Magellan: Charting 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

Shuqiao Zhao
Multimedia Development Group
UUSee, Inc.
zhaosq@uusee.com

Abstract

Live peer-to-peer (P2P) streaming applications have been successfully deployed in the Internet. With relatively simple peer selection protocol design, modern live P2P streaming applications are able to provide millions of concurrent users adequately satisfying viewing experiences. That said, few existing research has provided sufficient insights on the time-varying internal characteristics of P2P topologies in live streaming. With 120 *GB worth of traces in late 2006 from a commercial P2P live streaming system of UUSee Inc. in Beijing, this paper represents the first attempt in the research community to explore topological properties in practical P2P streaming, and how they behave over time. Starting from classical graph metrics, such as degree, clustering coefficient, and reciprocity, we explore and extend them in specific perspectives of streaming applications. We also compare our findings with existing insights from topological studies of P2P file sharing applications, which shed new and unique insights specific to streaming. Our characterization reveals the scalability of the commercial P2P streaming application even in case of large flash crowds, the clustering phenomenon of peers in each ISP, as well as the reciprocal behavior among peers, all of which play important roles in achieving its current success.*

1 Introduction

Based on the peer-to-peer (P2P) communication paradigm, live P2P multimedia streaming applications have been successfully deployed in the Internet with up to millions of users at any given time. Prominent examples that are better known to the research community include CoolStreaming [19], PPLive [10] and TVAnts [16]. 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 for most of the recent successful P2P live streaming applications, blocks of live media contents are being delivered over a mesh overlay topology, featuring reciprocal exchanges of useful content blocks among multiple peers. It is also interesting to observe that most current-generation P2P streaming applications employ relatively simple peer selection and mesh construction protocol designs. They typically use central tracking servers to gain initial knowledge of existing peers in the channels, and periodically exchange peer lists among peers themselves. As mesh-based streaming topologies play an important role towards the commercial success of P2P 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 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 topology has been characterized extensively at the IP layer, there has been relatively little literature on the discovery of application-level P2P topologies, and most of them are on previous-generation P2P file sharing applications [12, 15, 17]. Such P2P applications bear fundamental differences when compared to modern loosely-coupled P2P applications based on block exchanges and long-lived meshes, leading to possibly different topological properties. Furthermore, topological characteristics in block-exchanging P2P streaming applications may well be different from their file-sharing counterparts, due to its biased peer selection protocol design towards the timely delivery of media content.

With the objective of gaining in-depth insights and a complete understanding of the P2P streaming characteristics, we have launched the *Magellan* project, with collaborative efforts with UUSee Inc. [1], one of the main P2P live streaming solution providers in mainland China. This paper represents the first milestone of Magellan, presenting our comprehensive insights from exploring graph theoretical properties in actually formed live streaming topologies, based on over 120 GB of traces and 10 million unique IP addresses that we have collected over a two-month period.

---

*The completion of this research was made possible thanks to Bell Canada's support through its Bell University Laboratories R&D program.

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, and reciprocity — 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 peer-to-peer 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.

In what follows, we briefly review existing work in P2P measurements and topology characterization in Sec. 2, and outline the basics of UUSee peer selection protocol and trace collection methods in Sec. 3. In Sec. 4, we analyze the topological properties from a number of important perspectives, and discuss the implications of our discoveries. We summarize our findings and conclude the paper in Sec. 5.

2 Related Work

There have been a number of measurement studies on various P2P applications in recent years. For KaZaA overlay, Gummadi *et al.* [8] characterized its workload, and Liang *et al.* [13] studied its two-tier overlay structure and dynamics. For the previous-generation Gnutella network, earlier work [2, 12, 15] reported the discovery of power-law degree distributions and strong "small-world" properties, *i.e.*, small network diameter and peer clustering.

As one of the most pronounced work on unstructured overlay topology characterization, Stutzbach *et al.* [17] reported a detailed investigation of the topologies of modern Gnutella networks. Utilizing their fast Gnutella crawler, *Cruiser*, they captured back-to-back snapshots of Gnutella network, which reflected more accurate graph properties and dynamics of the topology. Comparing to earlier reports on Gnutella topologies, they believed that the discovery of power-law peer degree distribution is 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 a spike in between the two segments. They also reported that modern Gnutella exhibits the "small-world" phenomenon, but comparably less clustering than the previous-generation Gnutella network.

The focus of *Magellan* has been on modern P2P live streaming topologies, which is significantly different from existing work on P2P topology characterization. First, modern P2P live streaming applications are based on a BitTorrent-like block exchange mechanism, over mesh overlay topologies featuring dynamic reconfiguration. No previous work exists for topology characterization of such modern BitTorrent-like P2P applications. Second, mesh-based P2P streaming has fundamentally different requirements as compared to BitTorrent-like file-sharing applications, with respect to peer selection and topology construction.

Topology characterization aside, towards measurements of various modern P2P applications, Pouwelse [14], Izal [11] and Guo [9] *et al.* have investigated the performance of BitTorrent. Various aspects of *Skype*, a P2P VoIP application, have been explored as well [4, 7].

Targeting commercial P2P live streaming applications, Hei *et al.* [10] carried out detailed measurement studies of *PPLive*, from both global views of peer distributions and local views of media traffic characteristics. Their global property investigation includes the evolution of peer numbers and peer geographic distributions, but does not include other topological properties. In addition, Silverston *et al.* [16] studied the upload/download traffic generated by another P2P IPTV system, TVAnts, while users were watching the last FIFA World Cup. Ali *et al.* [3] presented an analysis of the control traffic, resource usage, locality, and stability of two P2P live streaming applications: PPLive and SOPCast. In addition, Cheng *et al.* [5] studied end user experience in a P2P video-on-demand system, GridCast.

To the best of our knowledge, this paper represents the first attempt at topology characterization in large-scale P2P streaming applications. Nevertheless, when applicable, we also make attempts to compare our new insights with those discovered with traditional P2P file sharing applications.

3 Background

3.1 UUSee P2P Streaming

Backed by venture capital funding from recognized investors, UUSee Inc. [1] 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. With a large collection of streaming servers around the world, it simultaneously broadcasts over 800 channels to millions of peers, mostly encoded to high quality streams around 400 Kbps.

Similar to all current-generation P2P streaming protocols, UUSee's streaming protocol design is based on the principle of allowing peers to serve each other by exchanging blocks of data in a *sliding window* of the media channel. After a new peer joins a channel in UUSee, the initial set of a number of *partners* (up to 50) is supplied by one of its tracking servers. The peer establishes TCP connections with these partners, and buffer maps are exchanged periodically. During this process, it measures the round-trip delay and TCP throughput of the connection, and then selects a number of most suitable peers (around 30) from which it actually requests media blocks.

In addition, the UUSee peer selection protocol incorporates a number of strategies to maximally utilize peer upload bandwidth. Each peer estimates its maximum upload capacity, and continuously monitors its aggregate instantaneous sending throughput to its partners during streaming. If its aggregate sending throughput is lower than its upload capacity for a sustained period of time, it will inform one of the tracking servers that it is able to receive new connections. Each tracking server keeps a list of such peers, and bootstraps new peers with peers randomly selected from this set. During streaming, neighboring peers also recommend known partners to each other, based on estimated availability for them to assist each other. As a last resort, a peer will contact a tracking server again to obtain additional partners, if its playback rate is not sustained for a certain period of time.

3.2 Collection of Traces

To discover and chart live P2P streaming topologies, we have implemented detailed measurements within the UUSee streaming protocol on each peer.

Each peer estimates its total download and upload bandwidth capacities, and continuously monitors its instantaneous aggregate receiving and sending throughput, from and to all its partners. For each active partner with which it has a live TCP connection, it actively measures the number of sent/received segments over the TCP connection.

A standalone trace server is responsible for the collection of measurement reports from existing peers. Each report includes basic information such as the peer's IP address, the channel it is watching, its buffer map, total download and upload capacities, as well as its instantaneous aggregate receiving and sending throughput. In addition, the report also includes a list of all its partners, with their corresponding IP addresses, TCP/UDP ports, and number of segments sent to or received from each partner.

Using UDP datagrams, a new peer sends its initial report to the trace server 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 known partners, 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. In a two-month span, we have collected over 120 GB of traces with more than 10 million unique IP addresses, constituting continuous-time snapshots of P2P streaming topologies throughout this period.

4 Charting Large-scale P2P Streaming Topologies

Taking advantage of the traces, we carry out an in-depth investigation of global topological properties of the P2P streaming overlay, and their evolutionary dynamics over time. While we have studied the entire trace period, we are constrained by space to show results obtained over such a long time. Therefore, we choose two representative weeks, from 12:00 a.m. October 1st, 2006 (GMT+8) to 11:50 p.m. October 14th, 2006 (GMT+8), and show results from these weeks in our figures. These selected periods, we believe, present all the typical scenarios that we shall discover.

4.1 Scale of UUSee Topologies

As a natural first step, we investigate the scale and general streaming performance of the UUSee application.

4.1.1 Overall number of simultaneous peers

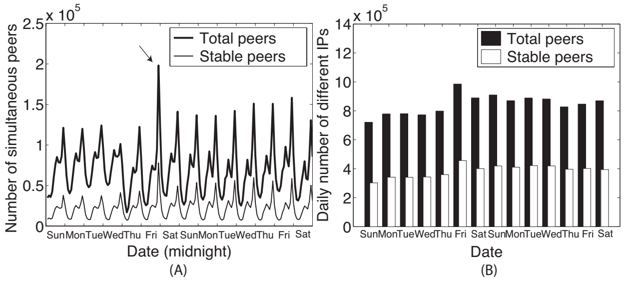

Figure 1. Daily peer number statistics.

To discover the number of concurrent online users in the UUSee streaming overlay and the percentage of stable peers (whose reports are received by the trace server) in the overall peer population, 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 peers that have reported and peers in their partner lists. The peer number statistics are shown in Fig. 1(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 similar daily peer number patterns as that presented in a previous PPLive study

[10]. Different from [10], we observe only a slight number increase over the weekend, considering the weekly variance trend. In addition, we clearly observe 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.[1] Comparing the number of stable peers to the total number of peers, we discover that the former is asymptotically $1/3$ of the later.

We further summarize the number of distinct IP addresses that appeared in the traces on a daily basis in Fig. 1(B). The statistics exhibit that UUSee serves up to 1 million different users each day.

4.1.2 Number of simultaneous peers in different ISPs

Throughout the paper, we also emphasize on the mapping of the abstract streaming topology to the real world scenario, with respect to the ISP 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. For each IP address in China, the database provides the China ISP it belongs to; for IP addresses out of China, it provides a general code indicating foreign ISPs.

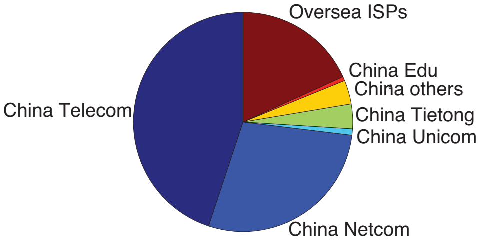

Figure 2. Peer number statistics for different ISPs.

Using this mapping database, we have determined the ISP membership of simultaneous peers at any time, and discovered that the ISP distributions of peers do not vary significantly over the two-month period. Therefore, we only show the averaged shares of peers in major ISPs over the trace period in Fig. 2. We observe that while users in China dominate the UUSee streaming network, there still exist a substantial number of peers from overseas. In our subsequent studies, when we explore characteristics of global topologies, we include all peers from all ISPs; when we investigate ISP-based topology properties, we will mainly focus on the ISPs inside China.

4.1.3 Streaming quality

To explore the streaming quality of UUSee, we make use of the aggregate instantaneous receiving throughput measured at each peer in different channels. We select two of the most popular channels broadcast by UUSee, CCTV1 and CCTV4 (both from the official Chinese television network). Fig. 3 shows the percentage of peers in both channels whose receiving throughput is higher than 90% of the channel streaming rate. Although we have found that the numbers of concurrent peers in CCTV1 and CCTV4 differ significantly[2], we see that around $3/4$ of all viewers in both channels can achieve satisfactory streaming rates, and the percentages are quite consistent over time. In addition, it is a bit surprising to find that, the percentages are generally a bit higher at the peak hours of a day. Especially, during the flash crowd scenario on October 6, 2006, the percentages of CCTV4 viewers with satisfactory streaming rates represent a sharp increase. We explain this phenomenon as that, as UUSee aims to maximally utilize peer upload capacity to assist each other, when more peers are online, they are able to effectively utilize the available upload capacities in the entire network. This shows that the P2P streaming protocol scales well to a large number of peers.

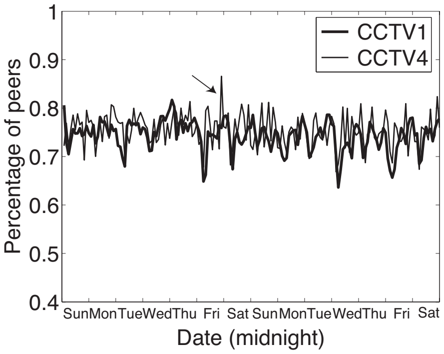

Figure 3. Percentage of peers with satisfactory streaming rates.

4.2 Degree Distribution

In our traces, each stable peer reports IP addresses in its list of partners, as well as the number of segments 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 partners, from which the number of received segments is larger than a certain threshold (10 segments); (2) active receiving partners, to which the number of sent segments is larger than the threshold; (3) nonactive partner, otherwise.

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

---

[1] In all the 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.

[2] CCTV1 concurrent viewers are 5 times more than those of CCTV4, with the number of the former around $30,000$ at any time and that for the latter around $6,000$.

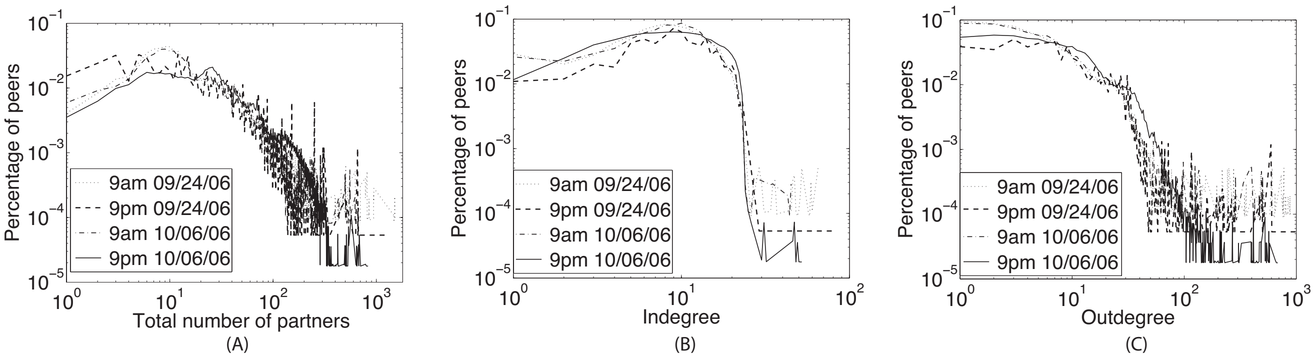

(A) (B) (C)

Figure 4. Degree distributions of stable peers in the global topology.

number of partners in the partner list including both active and nonactive partners. Note that in a mesh network, it is common for a partner to be both a supplying partner and a receiving parter 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 study for modern Gnutella topology, Stutzbach *et al.* [17] 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. 4(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, recall that each peer has an initial set of up to 50 partners upon initial joining. These observations exhibit that the number of partners decreases at most peers during streaming, due to partner departures or failures.

For the peer indegree distribution shown in Fig. 4(B), we observe spikes around 10, and the spike is at a slightly larger 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 bandwidth on each upload link is guaranteed. Besides, during streaming, the aggregated download rate is limited by the streaming rate in the UUSee application. All these explain the observation that the number of supplying partners at each peer, which guarantees a satisfactory streaming rate, is relatively small in the UUSee overlay, as compared to other file sharing applications.

The peer outdegree distributions in Fig. 4(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.

4.2.2 Degree evolution in the global topology

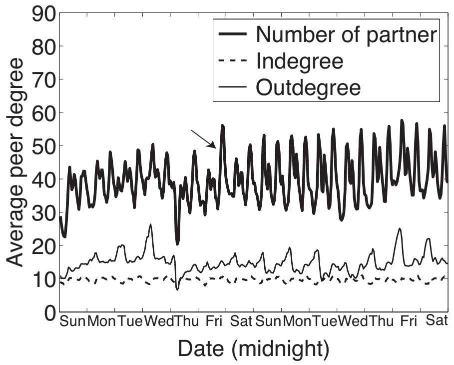

Figure 5. Evolution of average degrees for stable peers in the global topology.

We next show the evolution of average degrees of stable peers during the two-week period in Fig. 5. 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 the UUSee protocol does not explicitly impose such an upper bound for active incoming connections at each peer, we explain this discovery as follows: At peak times with a large number of peers in the network, there are abundant supplies of upload bandwidth, considering the streaming rate around 400 Kbps is lower than the upload capacity of most ADSL/cable modem peers, 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 become recognized 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.

4.2.3 Intra-ISP degree evolution

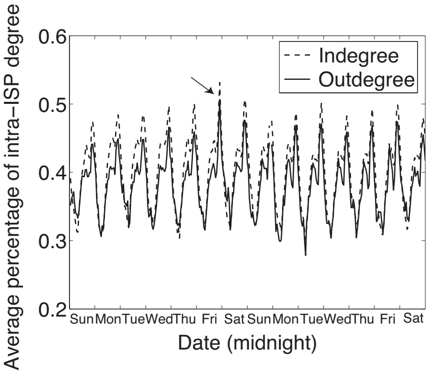

Figure 6. Evolution of average intra-ISP degrees for stable peers in the network.

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. 6 plots the evolution of the intra-ISP degree percentage, averaged over all stable peers in the network at each time. 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 supplying/receiving partners of each peer are within the same ISP. As the UUSee protocol does not take ISP membership into consideration when the tracking server assigns new partners to a peer and 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. 6 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 which are largely intra-ISP.

4.3 Clustering

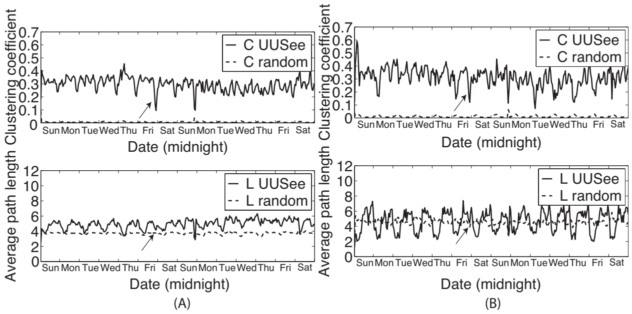

(A) (B)

Figure 7. (A) Small-world metrics for the entire stable-peer graph; (B) Small-world metrics for ISP — China Netcom — subgraph.

Studies on 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 random networks of similar peer numbers and link densities. To investigate whether a graph g is a small-world graph, a clustering coefficient is calculated as $C_g = \frac{1}{n}\sum_{i=1}^{n} C_i$, 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 [18]. A graph 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. 7(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 more relaxed choice of partners at each peer in larger networks at peak times.

In Sec. 4.2.3, we have observed ISP-based peer clustering. Here, we wish to further validate this finding by calculating the clustering coefficient and average 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. 7(B) with respect to a major China ISP — China Netcom. Comparing Fig. 7(B) with Fig. 7(A), we conclude 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. Similar properties were observed for sub topologies for other ISPs as well.

Another observation we can make from these figures 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 path lengths in the original UUSee topologies should be close to those in the stable-peer graphs. Therefore, the overall UUSee streaming network may represent a low network diameter, which facilitates the quick distribution of media segments throughout the entire topology.

4.4 Reciprocity

In a modern P2P streaming application such as UUSee, a mesh streaming topology is constructed and BitTorrent-like block distribution is employed over the mesh. However, as all media segments 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 segment 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 total number of edges in the graph:

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

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 links as well, due exclusively to random factors. Therefore, we utilize another more accurate edge reciprocity metric proposed by Garlaschelli *et al.* [6]:

$$\rho = \frac{r - \bar{a}}{1 - \bar{a}}, \tag{2}$$

where r is 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 probability of finding a link between any two nodes, $\bar{a}$ represents the reciprocity computed 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 topology 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 peers do help each other materially by exchanging media segments, the edge reciprocity should be larger than 0.

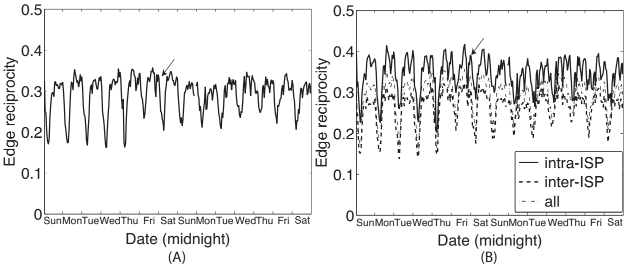

Figure 8. (A) Edge reciprocity for the entire topology; (B) Reciprocity for edges in the same ISP and across different ISPs.

Fig. 8(A) plots evolution of edge reciprocity in the entire UUSee topology. The consistent greater-than-zero edge reciprocity reveals significant reciprocal exchanges of available segments among pairs of peers in such mesh-based streaming. It also implies that the sets of supplying and receiving partners at each peer are strongly correlated, as compared to a purely random network. Furthermore, the reciprocity exhibits daily variance patterns with peaks at the peak hours as well.

We have discovered ISP-based clustering of the peers in previous sections, where the direction of P2P links is not considered. We now 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. 8(B) shows edge reciprocities of the two sub topologies. For the purpose of comparison, it also plots the edge reciprocities of the entire topology. We observe 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.

5 Conclusion

This paper presents *Magellan*, our first effort to characterize topologies of modern large-scale peer-to-peer streaming networks, with abundant traces from a successful commercial P2P streaming application, UUSee. Utilizing a number of meaningful graph metrics, we discover the structural properties of the streaming topologies at short time scales, as well as their dynamics over 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. In addition, the peers are reciprocal to each other to a great extent, which contributes to the stable performance of streaming in such mesh networks. 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 P2P streaming protocol design. Such improvements have indeed been part of our ongoing work and future plans.

References

[1] *UUSee Inc.*, http://www.uusee.com/.

[2] L. A. Adamic, R. M. Lukose, A. R. Puniyani, and B. A. Huberman. Search in Power-Law Networks. *Physical Review E 64(46135)*, 2001.

[3] A. Ali, A. Mathur, and H. Zhang. Measurement of Commercial Peer-To-Peer Live Video Streaming. In *Proc. of Workshop in Recent Advances in Peer-to-Peer Streaming*, August 2006.

[4] K.-T. Chen, C.-Y. Huang, P. Huang, and C.-L. Lei. Quantifying Skype User Satisfaction. In *Proc. of ACM SIGCOMM Conference 2006*, September 2006.

[5] B. Cheng, X. Liu, Z. Zhang, and H. Jin. 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)*, February 2007.

[6] D. Garlaschelli and M. I. Loffredo. Patterns of Link Reciprocity in Directed Networks. *Physical Review Letters 93(26):268701*, 2004.

[7] S. Guha, N. Daswani, and R. Jain. An Experimental Study of the Skype Peer-to-Peer VoIP System. In *Proc. of the 5th International Workshop on Peer-to-Peer Systems (IPTPS '06)*, February 2006.

[8] K. P. Gummadi, R. J. Dunn, S. Saroiu, S. D. Gribble, H. M. Levy, and J. Zahorjan. Measurement, Modeling and Analysis of a Peer-to-Peer File-Sharing Workload. In *Proc. of the 19th ACM Symposium of Operating Systems Principles (SOSP)*, October 2003.

[9] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang. Measurements, Analysis, and Modeling of BitTorrent-like Systems. In *Proc. of the Internet Measurement Conference (IMC)*, October 2005.

[10] X. Hei, C. Liang, J. Liang, Y. Liu, and K. W. Ross. A Measurement Study of a Large-Scale P2P IPTV System. Technical report, Department of Electrical and Computer Engineering, Polytechnic University, 2006.

[11] M. Izal, G. Urvoy-Keller, E. Biersack, P. Felber, A. A. Hamra, and L. Garces-Erice. Dissecting BitTorrent: Five Months in a Torrent's Lifetime. In *Proc. of the 5th Passive and Active Measurement Workshop (PAM '04)*, April 2004.

[12] M. Jovanovic, F. Annexstein, and K. Berman. Modeling Peer-to-Peer Network Topologies through Small-World Models and Power Laws. In *TELFOR*, November 2001.

[13] J. Liang, R. Kumar, and K. Ross. The FastTrack Overlay: A Measurement Study. *Computer Networks*, 50(6):842–858, April 2006.

[14] J. Pouwelse, P. Garbacki, D. Epema, and H. Sips. The BitTorrent P2P Filesharing System: Measurements and Analysis. In *Proc. of the 4th International Workshop on Peer-to-Peer Systems (IPTPS '05)*, February 2005.

[15] M. Ripeanu, I. Foster, and A. Iamnitchi. Mapping the Gnutella Network: Properties of Large-Scale Peer-to-Peer Systems and Implications for System Design. *IEEE Internet Computing Journal*, 6(1), 2002.

[16] T. Silverston and O. Fourmaux. P2P IPTV Measurement: A Case Study of TVants. In *Proc. of the 2nd Conference on Future Networking Technologies (CoNEXT '06)*, December 2006.

[17] D. Stutzbach, R. Rejaie, and S. Sen. Characterizing Unstructured Overlay Topologies in Modern P2P File-Sharing Systems. In *Proc. of Internet Measurement Conference (IMC)*, October 2005.

[18] D. J. Watts. *Six Degrees: the Science of a Connected Age*. ACM Press, 2003.

[19] X. Zhang, J. Liu, B. Li, and T. P. Yum. CoolStreaming/DONet: A Data-Driven Overlay Network for Live Media Streaming. In *Proc. of IEEE INFOCOM 2005*, March 2005.