

## P2P, THE GORILLA IN THE CABLE

Alexandre Gerber, Joseph Houle, Han Nguyen, Matthew Roughan, Subhabrata Sen  
AT&T Labs - Research

### *Abstract*

*There is considerable interest in peer-to-peer (P2P) traffic because of its remarkable increase over the last few years. By analyzing flow measurements at the border routers of a Tier-1 ISP backbone that carry broadband traffic, we are able to study its properties. P2P has become a large part of broadband traffic and its characteristics are different from older applications, such as the Web. It is a stable balanced traffic: the peak to valley ratio during a day is around two and the IN/OUT traffic balance is close to one. Although P2P protocols are based on a distributed architecture, they don't show strong signs of geographical locality. A broadband subscriber is not much more likely to download a file from a close region than from a far region.*

*It is clear that most of the traffic is generated by heavy hitters who "abuse" P2P (and other) applications, whereas most of the subscribers only use their broadband connections to browse the web, exchange emails or chat. However it is not easy to directly block or limit P2P traffic, because these applications adapt themselves to their environment: the users develop ways of eluding the traffic blocks. The traffic that could historically be identified with five port numbers is now spread over thousands of TCP ports, pushing port based identification to its limits. More complex methods to identify P2P traffic are not a long-term solution, the cable industry should opt for a "pay for what you use" model like the other utilities.*

### INTRODUCTION

P2P (peer-to-peer) file sharing applications have grown dramatically over the past few years and contribute a significant share of the total traffic in many networks. In this paper, we analyze flow-based measurements of broadband traffic spanning several months, gathered in the backbone of a large ISP network. We first develop an understanding of P2P traffic behavior from the viewpoint of broadband provider networks (earlier studies were based on a Tier-1 ISP backbone viewpoint [1] and on a University edge-network viewpoint [2]). The study then describes some key issues and challenges in handling/controlling this traffic, and presents a potential solution approach. We begin with a description of these P2P systems

#### File Sharing Applications

Many popular P2P applications such as KaZaA and Gnutella are organized as application-level overlay systems in which large numbers of computers (called peers) across the Internet link together in a decentralized manner via application-level connections. The predominant use of these systems is for sharing large data files (particularly music and video) among the connected users. The data files and associated metadata information (useful for searching content) are distributed across the different peers. A key difference with traditional client-server systems is that each host in a P2P system acts as both a client and a server of content. In contrast to the stable configurations of traditional distributed systems, the individual peers can frequently join and leave the P2P system.

The process of obtaining a file can be broadly divided into two phases – query search followed by object retrieval. First, a user specifies a query (e.g., a combination of name, genre, artist name etc.), and the P2P protocol searches for the existence of file(s) that match the query. The requesting peer receives one or more responses, and if the search is successful, identifies one or more target peers from which to download each file. The search queries as well as the responses are transmitted via the overlay connections. The details of how the search is propagated through the overlay is protocol-dependent. In earlier P2P protocols exemplified by Gnutella version 4.0, a peer initiates a query by flooding it to all its neighbors in the overlay. The neighboring peers in turn, flood to their neighbors, using a scoping mechanism to control the flood. In contrast, for newer protocols like KaZaA, as well as for newer versions of Gnutella, queries are forwarded to and handled by only a subset of special peers (called SuperNodes in KaZaA, and UltraPeers in Gnutella). A peer transmits an index of its content to the “special peer” to which it is connected. The special peer then uses the corresponding P2P protocol to forward the query to other such peers in the system.

Once the search results are in, the requesting peer directly contacts the target peer, typically using some variant of HTTP (the target peer has a HTTP server listening by default on a known protocol-specific port), to get the requested resource. Some new systems use *swarming download*-- a file is downloaded in chunks from multiple peers.

Although the earlier P2P systems mostly used default network ports for communication, there is strong evidence to suggest that substantial P2P traffic nowadays is transmitted over a large number of non-standard ports. This seems to be primarily

motivated by the desire to circumvent firewall restrictions as well as rate-limiting actions by ISPs targeted at such applications - we shall discuss this more later in the paper.

Another recent occurrence has been the development of tools that allow an end-user to explicitly select the SuperNode it connects to [3]. This appears to be an attempt to improve the quality of the best-effort search process in the P2P system, for files that are not widely distributed, but are geographically localized. For instance, connecting to a SuperNode in Brazil may increase the chances of locating Samba-related content.

### Data Collection

We have access to “flow-level” data about broadband traffic at the border routers of a large ISP. Flow-level data is considerably more detailed than data sets such as SNMP, and at least this level of detail is needed to perform application classification. When looking at these flows we can make a very educated guess about whether the flow is associated with a Broadband consumer and from which region it originates. A region typically ranges from an extended metropolitan area to a state. For the remainder of this paper we focus on traffic that appears to be associated with broadband subscribers.

By flow, we mean a sequence of packets exchanged by two applications. More precisely we define a flow to be a series of uni-directional packets with the same IP protocol, source and destination address, and source and destination ports (in the case of TCP and UDP traffic). The flow measurements used here are called Cisco Netflow [4]; they are implemented in many of Cisco’s routers. The data collected about a flow (apart from the information above) are the duration, the number of packets, and bytes transmitted, and which header flags (SYN, ACK, ...) were used

in the flow. Measured flows are also constrained in time (Cisco Netflow collection sends flows from the router at 30 minute intervals), so there is a need to reconstruct the actual traffic from a single “connection”. After reconstruction there will be one flow per connection – a potentially enormous volume of information.

In order to minimize any performance impact on the routers collecting the flow measurements the measurements are based on sampled packets collected on the routers, which then export the flows to aggregators. To reduce the huge data volume the aggregator further samples the flows using the smart sampling algorithm [5] that is better suited for heavy tailed distribution, such as typically found in Internet flows. In addition, there is also an uncontrolled sampling due to measurement packet losses. These three types of sampling can be estimated and corrected and don't affect our results that are based on the weekly or monthly average traffic generated by hundreds of thousands of broadband subscribers between May 2002 and February 2003.

### Identifying Applications

There are a number of ways one could go about identifying individual applications within IP traffic. However, as noted, Netflow only keeps data on some aspects of flows. The most useful of these for application breakdowns are the source and destination port numbers, and the IP protocol number. The protocol numbers used are well documented [6], with TCP being protocol 6, and UDP being 17. TCP, and UDP traffic also define (16 bit) source and destination port numbers intended (in part) for use by different applications. The port numbers are divided into three ranges: the Well Known Ports (0-1023),

the Registered Ports (1024-49,151), and the Dynamic and/or Private ports (49,152-65,535).

A typical TCP connection starts with a SYN/ACK handshake from a client to a server. The client addresses its initial SYN packet to the server port for a particular application, and uses a dynamic port as the source port for the SYN. The server listens on its port for connection. UDP uses ports similarly though without connections. All future packets in the TCP/UDP flow use the same pair of ports at the client and server ends. Therefore, in principle the server port number can be used to identify the higher layer application using TCP or UDP, by simply identifying which port is the server port (the one from the well-known, or registered port range) and mapping this to an application using the IANA list of registered port [7].

There are many barriers to determining applications from port numbers. For instance, well know and registered ports are not defined for all applications and this is typical of P2P applications. Further more, in some cases server ports are dynamically allocated as needed (for instance, one might have a control connection on which a data port is negotiated). Finally, the use of firewalls to block unauthorized and unknown applications from using a network has spawned work arounds that have made the mapping from port number to application ambiguous.

Despite this, a great deal can be said about the mapping of port to application, though obviously there will still be some ambiguity, and chance for errors. Note that both ports must be considered as possible candidates for the server port, unless other data is available to rule out one port.

The algorithm that we have adopted here chooses the server port by (1) looking for a

well known port, (2) a registered port, or (3) an unregistered port which is known (from reverse engineering of protocols) to be used by a particular (unregistered) application. If both source and destination port could be the server, then we choose the most likely one through ranking applications by how prevalent they are in detailed (packet level) traffic studies – for instance, WWW is considered a high ranking application, as are email, and P2P applications.

The result is a mapping from flows to applications, that while not perfect, has been shown to be reasonably effective. The biggest problem is that there are still a substantial number of flows which cannot be mapped to an application. We further classify these unknown flows by the size of the flows: the category of most interest here is “TCP-big”, which consists of unknown flows that transmit more than 100kB in less than 30 minutes.

We shall argue in this paper that the TCP-big traffic is primarily P2P traffic that is using unregistered ports unknown to us. P2P applications already use unregistered ports, and the structure of P2P protocols (with separate control and data traffic) allows data traffic to be assigned to arbitrary ports. In the past the major applications have typically used default ports (for instance 1214 for KaZaA) but in the recent past many efforts have been made to constrain P2P traffic through rate limiting single ports or by blocking some ports at firewalls, with the result that P2P users commonly use work-arounds. Where-ever we refer to P2P traffic we are using the traffic on the ports known to be directly associated with P2P applications: we shall keep this separate from TCP-big except where explicitly noted. Also note that some P2P traffic may be misclassified into other application classes and so our estimates of the total volumes of P2P traffic are conservative.

We should note that we are not collecting any information about URL's, or individual subscribers usage: IP addresses measured are not related to individual subscribers, and we only view the bulk properties of the traffic, such as its distributions.

## APPLICATION COMPOSITION

### Overview

Table 1 shows the application traffic composition for 2 broadband regions in May 2002 and January 2003. For each of these regions, we examine both the traffic coming from outside the region to some IP address within the region (referred to as IN) and the traffic sourced within the region and destined for outside the region (OUT). For each time period and region, we display the per-application traffic volume in each direction as a percentage of the total traffic in that direction. For a given application we also show the traffic normalized by dividing by its IN traffic volume for May 2002, in order to show the In/Out ratio, and the growth between the two periods.

We note that in either direction the P2P traffic forms a much smaller percentage of the overall traffic in January 2003 than in May 2002. TCP-big registered dramatic increases in traffic contribution in both directions (10.5 times for Outgoing and 6 times for Incoming) over the same period. The normalized figures show that the P2P incoming and outgoing traffic are very similar for either of the 2 months considered. Note also that the TCP-big traffic in the 2 directions becomes much more balanced recently than earlier. For example for broadband region X, the ratio between incoming and outgoing TCP-big traffic volumes changes from 1.94:1 in May 2002 to 1.12:1 in January 2003.

	Broadband Region X								Broadband Region Y							
	Applicationx Mix (percentage)				Normalized Consumption				Applicationx Mix (percentage)				Normalized Consumption			
	May 2002		January 2003		May 2002		January 2003		May 2002		January 2003		May 2002		January 2003	
	OUT	IN	OUT	IN	OUT	IN	OUT	IN	OUT	IN	OUT	IN	OUT	IN	OUT	IN
All	100.0%	100.0%	100.0%	100.0%	1	1.65	1.97	3.2	100.0%	100.0%	100.0%	100.0%	1	2.19	1.83	4.08
ESP/GRE	0.4%	0.5%	0.6%	0.5%	1	1.98	3.12	4.3	0.4%	0.5%	0.3%	0.4%	1	2.71	1.7	4.67
OTHER	4.4%	3.7%	5.7%	4.5%	1	1.37	2.54	3.23	4.6%	3.2%	5.4%	3.4%	1	1.53	2.16	2.97
TCP-BIG	8.9%	10.5%	47.5%	32.5%	1	1.94	10.5	11.68	9.5%	11.8%	45.3%	32.1%	1	2.71	8.71	13.72
AUDIO/VIDEO	0.2%	1.6%	0.2%	1.6%	1	16.61	2.77	32.64	0.1%	1.5%	0.2%	1.5%	1	23.71	3.1	44.29
CHAT	0.7%	1.3%	1.0%	1.7%	1	3.08	2.93	7.93	0.7%	1.2%	0.7%	1.4%	1	3.81	2.02	8.67
FTP	1.0%	1.3%	1.0%	0.7%	1	2.22	1.91	2.4	1.4%	1.4%	0.4%	0.9%	1	2.24	0.56	2.64
GAMES	1.6%	1.2%	3.6%	2.5%	1	1.29	4.54	5.15	1.3%	1.2%	3.4%	2.4%	1	1.92	4.73	7.43
MAIL	1.7%	0.6%	1.1%	0.7%	1	0.6	1.26	1.28	1.0%	0.5%	0.9%	0.5%	1	1.13	1.71	1.88
NEWS	0.3%	7.3%	0.2%	5.3%	1	38.52	1.51	54.55	0.7%	17.5%	0.7%	14.6%	1	54.99	1.76	85.33
P2P	75.2%	45.6%	32.9%	20.6%	1	1	0.86	0.87	75.1%	38.5%	36.7%	19.5%	1	1.12	0.9	1.06
WEB	5.6%	26.4%	6.2%	29.4%	1	7.8	2.2	16.88	5.2%	22.8%	5.9%	23.5%	1	9.53	2.06	18.27

Table 1: Application Composition of two broadband regions in May 2002 and January 2003.

### Time of Day Pattern

We next examine the diurnal behavior of P2P traffic. Fig. 1 plots the time series of the incoming and outgoing traffic volumes (P2P, web and TCP-big) for a given broadband region across a week in February 2003. For each application, all the data values are normalized by the mean per-hour incoming data volume for that application, averaged across that week.




Fig. 1: Time of day pattern of P2P and Web traffic.

All three applications exhibit similar diurnal behaviors with peak loads (in either direction) around 2.00 AM GMT (10.00 PM EST, 7.00 PM PST). The P2P traffic exhibits

less variability across a day than Web traffic. The peak load is about 2 times the minimum as opposed to 5 times for Web traffic. The smaller variance in P2P traffic across a day may be a function of the programmed download feature in P2P applications that allow users to specify multiple files in advance, that can be downloaded asynchronously by the P2P application.

For Web, the outgoing traffic is significantly smaller than (at most 20% of) the incoming traffic, suggesting that the broadband subscribers are mostly consumers of web data. In contrast, for P2P, the traffic in the 2 directions track each other much more closely, across a day and across the week. Another notable here is that the TCP-big traffic distribution across time is very similar to the P2P traffic. Also, just like P2P, the TCP-big traffic in the 2 directions are similar. These behavioral similarities are another indicator that the TCP-big traffic includes some P2P applications. Finally for all 3 applications, we do not see significant variations across days and between weekdays and weekends.

## P2P LOCALITY

One of the potential advantages of P2P applications is that by distributing content, they provide the ability to download this content from locations closer to a user. It is therefore interesting to consider whether this really happens, and moreover to consider the question of locality in P2P traffic in general.

We approach this question by considering the simplest possible counter examples to localized traffic: the simple gravity model [8]. In this model, a packet entering the network at S, makes its decision about its destination D independent of the arrival point. That is, the packet is drawn (as if by gravity) to destinations in proportion to the volume of traffic departing at those locations.

The gravity model can be used to make predictions of the traffic volumes between two regions based purely on the volumes entering and exiting at those two regions, by the formula

$$T^{S,D} = \frac{T_{in}^S T_{out}^D}{T}$$

where T is the total volume of traffic across the network,  $T_{in}^S$  is the traffic entering the network at region S, and  $T_{out}^D$  is the traffic exiting the network at region D. Fig. 2 below shows a comparison of the gravity model predictions for inter-regional traffic of a broadband ISP. The plot is based on Netflow traffic collected during one week in September 2002; it shows traffic traversing the backbone between regions. The figure shows a scatter plot of the real inter-regional traffic versus the gravity model prediction, for both P2P traffic, and the total traffic to the broadband regions. One can see that in both cases the gravity model predicts the true traffic within about  $\pm 20\%$ .

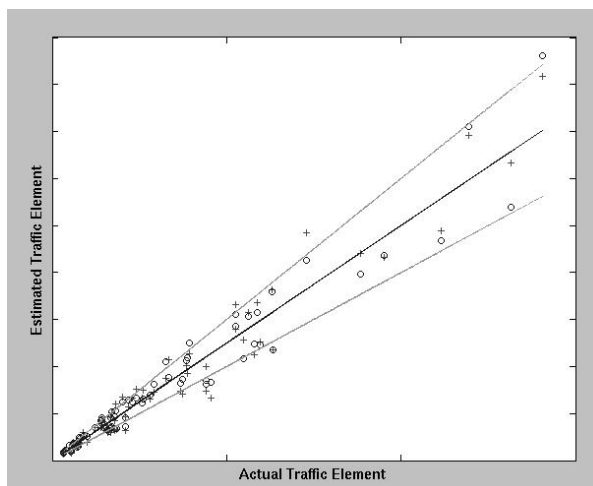


Fig. 2: Comparison of the real matrix elements to the estimated traffic matrix elements for a broadband ISP. The circles represent purely P2P traffic and pluses represents the total traffic. The blue solid diagonal line shows equality and the green dashed lines show  $\pm 20\%$ .

What does that tell us? Well the main point is that the gravity model above explicitly excludes any notion of geographic, or topological distance. Therefore, as the measured traffic fits this model to some extent, we may believe that neither P2P traffic nor the traffic overall exhibit strong locality at the regional level. A further, somewhat subjective conclusion one might draw from the graph is that P2P traffic actually seems to fit the gravity model slightly worse, and so we may hypothesize that P2P traffic shows more locality than other traffic sources.

To examine these hypothesis in more details we present Table 2, which shows the normalized traffic volumes between regions for the P2P traffic. The table shows the normalized probability that traffic originating from a particular region in one broadband network, will depart from each region in the same broadband ISP (given it stays on the same broadband network). Table 2 can be seen to have a number of almost identical rows (for instance the group of regions R1, R2, and R5 are very similar, as is the group R6, R7 and R8) indicating a complete lack of locality of traffic with reference to these regions. Other

regions (specifically R3 and R4) are not dramatically far away, but rather fall somewhere in between the other two groups.

However the table also shows some disparity between the groups of rows. This disparity is at its height when comparing the regions in the Eastern Standard Timezone (EST), with those in the Pacific Timezone (PST). This is an indication of some degree of weak locality in P2P traffic, at the “super-regional” level.

From/To	R1 (PST)	R2 (PST)	R3 (MST)	R4 (MST)	R5 (CST)	R6 (CST)	R7 (EST)	R8 (EST)
R1 (PST)	-	0.18	0.14	0.126	0.174	0.128	0.124	0.127
R2 (PST)	0.172	-	0.141	0.126	0.19	0.132	0.118	0.12
R3 (MST)	0.132	0.12	-	0.189	0.135	0.145	0.139	0.14
R4 (MST)	0.107	0.111	0.182	-	0.124	0.163	0.155	0.158
R5 (CST)	0.161	0.18	0.136	0.132	-	0.135	0.127	0.129
R6 (CST)	0.107	0.108	0.145	0.155	0.125	-	0.187	0.173
R7 (EST)	0.107	0.106	0.137	0.157	0.127	0.182	-	0.184
R8 (EST)	0.109	0.111	0.127	0.161	0.128	0.178	0.185	-

Table 2: Normalized inter-regional traffic matrix of broadband ISP X weighted by P2P+TCP-big traffic (Longitude defined by the Timezone).

This super-regional locality could arise for a couple of reasons (other than P2P applications explicitly taking advantage of content locality to improve performance). Firstly, because of usage patterns (specifically the times at which a user is connected to the P2P network), there is a slight increase in the likelihood that a search will find content in a local time zone. Secondly, there may be a group of people within a super-region with content that is slightly more relevant to the local super-region. However, the data so far suggests that both of these effects are not dominant, and certainly there is no strong locality influence such as might be seen if the main P2P applications exploited locality information.

In both of the above examples the monitoring location limits our data to seeing only inter-regional traffic. Thus, one might argue, we are missing the key component in any study of traffic locality: the intra-regional traffic. While the data limitations prevent us from seeing the intra-regional traffic on a

single broadband ISP, we can gain a good view of this data by considering the traffic between broadband ISPs. If locality were being exploited in P2P applications, then one would expect traffic from ISP Y, region R to prefer going to ISP X, region R, rather than the alternative regions.

Table 3 shows an example, giving the normalized probabilities that traffic from ISP Y to X will go from regions M to R. Although the regions for the two broadband ISPs are slightly different, regions M3 and R7 are very closely matched as are M4 and R8. However, we see only very minor bias towards traffic from M3 to R7 (compared to other EST regions), and similarly from M4 to R8.

From / To	R1 (PST)	R2 (PST)	R3 (MST)	R4 (MST)	R5 (CST)	R6 (CST)	R7 (EST)	R8 (EST)
M1 (MST)	0.133	0.121	0.157	0.125	0.118	0.111	0.089	0.146
M2 (CST)	0.121	0.095	0.114	0.158	0.117	0.145	0.094	0.156
M3 (EST)	0.12	0.114	0.12	0.138	0.119	0.128	0.14	0.122
M4 (EST)	0.11	0.115	0.109	0.137	0.135	0.119	0.133	0.142
M5 (EST)	0.117	0.115	0.133	0.135	0.129	0.12	0.121	0.129

Table 3: Normalized traffic matrix from broadband ISP Y to broadband ISP X weighted by P2P+TCP-big traffic.

Our conclusion is that, although there is some evidence for weak locality at a large spatial scale, P2P applications do not yet exploit such information on a large scale, and consequently, P2P traffic does not show strong signs of geographic locality. Recent developments such as the KazuperNode tool [3]) provide methods for selecting the super-node to which one connects. On the one hand this could potentially increase locality if users tend to connect to nearby supernodes. On the other hand, there could be less locality if users connect to supernodes in different locations in their attempts to locate content.

## HEAVY HITTERS AND P2P

It is well known in the broadband industry that some heavy hitters consume most of the bandwidth. We shall divide subscribers into classes by their total usage, and analyze their consumption characteristics such as the application composition and the traffic balance

per class. We define three groups of users: the heavy users who consume more than 1 Gbytes/day in average over a week, the medium users who consume between 50 Mbytes/Day and 1 Gbytes/Day and the light users who consume less than 50 Mbytes/Day.

### User Distribution

We first compare the distribution of traffic per subscriber. In order to see if there are consistent patterns we compare three regions, all at two different points in time: during the week ending June 26<sup>th</sup> 2002 and during the week ending February 9<sup>th</sup> 2003. By subscriber, we mean an active IP address. Even though the IP address is not statically assigned (the user

obtains an IP automatically via DHCP), in the networks we examined it is “sticky”. That is, over a week a subscriber maintains the same IP address in practice, because the DHCP lease expires only after 4 days and it is reassigned to him if it is still available. However, the IP address distribution doesn’t reflect exactly the subscriber distribution since it misses the inactive subscribers and the subscribers with a very low usage that may not be sampled.

The six distributions in Figure 3 and 4 are quite consistent. In each case, the top 1% of the IP addresses account for 18.6 — 24.4% of the total traffic and the top 20% of the active IP addresses account for slightly more than 80% of the traffic.

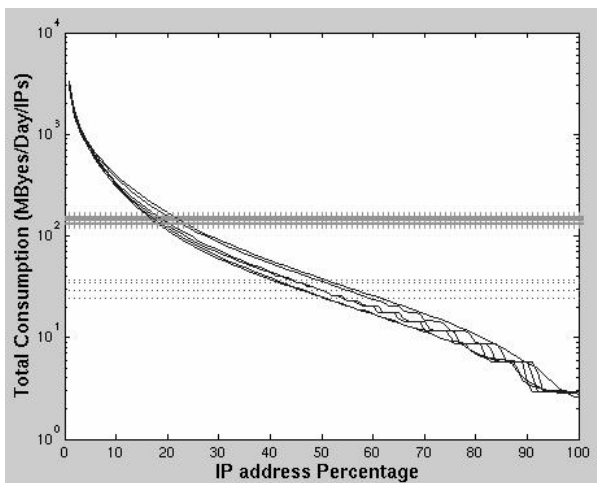


Fig. 3: Consumption per percentile of IP addresses of three regions during a week in June 2002 and a week in February 2003. The mean consumptions are around 140 Mbytes/Day/IP and the medians are roughly 30 Mbytes/Day/IP.

### Consumption Characteristics

Since the median consumption is 4 to 5 times smaller than the average consumption, it is clear that the average consumption doesn’t reflect the behavior of most of the subscribers. This still holds if we compare the application composition of each group of users, as defined earlier, with the average application composition that was studied earlier in this paper. Indeed, in a close look at one of these regions Table 4 shows that the light user group




Fig. 4: Cumulative Consumption of three broadband regions during a week in June 2002 and a week in February 2003.

(67% of the IP addresses) is still mainly browsing the web, exchanging email and chatting online. Its traffic balance – the IN/OUT ratio – is 4.8, which is far from the traffic balance of the heavy and medium user groups at 1.4-1.7 and 1.8, respectively. Table 5 makes it clear that this class of subscriber is not familiar with P2P or News since only 12.6 % of that group is lightly using one of these applications and it generates 1.1 % of the outgoing News traffic and 1.8 % of the outgoing P2P traffic.

User Type	Week ending June 26th 2002									Week ending February 9th 2003								
	Heavy			Medium			Light			Heavy			Medium			Light		
Direction	OUT	IN	OUT	IN	OUT	IN	IN/OUT	IN/OUT	IN/OUT	OUT	IN	OUT	IN	OUT	IN	IN/OUT	IN/OUT	IN/OUT
Normalized Traffic per Sub	266.8	445.5	27.0	48.9	1.0	4.8	1.7	1.8	4.8	288.3	415.1	26.1	47.8	1.1	5.2	1.4	1.8	4.8
AUDIO/VIDEO	0.1%	0.3%	0.1%	1.9%	0.4%	2.7%	3.2	26.4	29.8	0.1%	0.5%	0.2%	2.2%	0.4%	2.6%	4.9	17.3	28.4
CHAT	0.2%	0.4%	0.6%	0.8%	2.9%	2.0%	3.2	2.4	3.4	0.3%	0.6%	0.7%	1.2%	2.6%	2.3%	3.0	3.0	4.1
NEWS	1.1%	34.9%	0.5%	13.5%	0.2%	2.1%	53.6	54.1	55.1	1.0%	32.8%	0.4%	10.5%	0.1%	1.4%	49.6	46.6	46.2
MAIL	0.4%	0.1%	1.5%	0.4%	8.3%	2.3%	0.5	0.5	1.4	0.1%	0.3%	1.3%	0.7%	8.1%	2.7%	2.7	0.9	1.6
FTP	0.7%	0.9%	0.6%	1.1%	0.8%	0.3%	2.2	3.5	1.7	0.8%	0.7%	0.5%	0.8%	0.6%	0.2%	1.4	2.8	1.9
GAMES	0.4%	0.5%	1.5%	1.5%	2.8%	1.0%	2.0	1.7	1.7	3.3%	1.9%	4.1%	2.7%	2.9%	1.0%	0.8	1.2	1.7
ESP/GRE	0.0%	0.2%	0.7%	1.1%	5.3%	2.8%	6.9	3.0	2.6	0.1%	0.3%	1.0%	1.4%	6.0%	3.1%	5.6	2.5	2.5
P2P	87.4%	44.0%	82.3%	43.2%	18.5%	6.8%	0.8	1.0	1.8	37.7%	22.9%	29.5%	14.0%	7.0%	2.3%	0.9	0.9	1.6
TCP-BIG	6.9%	8.4%	3.3%	6.3%	2.4%	2.5%	2.0	3.4	5.1	51.2%	30.5%	47.6%	29.3%	13.1%	6.8%	0.9	1.1	2.5
WEB	0.9%	5.3%	5.1%	26.6%	46.2%	71.6%	10.1	9.5	7.5	1.6%	6.5%	6.4%	31.5%	46.7%	72.3%	5.7	9.0	7.5
OTHER	2.0%	5.1%	4.0%	3.7%	12.2%	5.7%	4.3	1.7	2.3	3.9%	3.1%	8.2%	5.8%	12.5%	5.3%	1.1	1.3	2.1

Table 4: Comparison of the application composition of the heavy, medium and light user groups of a typical region.

On the other hand the heavy user group is mainly generating file sharing traffic. That group is actually providing content to the rest of the P2P community since its P2P traffic balance is below 1. Even though that subscriber group accounts for only 2.9% of the subscriber population, it generates almost half of the P2P traffic (table 5). What is more surprising is that these P2P applications are not the only way for the heavy hitter class to download files. Only 83.6 % of that group of users installed one of these major P2P applications. This percentage goes up to 96.7% if we take also Netnews into account. Finally the remaining 3.3 % chose other solutions that include FTP and downloads from the Web. It is interesting to notice that Netnews and the Web are only means to download content but not to share it and so the traffic balance for these applications is very large: up to 50 bytes received for one byte sent.

Direction	Week ending June 26th 2002					
	OUT			IN		
User Class	Heavy	Medium	Light	Heavy	Medium	Light
IP address Percentage	2.9%	30.1%	67.0%	2.9%	30.1%	67.0%
Traffic Percentage	46.6%	49.4%	4.1%	41.6%	47.9%	10.5%
NEWS	68.6%	30.4%	1.0%	68.4%	30.5%	1.1%
P2P	49.6%	49.5%	0.9%	46.2%	52.1%	1.8%
TCP-BIG	64.9%	33.1%	2.0%	51.5%	44.5%	4.0%
WEB	8.5%	52.2%	39.3%	9.8%	56.6%	33.6%
P2P Users in that Class	83.6%	63.4%	10.1%	83.6%	63.4%	10.1%
News Users in that Class	25.8%	12.4%	2.6%	25.8%	12.4%	2.6%
News or P2P Users	96.7%	71.6%	12.6%	96.7%	71.6%	12.6%

Table 5: P2P and News Users in a region having more than 100 000 subscribers.

Looking at the evolution of the traffic balance of Web traffic of the heavy users also leads to the conclusion that a more complex phenomenon is happening. Indeed in June 2002, the web traffic balance of the heavy

users – 10.1 - was clearly higher than the web traffic balance of the light users whereas, in February 2003, that heavy hitter web traffic balance went down to 5.7, i.e. even lower than the one of the light users. This suggests that web traffic starts to be contaminated by a more balanced traffic, namely P2P applications. Furthermore, the traffic balance per application is another evidence that most of the traffic classified as TCP-big this year was actually what was classified as P2P last year. While the TCP-big traffic of the heavy hitters increased enormously, its traffic balance shifted from 2.0 to 0.9 and is now equal to the traffic balance of the P2P traffic that is still classified as P2P. It is now high time to understand why we are reaching the limits of port based identification of P2P traffic.

### LIMITING P2P TRAFFIC

The ability to accurately identify P2P traffic is a crucial requirement for appropriately handling this traffic in the network - through either traffic engineering, provisioning, rate-limiting or pricing. However, P2P applications have evolved rapidly in a direction which makes accurate accounting of the traffic more difficult. In particular, previously the applications used default TCP ports, and it was possible to account for the bulk of the P2P traffic by monitoring a relatively small number of ports. However, the current widespread use port-hopping makes such mapping exceedingly impractical. We next present specific evidence

of this trend and then discuss the implications for managing this traffic.

### KaZaA Rate Limiting Experiment

We first show an interesting case study which graphically illustrates how difficult it can be to limit P2P traffic. In Fall 2002, a particular broadband region began rate limiting traffic on port 1214 (the default port for KaZaA). Fig. 5 shows the IN traffic for web, p2p and TCP-big for that region before and after the rate limiting was initiated. Note that the P2P traffic decreases significantly after the rate-limitation was initiated. However, the TCP-big starts increasing and in 2 months has tripled compared to its value just before rate-limiting began. The web traffic (port 80, 8000, 8080) also increases over the same period. A reasonable explanation for the jump in the TCP-big traffic coincident with the rate limiting action on the KaZaA port is that the traffic spurt was caused by KaZaA traffic migrating to other ports that were mapped to TCP-big.

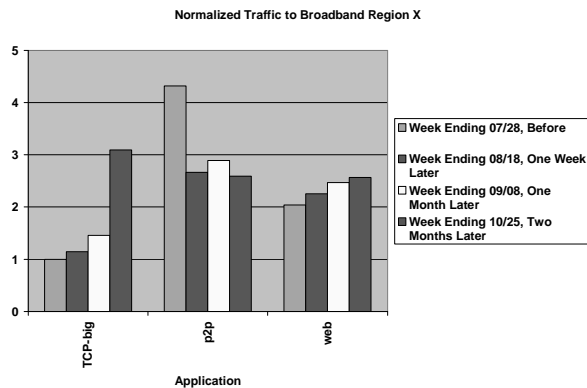


Figure 5: Mutation of P2P traffic into TCP-big traffic.

This conclusion is supported by the previous findings of this paper, but we shall investigate in even more detail. Fig. 6 plots the per-port traffic distribution for June 2002 and February 2003, for the P2P or TCP-big ports for the 2 time periods. Note that in 2002, 60 %

of that P2P and TCP-big traffic was contributed by only three ports. However, in February 2003, the traffic was much more uniformly distributed among a larger number of ports – the top 3 ports now account for only 20 % of the traffic. To get 60 % of the traffic we would need to monitor a larger number (1000) of ports.

Much more difficult is the task of mapping the traffic on these heavy-hitter ports to specific applications. Given the use of port-hopping by bandwidth-intensive applications like P2P, an important unanswered question is how much of the traffic on these ports can be attributed to the IANA-registered applications, and how much is P2P.

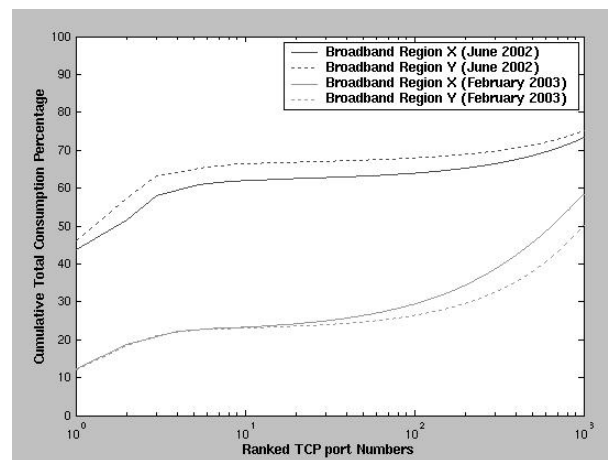


Fig. 6: Distribution of traffic by TCP port numbers classified as P2P or TCP-big

Given the limitation of port-based accounting, one might try to develop alternative techniques to accurately identify P2P applications. For example, additional information such as packet-level data, identification of SuperNodes etc. could help in developing signatures of P2P traffic. However, P2P applications have exhibited remarkable ability to rapidly evolve to evade detection and control. For example, many P2P applications now encrypt their communications, making it more difficult to reverse-engineer and/or monitor such systems at the application-level.

The above trends have important implications for port-based traffic control of P2P applications. If the rate control is targeted to a few well-known P2P ports, a significant fraction of the P2P traffic will evade the limit by hopping to other ports. The alternative is to track a larger number of ports that contribute significant traffic volumes and that are suspected to carry P2P traffic. The problem with this approach is that (i) it may not be feasible to track such a large and potentially dynamic set of ports, and (ii) such a widespread rate control may adversely affect the performance of many non-P2P users running valid applications on these other ports – this would be undesirable for the broadband providers.

#### SERVICE EVOLUTION TO TAME THE P2P GUERRILLA

There are an assortment of approaches to address the “problem” of P2P traffic. Let’s review a few that may be applicable to the cable industry.

Over the past few years many Multiple System Operators (MSOs) have incorporated “caps” into their service definition. These service caps tend to be implemented by controlling the rate at which data can flow into or out-of the network. The effect of these caps is to limit the instantaneous peaks of on-demand transactions. This has started us down the path of keeping bandwidth hogs in check. Some MSOs are now adding “tiered caps”. This allows the bandwidth hogs to identify themselves as such and pay a price for the enhanced service they are receiving.

Caps have been good to the industry and take us part of the way to where we want to go. However, P2P traffic is a relatively “passive” phenomena. The requester can queue-up a set of requests for files then walk away. The file

provider does not even need to be at the serving PC. In this situation rate capping will make the requests take longer, but will likely not change the behavior of the P2P participants. Fig. 1 enforces this point with the lower correlation between P2P traffic with the times users tend to be at their PCs.

Attempts to manage P2P traffic explicitly have met with little success. As illustrated in Figure 5, attempts to block standard ports of one P2P application only cause the user population to shift their behavior so that the traffic reappears on other ports. Devices inside the network to block or significantly throttle specific port numbers have questionable economic return given the “slipperiness” of ports that P2P applications use and the risk that valid applications also are using those ports.

Not that we should treat High Speed Data Services as a classic utility, but let’s look at how other “utilities” handle the problem of consumption hogs. Water, power, landline phone utilities all have a “pay for what you use” model. There is no attempt in these industries to limit the usage besides the economic consequence of paying for what is used. Cell phone providers put an additional twist on this model and provide usage bands. These bands allows a subscriber to sign-up for a usage band that best represents their need, but then gets charges for usage beyond what is included. With these revenue models consumption hogs are not “bad”, they are just big consumers.

User response to these revenue models may not be as bad as we may fear. Users will be concerned that this will raise their rates. Surveys suggest that many users, on the average, feel they themselves are heavy hitters. But Figure 4 suggests only 5% of the users are creating 50% of the traffic. With strategic selection of banding, the users will be pleasantly surprised to find that they can buy

one of the lower bands. There will be a small percentage of users (maybe the 1% that is causing the 20% of the traffic) that will not be happy with their new rates and will balk to other broadband services, but those are the ones that the cable industry can afford to lose.

### CONCLUSION

In this paper, we examined a large set of flow-based measurements of network traffic associated with broadband consumers, spanning several months. Our analysis reveals several interesting features. Firstly, they illustrate that broadband consumer traffic is dominated by P2P applications. We further look into the properties of the various application classes, in particular the traffic patterns, and IN/OUT ratios, noting that P2P traffic has a much more balanced traffic pattern and IN/OUT ratio than applications such as the web. In addition we show that geographic locality is not yet a dominant feature of P2P traffic.

The paper then considers the traffic patterns of user groups, showing that the well known 80-20 rule (80% of the traffic is generated by 20% of the users) applies here, but moreover that the group of heavy users actually tend to use different applications : they tend to generate more P2P and Netnews traffic, while the group of light users tend to use more web, email and chat applications.

Finally the paper considers how one might control the large volumes of P2P traffic that currently flood the broadband networks. The more obvious controls, such as rate limiting traffic on particular ports are shown to be ineffective, because they simply push the traffic onto alternate ports. A more practical

approach is to adopt a usage-based pricing approach, where the customers are billed for the resources they use.

### REFERENCES

- [1] "Analyzing Peer-to-Peer Traffic Across Large Networks", S. Sen and J. Wang, Proceedings of ACM SIGCOMM Internet Measurement Workshop, Marseilles, France, November 2002.
- [2] "An Analysis of Internet Content Delivery Systems", S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy, Proceedings of the 5th Symposium on Operating Systems Design and Implementation (OSDI), Boston, MA, December 2002.
- [3] KaZaA supernodes:  
<http://www.fasttrackhelp.com/kazupernodes/en>
- [4] Cisco Corp. NetFlow Services and Applications:  
[http://www.cisco.com/warp/public/cc/pd/iosw/ioft/neflct/tech/napps\\_wp.htm](http://www.cisco.com/warp/public/cc/pd/iosw/ioft/neflct/tech/napps_wp.htm)
- [5] "Charging from sampled network usage", N.G. Duffield, C. Lund, M. Thorup. ACM SIGCOMM Internet Measurement Workshop 2001, San Francisco, CA, November 1-2, 2001.
- [6] Internet Assigned Numbers Authority, protocol-numbers:  
<http://www.iana.org/assignments/protocol-numbers>
- [7] Internet Assigned Numbers Authority, port-numbers: <http://www.iana.org/assignments/port-numbers>
- [8] "Fast Accurate Computation of Large-Scale IP Traffic Matrices from Link Loads", Yin Zhang, Matthew Roughan, Nick Duffield and Albert Greenberg, to appear in ACM SIGMETRICS 2003.