Topic Peer-to-Peer Networks

Georgi Georgiev
(georgiev@in.tum.de)

Seminar "'Internet Measurement"' ,
Technical University Munich

SS /2005 (Version from July 24, 2005)

Abstract

This paper contains an analysis of the P2P(Peer-To-Peer) traffic of three popular P2P systems - FastTrack, Gnutella, and Direct-Connect, done by measuring the flow-level information collected at multiple border routers across a large ISP network. It also characterizes the query behaviour of the peers in the P2P network Gnutella. The P2P traffic is measured at three levels of granularity - IP address, prefix and AS (Autonomous System). To analyze the P2P traffic the host distribution, the volume of the traffic, the connectivity of the hosts, their connection duration and the bandwidth usage are being measured. The passive peers are also observed, although they produce no queries. They are important part of the traffic characterization, because they are part of the overlay network that forwards and respond to queries. The active peers' session duration is also analyzed. Its characteristic is composed of the number of queries issued in a session, the time between the establishment of the connection and the sending of the first query, the time between the sending of two successive queries and the time after sending the last query until termination of the connection. At the end the paper compares the P2P traffic with the web and the overall traffic.

1 Introduction

A peer-to-peer (or P2P) computer network is a network, that relies on computing power at the ends of a connection rather than in the network itself. A pure P2P file transfer network does not have the notion of clients or servers, but only equal peer nodes, that simultaneously function as both clients and servers to the other nodes on the network. P2P networks are used for sharing content like audio, video data or anything in digital format, but also for gaming and instant massaging, which makes them one of the most popular Internet applications recently. The analysis of the P2P traffic is done by measuring flow-level information collected at multiple border routers across a large ISP network.The research is based on the three popular P2P systems - FastTrack, Gnutella, and Direct-Connect. The characterization of the P2P traffic is done by observing a single ISP and its impact on the underlying network. The distribution of the traffic across the network at different levels of spatial aggregation (IP, prefix, AS) is very asymmetric. The P2P network behavior is dependent on the time of the day, the geographical region and other factors, which make it so dynamic. Still, the P2P traffic is more stable than the corresponding distribution of the Web traffic and the overall traffic. An important advantage of the P2P networks is that the bandwidth of all clients can be used, so the total bandwidth - and usually the available download bandwidth for the average user grows with the number of nodes, instead of all clients having to share the bandwidth of one server, where adding more clients could mean slower data transfer for all users.

2 P2P applications

FastTrack, Gnutella, and Direct-Connect are three of the most common P2P networks. Nullsoft's Gnutella is a P2P network without a central server. In this network every user is a node. Every time, when a node connects another one, the second node will send to the first one a list of working nodes. By searching the second node forwards the request to the other nodes, with which it is connected, and they in turn forward the request, and so on. In theory, the request will eventually find its way to every user on the Gnutella network. In practice, searching on the Gnutella network is often unreliable. Each node is a regular computer user, which are constantly connecting and disconnecting, so the network is never completely stable. Since individual users' bandwidth are likely to be limited, some search requests may be dropped before they reach the whole network. As a result most queries will never reach more than half of the network. The real benefit of Gnutella's node model is, that the network can't be shut down as long as there are at least two users. The most popular Gnutella clients are LimeWire, BearShare and Shareaza.

FastTrack is based on the Gnutella protocol and extends it with the addition of supernodes to improve scalability. The supernode functionality is built into the client; if a powerful computer with a fast network connection runs the client software, it will automatically become a supernode, effectively acting as a temporary server for other, slower clients. To allow downloading from multiple sources, FastTrack employs the UUHash hashing algorithm. As of early 2003, FastTrack was the most popular P2P network. It has such clients as Kazaa and iMesh.

The Direct-Connect network is not as decentralized as Gnutella or FastTrack, since it uses hubs, which connect a group of users.

3 Methodology

3.1 Measurement approach

According to [1], the measurements must be efficient, should catch the system dynamics in sufficient detail and should not affect the system. Because of the decentralization of the P2P network, the large number of hosts involved and the dynamic nature of the peer membership, it is a challenge to characterize and measure the P2P traffic and its complex behavior. Each P2P system is measured here at three levels of granularity - IP address, prefix and AS (Autonomous System). IP level information provides fine observing model of the load distribution across the network. It gives us information about the overall traffic distribution. Prefixes are the unit of routing at the IP layer, so understanding traffic at this level is important for ISP traffic engineering. Also, the prefix level aggregation, by grouping IP addresses that are topologically close together from a network routing viewpoint, enables capturing locality characteristics in the P2P system. At the AS level, we identify an AS by its unique public AS number.

3.2 Passive network measurement approach vs. active probing

The passive network measurement does not require knowledge about the P2P protocol, beyond port number information. Also it is a non-intrusive method and all the traffic data can be collected without interfering with or impacting the peers. This way we can get a more complete view of the P2P host distribution and their traffic patterns. The passive network measurement gathers information on both the P2P signaling traffic and the actual data download traffic. This is a major priority of the method over the active probing, because the P2P systems are used to download large files and it is important to be able to analyze the actual data traffic. The Passive network measurement allows localized analysis, which is important for the local traffic engineering and provisioning at an ISP. The issues of the passive measurement are first that the data is aggregated at the flow-level. There is no possibility to obtain application-level details, such as the actual P2P messages exchanged between peers, or the specific files, being requested and actually downloaded. It is also impossible to capture the complete flow of traffic and due to the asymmetric IP routing, we may see only one direction of the traffic between a given pair of hosts.

4 Characterization metrics

We want to characterize the P2P systems behavior, to understand how these systems affect the underlying network, and to gain ideas of developing P2P systems with superior performance. To reach our goal we measure the distribution of the P2P hosts across the network, the traffic volume transmitted between them, the connection duration, the mean bandwidth usage for each host and their behaviour during time.

4.1 Host distribution

We analyze the distribution of the P2P hosts across the network. Measuring the number of unique IP addresses, prefixes, and ASes participating in each P2P system in each one-day period across several weeks spread over several months, makes us able to indicate the trends in the size of the P2P systems. Comparing the measurement results at different levels of topological granularity, we can infer locality characteristics of the P2P hosts distribution, which can be used in traffic engineering and P2P architecture design.

4.2 Traffic volume

As the P2P systems are mainly used for sharing large audio/video files or software, we measure the traffic volume transmitted between P2P hosts, and compute the overall data transmitted or received by each IP address, prefix and AS per day.

4.3 Host connectivity

For each aggregation level, we compute the total number of unique entities, that it communicates with (either transmits to or receives data from). The resulting distribution is used to characterize the host connectivity in the P2P network.

4.4 Traffic pattern over time

We form a traffic pattern over time, analyzing the number of hosts participating in the P2P system at a given time and the traffic volume transferred among them during this time. For our goal we divide the entire data set into small time bins. For each bin, the number of unique entities (IPs, prefixes, or ASes) that are participating in the P2P system and the traffic volume transferred among them are computed.

4.5 Connection duration

Connection duration means, how long a host stays in the P2P system. The length of the host connection duration tells us how long a host stays in the P2P system, once it joins. The on-time of a host means how long a host stays in the P2P system during a certain period of time. It is computed as the sum of all the connection durations over a given time period.

StartTime(c) = $\min_{f \in c} StartTime(f)$
FinishTime(c) = $\max_{f \in c} FinishTime(f)$
Volume(c) = $\sum_{f \in c} Volume(f)$
c - our connection
StartTime() - the time, when the host starts to send or receive data
FinishTime() - the time, when the host finishes sending or receiving data
f - concurrent flows associated with the connection

4.6 Mean bandwidth usage

Finding the mean bandwidth usage would enable network administrators to understand the bandwidth demand that hosts, running P2P applications might impose on the network. The upstream and downstream bandwidths for each host are measured separately. For a host h we measure the mean bandwidth usage like this :

Bandwidth$_r(h) = \frac{Volume_r(h)}{OnTime_r(h)}$,

where r is upstream or downstream and OnTime is the total time of the observed transmission.

5 P2P Traffic

The Analysis of the P2P traffic is based on the flow records from multiple border router's interfaces across the ISP backbone. For each P2P system, records that matched the corresponding default application ports (source or destination), involving TCP traffic are extracted. To ensure the exact results the invalid IPs, no matched prefixes in routing tables, and invalid AS numbers are ignored. All the IP addresses that are in the following ranges : 10.0.0.0 - 10.255.255.255, 172.16.0.0 - 172.31.255.255, and 192.168.0.0 - 192.168.255.255, also all the AS numbers between 64513 and 65535 are considered as invalid. This way 4% of the captured flows are eliminated. As a result the final dataset consists of around 800 million flow records.

Date	Protocol	Number of records	Total number of unique IPs	Number of unique IPs per day	Total traffic (GBytes/day)	Traffic volume per IP (MBytes/day)
9/10/2001-9/15/2001	Gnutella	37,853,281	718,464	197,445	211	2.2
	FastTrack	110,533,024	3,403,900	998,669	773	1.6
	DirectConnect	595,606	22,852	6,244	48	15.4
10/9/2001-10/13/2001	Gnutella	49,649,348	823,532	247,114	272	2.2
	FastTrack	184,113,038	4,450,149	1,485,370	1,153	1.6
	DirectConnect	566,740	23,211	7,193	56	15.6
12/10/2001-12/16/2001	Gnutella	69,578,723	887,520	236,954	242	2.0
	FastTrack	340,690,074	5,924,072	1,934,460	1,776	1.8
	DirectConnect	701,712	29,925	7,213	71	19.6

Table 1: Netflow data set of P2P traffic over TCP

The Table above shows us that FastTrack, with clients like Kazaa and I-Mesh, is the most popular P2P system followed by the open-source Gnutella. The newer system DirectConnect has a smaller user base, but the DirectConnect hosts tend to stay active longer than FastTrack and Gnutella and also have higher average bandwidth than the other two concurrent networks. Of course the above table is out of date, as it shows the state of the three P2P networks at the end of 2001.

5.1 Fraction of passive peers

According to [2], passive peers are connected in the overlay network, but perform no queries. These peers constitute an important component of a realistic workload, because they dont generate any query load, but they form part of the overlay network that forwards and respond to queries. We count the number of peer sessions that begin in a 1-hour interval that issue no queries and calculate the ratio to all sessions that start in the same hour. The average for each 1-hour interval is computed over the entire measurement period.

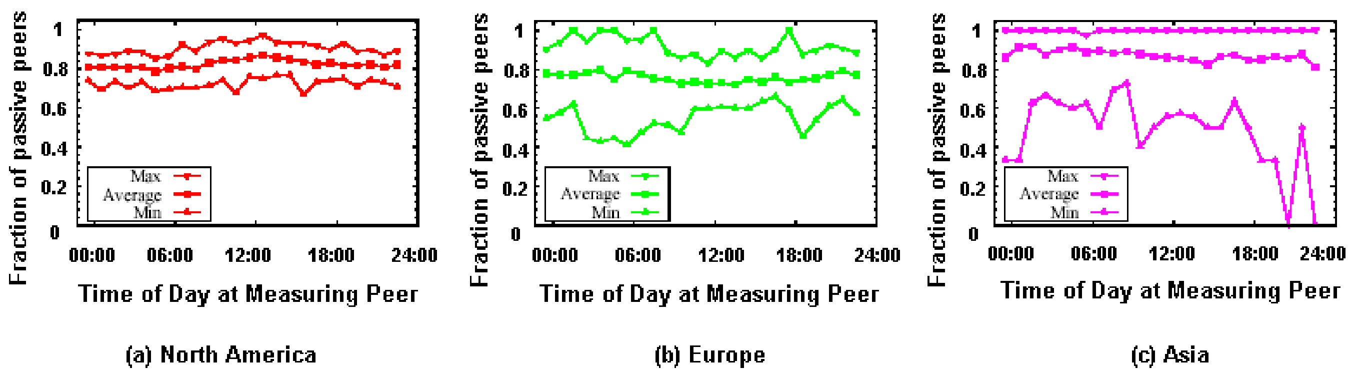

Figure 1: Fraction of connected peers that are passive.

We observe that the fraction is almost the same for each geographical region, with about 80% to 85% for North America, 75% to 80% for Europe, and 80% to 90% for Asia. The fraction of passive peers varies only by about 5% over time of day.

5.2 Active peer session characteristics

In [2] is mentioned, that the connected session duration for active peers is a measure, composed of the number of queries issued in a session, the time between the establishment of the connection and the sending of the first query, the time between the sending of two successive queries and the time after sending the last query until termination of the connection.

5.2.1 The number of queries in a session

The Fraction of peers that issue less than 5 queries is 92% for Asia, 80% for North America, and only 70% for Europe. 5% of the Asian, 8% of the North American 13% of the European peers issue 5 to 10 queries. The sessions with many queries comprise 3% of the session in Asia, 10% in North America and 13% in Europe. We can conclude that European peers issue significantly more queries in a session than peers from the other geographical regions.

5.2.2 The Distribution of time until first query for active sessions

Figure 2 (a) shows the time until the first query after the connection establishment of a session broken down by geographical region. The curves for North American and European peers are very similar, but significantly different from the Asian peers. The first query within a session from Asian peers is issued within 10 seconds for 10% of the peers, whereas this fraction constitutes 20% for North America and Europe. The fraction of peers that issue a query within 30 seconds stays with 40% almost equal for all regions. Another 50% of the Asian peers issue the first query within 30 and 90 seconds. The same fraction of peers issues the first query within 30 and 1,000 seconds for Europe, indicating a significant correlation to geographical region. Furthermore, a fraction of 1% of the sessions started in North America and Europe issue the first query after 80,000 seconds, i.e., after more than 20 hours.

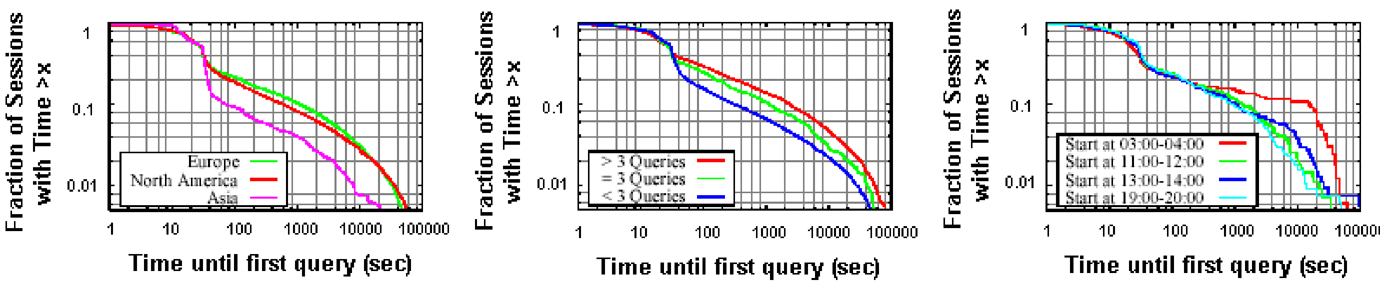

(a)Each geographic region (b)Conditioned on session length, peers in North America (c)Each Key period peers in Europe

Figure 2: Distribution of time until first query

To analyze correlations between the time until first query and the number of queries issued in a session, Figure 2 (b) shows sessions with less than 3 queries, exactly 3 queries and more than 3 queries for North American peers. This figure shows that the conditional distributions are equal for 50% of the sessions with an early first query, while there is a significant difference for the 50% of the session with late first query. In 90% of the sessions with less than 3 queries the first query is issued before 200 seconds, in the sessions with exactly 3 queries before 1000 seconds and in sessions with more than 3 queries before 2,000 seconds. We conclude from Figure 2 (b) that for North America the time until first query is correlated with the session length in number of queries.

We find that in sessions started in the non-peak hours in a significant fraction of the sessions the first query is sent 10000 seconds and more after a session starts. These fractions constitute 10% for Europe. The same trend can be observed for the other geographical regions. We conclude from Figure 2 (c) that there is a significant correlation between time of day and the time until first query.

5.2.3 Distribution of time between queries for active sessions

Figure 3 (a) plots the significant correlation between the query interarrival time and geographical region. This figure shows, that queries generated by European peers have shorter interarrival times than the peers from the other two regions. For instance, the fraction of interarrival times below 100 seconds constitutes 90% for Europe 80% for Asia and 70% for North America.

Sessions of European peers with many queries have smaller interarrival times than sessions with few queries as can be seen in Figure 3 (b). This shows that there is a difference in the environment of North American and European peers like the prizing model of the Internet service providers. Due to this difference sessions of North American peers with many queries tend to connect for a longer period compared to similar sessions of European peers. We conclude that the query interarrival time has to be conditioned on the number of queries per session for European peers but not for North American peers.

To analyze the correlation to time of day, Figures 3 (c) plots the dependence of query interarrival time on the important daily time periods for European peers. It show that queries issued in peak hours (all daily periods except 03:00-04:00) have longer interarrival times than queries issued in non-peak hours. For example, 94% of the queries issued in Europe between 3:00 and 4:00 have an interarrival time below 100 seconds, while this fraction is only 85% for sessions starting between 11:00 and 12:00. Results for the other geographical regions are identical. The query interarrival time shows a significant correlation also to the time of day.

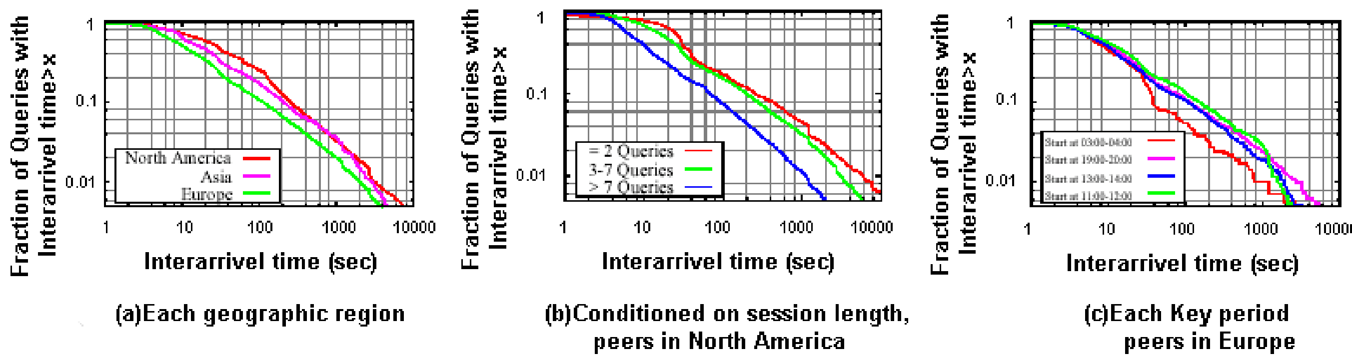

Figure 3: Distribution of time Between queries

5.2.4 Distribution of time after last query for active sessions

Figure 4 (a) shows, that there is a significant correlation between time after last query and geographical region. The distributions are very similar for North American and European peers, while Asian peers tend to close sessions much faster. The fraction of sessions with a time after last query of more than 1000 seconds is 20% for Europe and North America, while it is only 10% for Asia.

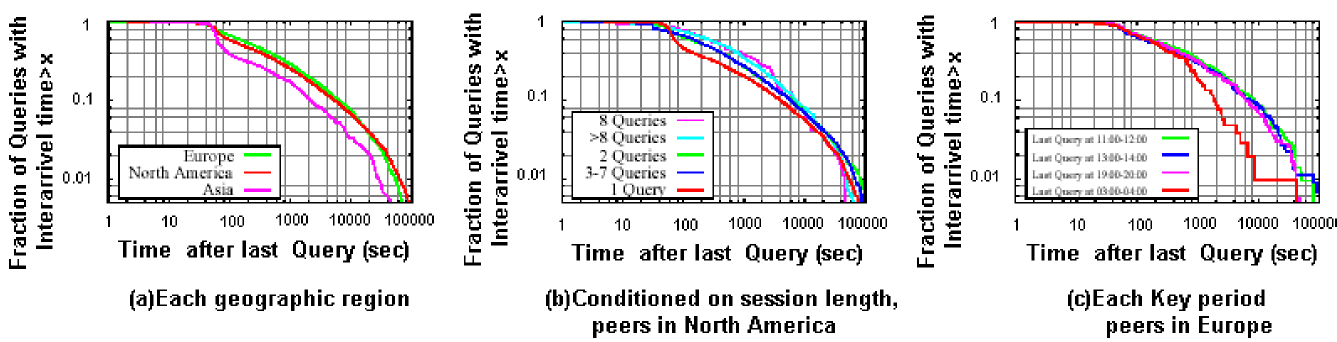

Figure 4: Distribution of time after last query

We conclude from Figure 4 (b) that the distribution of time after last query must be conditioned on number of queries per session. The correlation between time after last query and the number of queries per session is shown in Figure 4 (b). We observe the smallest and greatest values for the time after the last query for sessions with a single query, and with 8 and more queries, respectively. The conditional distributions for 2 queries and 3 to 7 queries are identical for 99% of the sessions, similar to the curves for exactly 8 and more than 8 queries. Combining both distributions of these pairs, we observe a positive correlation between time after last query and number of queries per session for 90% of the sessions.

Analyzing the correlation to time of day for European peers in Figure 4 (c), we find that sessions sending the last query in the nonpeak hours have a shorter time after last query than sessions sending the last query in peak hours. This trend is most noticeable in Europe, but we see it also in North American.In Europe the time after the last query for sessions sending the last query between 03:00 and 04:00 is below 10,000 seconds for more than 99% of the sessions, while it is below 91% of the sessions sending the last query at other times. We conclude from Figure 4 (c) that time after last query is significantly depends on the time of day.

5.3 Host distribution

The number of IP addresses participating in FastTrack each day ranges from 0.5 million to 2 million, according to [1]. The number of unique prefixes ranges from 17000 to 26000, and the number of unique ASes - from 4000 to 5500. The number of unique active IP addresses belonging to of prefixes in FastTrack is 4 times better that by Gnutella and 16 times than DirectConnect. That tells us, that FastTrack hosts have better potential to find nearby peers, and that most of the queries can be resolved locallv.

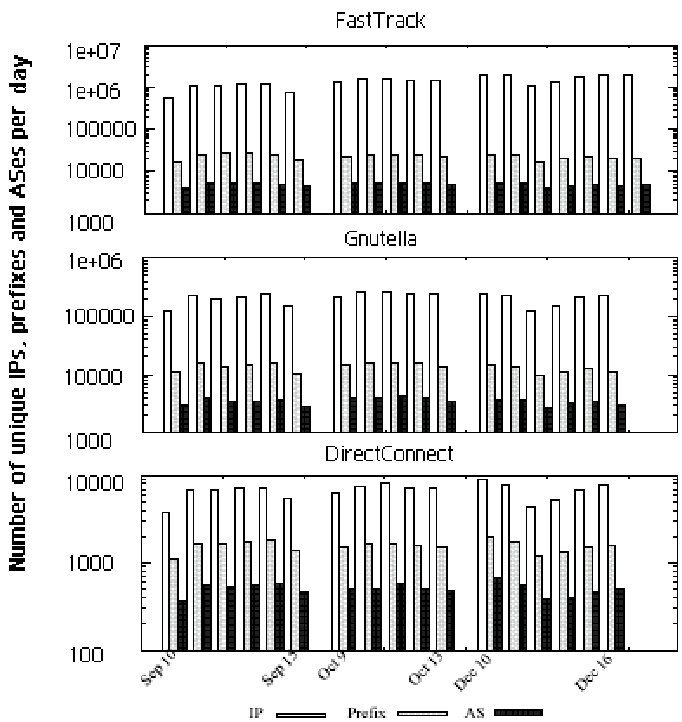

Figure 5: The distribution of the hosts participating in three P2P systems per day(the y-axis is in logscale).

5.4 Query popularity distribution

About 80% of the days the number of top 10 queries that is found in the top 100 on the subsequent day is not larger than 4, indicating a significant hot set drift. Even the top 100 queries change significantly from day to day. The query popularity distribution cannot be calculated over the entire trace, since the hot set drift must be considered.

In addition to temporal influences, we conjecture that the query popularity distribution depends on the geographical location of peers. To confirm this conjecture, we determine the set of distinct queries issued by North American, European and Asian peers, subsequently, for periods of length N=1,2, and 4 days. Furthermore, we determine the pair-wise intersection between the query sets and the intersection of all three sets. The cardinalities of the sets for typical periods are shown in Table 2.

Measure	4-Day Period	2-Day Period	1-Day Period
Number of Different Queries from North American Peers	6106	3588	1990
Number of Different Queries from European Peers	5382	3729	1934
Number of Queries in Intersection Set between North American an European Peers	323	114	56
Number of Queries in Intersection Set between North American and Asian Peers	41	15	5
Number of Queries in Intersection Set between European and Asian Peers	28	10	5
Number of Queries in Intersection Set between North American, European, and Asian Peer	17	4	2

Table 2: Query Class Sizes

We note that the cardinality of the intersection between the query sets of North American and European peers is about 2.8% of the cardinality of the North American set and the European set for a single day. Even for a 4-day period, the cardinality of the intersection is not larger than 6%. The relative cardinality of the intersection of the query sets from all three continents is about 0.001% and 0.02% for all geographical regions and periods. We conclude from Table 2 that peers from different geographical regions issue different queries. There is a small intersection, but it should be considered in an accurate workload model.

5.5 Host Connectivity

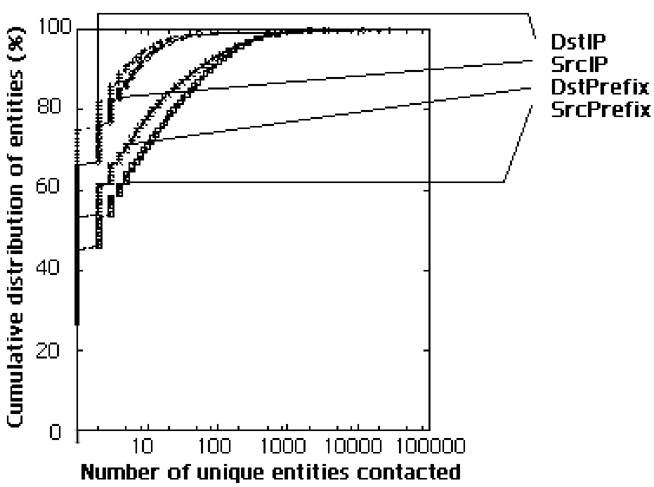

Figure 6: The cumulative distribution of network connectivity at the IP and network prefix (PR) levels, for hosts participating in FastTrack on September 14, 2001.

We observe here the FastTrack system. Because of its specific node structure, using supernodes, 48% of the individual IPs communicate with at most one IP, and 89% with at most 10 other IP addresses. Only 1% communicates with more than 80 IPs. The distribution is less skewed for the network prefix and AS level connectivity. 75% of the prefixes communicate with at least 2 prefixes, and the top 1% of the prefixes talk with at least 1000 prefixes. 80% of the ASes communicate with multiple ASes, and the top 1% of the ASes communicate with at least 476 other ASes.

6 P2P system dynamics

The P2P networks are highly dynamic systems. Their traffic is dependent on many factors, like the number of the connected hosts involved, the duration and the activity of their connection, the time of the day, even of the week, the geographical region, etc.

6.1 Traffic pattern over time

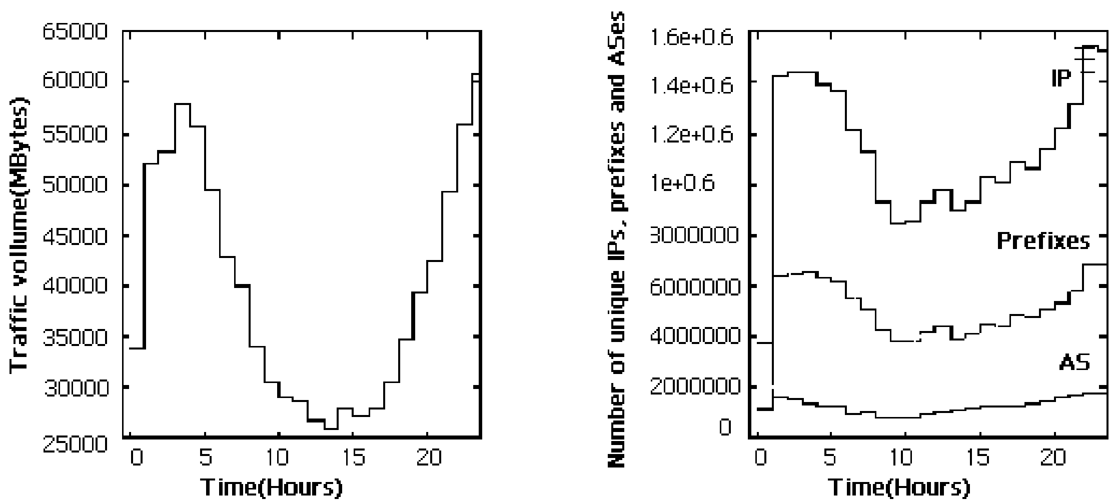

Figure 7: The distribution of number of IP addresses and traffic volume across hours in FastTrack on September 14, 2001. (a) The traffic volume transferred in each bin. (b) The number of unique IP addresses, network prefixes, and ASes that are active in each bin.

In [1] is measured how many hosts are active in the P2P systems during a certain time period. The traffic pattern is characterize, using discrete time bins. The bin size vary from 15 minutes to 1 hour. We compute the number of unique IP addresses, network prefixes, and ASes observed with each bin and the corresponding traffic volume transferred.

Observing Figure 7, we can conclude that the FastTrack hosts are less active in the early morning, but they are increasing their activity later in the day and reaching their peak between midnight and early morning. The traffic volume shows also the same correlation with the time of the day. The most of the traffic is concentrated in the hours between 0 and 5. It reaches its lowest point in the middle of the day and increases in the late hours.

6.2 Host connection duration and on-time

60% of the IP addresses, 40% of network prefixes and 30% of the ASes stay in FastTrack for 10 minutes or less per day. The P2P system is much less transient at the prefix and AS aggregation levels. 65% of the IP addresses join FastTrack only once. The most of the connections are very short, because the vast majority of the data transfer events are queries and responses. Although the distribution of the number of connections of the hosts in all three P2P systems are similar, hosts are staying longer in Direct-Connect than in the other two systems, which contributes to the more traffic volume, distributed by the individual Direct-Connect hosts.

6.3 Mean bandwidth usage for hosts

By FastTrack and Gnutella 1/3 of the IP addresses have mean downstream bandwidths of 56Kbits in second or less. Probably those are users with dial-up internet connections and the other 2/3 with broadband network connectivity.

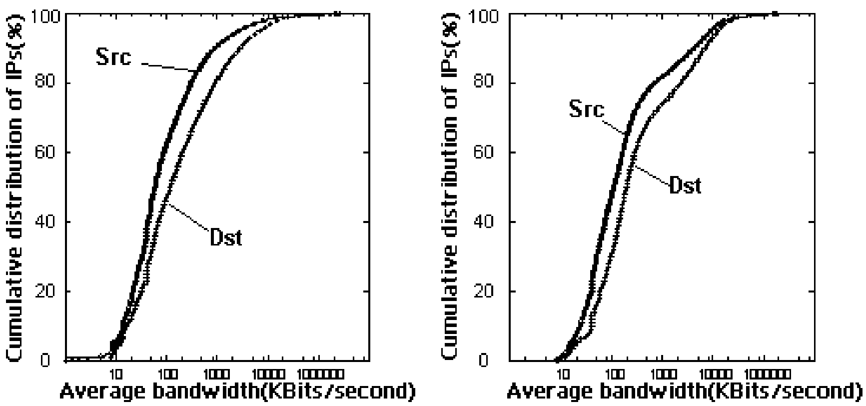

Figure 8: The cumulative distribution of the mean upstream and downstream bandwidth usage of hosts participating in FastTrack, and DirectConnect on September 14, 2001 (x-axis is in logscale). (a) FastTrack, (b) DirectConnect.

The observation also shows, that the average upstream bandwidth is smaller than the average downstream bandwidth. That is, because the presence of nodes with asymmetric bandwidth connectivity by DSL and cable modem and because the users often limit the upload traffic bandwidth. 50% of the IP addresses have average upstream bandwidth of 56Kbits/second or less.

For DirectConnect the bandwidth, downstream and also upstream, is higher than in the other two P2P networks. 20% of the IP addresses have a mean downstream bandwidth of 56 Kbits/second or less, while another 40% of the IP addresses have mean downstream bandwidth of 56 to 256 Kbits/second. 1/3 of the hosts have average upstream bandwidth of 56 Kbits/second, which also contributes to the more traffic, distributed by the DirectConnect hosts.

6.4 Geographic Distribution

Because of economical, political and legal reasons the provided traffic is different in the different geographic regions.

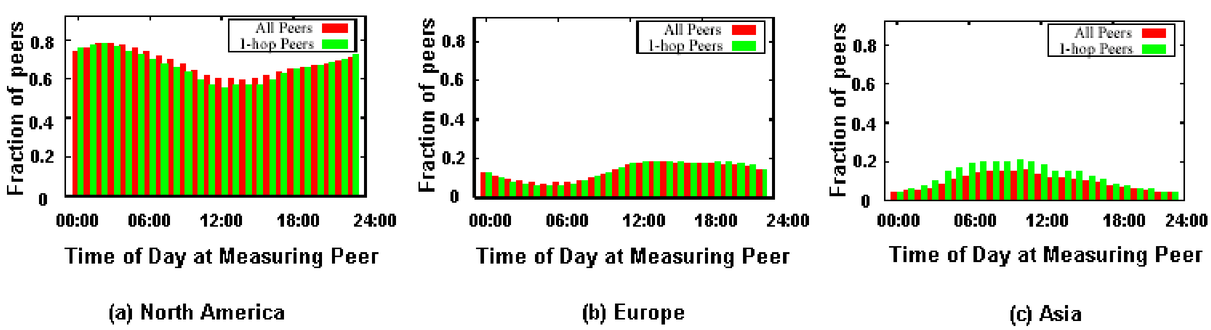

Figure 9: Geographic Distribution

Most of the traffic comes from North America. The traffic there decreases from 10pm to 6am from 80% to 60% and between 6am and 10am rises rapidly back to 80%. The European traffic reaches its peak 20% between noon and midnight. In the early morning the traffic fraction of Europe constitutes only of about 6%. The traffic in Asia is about 14% of the world traffic in the peak hours and about 4% at 6am. 5-10 % of the traffic comes from other geographic regions or from unknown locations.

7 P2P traffic vs. web traffic

According to [1], 97% of the prefixes contributing to P2P traffic also contribute Web traffic and also the heavy hitter P2P prefixes all tend to be heavy hitters in terms of web traffic. The top 0.01%, 0.1%, 1%, and 10% heavy hitter prefixes are responsible for 10%, 30%, 50%, and 90% of the corresponding monthly aggregated traffic volume. Although the P2P systems are very dynamic, the P2P traffic contributed by the top heavy hitter prefixes is more stable than that of the web traffic. This fact is due to long lasting data transfer flows, which contributes volume of the data. The stability of the traffic of the P2P networks makes them a preferred way of file exchange, and the most used system for downloading and uploading large audio and video files and software.

8 Conclusion

The Measuring of the traffic of P2P networks is a major factor for their development. By gathering information about the traffic volume, the bandwidth usage, the dependence of the time and geographic place, about the connection duration and the host distribution and connectivity, it is easier to find, manage and remove the weaknesses of the systems and to improve their efficiency. Because of the large amount of the hosts involved, the decentralization, the use in different geographical regions in different time periods the P2P networks are dynamic systems with highly skewed traffic distributions. Despite their asymmetric character FastTrack, Gnutella and Direct-Connect show stable traffic in comparison with the web systems, giving their users that way the chance to transfer large files, but also to chat or game. Only Kazaa claims to have 20 million downloads per month, and Morpheus has 115 million users worldwide. Despite their security and privacy issues, the P2P systems are and have the future of the most common file-transfer systems, been constantly developed and improved they increase their users number with the speed of increasing number of the internet users worldwide.Their future means the future of the free file communication among the different hosts in different continents in the whole network.

9 Bibliography

[1] Subhabrata Sen, Jia Wang: *Analyzing Peer-to-Peer Traffic across large Networks*; Second Internet Measurement Workshop IMW 2002

[2] Alexander Klemm, Christoph Lindemann, Mary Vernon, Oliver P. Waldhorst: *Characterizing the Query Behavior in Peer-to-Peer File Sharing Systems*; ACM Internet Measurement Conference IMC 2004