Distribution of traffic among applications as measured in the French METROPOLIS project

Répartition du trafic par application mesuré dans le projet METROPOLIS

Philippe Owezarski[1], Nicolas Larrieu[1], Laurent Bernaille[2], Walid Saddi[3], Fabrice Guillemin[3], Augustin Soule[2], Kavé Salamatian[2]

Abstract.

We investigate in this paper the time evolution and the composition in terms of applications of traffic in two different networks, namely the Renater network, dedicated to the French academic and research community, and the France Télécom backbone network supporting commercial traffic. For each network, we present the time evolution of traffic in terms of applications, the associated pie charts for global results, as well as, for each detected application, its flow size distribution, that should have an impact on the traffic nature (self-similarity or long range dependence due to the heavy tail of flow size distribution). Based on these results, this paper presents a discussion on the differences between academic and commercial traffic in terms of usage, as well as possible solutions against LRD and its associated degradation of network performance. For traffic analysis, we propose a new method of classifying traffic according to applications, which relies on applicative protocols recognition instead on the IANA ports numbers.

Résumé.

Nous étudions dans cet article l'évolution au cours du temps et la composition du trafic en termes d'applications dans les deux réseaux que sont Rénater, utilisé par la communauté académique et de recherche française, et le réseau France télécom qui transporte du trafic commercial. Pour chaque réseau, nous présentons l'évolution au cours du temps de la composition de trafic en termes d'applications, les camemberts associés pour présenter les résultats globaux ainsi que, pour chaque application observée, la distribution des tailles de ses flux qui ont un impact sur les caractéristiques du trafic (auto-similarité ou dépendance longue dues à des distributions de tailles de flux à décroissance lente). A partir de ces résultats, cet article étudie les différences entre les trafics académique et commercial en termes d'usages, ainsi que des solutions pour réduire la LRD est les baisses de performance réseau qu'elle induit. A noter également que pour l'analyse du trafic, nous proposons une nouvelle méthode de classification du trafic par application qui repose sur la reconnaissance des protocoles applicatifs plutôt que sur les numéros de port IANA.

1. Introduction

[1] LAAS-CNRS, 7, Avenue du Colonel Roche, 31077 Toulouse Cedex 4, France
[2] LIP 6, Université Pierre et Marie Curie, 4 place Jussieu, 75005 Paris, France
[3] France Télécom, 2 Avenue Pierre Marzin, 22300 Lannion, France

While the Internet was originally used by the scientific community for exchanging electronic mails, text or binary data without any strong delay requirements, the usage of this network has been rapidly changing for the past ten years. As a corollary of its immense commercial success, in particular through the massive deployment of residential broadband access via ADSL lines, the Internet has nowadays to support a variety of applications, which did not even exist when the basic IP protocols were designed. This is particularly true for many applications, which are currently utmost popular on the Internet, such as peer-to-peer (P2P) applications for exchanging music and movies, Video on Demand (VoD), distributed games, video conferencing, IP telephony, etc.

These new applications have very different requirements in terms of Quality of Service, especially with regard to information transfer delay through the network. Moreover, those applications give rise to various traffic patterns, which may have a great impact on the behavior and the performance of the network (burstiness, elasticity, streaming, etc.). Designing mechanisms, protocols and architectures able to meet the quality requirements of new multimedia applications is a very challenging issue in the evolution of the Internet. Yet, as a preliminary task, it is essential to exactly know which applications give rise to traffic in current networks. This is of great importance for network engineers, administrators and of course researchers when designing new traffic management policies.

Estimating the distribution of traffic among applications is one of the key contributions of monitoring and measurement systems. However, the applicative classification is nowadays often a tricky task as more and more applications use encryption (for instance in an encrypted IPSEC tunnel), or use dynamic port numbers. This greatly complicates the recognition of applications generating traffic, in particular through the analysis of layer-4 port numbers. The Metropolis project, funded by the French council for research in telecommunications (RNRT), has been addressing the above issue for the last three years, and has proposed a new methodology for accurately classifying traffic among the different applications, based on the recognition of the application protocol, thus making the classification process insensitive to dynamic port changes.

In this paper, we present results obtained, in the framework of the Metropolis project, for two different networks with different characteristics in terms of usage and users. The Metropolis project has been monitoring for three years several links of the Renater network (the French network for education and research)[4]. In parallel, France Télécom has carried out several measurement experiments in its different networks for learning the characteristics of IP traffic and the usage of customers in a commercial environment. In this paper, some results based on measurements from on a high speed link of the France Télécom backbone network are presented. This paper thus describes results on both an academic and a commercial network, enabling a comparison in terms of usage and traffic characteristics.

This paper is organized as follows: in Section 2, we describe in details the networks which are monitored as well as the measurement points where traffic has been captured. Section 2 also presents the new classification method designed for the Metropolis project and based on the recognition of the protocol used by an application. Measurement and classification results are then shown in Section 3, focusing on the difference between academic and commercial networks in terms of usage. Section 4 concludes this paper and presents future work.

2. Classification methodology and description of monitored networks

2.1. Experimental setting

[4] Renater interconnects all universities, public research labs, some schools, as well as some industrial partners (depending on the projects in relation with academia they are involved in).

The Metropolis project, started in 2001, gathers most of the public research in networking and metrology in France. In terms of measurement techniques, Metropolis aims at combining active and passive measurements in order to obtain the best of both techniques. The monitoring and measurement platform has been designed and deployed in the Renater network. This platform consists of:

- Passive microscopic monitoring devices equipped with the famous DAG card (designed and provided by the University of Waikato and Endace in New Zealand), which collects every packet transmitted on the monitored link, including a very accurate GPS timestamp [CLE00];
- Active measurement devices. Active probes rely on the RIPE boxes that have been extended to support the NIMI software as well as a self designed and developed measurement software environment called MetroMI.

The Metropolis monitoring and measurement platform is depicted in Figure 1. Even if active measurement results are also available, we focus in this paper on passive measurements only, specifically traffic traces captured in Paris (University of Paris 6) and Toulouse.

In addition, we present results obtained by analyzing traffic measurements from the France Télécom network. Traffic traces are captured on a high speed (1 Gbit/s) link connecting several ADSL areas to the France Télécom IP backbone network. Only traffic from the backbone in direction to the ADSL areas (downstream traffic) is analyzed. The reason for limiting the analysis to downstream traffic is that downlink bit rates are much more various and much higher than the uplink bit rates. Depending upon customer's subscription, the downstream bit rates vary from 128 Kbit/s up to 1 Mbit/s, the majority of customers having a 512 Kbit/s subscription. The upstream bit rates for the experiment described in this paper and carried out in 2005 are then most of the time limited to 128 Kbit/s. Note that with the introduction of new services such as video on demand, visiotelephony, etc., the access bit rates are rapidly changing through new commercial offers by France Télécom. The measurement device was installed in derivation of the transmission link and performed an on line analysis of the first 200 bytes of IP packets. We were thus able to determine to which micro-flow, identified by the 5-uple (source IP @, destination IP @, source port, destination port, protocol type), a packet belongs to as well as the application generating the IP packet.

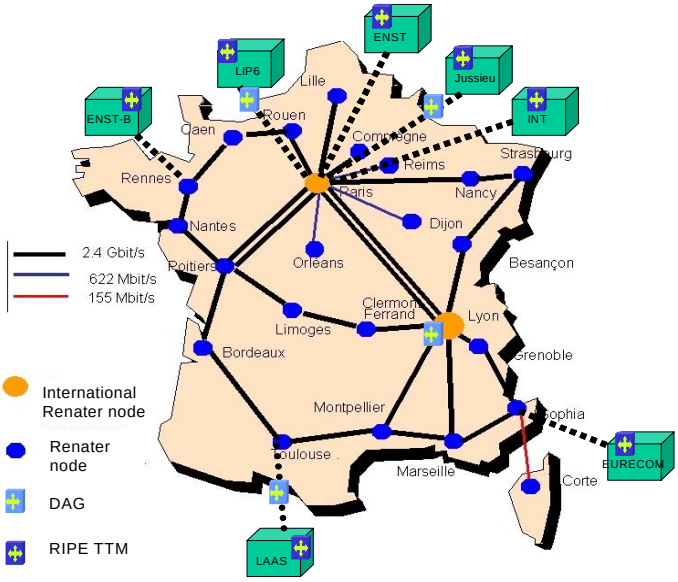

Figure 1. Metropolis monitoring and measurement platform.

2.2. Classification method

Our method for classifying packets and flows among applications is performed by using the QoSMOS[5] Traffic Designer (TD) box [QoS04]. Traffic Designer is a traffic monitoring box, namely a PC equipped with a bypass Ethernet card, which performs a real time analysis of ongoing packets. The TD software then provides macroscopic statistics of traffic. The monitoring box is also designed for re-shaping traffic according to measurements and the traffic policy defined by the network manager. This functionality is not used in the work presented in this paper. Only the monitoring facility is of interest and more specifically the capability of classifying traffic via the recognition of protocols used by applications. The recognition mechanism relies on the payload analysis of packets. First, flows are separated using the classical 5-tuple information (IP addresses, transport protocol, transport ports). Then the content of packets belonging to a given flow is analyzed until the flow is associated with an application. A flow is said to belong to a specific application only if the syntax of the data exchanged between the two TCP peers matches to the application syntax. For instance if a flow is transmitted on port 21 and contains something like "GET /index.html HTTP/1.1", Traffic Designer is going to classify it HTTP whereas a standard port classification would describe it as FTP.

Such a method proves much more efficient than the simple classification method based on standard IANA port numbers. Indeed, many applications are now using dynamic port negotiation. Thus, even if ports can be used to identify control connections, those flows with dynamically negotiated port numbers, used for data transfer, are not going to be recognized. FTP is a very good example of this kind of behavior. In this case, the content of control connections is analyzed in order to find out the negotiated port numbers and to be able to subsequently classify data flows. For instance, if a flow is recognized as an FTP control flow, the classification engine will search for packets containing the “PORT” command which gives information on the setup of the data connection. In addition, more and more applications are

[5] QoSMOS is a spin-off of LIP6 laboratory located in Paris.

not using standard ports at all. Finally, some applications may even use a well-known port for a different usage, for instance in order to bypass firewall rules.

The main advantage of using an applicative classification tool is that it significantly reduces the amount of non-classified traffic. Thus, for the traffic traces from the Renater network, 30 % of traffic would not be classifiable when using standard port numbers, compared to less than 5 % when using Traffic Designer.

Since the QoSMOS TD can be run only with Fast Ethernet cards, we solve the problem of monitoring Gbit/s transmission link by using another QoSMOS software called TD player which allows us to replay traffic collected with DAG passive monitoring boxes. DAG systems are able to collect all packets on gigabit links without skipping one of them, and then by replaying this traffic trace at an acceptable speed for the TD box, we can get a very accurate classification of traffic according to applications. The following section presents the results obtained via the above measurement method on traffic traces from Renater and the France Télécom IP network.

3. Classification results

This section presents the results obtained by monitoring and classifying traffic among applications. We give some temporal breakdowns of the evolution of traffic distribution, as well as pie charts for those readers who are only interested in global cumulative results. All these classification results are presented in terms of number of bytes, packets and flows (e.g., TCP connections). Also, this section gives the flow size distribution for all applications, which are detected in analyzed traffic. This is motivated by the fact that references [PAR96] and [PAR97] show that the change of the distribution of flow size, which is becoming more and more heavy tailed, has a strong impact on the nature of traffic, in particular by giving rise possibly to self-similarity and at least to long range dependence properties (LRD) [OWE04]. It is then of particular interest to check whether these observations hold for all traffic types.

3.1 Data from Renater POP in Jussieu (Paris)

For collecting the data described in this paper, we captured several traffic traces in the Jussieu campus network. All traces gave rise almost the same amounts of data, packets and flows, as well as the same distribution of traffic among applications. We then arbitrarily selected one of them, and this section provides the breakdowns and pie charts obtained via our applicative classification methodology and tools described in the previous section. The traffic trace has been captured by using a DAG card for both incoming and outgoing traffic. To determine the direction in which packets were sent (i.e., from the university network or towards this network), we have used the MAC addresses of the edge router. The measurement experiment for the Jussieu campus network is globally characterized as follows (see Table 1):

- **Capture date:** Monday October 11th, 2004
- **Location of the measurement device:** campus network of Jussieu
- **Start time:** 2:50 pm
- **Duration:** 3600 seconds
- **Total number of packets:** 80,437,378
- **Total number of flows:** 2,322,931

Table 1. Number of flows of main Internet protocols.

Name	Flows	Flow Ratio	Volume Ratio
TCP	1,745,870	75.2 %	92 %
UDP	492,732	21.2 %	7.4 %
ICMP	84,269	3.6 %	0 %
ESP	28	0 %	0 %
GRE	15	0 %	0.6 %
IPv6	7	0 %	0 %
IGMP	7	0 %	0 %
PIM	2	0 %	0 %
HOPOPT	1	0 %	0 %

Since we want to analyze full flows in order to get an accurate flow size distribution according to applications, we have to focus on TCP flows starting and finishing within the experiment time window. A flow is said to have ended if its last TCP flags matches a connection termination (either gracefully with a FIN – ACK sequence or abruptly with a RST) or if no packets have been seen for this flow for more than 4 minutes.

As an indication, flows started before the beginning of trace represent 0.5 % of TCP flows (20 % of the total volume), whereas flows still active after the last packet of the trace represent 0.6 % of the flows (10 % of the volume). While the difference in terms of flows is negligible, the gap in terms of volume is relatively important. This bias explains why the traffic rate seems to be small at the beginning of the traces in the graphs.

Figures 2 and 3 present the different breakdowns and pie charts obtained via the classification process of traffic among applications. Figure 2 gives for all applications the total amount of traffic in terms of bytes, packets and flows. The same is shown in Figure 3 by considering family of applications. We call an application family all applications that have the same purpose: for instance, Kazaa and E-donkey belong to the P2P family.

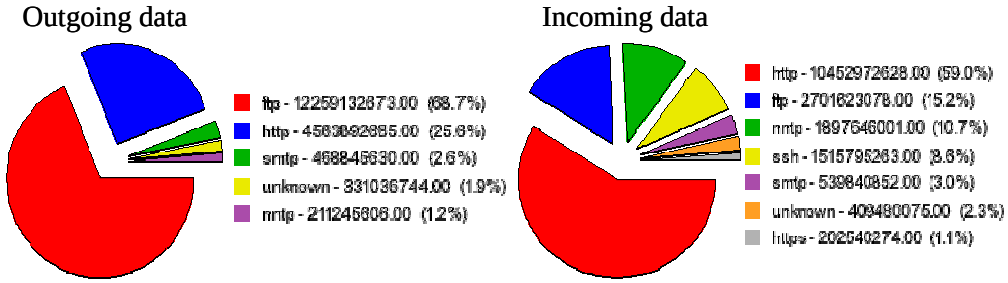

Figure 2. Global amount of traffic per application (outgoing and incoming data). These pie charts represent the distribution of traffic among applications in term of bytes

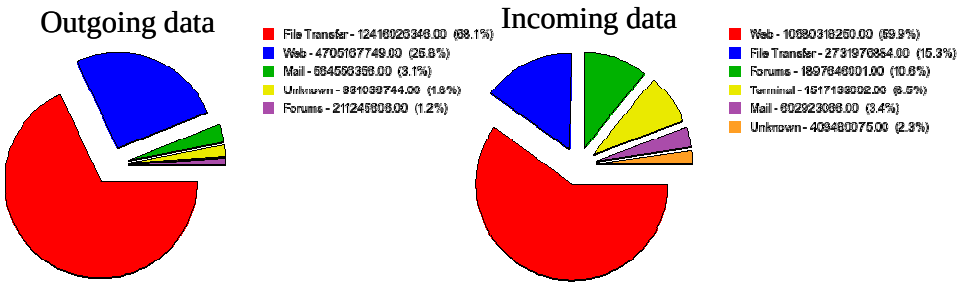

Figure 3. Global amount of data per applicationl families (outgoing and incoming data). These pie charts represent the distribution of traffic among family of applications in term of bytes

Figures 4 and 5 depict over the time evolution of the bit rate (in Byte/s) per application and family of applications, respectively. Figure 6 gives the bit rate in packet/s. Figure 7 shows the application flow rate in flow/s.

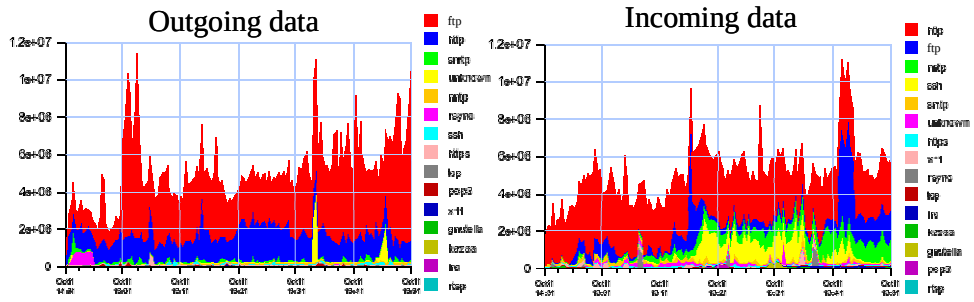

Figure 4. Application throughput breakdown in Bytes/s (outgoing and incoming data). These breakdowns represent the time evolution of the bit rate of the principal applications of the Jussieu's traffic

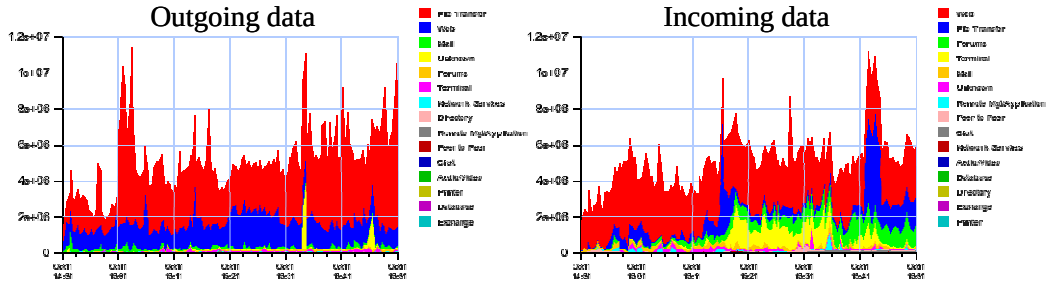

Figure 5. Application family throughput breakdown in Bytes /s (outgoing and incoming data). These breakdowns represent the time evolution of the bit rate of the principal contributing families of applications of the Jussieu's traffic.

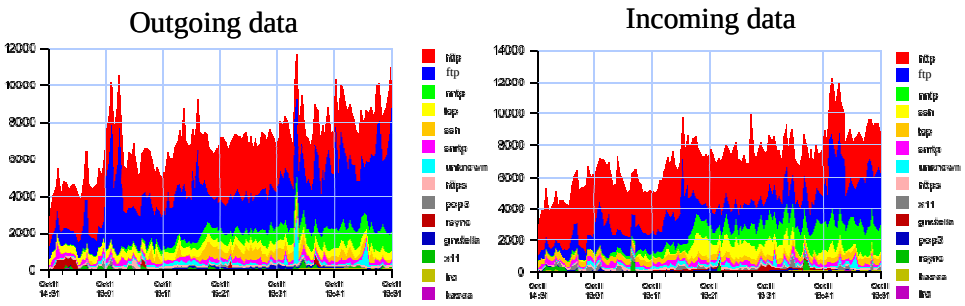

Figure 6: Application packet rate breakdown (in Packets/s) (outgoing and incoming data). These breakdowns represent the time evolution of the bit rate in packet/s due to the main contributing applications of the Jussieu's traffic

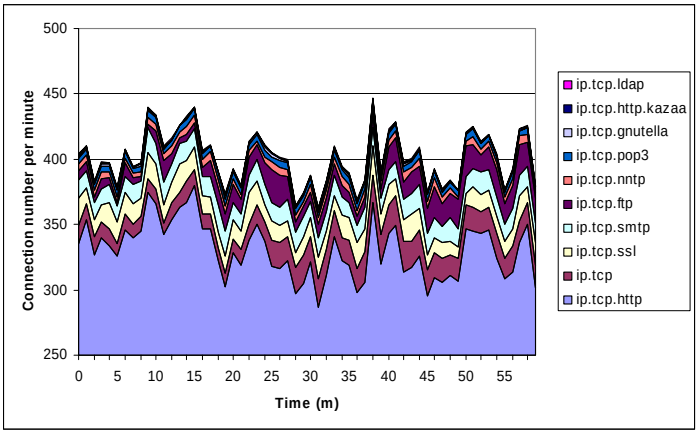

Figure 7. Application flow rate breakdown (in flows /s). This breakdown represents the time evolution of the number of new flows due to the main contributing applications of the Jussieu's traffic. ip.tcp flows are flows where no application data was exchanged (for instance connection attempts that did not go through).

Figure 8 represents the flow size distribution for each application with a log-log scale. The objective is precisely to evaluate whether flow size distributions are heavy tailed or not. Such a characteristic is important as heavy tailed distributions of flow sizes infer LRD and might reveal self-similarity properties of traffic. These two properties have been shown in the literature to impact many performance issues for the network. In fact, Figure 8 more or less confirms that applications which generate long flows (a.k.a. elephants) likely have a heavy tailed flow size distribution, as already demonstrated in [PAR96]. In particular, it is important to note that the most heavy-tailed distributions are not those of P2P or FTP applications, but those of streaming oriented applications (based on RTSP), terminal applications, and chat. Nevertheless, it seems that in spite of these differences in flow size distributions, P2P and File transfer applications have more impact on LRD than other applications (see table 2).

For more information on LRD (Long Range dependence), the reader can refer to the huge existing literature, for example [PAR00] [LEL93] [WIL95] To quantitatively characterize the LRD level of current Internet traffic, we have performed a wavelet based decomposition of traffic traces. Interested readers can refer to [ABR98] for details on the wavelet based method used in this paper, as well as on the

LDestimate tool that provides LRD diagrams. From these LRD diagrams it is then possible to compute the Hurst parameter (H), which characterizes the long range dependence of the traffic process. The parameter H has in general a value between 0.5 and 1: The value 0.5 indicates that there is no dependence in the process. The value 1 indicates very strong long range dependence. Sometimes, the LDestimate tool yields H values larger than 1. This is due to approximation and roud off erros when the number of values is not sufficient for having small confidence intervals. Nevertheless, even if the value is not very accurate, the qualitative indication for strong LRD still holds.

Results in table 3 show that even if Web distribution is not as heavy-tailed as the global distribution, the LRD level is really higher. Moreover, when we analyze Web traffic, we can notice that 0.00058 % of the total number of Web flows generate more than 37.3 % of the total amount of Web traffic (cf. table 3 for details). Hence, it is clear that Web clients observed in this trace have not only browsed Web pages. In fact, a more detailed analysis shows that Web clients have downloaded large files, which means that HTTP protocol is used as a FTP-like protocol. Table 4 generalizes this result for the whole set of flows (and not only web flows). In fact, LRD is introduced by the largest flows, independently of the associated application. As demonstrated in [PAR97], those flows are responsible for the most significant degradation of network performance.

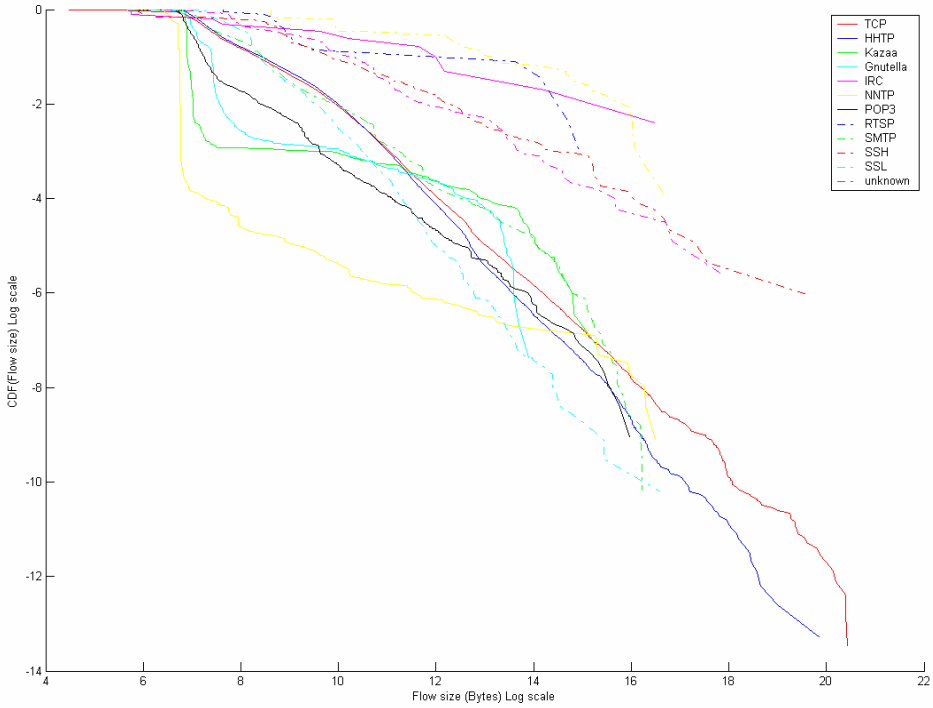

Figure 8. Flow size distribution for each application

Table 2. LRD inferred by some application families

Application family	LRD (Hurst factor)
Web traffic	0.905
P2P traffic	1.010
Terminal traffic	1.15
FTP traffic	1.245

Mis en forme : Paragraphes solidaires, Lignes solidaires

Mis en forme : Paragraphes solidaires, Lignes solidaires

Mis en forme : Paragraphes solidaires, Lignes solidaires

Table 3. Relation between file sizes and LRD for web traffic

	Volume (%)	Number of flows (%)	LRD Level (Hurst value)
Total Web traffic	14.4 GBytes (100%)	1 167 759 (100%)	H=0.905
Traffic due to large web flows (size> 1 MBytes)	5.31 GBytes (37.3%)	679 (0.00058%)	H=1.011
Traffic due to the 100 largest Web flows	3.99 GBytes (27.7%)	100 (0.000086%)	H=1.201

Table 4. Relation between file size and LRD for all flows

Flow sizes (MBytes)	Total volume (GBytes)	Total volume percentage (%)	LRD level (Hurst value)
> 0	43.4	100%	0.855
Between 1 and 10	6.56	15.1%	0.814
Between 10 and 50	5.34	12.3%	0.834
Between 50 and 100	1.91	4.4%	0.873
Between 100 and 300	5.39	12.4%	1.03
> 300	14.1	32.5%	1.07

3.2. Data from France Télécom's network (link connecting several ADSL areas in Paris)

This section presents the same analysis as that for Jussieu's traces, but for a France Télécom trace. The global characteristics, in particular the contribution in terms of flows of the different protocols (see Table 5), are as follows:

- **Capture date:** Thursday October 15th, 2004
- **Location of the measurement device:** in derivation of a high speed link connecting different ADSL areas in Paris; only downstream traffic is observed
- **Start time:** 19:01 PM
- **Duration:** 1300 seconds
- **Total number of packets:** 134,434,541
- **Total number of flows:** 9,636,105

Table 5: Number of flows of main Internet protocols:

Protocol	Number of Flows	Flow Ratio	Volume Ratio
ICMP	121,378	1.26 %	0.04 %
TCP	4,805,502	49.87 %	87.85 %
UDP	4,709,060	48.87 %	11.95 %
ESP	91	0 %	0.14 %
GRE	38	0 %	0.03 %
IPv6	26	0 %	0 %
AH	10	0 %	0 %

Figures 9 and 10 give the contributions to the global volume of the different applications and families of applications, as defined in the previous section. It clearly appears from these two figures that P2P applications have significant contribution to the global load. This is the key difference between commercial and campus traffic. In campus traffic, as observed in the traffic trace of the Jussieu campus network, P2P traffic is almost inexistent. In a commercial network, P2P has a major impact. We shall see that this point is important when examining flow size distributions.

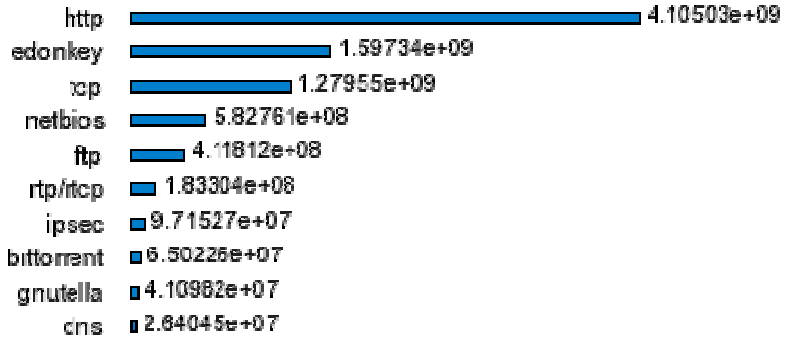

Figure 9. Global amount of data per application (incoming data). This histogram represents the distribution of traffic among applications in term of bytes

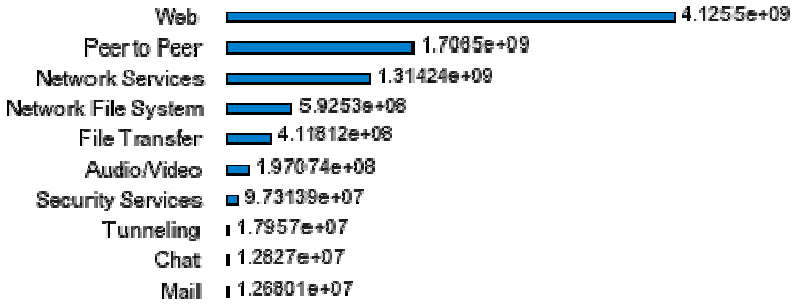

Figure 10. Global amount of data per protocol family (incoming data). This histrogram represents the distribution of traffic among family of applications in term of bytes

Figures 11, 12, and 13 depict the bit rate for applications and families of applications. These figures show that the global bit rate is roughly stationary; there is no evidence of a significant drift in the bit rate processes. It is also worth noting that the proportion of P2P traffic is increasing so that its contribution to the global load becomes more and more significant. Several traffic traces show that P2P traffic is dominating in the evening (say, between 9:00 pm and 12:00 pm). Figure 14 depicts the flow rate; it shows that the situation is stable during the observation period. It clearly appears from this figure that the number of TCP flows is very high and that a large number of flows do not lead to the establishment of TCP connections. This phenomenon is characteristic of P2P protocols. Those protocols indeed generate a huge amount of signaling in the form of small flows of a few packets (a.k.a. as mice), related to file search or maintenance in P2P networks. Moreover, since the allocation of IP addresses for ADSL customers is dynamic, many peers outside the observed ADSL areas try to contact ADSL peers, when they are not active or not connected.

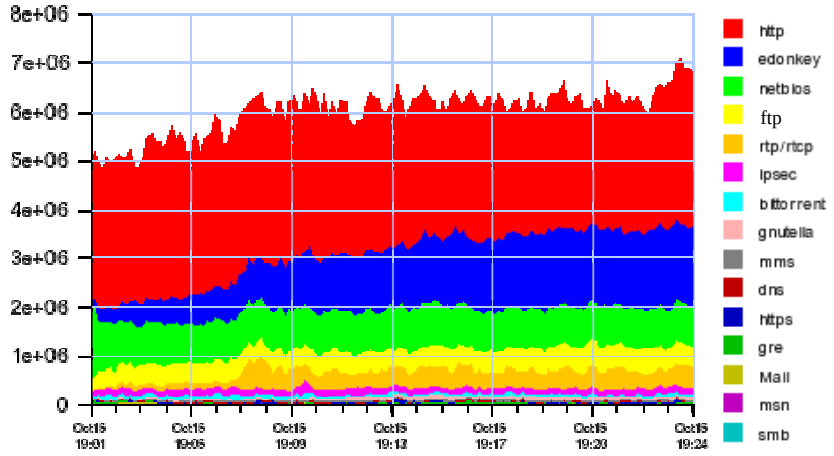

Figure 11. Application throughput breakdown in Bytes /s (incoming data). This breakdown represents the time evolution of the byte traffic due to the main contributing applications of the Fontenay's traffic

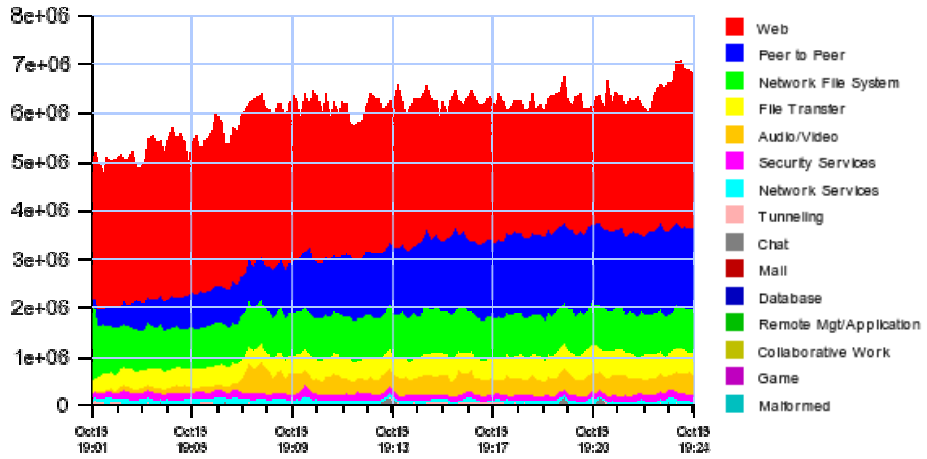

Figure 12. Application family throughput breakdown in Bytes/s (incoming data). This breakdown represents the evolution of the byte traffic during time due to the main contributing families of applications of the Fontenay's traffic

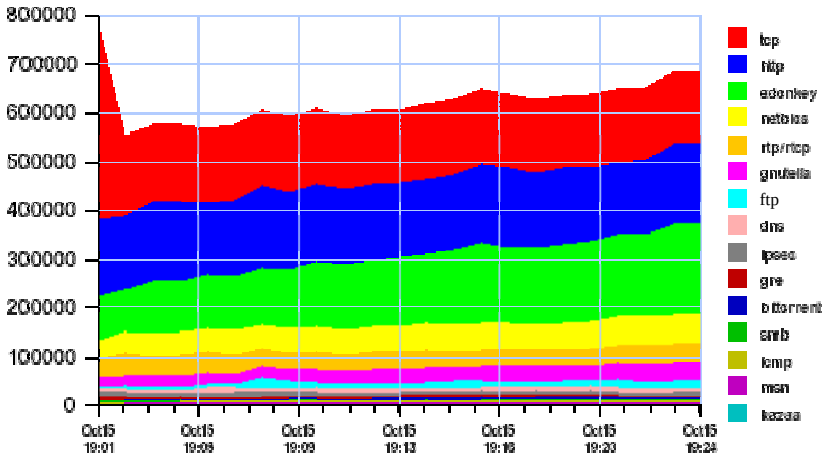

Figure 13: Application packet rate breakdown (in Packets/s) (incoming data). This breakdown represents the evolution of the packet traffic during time due to the main contributing applications of the Fontenay's traffic.

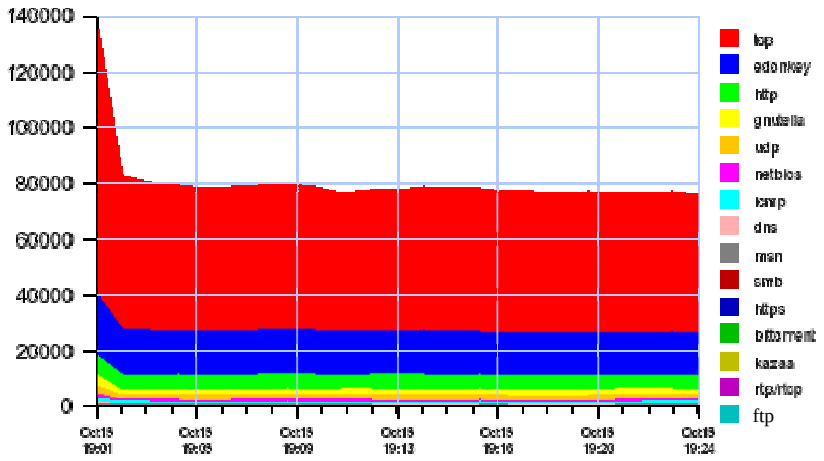

Figure 14. Application flow rate breakdown (in flows /s) (incoming data). This breakdown represents the evolution of the number of new flows during time due to the main contributing applications of the Fontenay's traffic. tcp flows are flows where no application data was exchanged (for instance connection attempts that did not go through).

Figure 15 displays the flow size distributions for the different applications. It is worth noting that the distributions are quite different from those obtained for the traffic trace from the Jussieu campus network. In log-log scale, the distributions cannot be easily approximated by straight lines. A finer analysis (see [BEN04]) shows that the tails of the distributions can be approximated by Weibullian distributions. In addition, the distributions of the flow size for P2P applications exhibit a bimodal behavior. By examining more precisely the composition of P2P traffic, we come up with the conclusion (see [BEN04]) that a huge number of flows are very small, comprising less than 8 packets. As mentioned above, these flows are due to signaling in P2P networks (file search or maintenance). In fact, for a finer analysis, we are led to separate small flows (mice) from long flows (elephants). Because most recent P2P protocols, in particular eDonkey, divide long files into smaller files (chunks), which can be downloaded asynchronously and in parallel by peers, the predominance of P2P traffic in commercial networks tends to smooth traffic, in particular to eliminate long range dependence and a fortiori self-similarity. This is illustrated on figure 16 were it clearly appears that the flow size distribution of eDonkey traffic is less heavy than the one of Kazaa. By measuring the Hurst parameter, we have obtained a value equal to 0.797 for eDonkey traffic and a value equal to 1.01 for Kazaa, which shows better performance of the network for recent P2P protocols than for previous ones.

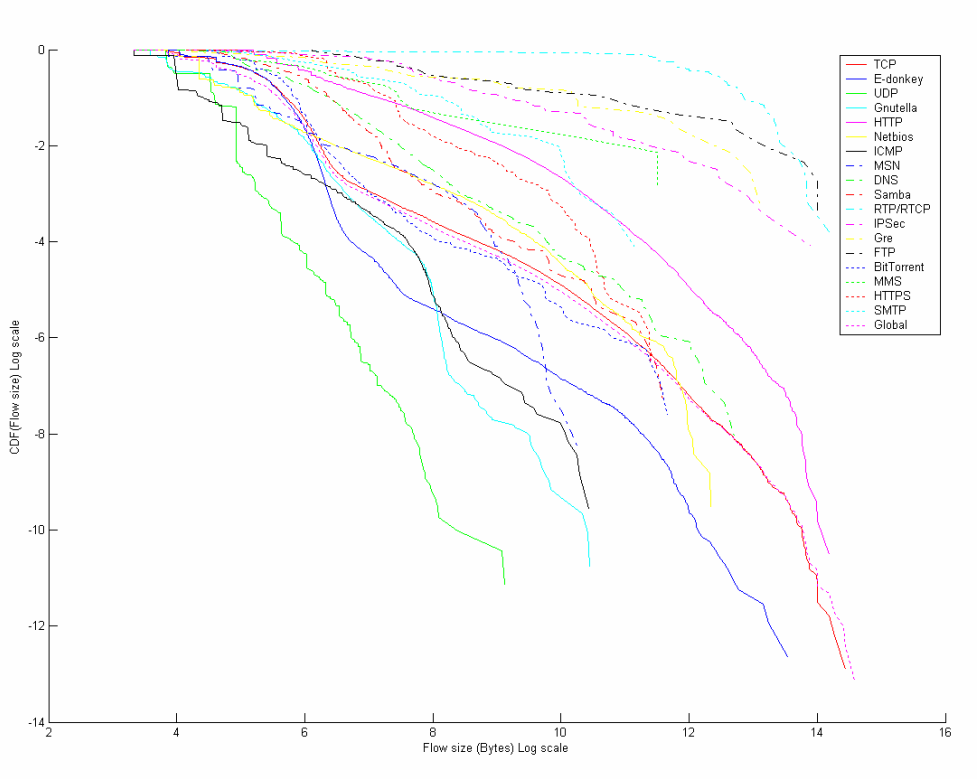

Figure 15. Flow size distribution for each application in France Télécom trace.

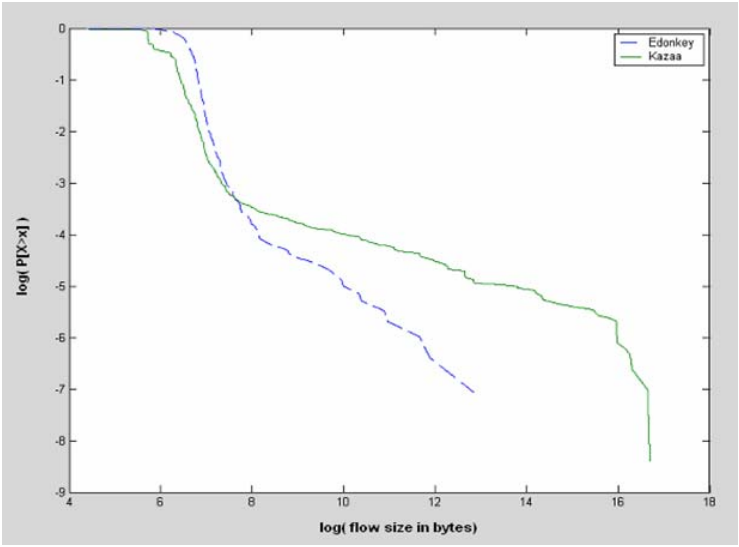

Figure 16. Flow size distribution for Kazaa vs. e-donkey flows

4. Conclusion

This paper investigates the distribution of traffic among applications in two different networks: a campus and a commercial network. For breakdowns of the distribution of traffic among applications, we have proposed a new methodology based on collecting traffic traces thanks to DAG equipments, and then on classifying traffic by using the QoSMOS TD box, which relies on applicative protocols recognition. We have exhibited some major difference in the usage of the two networks. The Renater network is devoted to education and research, excluding P2P exchanges that are not related to educational or research purposes. As far as we know, network administrators on the campus do their best to enforce this policy. In addition, since the campus of Jussieu in Paris is hosting one of the main FTP servers in France, mirroring many FTP servers in the world (GNU, Linux, etc.), the application generating the prevalent part of traffic is FTP. At the opposite, it clearly appears that the main part of traffic in commercial networks (such as France Télécom's networks) consists of P2P traffic.

It thus clearly appears from this paper that there does not exist a single type of IP traffic. While LRD and self-similarity were observed in LAN traffic of Bellcore in the mid 1990's, the composition of traffic in today's IP networks is quite different from the one in LANs in the mid 1990s. Hence, the conclusions valid for these kinds of networks at that time might not more be valid today. While LRD phenomena can still be observed in campus networks with a predominance of Web traffic, there is no evidence for this phenomenon in commercial networks because of the smoothing effect of P2P protocols.

The most important message of this work is to draw attention to the fact that traffic analysis should be based on flows and applications instead on the packet level. Analyzing traffic characteristics per application, such as correlation, LRD, self-similarity, etc. helps pointing out the applications and usage actually generating traffic as well as QoS and performance issues in the Internet. We expect, based on these results, to propose new solutions for improving network QoS and performance, as well as the way it is managed. The first result in this direction deals with segmenting large files before transmitting them onto the Internet. But other approaches are also investigated, for instance the ones proposing to replace TCP, which is responsible of performance degradation when it is used for transmitting large flows on high bandwidth networks [OWE04]. Some new protocols aiming at replacing TCP in the Internet are under study, such as DCCP at the IETF for instance, but describing such a protocol and its benefits is not in the scope of the present paper. It surely deserves a full paper for showing how DCCP succeeds in reducing LRD and performance degradation when it is used for transmitting large files.

5. References

[ABR98] P. Abry and D. Veitch, “Wavelet Analysis of Long Range Dependent Traffic”, Transaction on Information Theory, Vol.44, No.1, January 1998

[BEN04] N. Ben Azzouna and F. Guillemin. Experimental analysis of the impact of peer-to-peer applications on traffic in commercial IP networks. : European Transactions in Telecommunications, Special issue on P2P networking and P2P services, ETT 15(6), November-December 2004.

[CLE00] J. Cleary, S. Donnely, I. Graham, A. McGregor and M. Pearson, “Design principles for accurate passive measurement”, PAM (Passive and Active Measurements) Workshop, Hamilton, New Zealand, April 2000

[LEL93] W. Leland, M. Taqqu, W. Willinger and D. Wilson, “On the self-similar nature of Ethernet traffic”, ACM SIGCOM, September 1993

[OWE04] P. Owezarski, N. Larrieu, “Internet traffic characterization -- An analysis of traffic oscillations”, 7th IEEE International Conference on High Speed Networks and Multimedia Communications (HSNMC'2004), Toulouse, France, July 2004

[PAR96] K. Park, G. Kim, M. Crovella, “On the relationship between file sizes, transport protocols, and self-similar network traffic”, IEEE ICNP, 1996

[PAR97] K. Park, G. Kim and M. Crovella, “On the Effect of Traffic Self-similarity on Network Performance”, SPIE International Conference on Performance and Control of Network Systems, November, 1997

[PAR00] K. Park and W. Willinger, “Self-similar network traffic: an overview”, In Self-similar network traffic and performance evaluation, J.Wiley & Sons, 2000

[QoS04] QoSMOS, “Traffic Designer”, http://www.qosmos.fr/EN/home.htm

[WIL95] W. Willinger, M. Taqqu, R. Sherman and D. Wilson, “Self-similarity through high variability: statistical analysis of Ethernet LAN traffic at the source level”, In ACM Sigcomm’95, 1995.