Characterizing Peer-to-Peer Streaming Flows

Chuan Wu, Baochun Li, and Shuqiao Zhao

***Abstract*— The fundamental advantage of peer-to-peer (P2P) multimedia streaming applications is to leverage peer upload capacities to minimize bandwidth costs on dedicated streaming servers. The available bandwidth among peers is of pivotal importance to P2P streaming applications, especially as the number of peers in the streaming session reaches a very large scale. In this paper, we utilize more than 230 GB of traces collected from a commercial P2P streaming system, UUSee, over a four-month period of time. With such traces, we seek to thoroughly understand and characterize the achievable bandwidth of streaming flows among peers in large-scale real-world P2P live streaming sessions, in order to derive useful insights towards the improvement of current-generation P2P streaming protocols, such as peer selection. Using continuous traces over a long period of time, we explore evolutionary properties of inter-peer bandwidth. Focusing on representative snapshots of the entire topology at specific times, we investigate distributions of inter-peer bandwidth in various peer ISP/area/type categories, and statistically test and model the deciding factors that cause the variance of such inter-peer bandwidth. Our original discoveries in this study include: (1) The ISPs that peers belong to are more correlated to inter-peer bandwidth than their geographic locations; (2) There exist excellent linear correlations between peer last-mile bandwidth availability and inter-peer bandwidth within the same ISP, and between a subset of ISPs as well; and (3) The evolution of inter-peer bandwidth between two ISPs exhibits daily variation patterns. Based on these insights, we design a throughput expectation index that facilitates high-bandwidth peer selection without performing any measurements.**

***Index Terms*— Peer-to-Peer Streaming, Flow Characterization, Measurements.**

I. INTRODUCTION

THE FUNDAMENTAL advantage of peer-to-peer (P2P) live multimedia streaming is to allow peers to contribute their upload bandwidth, such that bandwidth costs may be saved on dedicated streaming servers. Server bandwidth cost savings are more substantial when participating peers contribute more bandwidth. It is therefore of pivotal importance for a peer to select other peers with high *inter-peer* bandwidth (*i.e.*, the available bandwidth between two peers) during a live streaming session, such that the media content can be timely retrieved to meet its playback deadline. As TCP is widely employed in P2P live streaming applications to guarantee reliability and to transverse NATs, the achievable TCP throughput is an essential metric when evaluating available inter-peer bandwidth.

Due to the inherent dynamic nature of peer arrivals and departures in a typical P2P streaming session, it is a daunting challenge to evaluate TCP throughput between two peers before data transmission begins. One may start a probing TCP connection to directly measure TCP throughput, but the time it takes for TCP to saturate available bandwidth leads to intrusive and expensive bandwidth usage, that can otherwise be available to stream actual media. A better approach would be to calculate TCP throughput based on flow sizes, maximum sender/receiver windows, and path characteristics such as delay and loss rate [1], [2], [3]. However, such calculations require the knowledge of TCP parameters or path characteristics, which may not be available without probing or new TCP connections, when a peer attempts to select high-bandwidth neighbors from a list of candidates. Yet another alternative may be to summarize historical TCP throughput using time series models, which may be utilized to forecast future TCP throughput [4], [5], [6]. Unfortunately, it is common for peers to come across neighbors with whom no historical TCP connections ever exist.

Though it is almost impossible to accurately predict TCP throughput between arbitrary peers without some probing or historical data, practical experiences show that it is *helpful* in the design of a peer selection protocol if the peer has only a *"rough idea"* about the available bandwidth between itself and a possible candidate, and such a "rough idea" can be used to *rank* the candidates based on available bandwidths. This paper represents the first step towards this objective, as we conduct a comprehensive statistical study of TCP throughputs, based on 230 GB of traces, 100 million unique IP addresses, and 370 million live streaming flows, collected as continuous-time snapshots over a four-month period in the entire network of a commercial P2P streaming system, operated by UUSee Inc. Backed by major venture capital investments, UUSee Inc. is a leading provider in mainland China for P2P live streaming solutions.

Our focus in this study is to statistically characterize TCP throughput distributions, investigate its application-layer deciding factors, and explore the correlation between these factors and TCP throughput, so as to derive useful insights towards the design of practical peer selection protocols based on inter-peer bandwidth availability. Unlike existing TCP throughput characterization which focuses on "microscopic" characteristics such as window sizes and delay, we explore "macroscopic" factors which we classify into two categories: (1) end-host characteristics, including peer last-mile upload/download capacity and the number of contingent sending/receiving TCP connections; and (2) membership factors, such as ISPs the peers reside in and geographic areas the peers locate at. We not only investigate evolutionary properties of TCP throughputs in the temporal dimension over a long period of time, but also zoom into snapshots at specific

Manuscript received on March 15, 2007. Revised on August 1st, 2007.
Chuan Wu and Baochun Li are with Department of Electrical and Computer Engineering, University of Toronto (e-mail: {chuanwu, bli}@eecg.toronto.edu).
Shuqiao Zhao is with Multimedia Development Group, UUSee, Inc. (e-mail: {zhaosq}@uusee.com).
Digital Object Identifier 10.1109/JSAC.2007.0712xx.

times, including both a representative regular time and special scenarios such as flash crowds. Our objective is to quantify the correlation between TCP throughput and its influential factors by modeling them into statistical regression models, and use such models to achieve practical peer selection based on available bandwidths.

Our original discoveries in this study include: (1) The ISPs that peers belong to are more correlated to inter-peer bandwidth than their geographic locations; (2) Inter-ISP peering does not always constitute bandwidth bottlenecks, which is ISP specific; (3) There exist excellent linear correlations between peer last-mile bandwidth availability and inter-peer bandwidth within the same ISP, and between a subset of ISPs as well; (4) The evolution of inter-peer bandwidth between two ISPs exhibits a daily variation pattern, although the level of throughput values shifts from time to time; (5) During a flash crowd scenario, the inter-peer bandwidth characteristics do not represent significant differences from those at regular times.

Based on these insights, we design a *throughput expectation index* that facilitates high-bandwidth peer selection without performing measurements. This index computes the relative *rank* of candidate peers based on the bandwidth availability between them and the receiver peer. We show that peers selected with this index are consistent with those with actual highest throughput values, based on cross-validation experiments over the trace period.

The remainder of this paper is organized as follows. In Sec. II, we briefly review existing work in flow rate characterization and measurements of P2P systems. In Sec. III, we outline trace collection methodologies in the UUSee streaming platform, and present the scale of our measurements. In Sec. IV, we investigate distributions of TCP throughput in different scenarios, within and across ISPs and areas, and for different peer types. We then zoom into a one-time snapshot of the entire network in Sec. V, and statistically characterize the correlation between TCP throughput and a number of influential factors. Next, in Sec. VI, evolutionary properties of inter-peer bandwidth availability are discovered over a long period of time. We develop the throughput expectation index to assist peer selection in Sec. VII. Finally, we conclude the paper in Sec. VIII.

II. Related Work

Significant research attention has been devoted to the characterization and prediction of TCP throughput. Existing research derives TCP throughput using either a *formula-based* approach in terms of key metrics such as window size, RTT or loss rate [1], [2], [3], or a *history-based* approach, based on historical throughput measurements on the same link, using standard time series forecasting techniques, *e.g.*, MA, EWMA, AR, ARMA, ARIMA [4], [5], [6], [7].

With respect to measurement based TCP throughput characterization, Balakrishnan *et al.* [8] suggested that end-to-end throughput, from an Olympic games website, can be modeled as a log-normal distribution. Zhang *et al.* [9] also exhibited a log-normal rate distribution of Internet flows between several sites on an ISP backbone. Compared to their study of unicast flow throughputs from a single web site or between a limited number of sites, our discoveries are more comprehensive and also original in the sense that we investigate various characteristics of flows between millions of pairs of peers that lie in a wide range of geographical regions and ISPs.

Little work exists for flow rate characterization in P2P networks. In their performance study of broadband hosts, Lakshminarayanan *et al.* [10] pointed out that the generally assumed inverse relationship between RTT and TCP throughput is masked by the wide range in last-hop peer bandwidth for broadband peers. In order to investigate the serving/downloading power of participating peers, Sen *et al.* [11] and Saroiu *et al*. [12] characterized upload and download bandwidth bottlenecks at the peers in P2P file sharing systems. Other than characterizing the overall upload/download capacities at each peer, our study investigates inter-peer available bandwidths, which take into consideration both the last-mile bandwidths at the end hosts and intermediate bandwidth bottleneck along the P2P links.

There have recently been a number of measurement studies on various P2P live streaming systems, *e.g.*, PPLive [13], [14], TVants [15], [16], and SOPCast [14], [16]. In terms of throughput, they studied the total upload/download throughput at peers in order to investigate network bandwidth utilization, but have not investigated the inter-peer bandwidths along P2P links.

With respect to general measurements on P2P file sharing and VoIP applications, a number of work have emerged in recent years, towards Kazza overlay [17], [18], Gnutella [19], [20], BitTorrent [21], [22], [23], and *Skype*, a P2P VoIP application [24], [25], [26]. However, none of them have systematically studied inter-peer bandwidth availability in the applications.

To the best of our knowledge, this paper represents the first comprehensive

study on flow characterization in large-scale P2P applications, towards the goal of better and practical peer selection protocol design, using a macroscopic approach. Nevertheless, when applicable, we will compare our forthcoming discoveries with the existing results.

III. Scale of Traces

We first present our methodology of trace collection, and then show the scale of the traces that this work is based on.

A. Collection methodology

Throughout this paper, we analyze traces from live P2P streaming sessions of a commercial P2P streaming company, UUSee Inc. [27], which is a leading P2P streaming solution provider in mainland China, with exclusive and legal rights to channels of CCTV, the official Chinese television network. With a large collection of streaming servers around the world, it simultaneously sustains over 800 channels, mostly encoded to 400 Kbps streams. Similar to all current-generation P2P streaming protocols, UUSee's streaming protocol design is based on the "pull-based" design 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 the channel in UUSee, the initial set of a number of *partners* is supplied by one of the tracking servers. The peer establishes TCP connections with these partners, and buffer maps are periodically exchanged.

To inspect the real-world bandwidth availability in UUSee P2P streaming, we have implemented related measurement and reporting capabilities within its P2P client application. Each peer in UUSee estimates its total download and upload bandwidth capacities. For each active partner with which it has a live TCP connection, it measures the maximum achievable sending or receiving throughput of the TCP connection every 5 minutes, and reports all measurements to a central trace server via UDP.

The download capacity of each peer is measured at the initial buffering stage of the peer, upon its first joining a streaming channel in the UUSee network. During this stage, the peer has no available blocks in its playback buffer, and can concurrently download from many supplying peers. In this case, its download bandwidth is largely saturated. Therefore, the download capacity of the peer is estimated as its maximum aggregate download rate at this initial buffering stage.

The upload capacity at each peer is measured upon its joining before the actual streaming starts, by setting up a temporary upload TCP connection with one of the nearest streaming servers. As we know, the upload bandwidth at each streaming server is mostly saturated due to its main upload functions, while the download bandwidth is largely idle. Therefore, we utilize the spare download capacity of the streaming servers, and have each peer send a randomly generated probing flow to a streaming server that is nearest to itself. The duration of the flow should be long enough for its TCP throughput to become stable, usually in seconds. The streaming server measures the stabilized TCP throughput on this connection, which is then estimated as the upload capacity of the respective peer.

The reported maximum throughput along a live TCP connection is measured in the following fashion: The time is divided into 30-second intervals. In each interval, the time that is actually used to transmit media blocks is summarized, excluding the idle TCP periods. An average throughput is calculated with the number of bytes sent in the block transmission time divided by the length of this duration. The maximum throughput is then derived as the maximum of all such average throughputs within 5 minutes. Taking the average transmission throughput within 30 seconds, we smooth out the periods of very bursty TCP throughput; deriving the maximum of all such 30-second measurements, we aim to obtain the maximally achievable TCP throughput on the link between two peers.

We further note that by only counting the time of actual block transmissions, our maximum throughput measurements essentially reflect the maximum availability of network bandwidths, eliminating any possible impact of media block availability. Therefore, such throughputs can be much larger than the streaming rate of the media channels, which is constrained by the block availability. We believe the investigation of such inter-peer bandwidth availability is important for any P2P live streaming protocol, as it is desirable to maximally download media blocks from the peers with large throughputs, in order to achieve timely delivery of the media streams before the playback deadline.

Each peer reports a collection of these measurements to the trace server every 5 minutes. Each report includes the IP address of the peer, its total download and upload bandwidth capacities, as well as a list of all its partners, with their corresponding IP addresses, TCP/UDP ports, number of segments sent to or received from each partner, and the current maximum sending/receiving throughput on each connection.

B. Trace summary

During a four month period from November 2006 to February 2007, we have collected more than 230 GB of traces with more than 100 million unique IP addresses and 370 million streaming flows, representing snapshots of the live P2P streaming network every five minutes throughout this period. In what follows, we illustrate the scale of the traces with respect to the numbers of simultaneous peers and P2P flows in live streaming sessions. Due to the large volume of the traces, in our figures, we will only depict results obtained over one representative regular week, 12:00am December 17th, 2006 (GMT+8) — 11:50pm December 23, 2006 (GMT+8).

1) Overall number: Fig. 1(A) shows that there are on average $100,000$ concurrent peers and $250,000$ active flows at any time in the UUSee streaming network. Both statistics show two daily peaks around 1pm and 10pm. Fig. 1(B) further summarizes the numbers of distinct IP addresses and P2P flows that appeared in the traces on a daily basis, which indicate that the traces contain information of up to 10 million different streaming flows among 2 million distinct peers each day. Such data abundance facilitates our forthcoming statistical investigations.

Besides the regular daily peer/flow numbers in Fig. 1, we have also observed a few flash crowd scenarios during the trace period. For example, a flash crowd scenario was observed around 23pm, February 17, 2007, caused by the broadcast of the celebration TV show on Chinese New Year Eve, with $871,000$ peers online in the UUSee streaming network and $2,271,000$ streaming flows among the peers.

2) Different ISPs: Using a mapping database obtained from UUSee Inc. that translates ranges of IP addresses to ISPs and geographic areas, we summarized the average ISP distribution of peers and P2P flows during the trace period. For each IP address inside China, the database provides the ISP it belongs to and province the user locates at; for each IP address out of China, it provides coarse geographic information of the continent the address lies in, but not detailed ISP information. Fig. 2(A) depicts the distribution of peers across major China ISPs and overseas. It exhibits that the two largest nationwide ISPs in China, Netcom and Telecom, own the largest user shares in the UUSee P2P network. While the majority of UUSee users are in China, peers from overseas also take a significant 20%, and their percentage shows a rising trend as we observed in our investigation. In our current study, we will mainly focus on streaming flows within China, and believe our discoveries will also bring useful insights towards global networks.

In addition, Fig. 2(B) summarizes the average number of concurrent streaming flows inside each major China ISP, and

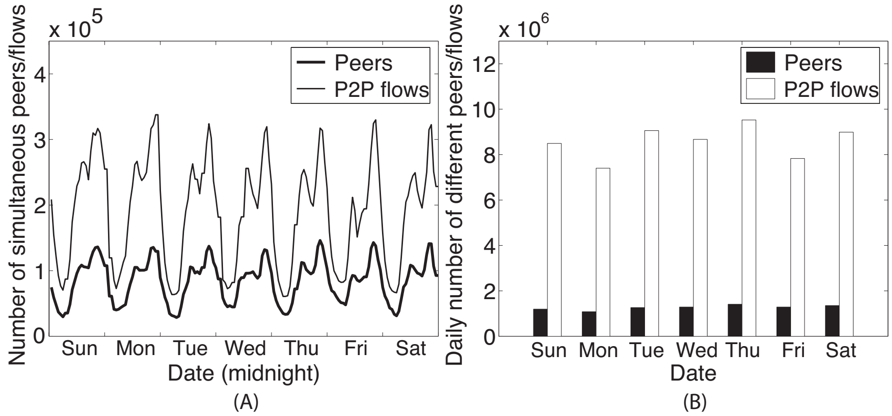

Fig. 1. Daily peer/P2P flow number statistics.

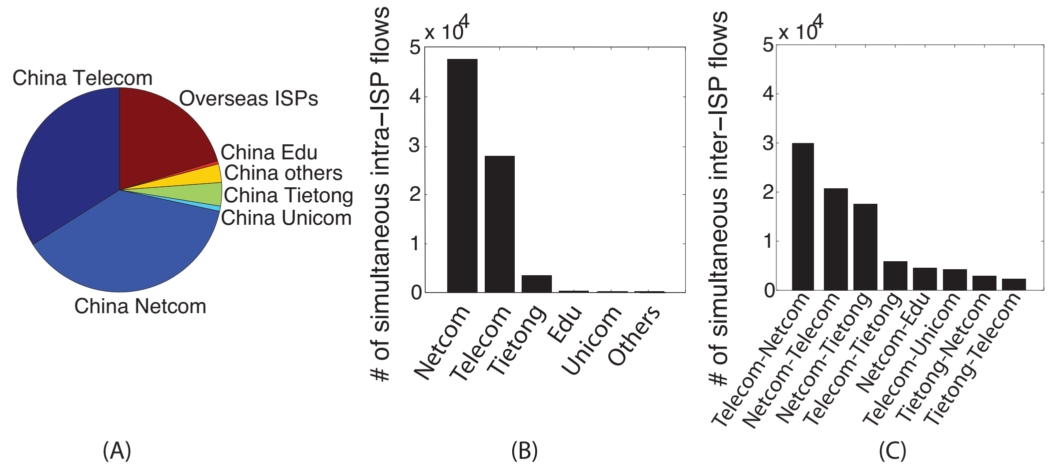

Fig. 2. Peer/P2P flow number statistics for different ISPs.

Fig. 2(C) illustrates the number of inter-ISP flows for ISP pairs that have more than 1000 concurrent flows in between. Again, the numbers of flows inside China Netcom and Telecom, and those for flows to and from these two ISPs dominate their respective categories.

3) Different peer types: We next categorize peers into two classes based on their download capacities in the traces, and the fact that the download bandwidth of the fastest ADSL connection in China is at most 3 Mbps: (1) Ethernet peers, for those with download capacities higher than 375 KBps; (2) ADSL/cable modem peers, for the remainder. While Fig. 3(A) exhibits the domination of ADSL/cable modem peer population, Fig. 3(B) shows comparable shares of P2P flows in each category, which demonstrates the contribution of the limited number of Ethernet peers in uploading to many other peers.

IV. Throughput Distributions

We start our P2P streaming flow characterization by analyzing the distributions of TCP throughput at representative times, across or within different ISPs/areas, and among different peer types. We note that in the rest of the paper, all our flow characterizations are based on the 5-minute maximum throughput measurements from the traces.

A. Overall throughput distribution at different times

Fig. 4 shows the throughput distribution[1] over the entire network at four representative regular times: Monday morning (9am 12/18/06), Monday evening (9pm 12/18/06), Friday morning (9am 12/22/06) and Friday evening (9pm 12/22/06). With throughput depicted in the log scale, the plots represent the shapes of normal distributions, corresponding to the original throughputs having the analytic distributions of log-normal [8], [28]. This finding is consistent with existing work of Balakrishnan *et al.* [8] and Zhang *et al.* [9], who also discovered log-normal rate distributions within their Internet flow sets.

The throughput distributions at the four times peak at 15KBps, 7KBps, 13KBps, 7KBps, respectively, with an 80th percentile of 280KBps, 96KBps, 275KBps, and 90KBps, respectively. We observe that the mean throughputs in the mornings, which are daily off-peak hours for the streaming application, are $2-3$ times higher than those at evening peak hours, and the variance of the throughputs in the mornings is larger than that in the evenings as well. For the same time at different days in a week, however, there does not exist an apparent throughput difference.

We further validate the above observations statistically using one-way analysis of variance (ANOVA) [29], [30]. The one-

[1]The bin size used in all our throughput distribution plots in this section is 1KBps.

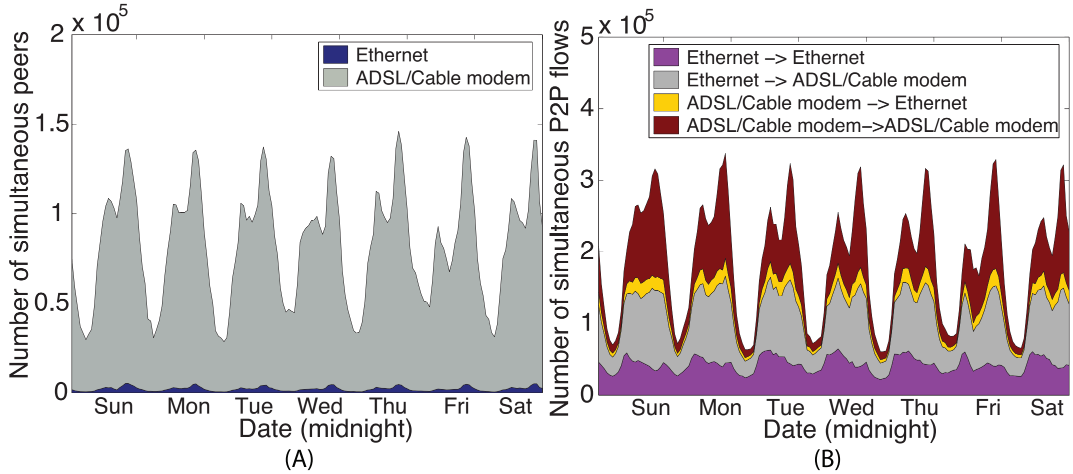

Fig. 3. Peer/P2P flow number statistics in two peer type categories.

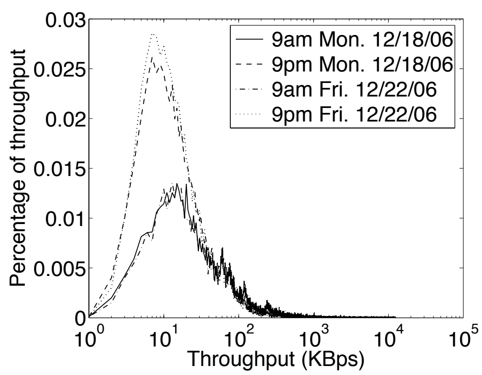

Fig. 4. Overall throughput distribution at different times.

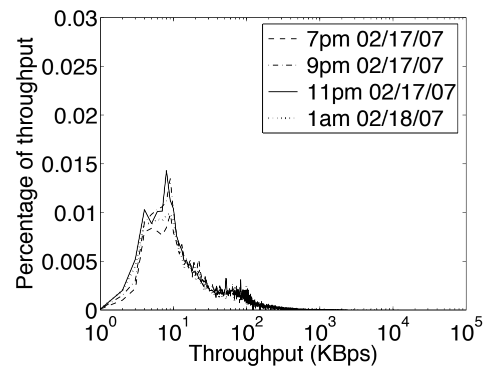

Fig. 5. Overall throughput distribution at Chinese New Year Eve.

TABLE I

KRUSKAL-WALLIS ANOVA FOR THROUGHPUTS AT DIFFERENT TIMES

Null Hypothesis	Throughput Sets	p-value
The two sets of throughputs have the same distribution	9am Mon. vs. 9am Fri.	0.8699
	9pm Mon. vs. 9pm Fri.	0.0684
	9am Mon. vs. 9pm Mon.	0
	9am Fri. vs. 9pm Fri.	0

way ANOVA is used to test the null hypothesis that different sets of samples for an independent variable have all been drawn indifferently from the same underlying distribution. In our case, we use ANOVA to examine whether the throughput distributions at different times on a same regular day are statistically equivalent, and whether those at the same time on different days are significantly different. As the numbers of throughput samples in the four sets are different, we conduct ANOVA by using the non-parametric Kruskal-Wallis test [30]. The comparisons and reported p-values are listed in Table I.

In our hypothesis test, if a result p-value is lower than the significance level of 0.05, the difference between the corresponding distributions are statistically significant, and the null hypothesis is rejected; otherwise there is insufficient evidence to reject the null hypothesis. The 0 p-values reported for the latter two tests strongly suggest the difference between throughputs at different times of a day, while the other large p-values validate the large similarity among morning/evening throughput sets on different days.

While the above observations generally apply for throughput sets on regular days, we have also investigated throughput distributions during a flash crowd scenario on Chinese New Year Eve (Feb. 17th, 2007), as shown in Fig. 5. Four representative snapshots are plotted: 7pm on the Eve, before the celebration TV broadcast started; 9pm, when the flash crowd started to gather as more and more viewers tuned in to the channel; 11pm, when the flash crowd reached its largest size as the Chinese New Year approached; and 1am on the next morning, when the crowd dismissed itself after the show ended. With ANOVA tests, we detected that the distributions are statistically different, with throughputs at 7pm statistically larger than those at 1am, followed by those around 9pm, and then those at 11pm. This reflects that inter-peer bandwidths became tight as the size of flash crowd increased and turned loose again when the crowd dismissed. Nevertheless, there does not exist a "crash" scenario with abrupt drop of throughput over the network, and the throughputs follow similar log-normal distributions as those at the same time on a regular day.

B. Intra/inter ISP throughput distribution

We next categorize the P2P streaming flows into two classes and investigate their respective throughput distributions: (1) intra-ISP flows, for which the sender and receiver are in the same ISP, and (2) inter-ISP flows, where they belong to different ISPs. Fig. 6 exhibits that, while they still follow log-normal distributions in each category, intra-ISP throughputs

TABLE II
KRUSKAL-WALLIS ANOVA FOR THROUGHPUTS ACROSS DIFFERENT ISPS AT 9PM, DEC. 18, 2006

Null Hypothesis	Throughput Sets	p-value	Multiple Comparison Test Result
Throughput sets within an ISP and from different ISPs to this ISP have the same distribution	(1) TC→TC, NC→TC, UC→TC, TT→TC, EDU→TC	0	$\text{Throughput}_{TC\to TC} \approx \text{Throughput}_{UC\to TC} \approx \text{Throughput}_{EDU\to TC} > \text{Throughput}_{TT\to TC} > \text{Throughput}_{NC\to TC}$
	(2) NC→NC, TC→NC, UC→NC, TT→NC, EDU→NC	0	$\text{Throughput}_{NC\to NC} \approx \text{Throughput}_{UC\to NC} \approx \text{Throughput}_{TT\to NC} > \text{Throughput}_{EDU\to NC} \approx \text{Throughput}_{TC\to NC}$
	(3) UC→UC, TC→UC, NC→UC, TT→UC, EDU→UC	0.062	
	(4) TT→TT, TC→TT, NC→TT, UC→TT, EDU→TT	0.081	
Throughput set from ISP1 to ISP2 and throughput set from ISP2 to ISP1 have the same distribution	(1) TC→NC, NC→TC	0.032	$\text{Throughput}_{TC\to NC} > \text{Throughput}_{NC\to TC}$
	(2) TC→UC, UC→TC	0.023	$\text{Throughput}_{UC\to TC} > \text{Throughput}_{TC\to UC}$
	(3) NC→TT, TT→NC	0.029	$\text{Throughput}_{TT\to NC} > \text{Throughput}_{NC\to TT}$
	(4) UC→TT, TT→UC	0.396	
	(5) EDU→UC, UC→EDU	0.153	

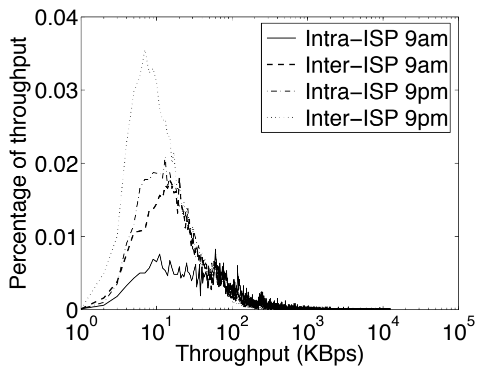

Fig. 6. Intra/inter ISP throughput distribution on Dec. 18, 2006.

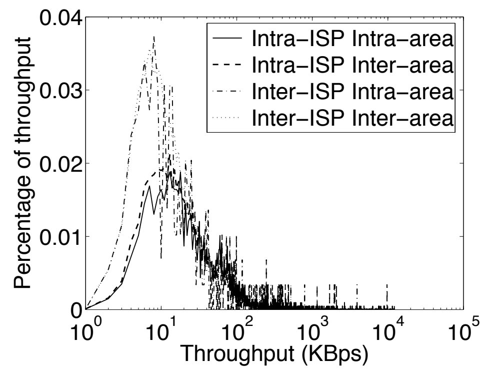

Fig. 7. Intra/inter area throughput distribution at 9pm, Dec. 18, 2006.

are generally larger than their inter-ISP counterparts measured at the same time: the former have peaks at 60KBps (9am) and 13KBps (9pm), while those of the latter are 20KBps (9am) and 7KBps (9pm), respectively. Also observed is that intra-ISP throughputs at peak hours are generally smaller than inter-ISP throughputs at off-peak hours on the same day. Within each intra-ISP or inter-ISP category, the throughput distributions show a similar diurnal pattern as that revealed by the overall throughput distributions in the previous section: both the mean and variance of the throughput distributions in the mornings are larger than those in the evenings.

While these observations meet our general expectation that bandwidth is more abundant within each ISP, we also notice many large inter-ISP throughput values and the large span for both inter-ISP and intra-ISP throughputs. This inspires us to further investigate: Are throughputs for flows within an ISP always statistically larger than those for flows to and from this ISP? Is there significant throughput difference across different pairs of ISPs? To answer these questions, we again conduct Kruskal-Wallis ANOVA tests to various throughput sets categorized based on contingent ISPs of the flows. If 3 or more throughput sets are compared in one test and significant difference is reported, we further perform the *multiple comparison test (or procedure)* [29], [30] to investigate the difference between each pair of sets. The representative tests and their results are given in Table II[2].

Again, taking 0.05 as the p-value threshold to determine if we should reject the null hypothesis, our discoveries from the ANOVA are the following. *First*, inter-ISP throughputs are not necessarily smaller than their intra-ISP counterparts. For the two largest China ISPs, Netcom and Telecom, the throughputs of their inbound flows are generally smaller than those of their internal flows. Throughputs are especially small between the two ISPs themselves. For every other ISP, there is no significant throughput difference among the internal flows and inbound flows. This validates the fact that there is a stringent bandwidth constraint between Netcom and Telecom, as two major ISP competitors in China, while no such caps exist across the other small ISPs and between those two and the small ISPs. *Second*, throughput asymmetry is exhibited from one direction to the other across the two largest ISPs, as well as between them and the other ISPs. The observation that throughput from large ISPs to small ISPs are smaller than those in the other direction may reveal possible bandwidth caps placed by large ISPs on such relay traffic.

C. Intra/inter area throughput distribution

To characterize the P2P streaming flow at finer granularity below the ISP level, we next compare throughput distributions

[2]To conserve space, we use the following abbreviations for ISPs: TC (Telecom), NC (Netcom), UC (Unicom), TT (Tietong), Edu (Education Network).

TABLE III

KRUSKAL-WALLIS ANOVA FOR INTER/INTRA AREA THROUGHPUTS BETWEEN DIFFERENT ISPS AT 9PM, DEC. 18, 2006

Null Hypothesis	Throughput Sets	p-value
Inside the same ISP, intra-area throughput set and inter-area throughput set have the same distribution	(1) intra-TC: intra-area set v.s. inter-area set	0.2396
	(2) intra-NC: intra-area set v.s. inter-area set	0.0701
	(3) intra-TT: intra-area set v.s. inter-area set	0.6228
	(4) intra-UC: intra-area set v.s. inter-area set	0.5751
Across two different ISPs, intra-area throughput set and inter-area throughput set have the same distribution	(1) TC→NC: intra-area set v.s. inter-area set	0.117
	(2) NC→TC: intra-area set v.s. inter-area set	0.179
	(3) NC→TT: intra-area set v.s. inter-area set	0.3105
	(4) UC→TT: intra-area set v.s. inter-area set	0.4575
Inside the same area, throughput sets within one ISP and from different ISPs to this ISP have the same distribution	(1) TC→TC, NC→TC, UC→TC, TT→TC, EDU→TC	0.0015
	(2) NC→NC, TC→NC, UC→NC, TT→NC, EDU→NC	0.0448
	(3) UC→UC, TC→UC, NC→UC, TT→UC, EDU→UC	0.5846
	(4) TT→TT, TC→TT, NC→TT, UC→TT, EDU→TT	0.5511
Across two different areas, throughput sets within one ISP and from different ISPs to this ISP have the same distribution	(1) TC→TC, NC→TC, UC→TC, TT→TC, EDU→TC	0
	(2) NC→NC, TC→NC, UC→NC, TT→NC, EDU→NC	0
	(3) UC→UC, TC→UC, NC→UC, TT→UC, EDU→UC	0.052
	(4) TT→TT, TC→TT, NC→TT, UC→TT, EDU→TT	0.2929

in the cases that the sender and receiver are located within or not within the same geographic area (intra-area vs. inter-area). Here, the peers are in the same area if they are in the same province of China. As we have concluded that ISP memberships of the peers may significantly affect the inter-peer bandwidth, we investigate four cases, as shown in Fig. 7. When ISP memberships are fixed, we observe no significant difference between the distributions of intra-area throughputs and inter-area throughputs; in either area case, intra-ISP throughputs are always larger than inter-ISP throughputs. To validate these observations, we again perform ANOVA to test the difference between the intra-area throughput set and inter-area throughput set for each specific ISP pair, and the difference among throughput sets from different ISPs to one ISP in both the intra-area and inter-area cases. The representative tests and their results are given in Table III.

Comparing the p-values with threshold 0.05, we first find that, in both cases that the sender and receiver do and do not belong to the same ISP, there does not exist a significant throughput difference when the sender and receiver are further in or not in the same area (province). For the two nationwide ISPs, Telecom and Netcom, considering the fact that they are organized on the provincial basis, our discovery shows that within each of them, the inter-province bandwidth constraints do not have apparent negative impact on inter-province P2P flow throughputs. In addition, across the two ISPs, a same provincial locality of two peers does not help in improving the inter-peer bandwidth. This can be explained by the facts that the two ISPs have only 4-6 fixed peering points across China, and even if two peers are in the same province, the underlying links in between them may well go via a peering point that is thousands of kilometers away. Second, in both cases that the sender and receiver are in and not in the same area (province), the comparisons of throughputs from different ISPs (including itself) to the same ISP exhibit similar results as those we have shown in Table II. While area information is included in the comparisons in Table III but not in Table II, they both show that, for large ISPs, there exist differences between their internal throughputs and those from other ISPs to them; for small ISPs, no difference is identified among the different throughput sets.

All these results lead to the conclusion that ISP membership has more significant impact on inter-peer bandwidths, as compared to geographic locations. In what follows, we mainly focus on ISP memberships when we discuss deciding factors that affect bandwidth availability in the middle of a P2P link.

D. Throughput distribution for different peer types

To discover the impact of peer types (*i.e.*, peer last-mile bandwidths) on inter-peer bandwidth, we further categorize intra-ISP and inter-ISP flows based on types of their incident peers, and plot the CDF of throughput in each category in Fig. 8. The plots and accompanying ANOVA tests exhibit that: throughputs are significantly larger when the sender is an Ethernet peer, in both the intra-ISP and inter-ISP cases; for the same sender type, flows with Ethernet receivers achieve higher throughput in most cases.

The above results reveal a coarse positive correlation between inter-peer bandwidth and the last-mile bandwidths at the peers. It inspires us to further consider the following questions: Is the peer last-mile download/upload capacity the key factor that decides inter-peer bandwidth, both when the peers are in the same ISP and when they are across any pair of ISPs? Or is inter-ISP peering the most important factor that affects throughput between some ISPs? In the following section, we seek to answer these questions with regression modeling of the throughputs.

V. THROUGHPUT REGRESSION: FOCUSING ON ONE SNAPSHOT

Focusing on one representative regular snapshot of the UUSee streaming network at 9pm December 18 2006, we investigate the impact of the following factors on inter-peer bandwidths: (1) ISP memberships of the peers, and (2) end-host characteristics, including upload/download capacities and the number of contingent sending/receiving TCP connections at the sender/receiver. We divide our discussions into two cases, intra-ISP case and inter-ISP case, and perform regression analysis on TCP throughputs and the respective end-host characteristics in each case.

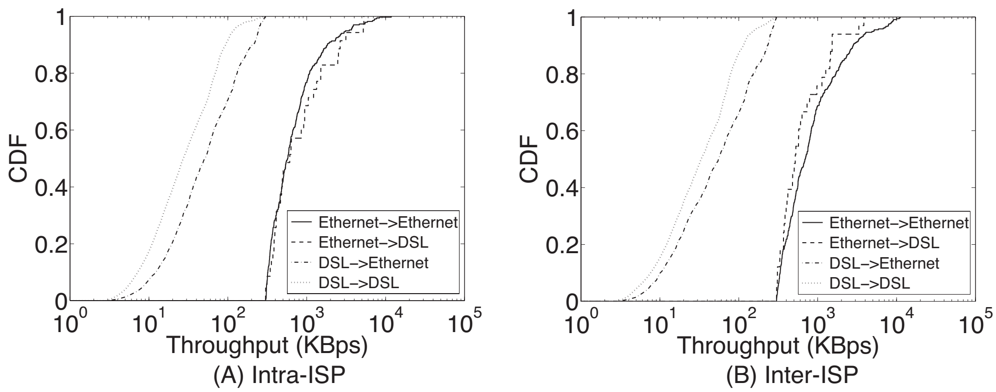

Fig. 8. Throughput CDF for different peer types at 9pm, Dec. 18, 2006.

A. *Intra-ISP throughput regression*

With the example of China Netcom, we check the correlation between flow throughputs on its internal P2P links and various end-host characteristics at the peers. To eliminate outliers and better capture the correlation, we divide the values of each capacity factor into small bins with a width of 5KBps, and plot the median throughput of flows falling into each bin at different levels of capacities in Fig. 9. The calculated Pearson product-moment correlation coefficient between throughput and a respective factor, *rho*, is marked at the upper right corner in each plot.

Fig. 9(A) exhibits that no significant linear correlation exists between throughput and the upload capacity at the sender peer, especially when the latter is large, *i.e.*, the Ethernet peer case. On the other hand, throughput and download capacity at the receiver peer is better correlated, as shown in Fig. 9(B). Nevertheless, there exist many cases in which the throughput is small when the capacity is large. Such unsatisfactory correlations inspire us to consider: Is the number of contingent upload/download flows high when the upload/download capacity is large, such that the bandwidth share for each flow is small? To answer this question, Fig. 9(C) shows a positive correlation between the upload capacity at the senders and the number of their concurrent upload flows, while no significant correlation is exhibited between receiver download capacities and their numbers of concurrent download flows in Fig. 9(D). The positive correlation in the former case can be explained by the UUSee streaming protocol design, which maximally utilizes upload capacity at each peer to serve more neighbors.

Naturally, we then investigate the correlation between throughput and per-flow upload/download bandwidth availability at the sender/receiver, defined as:

$$\textit{per-flow sender capacity} = \frac{\textit{sender upload capacity}}{\textit{no. of concurrent upload flows}},$$

$$\textit{per-flow receiver capacity} = \frac{\textit{receiver download capacity}}{\textit{no. of concurrent download flows}}.$$

Fig. 9(E) and (F) exhibit that these two characteristics constitute better explanatory variables towards the throughput regression.

When we take the minimum of per-flow sender capacity and per-flow receiver capacity, we obtain the best deciding factor of throughput, referred to as *per-flow end capacity (PEC)*:

$$PEC = \min(\textit{per-flow sender capacity, per-flow receiver capacity}).$$

Its excellent positive correlation with throughput, with a correlation coefficient of 0.81, is plotted in Fig. 10. We next fit PEC and throughput into a linear regression model:

$$\textit{Throughput} = \beta_0 + \beta_1 \times PEC + \epsilon, \quad (1)$$

where *PEC* is the explanatory variable, *Throughput* is the response variable, y-intercept β_0 and slope β_1 are regression coefficients to be estimated, and ϵ denotes the error term.

TABLE IV
ROBUST LINEAR REGRESSION STATISTICS FOR INTRA-NETCOM THROUGHPUTS AT 9PM, DEC. 18, 2006

β_0 (y-intercept)	β_1 (slope)	p-value for testing significance of β_0	p-value for testing significance of β_1
20.4228	1.1499	0	0

The basic assumption for least-squares based linear regression analysis is that the response variable is normally distributed. However, as we have shown in Sec. IV, throughputs follow approximate log-normal distributions, in which the few large tail values tend to have a strong influence on the regression model. Therefore, we employ *robust linear regression* [31], [32], which uses an iteratively re-weighted least-squares algorithm and is less sensitive to outliers by giving them lower weights. The derived regression statistics are given in Table IV:

The two p-values are results from tests of the following two null hypotheses, respectively: (1) the y-intercept is "0" (*i.e.*, the y-intercept is non-significant), and (2) the slope is "0" (*i.e.*, the slope is non-significant). As a 0 p-value rejects a corresponding null hypothesis and confirms the significance of regression, the statistics in Table IV further establish the linear correlation between PEC and throughput on intra-ISP flows. In addition, theoretically we expect the regression line to pass through the origin and the slope to be approximately at 45°, and these are validated by the small y-intercept value (as compared to peer last-mile capacities) and near-1 slope value.

We have conducted the same regression analysis to throughputs within other ISPs and observed similar correlations. Therefore, we may conclude that, within each ISP, inter-peer bandwidth bottleneck mainly lies at the end hosts,

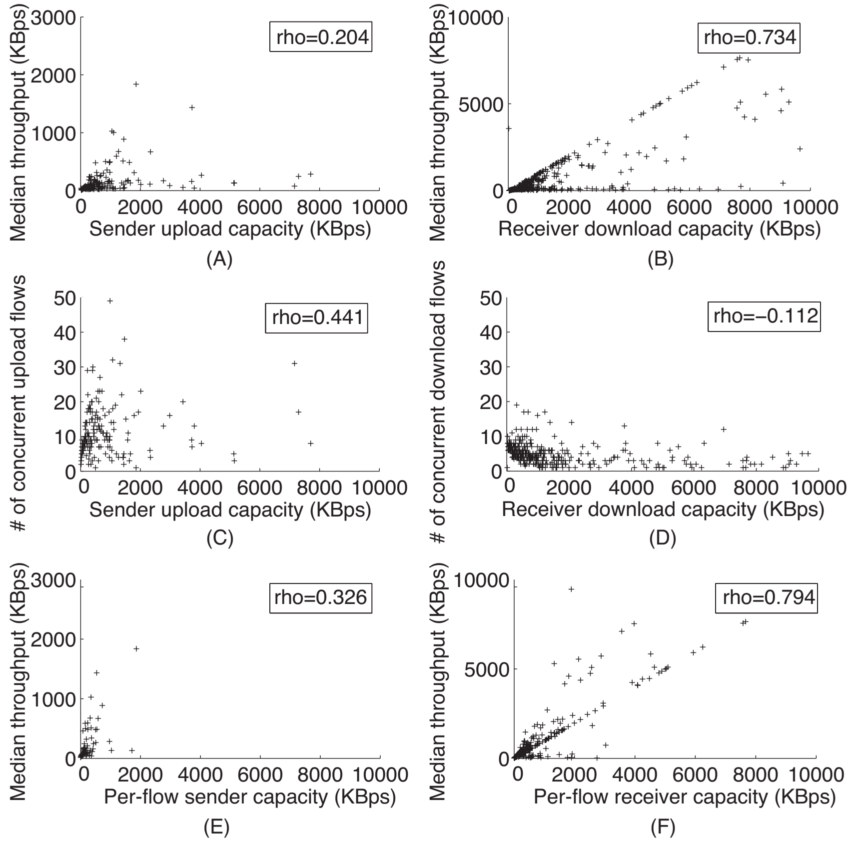

Fig. 9. Correlation of throughput with end-host characteristics for intra-Netcom flows at 9pm, Dec. 18, 2006.

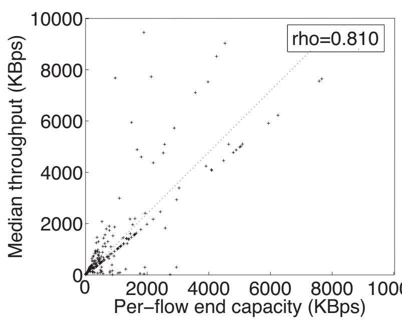

Fig. 10. Correlation of throughput with per-flow end capacity for intra-Netcom flows at 9pm, Dec. 18, 2006.

decided by peer last-mile capacities and their concurrent upload/download load: if the capacity bottleneck occurs at the upstream peer, *i.e.*, $\frac{\textit{sender upload capacity}}{\textit{no. of concurrent upload flows}} < \frac{\textit{receiver download capacity}}{\textit{no. of concurrent download flows}}$, the throughput is limited by the per-flow sender capacity; otherwise, the throughput is decided by the per-flow receiver capacity.

B. Inter-ISP throughput regression

When it comes to the inter-ISP case, we are interested to explore whether per-flow end capacity still poses a significant impact on inter-peer bandwidth, or it is shadowed by the inter-ISP peering bandwidth bottlenecks. We find the answer is different towards different ISP pairs.

Fig. 11(A) exhibits that no significant correlation exists between PEC and throughput for flows from Netcom to Telecom. This is further confirmed by its robust regression analysis statistics in Table V: a p-value of 0.6932 reveals the non-significance of the slope at the value of 0.0005. Nevertheless, when we investigate streaming flows from Netcom to Tietong, Fig. 11(B) shows a different result: throughput is linearly correlated with PEC with a slope of 0.4355, and a corresponding p-value of 0 indicates its significance. The regression statistics for representative flow groups between other ISPs are also listed in Table V.

TABLE V

ROBUST LINEAR REGRESSION STATISTICS FOR INTER-ISP THROUGHPUTS AT 9PM, DEC. 18, 2006

Throughput Set	β_0 (y-intercept)	β_1 (slope)	p-value for testing significance of β_0	p-value for testing significance of β_1
NC→TC	9.9526	0.0005	0	0.6932
TC→NC	20.4998	−0.0023	0	0.6585
NC→TT	30.5784	0.4355	0	0
TT→NC	39.094	0.316	0	0
TC→TT	24.1774	0.1109	0	0.0480
TT→TC	27.3144	0.5421	0	0
UC→Edu	20.2793	0.7098	0	0
Edu→UC	25.0535	0.4576	0	0

The statistics in Table V exhibit that: between the two largest ISPs, Netcom and Telecom, throughput is not contingent upon PEC, but limited by their peering bandwidth bottleneck; across other small ISPs and between other ISPs and the two, flow throughput is more or less decided by the peer last-mile bandwidth availability. In addition, in the latter cases, the regression slopes are generally smaller than

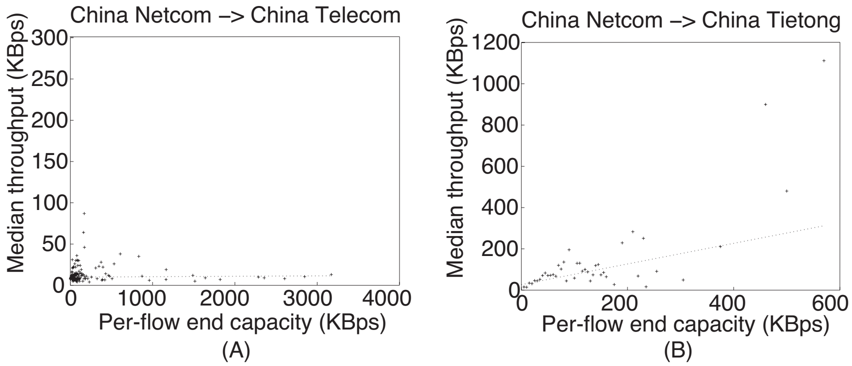

Fig. 11. Correlation of throughput with per-flow end capacity for inter-ISP flows at 9pm, Dec. 18, 2006.

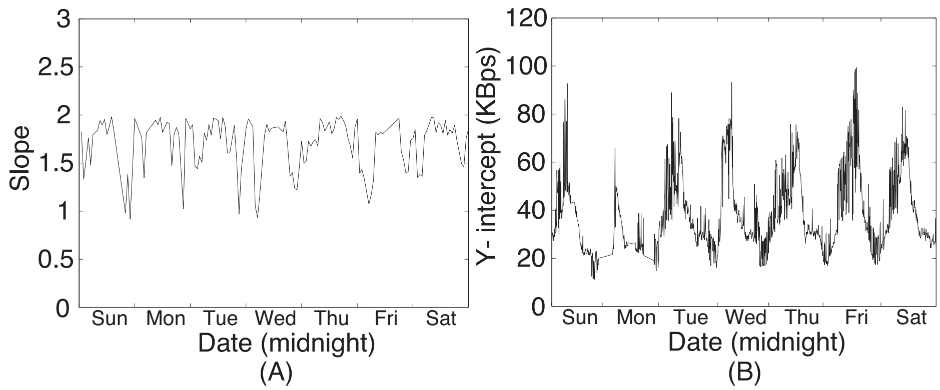

Fig. 12. Evolution of regression coefficients for intra-Netcom flows in the week of Dec. 17 — 23, 2006.

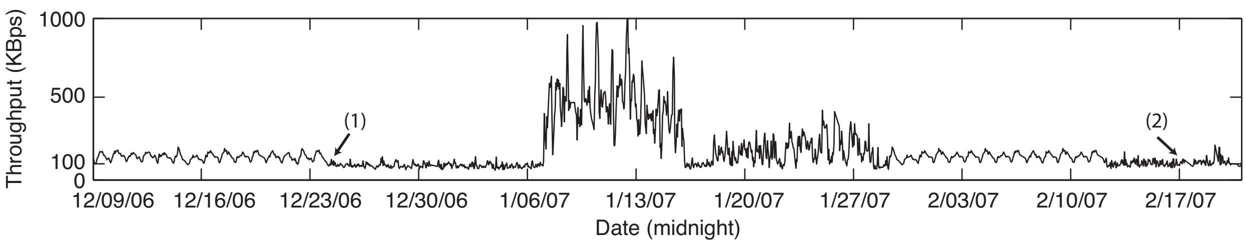

Fig. 13. Mean throughput evolution for intra-Netcom flows: (1) Taiwan earthquake, (2) Chinese New Year Eve.

those obtained for intra-ISP flows, revealing impact of inter-ISP peering. The unexpected discovery, that no apparent bandwidth limitations exist between the large and small ISPs, is quite interesting, especially if we consider the fact that large ISPs levy expensive bandwidth charges on small ones for relaying their traffic in both directions. This may be explained by that small regional ISPs have to rely on the large nationwide ISPs to deliver both their outbound and inbound traffic to and from the Internet.

Before we conclude this section, we add that besides regression study on the above snapshot on a regular day, we have also conducted regression analysis on snapshots during the Chinese New Year flash crowd scenario, and have observed similar correlations.

VI. Throughput Evolution: Time Series Characterization

With the one-time regression model derived, we now switch our focus to evolutionary characteristics of inter-peer bandwidth over time. Such an evolution of bandwidth is due to (1) the variation of the number of concurrent upload/download flows at the sender/receiver; and (2) the dynamics of cross traffic over the P2P links. Here, we are more concerned about the inter-peer bandwidth evolution caused by the latter. Based on the regression model we summarized in Sec. V, we are able to separate effects of the two causes, as the variation of coefficients in the linear models reflects the evolution of bandwidth availability over the P2P links when the per-flow end capacity is kept the same.

A. Intra-ISP throughput evolution

We now inspect the evolution of bandwidth availability over the internal P2P links of each ISP, by first studying the evolution of coefficients in the linear models, summarized with each of the continuous-time snapshots. Fig. 12 plots the evolution of regression coefficients during the week of December 17 — 23, 2006, again with the example of intra-Netcom flows.

Fig. 12 exhibits an apparent daily evolutionary pattern for y-intercept, whose value gradually increases at early hours of a day, peaks around $9-10$ am, and then drops and reaches the lowest around $10-11$ pm. The value of slope, although not as apparent, also shows a similar evolutionary pattern. Not shown in the figures is that p-values for testing the significance of slopes and y-intercepts are always asymptotically zero, exhibiting the significance of throughput regression for inter-Netcom flows at any time.

Though illustrated with the representative week only, the daily evolutionary pattern of regression coefficients — thus bandwidth availability on intra-Netcom P2P links — generally exists during the entire period of the traces. To validate this, we plot in Fig. 13 the evolution of mean throughput of intra-Netcom flows over more than 10 weeks of time, from December 10, 2006 to February 21, 2007.[3] To eliminate the effect of varying numbers of concurrent flows at the peers, the mean throughput at each time is calculated as the average of those flow throughputs with *PEC* in the range of $50-100$ KBps at that time. We observe a daily evolutionary pattern throughout the period, although it is more apparent on some days than others. Daily pattern aside, we also observe a few abrupt changes of the mean throughput *level* during this period: one around December 26, 2006, the date when an earthquake occurred in the Strait southwest of Taiwan, and another around January 8th, 2007. As the Taiwan earthquake damaged several undersea cables and disrupted some major overseas network connections of China, we conjecture that the first abrupt downgrade of bandwidth is caused by re-routing of traffic, which was originally directed towards overseas servers, to local servers, and the resulting tightness of bandwidths on local connections. We are not quite sure about the reason for the increase of throughput level around mid-January, while we conjecture that it might be caused by ISP upgrades, or measures taken by the ISP to counter the impact of the earlier earthquake. In addition, during the flash crowd scenario on Chinese New Year Eve, we observe no significant throughput downgrade.

Similar observations have been made during investigations of throughput evolution inside other ISPs. All these reveal that although the level of mean throughput may shift, the bandwidth availability on internal P2P links of an ISP statistically evolves following a daily pattern, which persists throughout the trace period.

B. Inter-ISP throughput evolution

In the inter-ISP case, we seek to answer the following questions: First, between Netcom and Telecom, does their inter-ISP peering always limit their inter-ISP P2P flow throughput? If so, is the bottleneck bandwidth availability varying at different times? Second, between the other ISP pairs, how does the bandwidth availability evolve over their inter-ISP links, and does per-flow end capacity always critically decide the throughput at any time?

With the example of the representative week, Fig. 14(A) reveals that, at most times of a day between Netcom and Telecom, P2P flow throughput is capped and is not correlated with per-flow end capacity, with a slope of approximately 0 and a corresponding p-value above 0.05. However, there does exist a certain period of time each day when throughputs are significantly correlated with PEC, usually in the early mornings, with slope values around 1 and corresponding p-values below 0.05. In addition, Fig. 14(B) exhibits daily evolutionary pattern for the y-intercept, which peaks at the time when the slope is well above 0 on a daily basis.

Based on the estimation algorithm of regression coefficients, we note that when the slope is non-significant, the y-intercept represents the mean throughput of the flows between the two ISPs; when the slope is significant, the throughput is decided by peer last-mile bandwidth availability, and does not show apparent inter-ISP peering caps. Therefore, the above observations reveal that: between the two largest ISPs, the limited bandwidth availability gradually improves at early times of a day, peaks in the morning when peer last-mile bandwidths come into play to decide the throughput, then drops and represents the lowest values in the evening.

Next, we inspect the throughput evolution between large ISPs and small ISPs. Fig. 15(A) exhibits that, for most of the time between Netcom and Tietong, the inter-ISP throughputs are significantly correlated with PEC, with non-zero slopes and near-zero p-values. Only occasionally at certain moments, there are observed drops of bandwidth availability, when the inter-ISP throughputs are limited regardless of the peer last-mile bandwidth availability. There also exists a daily evolutionary pattern for both the slope and y-intercept, similar to those in the previous cases, although not as apparent.

To validate the above observations in a longer period of time, we again plot the mean throughput evolution for Netcom→Telecom flows with PEC in the range of $10-60$ KBps in Fig. 16, and that for Netcom→Tietong flows with PEC in the range of $50-100$ KBps in Fig. 17. We also observe the rise of throughput levels in mid-January, but no apparent bandwidth downgrades around the earthquake scenario or flash crowd scenario on Chinese New Year Eve. Nevertheless, the daily evolutionary pattern of throughput persists at all times.

Similar observations have been made in investigations for other ISP pairs. Besides the daily throughput pattern, these observations also reflect that, no apparent inter-ISP bandwidth bottlenecks exist between a large ISP and a small one, and across small ISPs at most times. This again confirms that small ISPs do not usually impose low bandwidth caps at their peering point with large ISPs, in order to facilitate their traffic in both directions.

VII. Throughput Expectation Index: Application of Throughput Characteristics

The throughput characteristics we have derived in previous sections bring useful insights towards the improvement of current P2P streaming protocols. As an important application, we propose a *Throughput Expectation Index (TEI)*, to facilitate the selection of high-bandwidth serving peers.

For each pair of ISPs, as there exists a daily evolutionary pattern for each of the regression coefficients in its throughput model, we summarize a daily intercept function (y-intercept) and a daily slope function, $\beta_0(t)$ and $\beta_1(t)$, respectively, where

[3]Due to space limit, we are not able to show results obtained throughout the entire trace period, but choose this sufficiently long period of time which is representative of several typical scenarios.

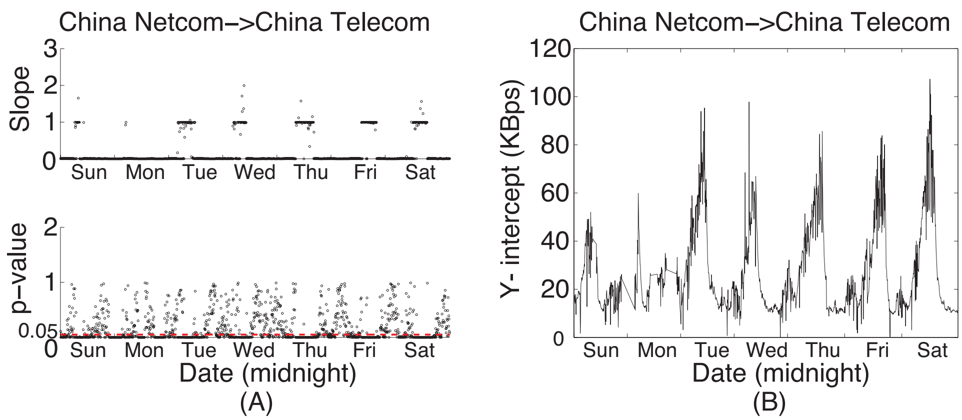

Fig. 14. Evolution of regression coefficients for Netcom→Telecom flows in the week of Dec. 17 — 23, 2006.

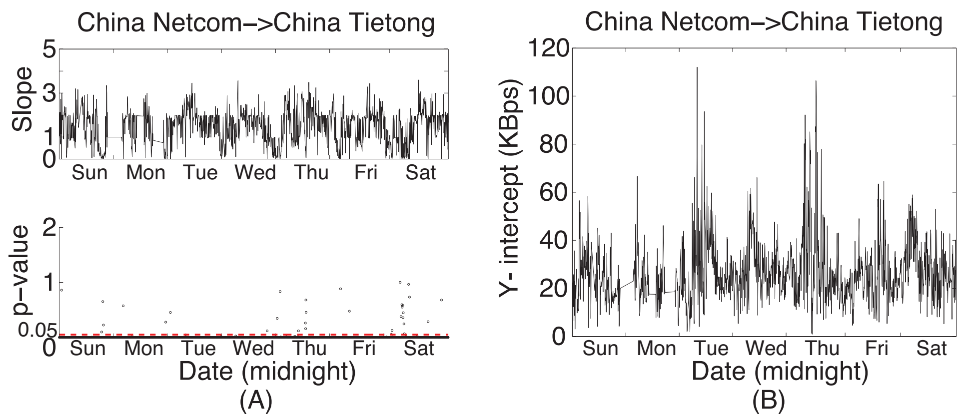

Fig. 15. Evolution of regression coefficients for Netcom→Tietong flows in the week of Dec. 17 — 23, 2006.

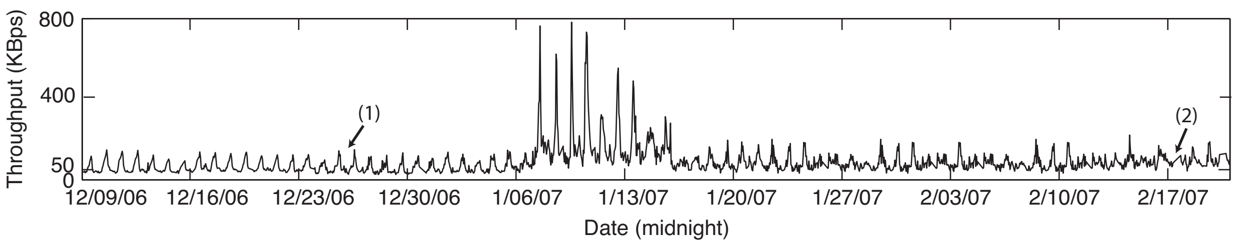

Fig. 16. Mean throughput evolution for Netcom→Telecom flows: (1) Taiwan earthquake, (2) Chinese New Year Eve.

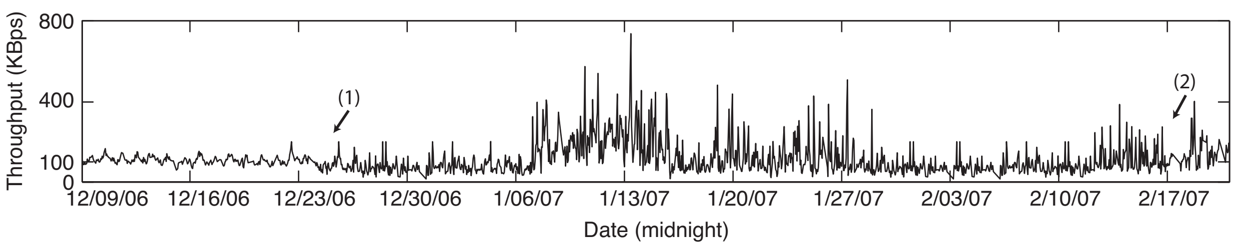

Fig. 17. Mean throughput evolution for Netcom→Tietong flows: (1) Taiwan earthquake, (2) Chinese New Year Eve.

t represents different times in a day, by taking the average of coefficient values at the same time on different days. For example, Fig. 18 depicts the daily intercept and slope functions for intra-Telecom flows, summarized by averaging coefficients at the same hour during the week of December 17 — 23, 2006. We then define the following throughput expectation index:

$$TEI = \beta_0(t) + \beta_1(t) \times PEC(t). \tag{2}$$

TEI approximates the achievable inter-peer bandwidth between two peers across two ISPs (including two identical ISPs) at a specified time of a day. The computation of TEI not only captures all the deciding factors of inter-peer bandwidth — upload/download capacities at the upstream/downstream peer, concurrent upload/download load at the upstream/downstream peer, and the ISPs both peers belong to — but also considers the temporal evolution of bandwidth availability at different times of a day. Therefore, it can be effectively utilized in peer selection at each peer, by *ranking* the candidate serving peers

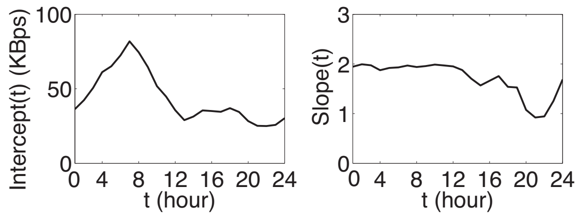

Fig. 18. Daily intercept/slope functions for intra-Telecom flows.

based on the computed TEI towards each of them. In more details, the TEI-assisted peer selection proceeds as follows:

The P2P streaming service provider derives the intercept and slope functions for each pair of ISPs, using the collected peer reports over a certain number of days (*e.g.*, one week). Upon bootstrapping a new peer, the intercept and slope functions of relevant ISP pairs, from each of the other ISPs to the ISP the peer belongs to, are loaded onto the peer. During the peer selection process, the peer obtains the following information from each of its candidate serving peers: IP address, upload capacity and the number of current upload flows. Then the peer calculates the per-flow end capacity of the potential P2P link from the candidate to itself, decides the intercept and slope functions to use (from its pre-loaded functions) by mapping the IP addresses of the candidate and itself to ISPs, and computes the TEI towards this candidate with y-intercept and slope values at the current moment. The peer ranks all candidate peers based on their derived TEIs. Then when the peer is deciding which media block to download from which candidate peer based on the exchanged buffer maps, it maximally retrieves available media blocks from the peers with the highest ranks.

Similar usage of TEI can be applied at a tracking server to select the best serving peers for a requesting peer. Note that such peer selections are performed without any intrusive measurements. Only a small number of intercept and slope functions need to be pre-loaded onto the peer, and a limited amount of information needs to be acquired from neighboring peers.

We further emphasize that in TEI-assisted peer selection, it is the relative *ranks* of peers computed by TEIs that are being used, instead of the absolute throughput values estimated with TEIs. This is because throughput levels may vary from day to day, but the daily throughput pattern persists for each ISP pair, and therefore the relative rank of peers may persist as well at a specified time on different days. This allows us to use the summarized intercept/slope functions and PEC values of end peers at a specified time to calculate the relative throughput ranks at the time.

To investigate the accuracy of the proposed TEI, we conduct a number of cross-validation experiments, by using intercept/slope functions summarized from the representative week (December 17 — 23, 2006) in TEI-assisted peer selection throughout the trace period. At each peer that appeared in the traces, we calculate the TEI towards each of its sending partners, and then compare their ranks computed by the TEIs with their *true ranks* based on the actual TCP throughput over the links. The experiments are divided into two parts.

First, we investigate the true rank of the best sending peer selected with TEI at each peer. Focusing on one snapshot at 9pm, December 18, 2006, Fig. 19(A) shows the distribution of this true rank at all the existing peers. At 70% of the peers, the TEI best peer coincides with the actual best sending peer with the largest throughput; at the majority of all peers, the TEI best peer ranks among top 3. Fig. 19(B) plots the evolution of the percentages of peers, at which the TEI best peer has a true rank no larger than 2 or 3, over the 10-week period of time. We observe that the former case achieves a peer percentage higher than 80% at all times, and the latter is consistently around 93%. During the throughput level shift around the earthquake scenario and the flash crowd scenario near Chinese New Year, the percentages represent larger fluctuations, but are nevertheless quite satisfactory as well.

Next, we compare the sum of throughputs on P2P links from two peer groups at each receiving peer: (1) the 5 top sending peers selected with TEI, and (2) the true top 5 peers with largest throughputs. Fig. 20(A) shows the distribution of throughput difference between the two groups, at peers that existed at 9pm, December 18, 2006. The TEI selection achieves less than 10 KBps throughput difference at more than 70% peers. Furthermore, the difference is shown to be larger around the time of the two special scenarios in Fig. 20(B), and nevertheless, it is minor at most peers at most regular times, *i.e.*, 75% peers are subjected to a difference less than 20 KBps.

The above results exhibit that, while we are using intercept/slope

functions summarized from only one week, the peer ranking mechanism of TEI

works quite well throughout the trace period. This reflects the practical usefulness of TEI in capturing the persistent *relative* ranks of inter-peer bandwidths at each specific time on different days, without the need of intensive training using a large amount of historical data. A more elaborate usage of TEI may involve the retraining of the intercept/slope functions over time at the P2P streaming service provider, based on feedbacks from the peers about the accuracy of the current functions in evaluating high-bandwidth peers. As our goal is to show one effective application of our derived P2P streaming flow characteristics, we choose not to go into details of such practical protocol design.

VIII. Concluding remarks

This paper represents the first attempt to characterize inter-peer bandwidth availability in modern large-scale peer-to-peer streaming networks. With abundant traces from a successful commercial P2P streaming application, UUSee, and using statistical means, we explore the critical factors that may determine such achievable bandwidth, from both the end-host and ISP/area perspectives. In addition, we also explore the evolution of such bandwidth over time. Among our many discoveries, we point out that the inter-ISP peering bandwidth constraints significantly limit the achievable streaming bandwidth only between the two largest China ISPs, and are less perceivable across other ISPs. In the latter case and in the case that both peers are in the same ISP, the streaming bandwidth

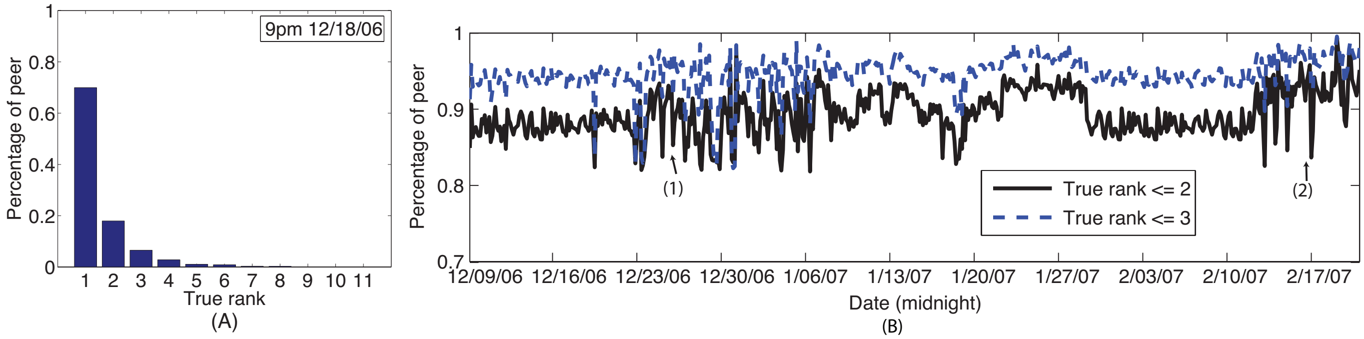

Fig. 19. True rank distribution of the best sending peer selected with TEI. (1) Taiwan earthquake, (2) Chinese New Year Eve.

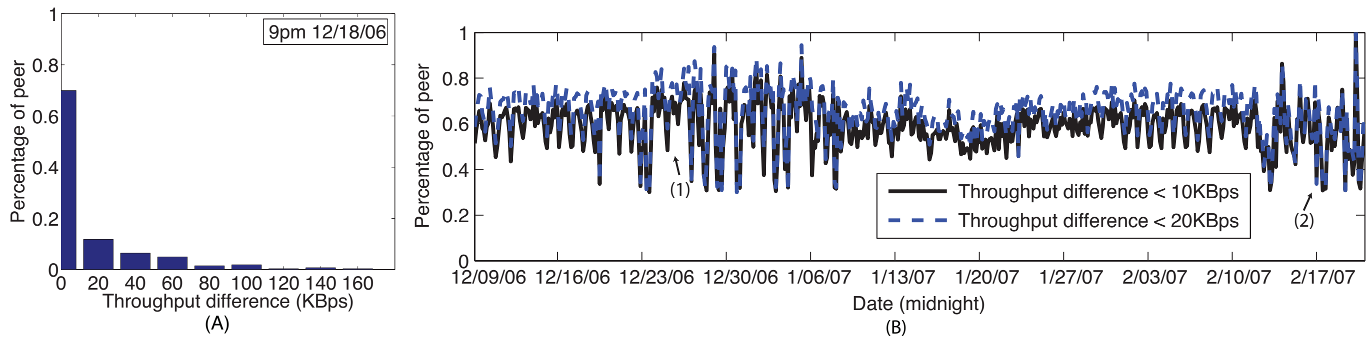

Fig. 20. Distribution of throughput difference between flows from 5 top peers selected with TEI and flows from true top 5 peers. (1) Taiwan earthquake, (2) Chinese New Year Eve.

is significantly correlated with last-mile bandwidth availability at the end peers. Another useful discovery is that inter-peer bandwidth exhibits an excellent daily evolutionary pattern within or across most ISPs, which we make use of in designing a throughput expectation index to achieve efficient peer selection based on bandwidth. We believe that our findings bring important insights towards a complete understanding of achievable bandwidth in practical P2P streaming applications, and will be instrumental towards further improvements of P2P streaming protocol design without active and intrusive measurements.

REFERENCES

[1] M. Mathis, J. Semke, and J. Mahdavi, "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm," *Computer Communications Review*, vol. 27, no. 3, 1997.

[2] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP Throughput: A Simple Model and its Empirical Validation," in *Proc. of ACM SIGCOMM 1998*, September 1998.

[3] B. Sikdar, S. Kalyanaraman, and K. Vastol, "Analytic Models for the Latency and Steady-State Throughput of TCP Tahoe, Reno and SACK," *IEEE/ACM Transactions on Networking*, vol. 11, no. 6, pp. 959–971, 2003.

[4] R. Wolski, "Dynamically Forecasting Network Performance using the Network Weather Service," *Cluster Computing*, vol. 1, no. 1, pp. 119–132, 1998.

[5] S. Vazhkudai, J. Schopf, and I. Foster, "Predicting the Performance of Wide Area Data Transfers," in *Proc. of the 16th International Parallel and Distributed Processing Symposium (IPDPS 2002)*, 2002.

[6] Q. He, C. Dovrolis, and M. Ammar, "On the Predictability of Large Transfer TCP Throughput," in *Proc. of ACM SIGCOMM 2005*, August 2005.

[7] D. Lu, Y. Qiao, P. A. Dinda, and F. E. Bustamante, "Characterizing and Predicting TCP throughput on the Wide Area Network," in *Proc. of the 25th International Conference on Distributed Computing Systems (ICDCS 2005)*, June 2005.

[8] H. Balakrishnan, M. Stemm, S. Seshan, and R. H. Katz, "Analyzing Stability in Wide-Area Network Performance," in *Proc. of ACM SIGMETRICS 1997*, 1997.

[9] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker, "On the Characteristics and Origins of Internet Flow Rates," in *Proc. of ACM SIGCOMM 2002*, 2002.

[10] K. Lakshminarayanan and V. N. Padmanabhan, "Some Findings on the Network Performance of Broadband Hosts," in *Proc. of Internet Measurement Conference (IMC)*, October 2003.

[11] S. Sen and J. Wang, "Analyzing Peer-to-Peer Traffic Across Large Networks," *IEEE/ACM Transactions on Networking*, vol. 12, no. 2, pp. 219–232, 2004.

[12] S. Saroiu, K. Gummadi, and S. Gribble, "Measuring and Analyzing the Characteristics of Napster and Gnutella Hosts," *ACM/Kluwer Multimedia Systems Journal*, vol. 9, no. 2, pp. 170–184, 2003.

[13] X. Hei, C. Liang, J. Liang, Y. Liu, and K. W. Ross, "A Measurement Study of a Large-Scale P2P IPTV System," *IEEE Trans. on Multimedia (to appear)*.

[14] 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.

[15] 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.

[16] ——, "Measuring P2P IPTV Systems," in *Proc. of the 17th International workshop on Network and Operating Systems Support for Digital Audio & Video (NOSSDAV'07)*, June 2007.

[17] J. Liang, R. Kumar, and K. Ross, "The FastTrack Overlay: A Measurement Study," *Computer Networks*, vol. 50, no. 6, pp. 842–858, April 2006.

[18] 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.

[19] 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.

[20] 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*, vol. 6, no. 1, 2002.

[21] 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.

[22] 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.
[23] 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.
[24] S. A. Baset and H. Schulzrinne, "An Analysis of the Skype Peer-to-Peer Internet Telephony Protocol," in *Proc. of IEEE INFOCOM*, April 2006.
[25] 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.
[26] 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.
[27] "UUSee Inc." http://www.uusee.com/.
[28] "Log-normal distribution." http://mathworld.wolfram.com/LogNormalDistribution.html.
[29] R. Christensen, *Analysis of Variance, Design, and Regression: Applied Statistical Methods*. Chapman and Hall/CRC, 1996.
[30] J. D. Gibbons and S. Chakraborti, *Nonparametric Statistical Inference, 4nd edition*. Marcel Dekker, 2003.
[31] D. C. Montgomery, e. A. Peck, and G. G. Vining, *Introduction to Linear Regression Analysis, Third Edition*. John Wiley and Sons, Inc., 2001.
[32] T. P. Ryan, *Modern Regression Methods*. John Wiley and Sons, Inc., 1997.

Chuan Wu. Chuan Wu received her B.Engr. and M.Engr. degrees from Department of Computer Science and Technology, Tsinghua University, China, in 2000 and 2002, respectively. She is currently a Ph.D. candidate at the Department of Electrical and Computer Engineering, University of Toronto, Canada. Her research interest is distributed algorithm design to improve Quality of Service of overlay multicast applications. She is particularly interested in applying theories (optimization, game theory, statistics) to guide practical protocol design.

Baochun Li. Baochun Li received his B.Engr. degree in 1995 from Department of Computer Science and Technology, Tsinghua University, China, and his M.S. and Ph.D. degrees in 1997 and 2000 from the Department of Computer Science, University of Illinois at Urbana-Champaign. Since 2000, he has been with the Department of Electrical and Computer Engineering at the University of Toronto, where he is currently an Associate Professor. He holds the Nortel Networks Junior Chair in Network Architecture and Services since October 2003, and the Bell University Laboratories Chair in Computer Engineering since July 2005. In 2000, he was the recipient of the IEEE Communications Society Leonard G. Abraham Award in the Field of Communications Systems. His research interests include application-level Quality of Service provisioning, wireless and overlay networks. He is a senior member of IEEE, and a member of ACM.