

Performance Limitations of ADSL Users: A Case Study*

Matti Siekkinen^{1,**}, Denis Collange², Guillaume Urvoy-Keller³,
and Ernst W. Biersack³

¹ University of Oslo, Dept. of Informatics, Postbox 1080 Blindern, 0316 Oslo, Norway
`siekkine@ifi.uio.no`

² France Télécom R&D, 905, rue Albert Einstein, 06921 Sophia-Antipolis, France
`denis.collange@orange-ftgroup.com`

³ Institut Eurecom, 2229, route des crêtes, 06904 Sophia-Antipolis, France
`{urvoy,erbi}@eurecom.fr`

Abstract. We report results from the analysis of a 24-hour packet trace containing TCP traffic of approximately 1300 residential ADSL clients. Some of our observations confirm earlier studies: the major fraction of the total traffic originates from P2P applications and small fractions of connections and clients are responsible for the vast majority of the traffic. However, our main contribution is a throughput performance analysis of the clients. We observe surprisingly low utilizations of upload and download capacity for most of the clients. Furthermore, by using our TCP root cause analysis tool, we obtain a striking result: in over 90% of the cases, the low utilization is mostly due to the (P2P) applications clients use, which limit the transmission rate and not due to network congestion, for instance. P2P applications typically impose upload rate limits to avoid uplink saturation that hurt download performance. Our analysis shows that these rate limits are very low and, as a consequence, the aggregate download rates for these applications are low.

1 Introduction

We analyze a large packet trace of clients connected to the Internet via ADSL to investigate the causes of throughput limitations experienced by the end users. For this purpose we use a TCP root cause analysis tool that we apply to TCP connections. We consider throughput as the performance metric. The cause that limits the performance of a particular connection can be located either at the edge (sender or receiver) of a connection or inside the network. Limitations at edge comprise the application not providing data fast enough to the TCP sender or the TCP receiver window being too small. A network limitation may be due to the presence of a bottleneck anywhere along the end-to-end path. We perform root cause analysis of performance both at connection level and at client level. Based on a packet level trace that captures the activity of over one thousand ADSL clients during 24 hours we see that

* This work has been partly supported by France Telecom, project CRE-46126878.

** Work mostly done while at Institut Eurecom.

- The distribution of the client activity in terms of volume and duration is highly skewed. Most clients are active only during a short period of time. Also, most clients generate a limited amount of traffic in the order of several tens of MB, while a small number of (heavy hitter) clients upload and download hundreds of MB each.
- The utilization of the uplink and downlink is very low for most of the clients. Even heavy hitters are far from saturating their access link.
- The low utilization is mainly due to the applications that limit their rate of transfer, which is now very common for P2P applications such as eDonkey.

2 Dataset

We collected one full day (Friday March 10, 2006) of traffic generated by approximately 3000 ADSL users identified by IP addresses. We captured all IP, TCP and UDP headers of packets without any sampling or loss. The data collected on this day represents approximately 290 GB of TCP traffic in total, out of which 64% is downstream and 36% upstream. This day can be considered as a typical day in terms of volumes uploaded and downloaded by clients. Out of those 3000 clients, 1335 generated enough data to enable any root cause analysis. We consider only those clients in further analysis. In addition to the packet trace, we have a list of IP addresses that belong to local clients, which allows us to distinguish the upstream traffic from the downstream traffic. However, we do not know the clients subscription rates, i.e., their uplink and downlink capacities. The offered subscriptions were (down/up): 128/64, 512/128, 1024/256, 1024/128, 2048/128, 2048/256, 3072-4096/160, 4096-5120/192, 5120-6144/224, 6144-8640/256, and 18500/840.

We first analyzed the overall characteristics of the trace. Due to space constraints, we report here only the main findings from this study. For further details, we refer the reader to our technical report[1] which is an extended version of this paper.

The average volume of data uploaded is quite constant during the whole day, around 2GB per 30 minute period. The volume of downloaded data is less constant, around 3 GB per 30 minute period from midnight to 6 am and around 4 GB per 30 minute period for the rest of the day.

Only 5 applications generated more than 5% of the total amount of bytes: E-donkey, applications using port 80/8080, BitTorrent, email and telnet (due to a couple of hosts that generated large amount of traffic using telnet for some unknown reasons). We identify applications using port numbers and associated the TCP port range 4660-4669 to eDonkey, the ports 6880-6889 and 6969 (tracker) to BitTorrent. We do not want to declare the traffic seen on ports 80 and 8080 as Web traffic since it is likely to include also P2P traffic. The dominant category of traffic, however, is the traffic from unidentified applications, referred to as “other” traffic in the rest of the paper. Since much of today’s traffic is not using fixed ports but “hiding” [2], we are not able with our port-based method to classify much of the traffic seen. Therefore, the “other” traffic represents about 50% of the total

traffic. However, our root cause analysis (see next section) does not rely on the identification of the application to infer the causes for throughput limitation.

Distributions of the traffic per connection and per client are heavily skewed. Consequently, clients can be classified into two classes: heavy hitters and non heavy hitters. We identified heavy hitters as the 15% of clients that generated 85-90% of the bytes both upstream and downstream. They represent 200 clients. Those results are in line with the ones of a recent study performed on a much larger scale for Japan's residential user traffic [3]. The average amount of bytes uploaded and downloaded by a heavy-hitter client is approximately 470 MB and 760 MB, respectively, while for a non heavy-hitter those average values are 9 MB and 27 MB.

Heavy hitters also differ from non heavy hitters in terms of the set of applications they use. Overall, heavy hitters tend to use P2P applications more extensively, which is visible when looking at the identified applications (heavy hitters heavily use eDonkey) and also when merely looking at the volumes uploaded and downloaded (see above), which are significantly more symmetric for a heavy hitter than for a non heavy hitter.

*Access link utilizations*¹, *uplink and downlink*, are in general very low. We observed that 80% of the clients have a downlink utilization of less than 20% and uplink utilization of less than 40% for a given 30 minute period.

Having seen that most clients achieve very low link utilization, we will now set out to investigate the causes. For this purpose, we will use some techniques referred to as root cause analysis (RCA) that has been originally proposed by Zhang et al. [4] and further refined by Siekkinen et al. [5].

3 Performance Analysis of Clients

3.1 Connection-Level Root Cause Analysis

To apply RCA, we need TCP connections that carry at least 130 data packets, which is equivalent to about 190 kB of data, if we assume MSS to be 1450 B. As pointed out in Section 2, most connections are quite small, but most of the bytes are carried in a tiny fraction of the largest connections. As a consequence, our RCA will only be able to analyze the 1% of the largest connections, which however carry more than 85% of total bytes.

We classify in a first step the packets of a connection into two groups. Each packet is either part of an **application limited period** (ALP) or a **bulk data transfer period** (BTP). Roughly speaking, the throughput of packets that are part of an ALP is *limited by the behavior of the application*. For example, an IP telephony application that produces packets at a fixed rate clearly determines (and limits) the throughput achieved. Therefore, the packets of the TCP connection carrying these data should all be put into an ALP.

¹ Due to lack of knowledge about clients' access link capacities, we estimated a lower bound for the capacity and, thus, obtain an upper bound for the utilization. Details are presented in the extended version[1].

The packets that are not part of an ALP will be part of a BTP. For the details on how packets get classified into ALPs and BTPs, we refer to our technical report [6].

For packets that are part of a BTP, there can be a number of causes that limit the throughput achieved, such as:

- **Network limitation:** A bottleneck limits the observed throughput. We distinguish between two types of network limitation. One is called **un-shared bottleneck** and corresponds to the case where a single connection uses the full capacity of the bottleneck link. The other type, called **shared bottleneck**, occurs when several connections share the capacity of the bottleneck link.
- **TCP end-point limitation:** The advertised receiver window is too small as compared to the bandwidth-delay product of the path, which prevents the sender to achieve a higher throughput. Note that in practice, the sender buffer size is rarely too small[5]. We count into this category also *transport limitation* which relates to the time spent for TCP ramp up[4].

The choice of the most likely limitation is based on a set of metrics computed from the packet header trace of the connection and a threshold-based classification scheme. For details, see [5] and esp. Chapter 7 of [7].

3.2 Client-Level Root Cause Analysis

We are interested in doing RCA not only at connection level but also at *client* level. We identify four types of limitations for clients, which are: (i) Applications, (ii) Access link saturation, (iii) Network limitation due to a distant bottleneck, and (iv) TCP end-point limitation. Our analysis showed the TCP end-point limitation category (described above) to be marginal in our data set. Hence, we exclude this limitation category from further discussions.

In this analysis, we focus on *active* clients. We define a client to be *active* during a period of 30 minutes if it transferred at least 100 kB during that period. For each active client we consider all the bytes transferred by all the connections of the client within a given 30-minute period. We then associate these bytes into the three considered client-level limitations. To do this association, we use the connection-level RCA as follows: All the bytes carried by the ALPs of all the connections of the client are associated to application limitation. All the bytes carried by all the BTPs that are labeled network limited (unshared or shared bottleneck) by connection-level RCA and during which the utilization is above 90% of the maximum are associated to access link saturation. All the bytes carried by the rest of the network limited BTPs during which the utilization is below 90% of the maximum are associated to network limitation due to a distant bottleneck. All the rest of the bytes transferred by the client, and not covered by these three limitations, are associated to “other” (unknown) client limitation. The amount of bytes associated with a limitation serves as a quantitative metric of the degree of that limitation for a given client during a 30-minute period.

We know from our previous work on RCA that for a single, possibly very long connection, the limitation cause may vary over time. Also, a single client may run one or more applications that will originate multiple connections. Assigning a single limitation cause to each client is therefore tricky. For this reason, we distinguish for each client between “main limitation” and “limitations experienced”. As **main limitation**, we mean the limitation that affects the most number of bytes for this client. This classification is exclusive. i.e. each client belongs to a single limitation category.

On the other hand, under **limitations experienced** a single client will be considered in all the categories whose limitation causes it has experienced. Therefore, this classification is not exclusive. The results are presented in Table 1 for two 30-minute periods of the day: 4-4:30am and 3-3:30pm, which are representative for the different periods of the day. During the night time, heavy hitters dominate (70 out of 77 active uploading clients and 61 out of 83 active downloading clients), which is not surprising if one considers that heavy hitters heavily use P2P applications and P2P file transfer that can run for several hours [8]. If we look at the absolute number of clients, we see that only a small fraction of 1335 clients is active in either 30-minute period. We show only the results for the upstream direction, the ones for the downstream direction being very similar and are given in our technical report [6].

Table 1. Number of active clients limited by different causes

			Upstream					
limitation cause			Total active #	application	access link	other link	other cause	
main limitation	all clients	4am	77	95%	0%	4%	1%	
		3pm	205	86%	6%	4%	4%	
	heavy hitters	4am	70	94%	0%	4%	2%	
		3pm	4am	111	92%	2%	3%	3%
			3pm	111	92%	5%	64%	–
limitations experienced	all clients	4am	77	100%	0%	60%	–	
		3pm	205	100%	7%	39%	–	
	heavy hitters	4am	70	90%	0%	66%	–	
		3pm	4am	111	92%	5%	64%	–
			3pm	111	92%	5%	64%	–

Main Limitation. If we look at the main limitation cause experienced by the clients, we see that *almost all clients see their throughput performance mainly limited by the application*. This holds irrespective of the direction of the stream (upstream or downstream), of the type of client, average client or heavy hitter, and of the period of the day.

The clients that are not application limited see their throughput either limited by the capacity of the access link or the capacity of another link along the end-to-end path. Capacity limitations occur more frequently during the daytime than at night. The very limited number of cases where we observe a saturation of the access link complies with the low access link utilization observed in the preliminary analysis (Section 2).

Limitations Experienced. Besides the main limitation, we also consider *all the limitation causes* experienced by a single client. The most striking result is the difference between main limitation and limitations experienced for the "other link" limitation. As we have seen, this limitation is rarely the main limitation, while the percentage of clients that experience such limitation is between 40% and 60%, which means that while approximately half of the clients experience such network limitation, this limitation cause is not dominant. Moreover, we checked that for a given client, the amount of bytes transferred while limited by the network is generally clearly less than the amount of bytes transferred while limited by the dominant cause, i.e. the application in almost all of the cases.

3.3 Throughput Limitations Causes Experienced by Major Applications

Having done the root cause analysis on a per-client basis, we now perform application-level RCA, i.e. we investigate what are the most important applications that experience the different limitation causes, namely (i) application limited, (ii) saturated access link, and (iii) bottleneck at distant link. For each 30-minute period, we associate bytes flagged with limitations by client-level RCA to different applications based on the used TCP ports (as in Section 2).

Figure 1(a) shows the main applications that generate traffic that is application limited. If we look at the evolution of the total volume of traffic that is application limited we see very little variation in time and an upload volume almost as big as the download volume, both being around 2 GB per 30 minutes. The largest single application that generates application limited traffic is, as expected, eDonkey. However, if we look by volume, the largest category is "other", i.e. the one where we were not able to identify the application generating the traffic. The overall symmetry of upload and download volumes for the "other" category as well as a manual analysis of the traffic of some heavy hitters strongly suggest that the "other" category contains of a significant fraction of P2P traffic.

Figure 1(b) shows the main applications that saturate the access link. For this cause, no traffic originating from recognized P2P applications was seen. Instead, a significant portion of traffic saturating the uplink is e-mail. For the downlink it is mainly traffic on ports 80 and 8080 and traffic for which the application could not be identified. The fact that the traffic using ports 80 and 8080 primarily saturates only downlink suggests that it could be real Web traffic that consists of small upstream requests and larger downstream replies from the server, as opposed to P2P traffic which is typically more symmetric. If we look at the absolute volumes, we see that most of the activity is concentrated to day time, with the peak being in the early afternoon.

Figure 1(c) shows the main applications that see their throughput limited by a link that is not the client's access link. The category of other applications is clearly dominating in terms of volume. Otherwise, we observe a mixture of applications. It is expected that the set of applications is diverse since this type of network limitation can occur at any point of the network regardless of the application behavior at the particular client experiencing that limitation.

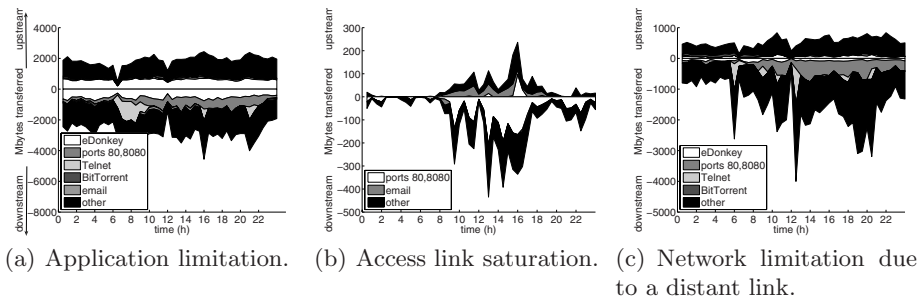


Fig. 1. Amount of bytes experiencing a particular root cause. Note the different scales.

In the download direction, the total traffic that is limited by a distant bottleneck reaches in the late afternoon a proportion that, in terms of volume, is almost as important as the download traffic that is application limited. The fact that this traffic peaks late afternoon² may be an indication of higher overall network utilization just after working hours, not only within the France Telecom network but in wider scale, that causes more cross traffic in aggregating links. Note that at the same time, the amount of traffic limited by the access link is very low (Figure 1(b)), which could indicate that these two groups represent different types of clients.

Finally, we would like to point out that a comparison of the absolute traffic volumes of Figures 1(a) – 1(c) reveal that the application limitation category represents the vast majority of the total number of transmitted bytes.

3.4 Impact of the Root Causes on Access Link Utilization

Now, we want to know how the three main root causes of throughput impact the access link utilization of the clients. We focus on link utilization and not on absolute throughput, because clients have different link capacities and we want to understand how far we are from the maximum, i.e. access link saturation.

As before, we included in the analysis for each client only the traffic of the 30-minute period for which that client achieved its highest instantaneous throughput. We computed client’s link utilization during ALPs and BTPs limited by different causes. In this way, we can quantify the impact of different limitation causes on the performance. Figure 2 shows CDF plots of the results.

We focus first on uplink utilization: We see that for the case of an unshared bottleneck, the utilization is in approximately 70% of the cases very close to one, which means that in these cases the uplink of the client is the bottleneck. In the remaining 30% of cases where we observe an unshared bottleneck, we see a link utilization between 0.4 and 0.85 that can be due to a distant access downlink,

² An analysis of the IP addresses using Maxmind (<http://www.maxmind.com/>) revealed that most of the local clients exchange data primarily with peers/servers located in France or surrounding countries.

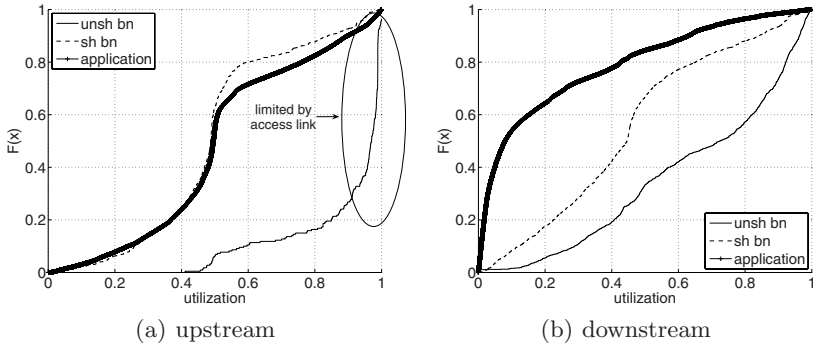


Fig. 2. CDF plot of access link utilization for the different root causes. For each client, we consider only traffic of the 30 min period during which that client achieved the highest instantaneous throughput of the day.

e.g. a peer that has lower downlink capacity than the uplink capacity of the local peer, or due to simply misclassification. For the two other root causes, application limitation and shared bottleneck, the clients achieve in about 60% of the cases a link utilization of less than half the uplink capacity.

If we look at the utilization of the downlink, we see that application limited traffic results most of the time in a very poor downlink utilization. Given that most of the application limited traffic is eDonkey traffic (cf. Figure 1(a)), one might be tempted to explain this low utilization by that fact that most likely the peer that sources the data has an asymmetric connection with the uplink capacity being much lower than the downlink capacity of the receiving peer³. However, a downloading peer has usually multiple parallel download connections, which in aggregation should be able to fully utilize the downlink capacity. The fact that this is not the case seems to indicate that many users of eDonkey use the possibility offered by today's P2P applications to limit their upload rate. Figure 3, which plots the maximum instantaneous aggregate download rates achieved per-client for different applications, further supports this hypothesis. We see that the maximum aggregate download rates of P2P applications, eDonkey and BitTorrent, are clearly below the maximum download rates of FTP and port 80/8080 traffic.

A recent study of eDonkey transfers by ADSL clients [8] found that the average file download speed achieved was only a few kB/sec. Our findings seem to indicate that such a poor performance is not due to network or access link saturation but rather due to *eDonkey users drastically limiting the upload rate of their application*.

3.5 Comparison with Other RCA Work

In [4], Zhang et al. performed flow-level root cause analysis of TCP throughput using a tool called T-RAT. They analyzed packet traces collected at high speed

³ Maxmind also reported that a clear majority of the distant IPs that the heavy-hitters communicated with were clients of ISPs providing residential services.

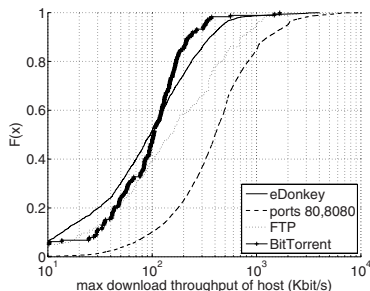


Fig. 3. CDF plot of maximum aggregate per-host download throughput computed over five second intervals

access links connecting two sites to the Internet, at a peering link between two Tier 1 providers, and at two sites on a backbone network. As results, the authors reported that, in terms of traffic volume affected, congestion (similar to network limitation in our vocabulary) was the most common limiting factor followed by host window limitation (TCP end-point in our vocabulary). It is important to notice that the data used were collected in 2001-2002. At that time, the popularity of P2P applications such as eDonkey was far from what it is today.

In order to understand whether or not our results are specific to this particular access network, we applied our RCA tool also to other publicly available packet traces collected at an ADSL access network in Netherlands (<http://m2c-a.cs.utwente.nl/repository/>). We looked at two 15-minute traces: one captured in 2002 and another one in 2004. A port based study similar to the one in Section 2 showed that in the 2002 trace, the applications generating most traffic were FTP and applications using ports 80 and 8080, while eDonkey and BitTorrent were dominating in the 2004 trace. We were unable to perform similar client-level study due to lack of knowledge about local client IP addresses and limited capture durations. However, simple connection-level RCA revealed that in the 2002 trace around 40% of bytes were throughput limited by the application, while this percentage was already roughly 65% in the 2004 trace, which demonstrates the impact of the increase in P2P application traffic.

4 Conclusions

We analysed one day ADSL traffic generated by more than one thousand clients. We observed that most of the clients never utilize more than a very small fraction of their upload and download capacity. TCP root cause analysis revealed that most of the user traffic is in fact *application* limited, which means that the users of P2P applications impose upload rate limits that are chosen to be very low. Other root causes that were typically observed in previous studies [4] play only a minor role: We saw some occurrences of network limitation, as well as

rare occurrences of limitations by TCP configuration issues such as too small a receiver window, but the amount of bytes transferred and affected by these causes were very small in comparison.

By severely limiting the aggregate upload rate of their P2P applications, clients certainly make sure that their P2P traffic does not interfere with concurrent activities such as Web serving or IP telephony. However, this comes at the price of very long download times, which makes the current rate limitation strategies used by P2P clients very inefficient from a users point of view.

The implication of such a low access link utilization is naturally low utilization of the entire access network, which is beneficial for the service provider. However, the utilization and traffic volumes can change dramatically in case a new type of popular P2P application is deployed or an already existing one is upgraded to utilize the uplink in a different, more effective way.

References

1. Siekkinen, M., Collange, D., Urvoy-Keller, G., Biersack, E.W.: Application-level performance of ADSL users. Technical report (2006) <http://www.ifi.uio.no/~siekkine/pub.html>.
2. Karagiannis, T., et al.: Is P2P dying or just hiding? In: Proc. Globecom 200. (2004)
3. Cho, K., Fukuda, K., Esaki, H., Kato, A.: The impact and implications of the growth in residential user-to-user traffic. In: Proceedings of ACM SIGCOMM '06, New York, NY, USA, ACM Press (2006) 207–218
4. Zhang, Y., Breslau, L., Paxson, V., Shenker, S.: On the characteristics and origins of internet flow rates. In: Proceedings of ACM SIGCOMM '02, Pittsburgh, PA, USA (2002)
5. Siekkinen, M., Urvoy-Keller, G., Biersack, E., En-Najjary, T.: Root cause analysis for long-lived tcp connections. In: Proceedings of CoNEXT. (2005)
6. Siekkinen, M., Urvoy-Keller, G., Biersack, E.W.: On the impact of applications on tcp transfers. Technical report, Institut Eurecom (2005) <http://www.eurecom.fr/~siekkine/pub/RR-05-147.pdf>.
7. Siekkinen, M.: Root Cause Analysis of TCP Throughput: Methodology, Techniques, and Applications. PhD thesis, Institut Eurécom / Université de Nice-Sophia Antipolis, Sophia Antipolis, France (<http://www.ifi.uio.no/~siekkine/pub.html>) (2006)
8. Plissonneau, L., Costeux, J.L., Brown, P.: Detailed analysis of edonkey transfers on adsl. In: Proceedings of EuroNGI. (2006)