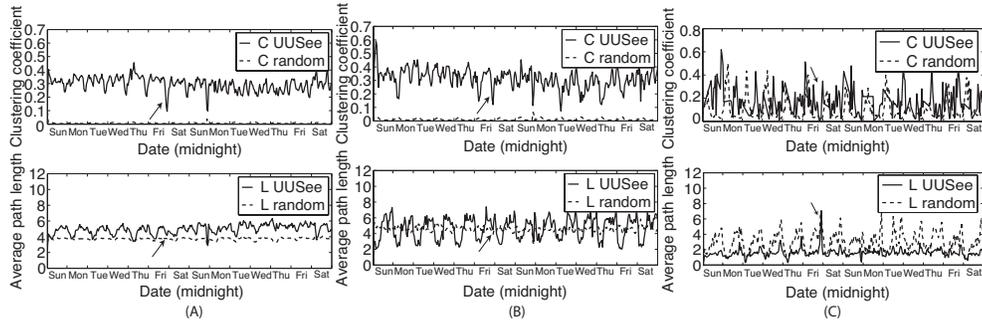


Fig. 11. Small-world property from October 1, 2006 to October 14, 2006. (A) Small-world metrics for the entire stable-peer graph; (B) Small-world metrics for an ISP subgraph (China Netcom); (C) Small-world metrics for an area subgraph (Zhejiang Province).

$C_g = \frac{1}{N} \sum_{i=1}^N C_i$, where N is the total number of vertices in the graph, and C_i is the clustering coefficient for vertex i , calculated as the proportion of edges between vertices within its neighborhood to the number of edges that could possibly exist between them [Watts 2003]. Therefore, a larger clustering coefficient represents more clustering at nodes in the graph. A graph g is identified as a small world if (1) it has a small average pairwise shortest path length l_g , close to that of a corresponding random graph l_r ; and (2) a large clustering coefficient C_g , which is orders of magnitude larger than that of the corresponding random graph C_r .

Based on the traces, we construct a subgraph of the entire UUSee topology at each time, by only including the stable peers and the active links among them. We investigate small-world properties of such stable-peer graphs, and believe they may reveal the connectivity of the original topologies as well.

Fig. 11(A) plots the clustering coefficients and average pairwise shortest path lengths of the stable-peer graph over the two-week period. We observe that its clustering coefficients are consistently more than an order of magnitude larger than those of a corresponding random graph, while their average path lengths are similar. This implies that the stable-peer graph does exhibit small-world properties. Besides, we observe slight decreases of clustering coefficients and slight increases of path lengths at peak hours of each day, which may be explained by the broader choice of partners at each peer in larger networks with significantly more peers at the peak times. In Vu *et al.*'s PPLive study, using the same clustering coefficient metric, they show that a less popular channel with 500 nodes is similar to a random graph, while the larger the channel popularity is, the more clustering it becomes. Our study focuses on the entire stable-peer topology, which is composed of tens of thousands of peers at each time, and thus represents the more clustering case discussed in the PPLive study.

Another observation we can make from Fig. 11(A) is that, the average pairwise shortest path length is quite steady, consistently around 5 at all times. This implies low network diameters in such stable-peer topologies. Considering that transient peers are connected to one or more stable peers with high probability, we conjecture that the pairwise shortest path lengths in the original UUSee topologies should be close to those in the stable-peer graphs. Therefore, the overall UUSee streaming

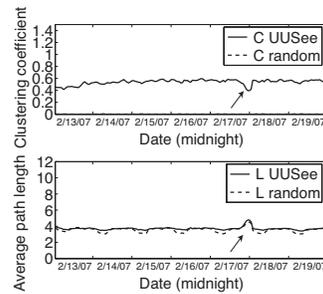


Fig. 12. Small-world metrics for the entire stable-peer graph from February 13, 2007 to February 19, 2007.

network may represent a low network diameter, which facilitates the quick distribution of media blocks throughout the entire topology.

In Sec. 4.2.3, we have observed the phenomenon of ISP-based peer clustering. Here, we wish to further validate this finding by calculating the clustering coefficient and average pairwise path length for the subgraph composed of stable peers in the same ISP and active links among them. A representative result is shown in Fig. 11(B) with respect to a major China ISP — China Netcom. Comparing Fig. 11(B) with Fig. 11(A), we have observed that the ISP subgraph has more clustering than the complete topology of stable peers, with (1) closer average path lengths to those of the random graphs, and (2) larger clustering coefficient difference from those of the random graphs. In our study, similar properties were observed for sub-topologies of other ISPs as well.

With the example of the sub streaming topology inside Zhejiang Province in China, we again investigate the clustering coefficient and average pairwise path length over the area sub-topology. Fig. 11(C) clearly demonstrates that there is no significant difference between the magnitudes of clustering coefficients for the area sub-topology and those of the corresponding random networks. Together with similar results from small-world metric evaluations of other geographic areas, it validates that there does not exist area-based clustering in the UUSee streaming network.

To examine whether such small-world properties may be different over a longer period of time, we have further investigated the clustering coefficient and average shortest path length in the entire stable-peer graph using the traces in February, 2007. Comparing the results in Fig. 12 to those in Fig. 11(A), we are able to identify a higher level of clustering in the UUSee streaming network at the later trace period. Again, decrease of the clustering coefficient and increase of the path length are observed during the flash crowd scenario on February 17, 2007. We omit the results with respect to the existence of ISP-based clustering and non-existence of area-based clustering observed in this period, which are nevertheless similar to those given in Fig. 11(B) and (C).

4.4 Reciprocity

In a modern P2P streaming application such as UUSee, a mesh streaming topology is constructed and a BitTorrent-like block distribution protocol is employed over

the mesh. However, as all media blocks originate from a collection of dedicated streaming servers and then propagate throughout the network, one may wonder: Is the media content propagating in a tree-like fashion, *i.e.*, a peer retrieves from a set of peers closer to the servers and further serves another group of peers farther away from servers? Or does such mesh-based streaming really benefit from reciprocal media block exchanges between pairs of peers? If it is the latter case, to what extent are the peers reciprocal to each other?

To answer these questions, we investigate another graph property on the P2P streaming topology: *edge reciprocity*. In a directed graph g , an edge (i, j) is reciprocal if vertex j is also linked to vertex i in the reverse direction, *i.e.*, (j, i) is also an edge in the graph. A simple way to obtain reciprocity of a graph is to compute the fraction of bilateral edges over the total number of edges in the graph:

$$r = \frac{\sum_{i \neq j} a_{ij} a_{ji}}{M}, \quad (1)$$

where a_{ij} 's are entries of the adjacency matrix of graph g ($a_{ij} = 1$ if an edge exists from i to j , and $a_{ij} = 0$ if not), and M is the total number of edges in the graph. However, this simple reciprocity metric cannot distinguish between networks with high reciprocity and random networks with high link density, which tend to have a large number of reciprocal edges as well, due exclusively to random factors. Therefore, we utilize another more accurate edge reciprocity metric proposed by Garlaschelli *et al.* [Garlaschelli and Loffredo 2004]:

$$\rho = \frac{r - \bar{a}}{1 - \bar{a}}, \quad (2)$$

where r is as defined in (1), and \bar{a} is the ratio of existing to possible directed links in the graph, *i.e.*, $\bar{a} = \frac{M}{N(N-1)} = \frac{\sum_{i \neq j} a_{ij}}{N(N-1)}$ with N being the total number of vertices. Since in a random network, the probability of finding a reciprocal link between two connected nodes is equal to the average probability of finding a link between any two nodes, \bar{a} actually represents the reciprocity calculated with (1), of a random graph with the same number of vertices and edges as g . Therefore, the edge reciprocity defined in (2) is an absolute quantity, in the sense that: if $\rho > 0$, the graph has larger reciprocity than a corresponding random graph, *i.e.*, it is a reciprocal graph; if $\rho < 0$, the network has smaller reciprocity than its random version, *i.e.*, it is an antireciprocal graph.

To compute the reciprocity among all the peers in the UUSee network at one time, we use all the directed active links among peers that appeared in the trace at the time. If streaming in the UUSee network takes place in a tree-like fashion, the computed edge reciprocity should be negative, as its $r = 0$ and $\rho = -\frac{\bar{a}}{1-\bar{a}} < 0$. If there is no strong correlation between the sets of supplying and receiving partners at each peer, the edge reciprocity will be around 0, *i.e.*, the case of a random network. If the peers do help each other materially by exchanging media blocks, the edge reciprocity should be greater than 0, and can be as large as 1.

Fig. 13(A) plots the evolution of edge reciprocity in the entire UUSee topology. The consistent greater-than-zero edge reciprocity reveals significant reciprocal exchanges of available blocks among pairs of peers in such mesh-based streaming. It also implies that the sets of supplying and receiving partners at each peer are

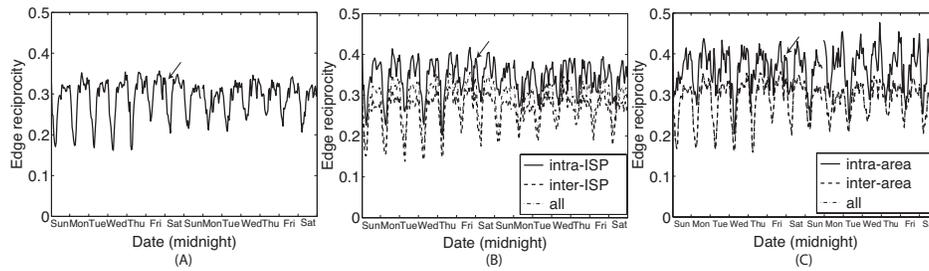


Fig. 13. Edge reciprocity from October 1, 2006 to October 14, 2006. (A) Edge reciprocity for the entire network; (B) Reciprocity for edges in the same ISP and across different ISPs; (C) Reciprocity for edges in the same area and across different areas.

strongly correlated, as compared to a purely random network. Furthermore, the reciprocity exhibits daily evolution patterns with peaks at the peak hours as well.

We have discovered ISP-based clustering of the peers in the previous sections, where the direction of P2P links is not essentially utilized. Now taking connection directions into consideration, we further investigate the reciprocity of links connecting peers in the same ISP and those among peers across different ISPs. For this purpose, we derive two sub-topologies from each topology we used in the previous reciprocity investigation: one contains links among peers in the same ISPs and their incident peers, and the other consists of links across different ISPs and the incident peers. Fig. 13(B) shows edge reciprocities for the two sub-topologies. For the purpose of comparison, it also plots the edge reciprocities for the entire topology. We have observed a higher reciprocity for the intra-ISP sub-topology and a lower reciprocity for the inter-ISP sub-topology, as compared to that of the complete topology. This implies that the streaming topology in each ISP is a densely connected cluster with large portions of bilateral links among the peers.

Similarly, we investigate the edge reciprocity for links connecting peers in the same area and those among peers across different areas, using the intra-area sub-topology and the inter-area sub-topology at each time. While the evolution curve for the inter-area sub-topology mostly overlaps with that of the entire topology, Fig. 13(C) reveals high reciprocities in the intra-area sub-topology. This is an interesting discovery, since even though we have observed no clustering for the streaming topology in one area, there does exist a high level of reciprocity over the limited number of intra-area links, *i.e.*, peers in the same area are largely reciprocal to each other.

In addition, the edge reciprocities derived using traces in February, 2007, are given in Fig. 14. Comparing Fig. 14 to Fig. 13, we observe generally smaller edge reciprocities in February, 2007, which we attribute to the expansion of UUSEE streaming network over the months, such that each peer has a broader choice of partners. Nevertheless, the other properties, such as the daily evolution pattern and the better reciprocity inside the same ISP or area, remain.

4.5 Supernode connectivity

In Sec. 4.2.5, we have observed the existence of supernodes in the P2P streaming overlay, *i.e.*, the peers with high sending throughput. Next, we start our investi-

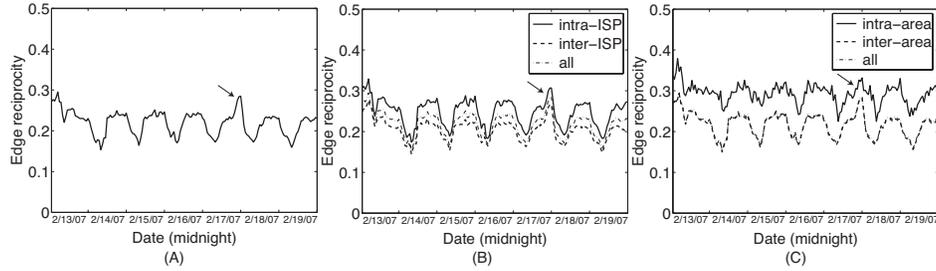


Fig. 14. Edge reciprocity from February 13, 2007 to February 19, 2007. (A) Edge reciprocity for the entire network; (B) Reciprocity for edges in the same ISP and across different ISPs; (C) Reciprocity for edges in the same area and across different areas.

gations on the connectivity among supernodes: Do they tend to connect to each other and constitute a hub in such a practical streaming network? Do they tend to exchange media blocks among each other in a reciprocal way?

To address these questions, we utilize the *likelihood* metric proposed by Li *et al.* in [Li et al. 2004], which is also linearly related to the *assortativity coefficient* discussed by Newman [Newman 2002]. These metrics suggest the connectivity of nodes with similar degrees in a graph, *i.e.*, whether they tend to be tightly interconnected or not. The likelihood metric for undirected graph g is defined as follows by Li *et al.* [Li et al. 2004]: Let $L(g)$ be the sum of products of degrees of adjacent nodes, *i.e.*, $L(g) = \sum_{(i,j) \in E(g)} d_i d_j$, where $E(g)$ is the set of edges in graph g and d_i is the degree of vertex i . Let L_{\max} and L_{\min} denote the maximum and minimum values of $L(g)$ among all simple connected graphs with the same number of vertices and the same node degree sequence as graph g . The likelihood is defined as:

$$\text{likelihood}(g) = \frac{L(g) - L_{\min}}{L_{\max} - L_{\min}}. \quad (3)$$

In order to compute L_{\max} and L_{\min} , we need to first generate graphs that have these likelihood values. A L_{\max} (L_{\min}) graph can be generated by the following simple heuristics: Sort nodes in graph g from the highest to the lowest degree. To generate the L_{\max} (L_{\min}) graph, connect the highest degree node successively to other high (low) degree nodes in decreasing (increasing) order of their degrees until it satisfies its degree requirement, and then connect the second highest degree node successively to other nodes in decreasing (increasing) degree order (which have not saturated their degree requirements) to achieve its degree, etc. This process is repeated for all nodes in descending degree order. In this way, the likelihood computed with (3) is a normalized quantity in the range of $[0, 1]$. A close-to-0 likelihood indicates that high degree nodes tend to connect to low degree nodes in the graph, while a close-to-1 likelihood reveals more clustering of the nodes with similar degrees.

To derive the connectivity among supernodes in the UUSee network, instead of using peer degrees in the computation of likelihood, we use the sending throughput of each peer, as we believe it is a better indication of the resource availability of the peers as supernodes. Besides, we have shown in Sec. 4.2.5 that peer sending throughput is positively correlated with outdegree in the UUSee streaming overlay.

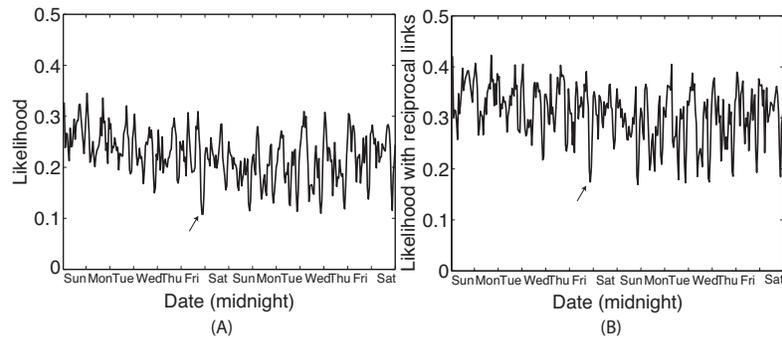


Fig. 15. Likelihood from October 1, 2006 to October 14, 2006. (A) Likelihood computed with all the links in the UUSee network; (B) Likelihood computed with only reciprocal links in the UUSee network.

We first compute the likelihood using all the active links included in the instantaneous UUSee streaming topologies. Fig. 15(A) shows that the likelihood values are below 0.3 at most times. These represent quite low likelihood in its value range of $[0, 1]$, *i.e.*, below the average likelihood among graphs with the same number of nodes and node sending throughput sequence. This observation indicates that supernodes do not tend to be tightly connected to each other in the UUSee streaming overlay, but are surrounded largely by nodes with lower sending abilities.

Motivated by the reciprocal properties that we have discovered in the previous section, we further explore the following question: Are supernodes more inclined to exchange media content with other peers with comparably high sending throughputs? To answer this question, we compute the likelihood of UUSee streaming network again with only reciprocal links in each instantaneous topology. Comparing Fig. 15(B) with Fig. 15(A), we find that the likelihood with respect to reciprocal links is generally larger than that in the entire topology, implying that the reciprocal links are relatively more likely to be connecting nodes with comparable sending throughputs. Nevertheless, the likelihood values are still lower than average among the set of graphs with the same number of nodes and the same sequence of node sending throughput.

Investigations of the likelihood property using traces in February, 2007 further validate our above discoveries. As illustrated in Fig. 16, the likelihood is at an even lower level in February, 2007. All these observations exhibit that, in such a practical streaming network, peers are not clustered based on their resource availability, and there does not exist a supernode hub at any given time in the overlay.

5. CONCLUDING REMARKS

This paper presents *Magellan*, the first effort in the research community to characterize topologies of modern large-scale peer-to-peer streaming networks, with abundant traces from a successful commercial P2P streaming application, UUSee. By collaborating with the streaming solution provider, we obtain instantaneous snapshots of the active topology of the P2P streaming network over a long period of time. Utilizing a number of meaningful graph metrics, we seek to discover structural

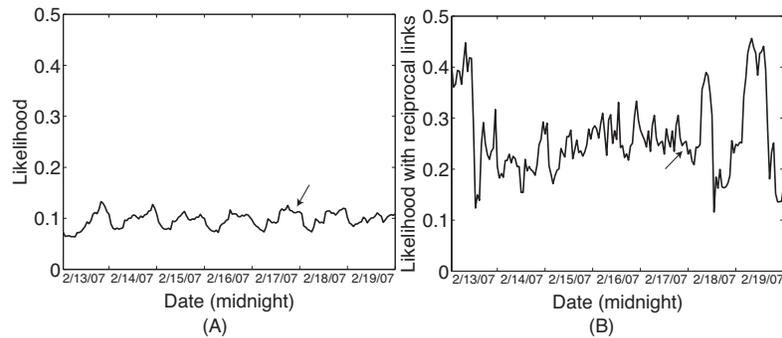


Fig. 16. Likelihood from February 13, 2007 to February 19, 2007. (A) Likelihood computed with all the links in the UUsee network; (B) Likelihood computed with only reciprocal links in the UUsee network.

properties of the streaming topologies at short time scales, as well as their evolutionary dynamics over longer periods of time. We have found that, even with a simple peer selection protocol, modern P2P streaming applications are able to sustain an acceptable streaming performance, even in the case of large flash crowds. We also discover that the topologies of modern P2P streaming overlays do not possess similar properties as those obtained from early Internet or AS-level topological studies, such as power-law degree distributions. Nevertheless, an interesting discovery is that the streaming topologies naturally evolve into clusters inside each ISP, but not within geographically adjacent areas. In addition, peers are reciprocal to each other to a great extent, which contributes to the stable performance of streaming in such mesh networks. There exist a small portion of high-throughput supernodes in the streaming overlay, each assisting a large number of peers with lower bandwidth availabilities, but not tightly connecting to each other in a hub-like fashion. We believe that our findings bring important insights towards a complete understanding of large-scale practical P2P streaming applications, and will be instrumental towards further improvements of modern P2P streaming protocol design. Such improvements have indeed constituted intriguing research topics in the community.

REFERENCES

- ADAMIC, L. A., LUKOSE, R. M., PUNIYANI, A. R., AND HUBERMAN, B. A. 2001. Search in Power-Law Networks. *Physical Review E* 64(46135).
- ALI, A., MATHUR, A., AND ZHANG, H. 2006. Measurement of Commercial Peer-To-Peer Live Video Streaming. In *Proc. of Workshop in Recent Advances in Peer-to-Peer Streaming*.
- BASET, S. A. AND SCHULZRINNE, H. 2006. An Analysis of the Skype Peer-to-Peer Internet Telephony Protocol. In *Proc. of IEEE INFOCOM*.
- CHEN, K.-T., HUANG, C.-Y., HUANG, P., AND LEI, C.-L. 2006. Quantifying Skype User Satisfaction. In *Proc. of ACM SIGCOMM Conference 2006*.
- CHENG, B., LIU, X., ZHANG, Z., AND JIN, H. 2007. A Measurement Study of a Peer-to-Peer Video-on-Demand System. In *Proc. of the 6th International Workshop on Peer-to-Peer Systems (IPTPS '07)*.
- GARLASCHELLI, D. AND LOFFREDO, M. I. 2004. Patterns of Link Reciprocity in Directed Networks. *Physical Review Letters* 93(26):268701.
- GUHA, S., DASWANI, N., AND JAIN, R. 2006. An Experimental Study of the Skype Peer-to-Peer
- ACM Transactions on Multimedia Computing, Communications and Applications, Vol. V, No. N, October 2007.

- VoIP System. In *Proc. of the 5th International Workshop on Peer-to-Peer Systems (IPTPS '06)*.
- GUMMADI, K. P., DUNN, R. J., SAROIU, S., GRIBBLE, S. D., LEVY, H. M., AND ZAHORJAN, J. 2003. Measurement, Modeling and Analysis of a Peer-to-Peer File-Sharing Workload. In *Proc. of the 19th ACM Symposium of Operating Systems Principles (SOSP)*.
- GUO, L., CHEN, S., XIAO, Z., TAN, E., DING, X., AND ZHANG, X. 2005. Measurements, Analysis, and Modeling of BitTorrent-like Systems. In *Proc. of the Internet Measurement Conference (IMC)*.
- HEI, X., LIANG, C., LIANG, J., LIU, Y., AND ROSS, K. W. 2006. Insight into PPLive: Measurement Study of a Large Scale P2P IPTV System. In *Workshop on Internet Protocol TV (IPTV) Services over World Wide Web, in conjunction with WWW 2006*.
- HEI, X., LIANG, C., LIANG, J., LIU, Y., AND ROSS, K. W. 2007. A Measurement Study of a Large-Scale P2P IPTV System. *IEEE Trans. on Multimedia (to appear)*.
- IZAL, M., URVOY-KELLER, G., BIERSACK, E., FELBER, P., HAMRA, A. A., AND GARCES-ERICE, L. 2004. Dissecting BitTorrent: Five Months in a Torrent's Lifetime. In *Proc. of the 5th Passive and Active Measurement Workshop (PAM '04)*.
- JOVANOVIĆ, M., ANNEXSTEIN, F., AND BERMAN, K. 2001. Modeling Peer-to-Peer Network Topologies through Small-World Models and Power Laws. In *TELFOR*.
- LI, L., ALDERSON, D., WILLINGER, W., AND J. DOYLE. 2004. A First-Principles Approach to Understanding the Internet's Router-level Topology. In *Proc. of ACM SIGCOMM Conference 2004*.
- LIANG, J., KUMAR, R., AND ROSS, K. 2006. The FastTrack Overlay: A Measurement Study. *Computer Networks* 50, 6 (April), 842–858.
- NEWMAN, M. E. J. 2002. Assortative Mixing in Networks. *Physical Review Letters* 89(208701).
- POUWELSE, J., GARBACKI, P., EPEMA, D., AND SIPS, H. 2005. The BitTorrent P2P File Sharing System: Measurements and Analysis. In *Proc. of the 4th International Workshop on Peer-to-Peer Systems (IPTPS '05)*.
- RIPEANU, M., FOSTER, I., AND IAMNITCHI, A. 2002. Mapping the Gnutella Network: Properties of Large-Scale Peer-to-Peer Systems and Implications for System Design. *IEEE Internet Computing Journal* 6, 1.
- SILVERSTON, T. AND FOURMAUX, O. 2006. P2P IPTV Measurement: A Case Study of TVants. In *Proc. of the 2nd Conference on Future Networking Technologies (CoNEXT '06)*.
- SILVERSTON, T. AND FOURMAUX, O. 2007. Measuring P2P IPTV Systems. In *Proc. of the 17th International workshop on Network and Operating Systems Support for Digital Audio & Video (NOSSDAV'07) (to appear)*.
- STEINER, M., BIERSACK, E. W., AND ENNAJJARY, T. 2007. Actively Monitoring Peers in KAD. In *Proc. of the 6th International Workshop on Peer-to-Peer Systems (IPTPS '07)*.
- STUTZBACH, D. AND REJAIE, R. 2006. Understanding Churn in Peer-to-Peer Networks. In *Proc. of the Internet Measurement Conference (IMC)*.
- STUTZBACH, D., REJAIE, R., AND SEN, S. 2005. Characterizing Unstructured Overlay Topologies in Modern P2P File-Sharing Systems. In *Proc. of Internet Measurement Conference (IMC)*. UUSEE INC. <http://www.uusee.com/>.
- VU, L., GUPTA, I., LIANG, J., AND NAHRSTEDT, K. 2007. Measurement and Modeling a Large-scale Overlay for Multimedia Streaming. In *Proc. of the Fourth International Conference on Heterogeneous Networking for Quality, Reliability, Security and Robustness (QShine 2007)*.
- WATTS, D. J. 2003. *Six Degrees: the Science of a Connected Age*. ACM Press.
- WU, C., LI, B., AND ZHAO, S. 2007. Magellan: Charting Large-Scale Peer-to-Peer Live Streaming Topologies. In *Proc. of the 27th International Conference on Distributed Computing Systems (ICDCS 2007)*.
- ZHANG, X., LIU, J., LI, B., AND YUM, T. P. 2005. CoolStreaming/DONet: A Data-Driven Overlay Network for Live Media Streaming. In *Proc. of IEEE INFOCOM 2005*.