Internet Traffic Measurements over the Spanish R&D IP/ATM Network Backbone*

M. Alvarez-Campana[1], A.Azcorra[2], J.Berrocal[1], J.R. Pérez[1]

[1] Departamento de Ingeniería de Sistemas Telemáticos
Universidad Politécnica de Madrid.
ETSI Telecomunicación, Ciudad Universitaria, 28040 Madrid, Spain
E-mail: mac@dit.upm.es Tel: +34 91 5495700 Fax: +34 91 3367333
(http://www.dit.upm.es)

[2] Area de Ingeniería Telemática
Universidad Carlos III de Madrid, Spain
28911 Leganés (Madrid)
(http://www.it.uc3m.es)

Abstract

The increasing adoption of ATM technology in the main Internet backbones is introducing new factors of complexity in the network dimensioning, management, and operation fields. On the one hand, the behavior of the IP/ATM protocol stack is still not well understood. On the other hand, the aggregated traffic supported by high-speed Internet backbones responds to complex patterns which are highly variable and unpredictable.

Most of the works on Internet traffic characterization over ATM networks rely on theoretical studies. They are however slightly underwritten by measured results, which often explains why the real behavior differs from the expected one. This paper provides a summary of the traffic measurements collected during approximately on year over the Spanish R&D IP/ATM Internet Backbone (RedIris). This data collection has served us to evaluate the performance of the IP/ATM protocol architecture under a real network scenario. We have also performed traffic characterization study identifying the most frequent Internet services and applications. The results have shown that some of them (e.g. the WWW) are not efficiently transported over ATM. In addition, we have obtained several other interesting data and conclusions that characterize the Spanish R&D Internet backbone. Some of them could be easily extrapolated to other Internet ATM backbones.

---

* This work has been partially supported by the CICYT project CASTBA (TEL-2608-E) within the Spanish R&D Program.

1 Introduction

Variability and unpredictability are probably the only well known characteristics of Internet traffic. These two aspects extraordinarily complicate the dimensioning, management, and operation of the underlying network resources. Such tasks cannot be successfully performed without a precise and up-to-date knowledge about what is going on in the network. This knowledge is especially critical in the case of the new high-speed ATM Internet backbones.

The traffic supported by Internet backbones responds to complex patterns [1], which are not well understood. This complexity is mainly caused by the aggregation of thousands of individual traffic flows generated by different types of users and applications [2]. The phenomenon becomes more pronounced as the bandwidth availability increases. This is the precisely the case of the ATM links that are being gradually introduced in the main Internet backbones. An additional source of uncertainty arises when considering the relatively new protocol architectures [3,4] used in these backbones.

To date, most of the studies on Internet traffic characterization over broadband networks rely on analytical models [5]. In practice, however, analytical results differ considerably from the real behavior observed in the network [6,7,8]. Once again, the reason must be found in the extraordinarily complex nature of Internet traffic, which hinders any serious attempt to obtain a reliable traffic model.

The above considerations have led researchers to perform Internet traffic studies based on measurements [9,10]. In this sense, it is worth to mention some recent works on Internet traffic characterization over IP/ATM backbones [11,12].

This paper provides a summary of the traffic measurements collected during approximately on year over the Spanish R&D IP/ATM Internet Backbone (RedIris [13]), within the framework of the CASTBA [14] project. This data collection has served us to evaluate the performance of the IP/ATM protocol architecture under a real network scenario. We have also carried out a traffic characterization study identifying the most frequent Internet services and applications. The results have shown that some of them (e.g. the WWW) are not efficiently transported over ATM. In addition, we have obtained several other interesting data and conclusions that characterize the Spanish R&D Internet backbone. Some of the results obtained in our study can be easily extrapolated to other Internet ATM backbones.

The rest of this paper is structured as follows. Next section describes the architecture of the RedIris backbone and the methodology employed to take the measurements. Sections 3 to 5 provide a summary of the most significant statistics resulting from the analysis of the different types of measurements carried out in the CASTBA project. Finally, the last section summarizes the main conclusions of our study and points out some future works.

2 Network Scenario and Measurement Methodology

RedIris is the Spanish R&D network that provides Internet access to the main academic and research organizations in Spain. This Internet backbone has recently experimented a technological transition from a classical IP architecture based on dedicated point-to-point links, to a network infrastructure based on IP/ATM connections.

Figure 1 shows the topology of the RedIris IP/ATM backbone consisting of 17 regional links interconnected to a central node located in Madrid. The regional links are supported by asymmetrical ATM circuits (with capacities between 2 Mbit/s and 8 Mbit/s) provided by Telefónica's GigaCom ATM public network. RedIris users (research and academic centers) access the network through the corresponding regional nodes via 64kbit/s-2Mbit/s dedicated links. In addition, there are three external links (not shown in the figure) that provide access to USA, Europe and the commercial

Internet. The traffic injected in the backbone is 100% IP, which is transported over the ATM links by encapsulating the IP packets over AAL5[3] frames according to the RFC 1483[4].

Figure 1. Topology of the RedIris IP/ATM backbone

Taking into account the star topology of the RedIris backbone, and in order to simultaneously take measurements from the 17 regional links, we decided to locate the monitor point at the RedIris central node in Madrid. Figure 2 shows the configuration of the measurement equipment. A Chameleon ATM protocol analyzer [15] with two ATM interfaces (one for each transmission direction, that is, input and output to/from each regional link) was inserted between the central ATM switch and the RedIris central router. To avoid interfering with the normal function of the network, two optical splitters were inserted so that the traffic is replicated to the ATM analyzer. In order to achieve a high traffic capture ratio, it was decided to use a separate machine (a standard PC) for data analysis. This allows the ATM analyzer to concentrate on the acquisition of traffic samples.

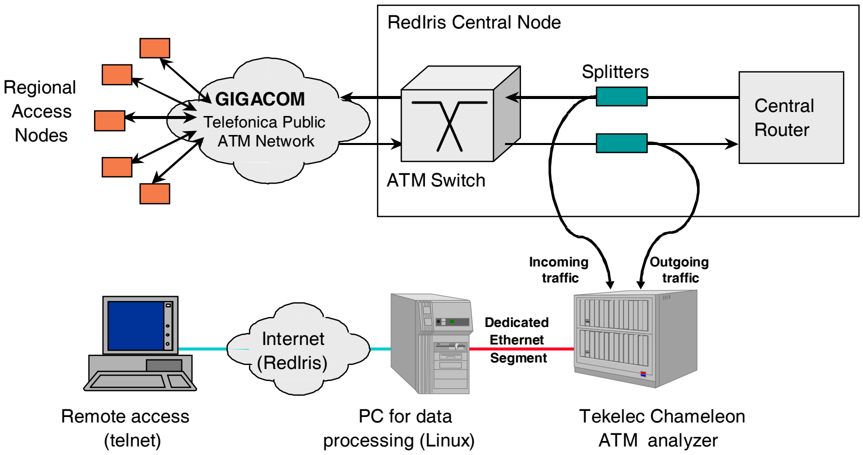

Figure 2. Configuration of measurement equipment

The data collected by the ATM analyzer are periodically transferred to the processing machine via a dedicated Ethernet segment. To automate the capture and analysis procedures, we developed a number of applications both in the processing machine and the ATM analyzer [16]. Note an additional link providing remote access to the processing machine. This way, the system can be remotely operated and configured, and the results can be downloaded to a local machine for further analysis and interpretation.

The following three sections provide a summary of the main statistics collected within the CASTBA project by using the measurement infrastructure above described. The results presented correspond to average values for different measurement periods from March 1997 to November 1997. A more exhaustive collection of results can be found in [17].

3 Traffic load and network performance measurements

Figure 3 show the overall average utilization of the RedIris backbone during a working day (graph 3a) and during a week (graph 3b). Note that we have represented separately the results corresponding to the incoming traffic (from the central node to the regional node) and the outgoing traffic. As a reference, a utilization of 100% corresponds to 68Mbit/s and 40Mbit/s of aggregated capacity for the incoming and outgoing links, respectively.

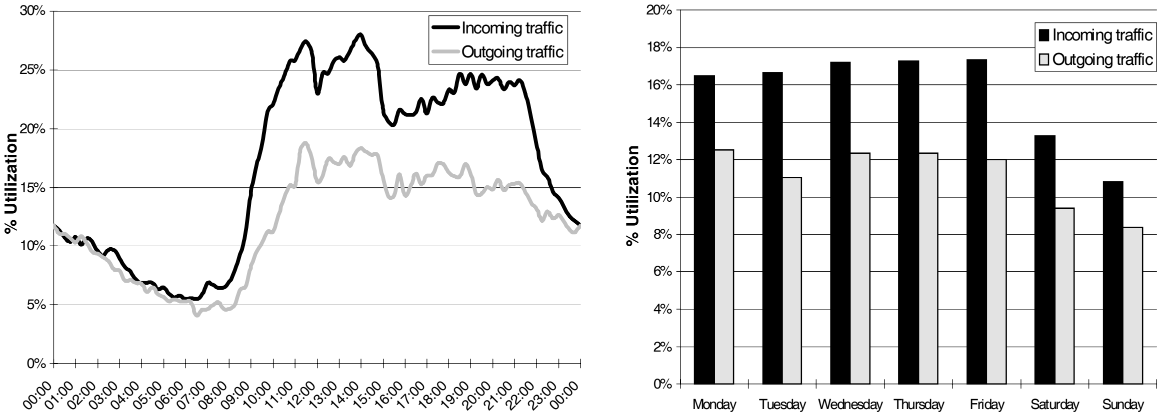

(a) **(b)**

Figure 3. Average utilization of the RedIris backbone a) during a working day b) during a week

Graph 3a shows that the maximum average utilization of the RedIris ATM links is approximately 25% (approx. 17 Mbit/s) for the incoming direction and 17% (6.8Mbit/s) for the outgoing direction. These occupancy levels correspond to the busy-hour period, which can be established from 9am to 9pm. It should be pointed out that we did not observed peaks of traffic that surpassed significantly the average utilization values shown in the graph.

The graphs in figure 4 show the average daily traffic volumes and occupancy levels for each of the 17 regional links[1] during a working day. In figure 4a, we observe that the incoming traffic (146 Gbytes, in total) nearly triple the outgoing traffic (59 Gbytes). This suggests that the RedIris regional nodes behave as sinks of information coming from the external links to USA, Europe and the commercial Internet.

---

[1] It should be mentioned that some of the links (e.g. CLM and RIO) were not fully operational during the sampling periods.

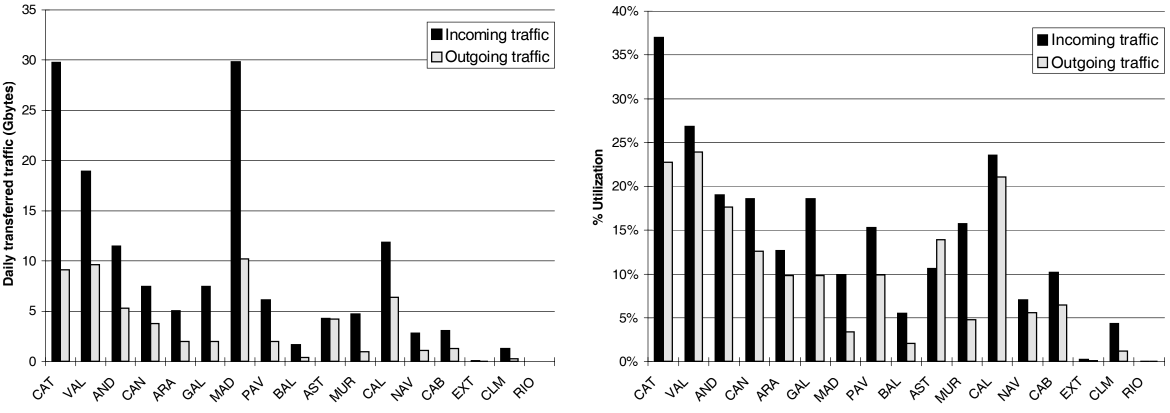

(a) **(b)**

Figure 4. a) Daily traffic volume by regional link b) Average link utilization

Regarding the network performance, we obtained several types of measurements to evaluate the quality of service provided by the RedIris backbone. We measured the end-to-end packet delay and the packet loss ratio. Both parameters resulted quite low and stable (average delays of 26ms for packets of 1Kbyte, and loss ratios below 0.01%). This is easily explained if we consider the moderate occupancy levels of the ATM links (see figure 4b).

Figure 5 shows one of the most interesting results obtained in the CASTBA project. The graphs prove the relatively high inefficiency of the IP/ATM protocol stack for transporting Internet traffic. This inefficiency is measured in terms of the overload added by the AAL5 and ATM layers.

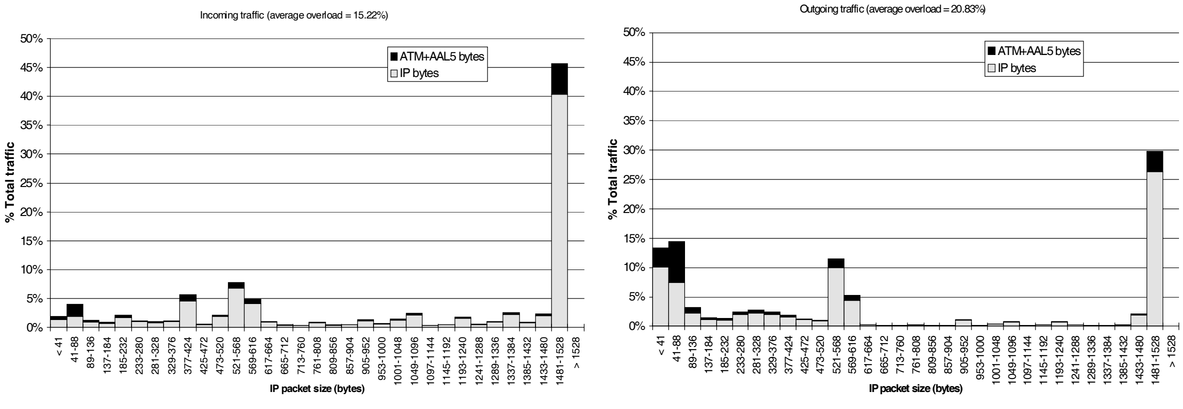

(a) **(b)**

Figure 5. IP/ATM overload a) incoming traffic b) outgoing traffic

The graphs in figure 5 shows the percentages of bytes corresponding to useful traffic (IP) and overload traffic (ATM and AAL5), as a function of the IP packet size. Each column in the histograms corresponds to the set of IP packet sizes that give rise to the same number of ATM cells. The results indicate a bandwidth waste of approximately 15% for the incoming direction (figure 5a), and 21% for the outgoing direction (figure 5b). Fortunately, RedIris makes use of the VC based multiplexing method for encapsulating IP/ATM instead of the LLC encapsulation method. Otherwise, the inefficiency would be even higher.

As it will be seen in the following section, the packet size distributions for incoming and outgoing traffic are different. The latter has a higher proportion of short packets (TCP control datagrams and request packets) which precisely support the higher overloads. This explains why the IP/ATM inefficiency is higher for the outgoing traffic than for the incoming one.

4 Analysis of IP packet size distributions

A second group of measurements of CASTBA focused on the characterization of Internet services by analyzing the packet size distributions. Figure 6a shows the distribution of IP packet sizes (grouped by AAL5 frame sizes) for one of the most loaded RedIris links. Figure 6b shows a zoom for packet sizes up to 100 bytes. In both graphs, the results are expressed as the percentage of packets over the total of captured packets. It should be pointed out that the values represented do not include fragmented IP packets. Nevertheless, we checked that these type of packets account for less than 0.05% of the total.

In figure 6a, we can observe three main groups of packet sizes: small packets (1 to 3 ATM cells), large packets (32 cells), and medium sized packets (9, 12 and 13 cells). Note that the proportion of small packets is quite high. More than 40% of the incoming packets and more than 75% of the outgoing packets are smaller than 100 bytes. As we anticipated in the previous section, this is the reason why the IP/ATM inefficiency results so high.

When analyzing the size distribution for small packets (figure 6b), we appreciate a dispersion caused by the mix of IP services and applications supported by the RedIris backbone and their different control PDUs. We can mention, for instance, that the typical size of TCP control packets is 40 bytes, the telnet protocol is the main responsible for the value 41, and the http protocol (web) for the value 44.

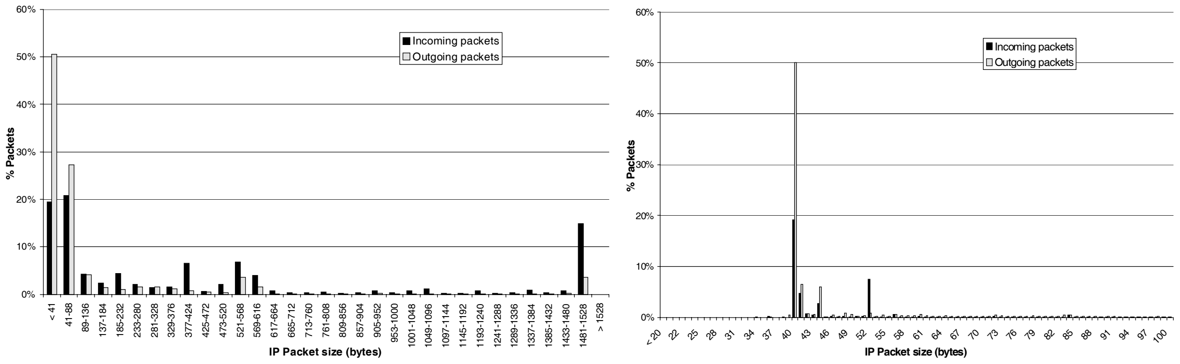

(a) **(b)**

Figure 6. IP packet size distribution a) grouped by AAL5 frame sizes b) zoom for packets up to 100 bytes

Medium sized packets are due to the mechanisms that TCP implementations usually apply to minimize packet fragmentation. For example, many BSD implementations require a MSS (maximum segment size) multiple of 512 bytes. If we add to this value the 40 bytes added by the IP and TCP headers, we obtain 552 bytes, which falls under the column 521-568 of figure 6a. On the other hand, if after the establishment of a TCP connection the peer entity does not specify a MSS, the default value 536 bytes. If we add the 40 bytes of IP/TCP, we obtain 576 bytes (column 569-616 in graph 6a) which correspond to the minimum MTU (maximum transmission unit) recommended for IP. This explains the low proportion of fragmented datagrams found in the captured traffic. We noted also a non-negligible percentage of packets in column 373-424, which are due to multicast traffic.

Large packets (1481-1528 bytes) correspond to the 1500 byte MTU size of Ethernet interfaces, one of the most popular types of LAN. The corresponding MSS is 1460 bytes, which is the default value for local address announced by many systems (SunOS, Solaris, AIX,...)

Figure 7 shows the average packet sizes for some of the main Internet services used in RedIris. The percentages of packets captured for each of the services are represented in figure 7b. As we see in graph 7a, the larger packets correspond to services that imply a massive transfer of data: nntp, ftp, and http. Note that although these services have some common characteristics (they use the same transport protocol and give rise to bulk data transfers), their average packet sizes are significantly different. The reason must be found in the particular configurations used by the servers.

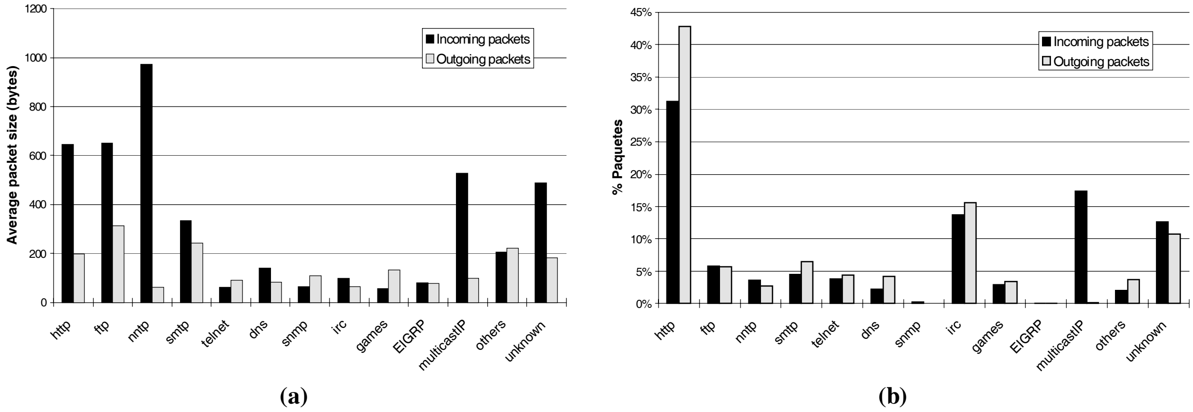

(a) **(b)**

Figure 7. a) Average packet size by service b) Percentage of captured packets by service

The graphs in figure 8 show the packet size distribution for some of the most popular services: http, ftp, telnet, smtp, nntp. The global packet size distribution is shown again in graph 8f in order to make easier the comparison. Note that all the services have a relatively high proportion of small packets, which are not efficiently transported by the IP/ATM stack. In particular, observe that the percentages of small and large packets for the http traffic are comparable. This case is especially critical, as the http traffic is the dominant in the RedIris backbone (see next section).

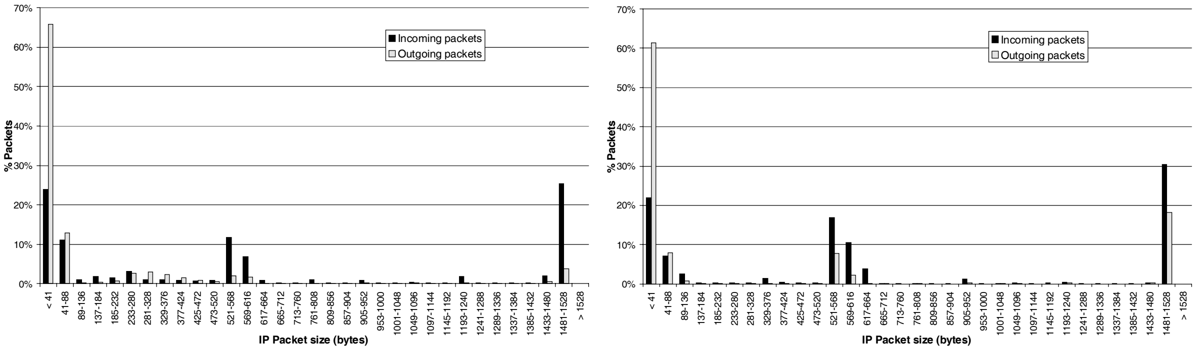

(a) http **(b) ftp**

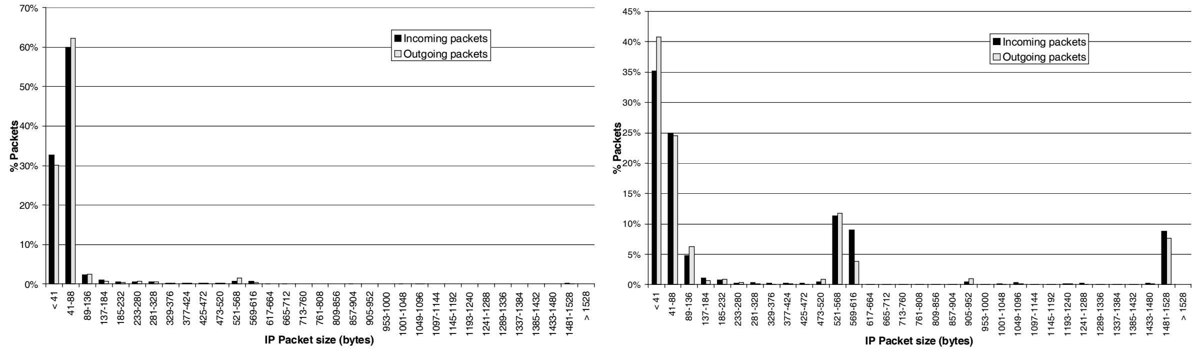

(c) telnet **(d) smtp**

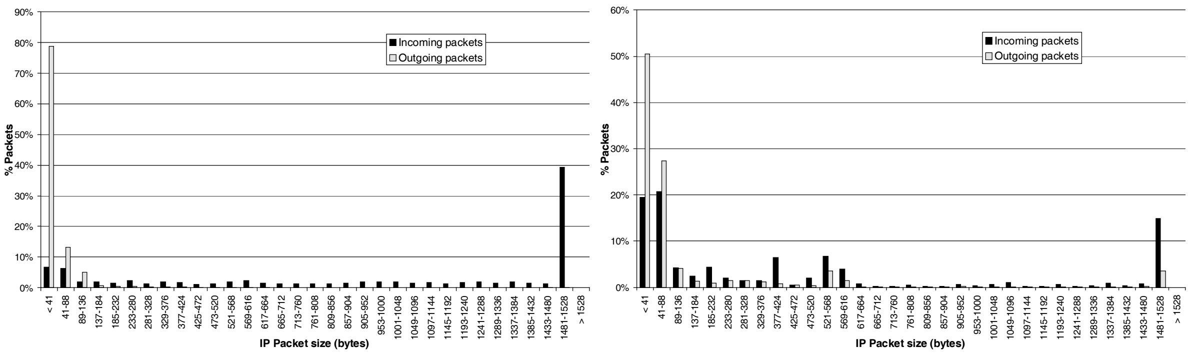

(e) nttp **(f) global**

Figure 8. Packet length distributions by services

5 Traffic characterization by services

This third block of results correspond to our first attempt to characterize the RedIris traffic by trying to evaluate the traffic load generated by the different services and applications. The methodology used in CASTBA for obtaining these data is based on the analysis of Internet protocol numbers and TCP/UDP ports. However, as it will be explained later, the results provided by this method can be regarded only as an approximation.

Figure 10 show the traffic load distribution by services during a working day. Graph 10a corresponds to incoming traffic and graph 10b to outgoing traffic. The results are expressed in Mbit/s. The measurements correspond to one of the most loaded regional links. Therefore, they can be considered sufficiently representative for the global characterization of the RedIris traffic.

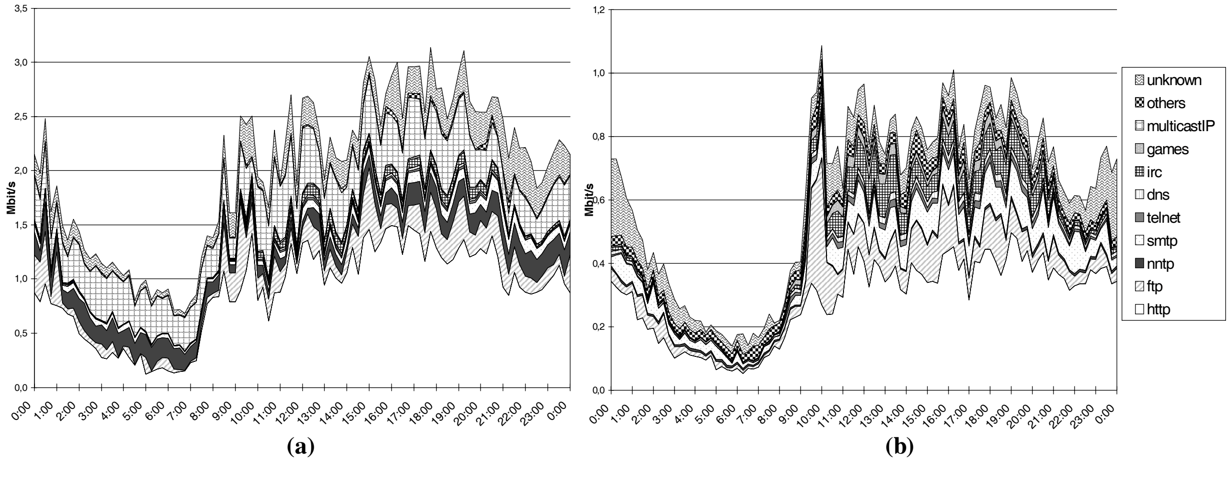

(a) **(b)**

Figure 10. Traffic load distribution by service a) Incoming traffic b) Outgoing traffic

The graphs in figure 10 indicate that the main service in terms of generated traffic is *http* (web) which accounts for 45% of the total. The second place is for *multicast traffic* (18% of incoming bytes). The third position is shared by *ftp* and *unknown* (see below) traffic which represent approximately a 10% of the traffic load each. Another relevant services by traffic load are: *smtp* (12% of outgoing traffic), *nntp* (7% of incoming traffic), and *irc* (6% of outgoing traffic).

As mentioned at the beginning of this section, the traffic characterization by services carried out in the CASTBA project cannot be considered 100% reliable. The reason is the existence of services making use of not registered ports and registered ports used for purposes other that the intended. The first aspect gives rise to an important fraction of traffic (10-20%) whose nature cannot be determined by a simple port analysis criteria. This is what we have called *unknown traffic*. The second aspect makes us suspect the existence of an undetermined amount of traffic classified under the wrong categories. These type of problems require a more refined traffic analysis methodologies, as the one considered in the MEHARI project [18], conceived as the continuation of the CASTBA project.

5 Conclusions and future work

Within the framework of the CASTBA project, we have obtained an exhaustive collection of traffic measurements over RedIris, the Spanish R&D Internet backbone. These data has served us to evaluate the performance of the new IP/ATM network infrastructure and to characterize the Internet traffic supported by RedIris.

The network performance measurements have proved the excellent quality of service provided by the RedIris backbone, which manifests in very low packet delay and packet loss values. However, we have observed a relative inefficiency of the IP/ATM protocol stack for transporting the Internet traffic. The results reveal that approximately 17% of the traffic on RedIris corresponds to overload bytes added by the AAL5 and ATM layers.

The traffic analysis by services has allowed us to identify the main Internet applications utilized by RedIris users and their importance in terms of traffic volumes. We have obtained also the individual packet size distributions for these services, which can be used as a starting point for defining Internet traffic models.

During the measurement period, the most frequent applications reported were the web (45% of the traffic), multicast traffic (18%), mail (12%) and news (7%). We must emphasize that this traffic distribution corresponds to a particular period of time. The constant proliferation of new Internet services will probably make obsolete these results in the short term. This consideration leads us to insist on the convenience of carrying on with this type of traffic measurements.

Finally, mention that the studies carried out in the CASTBA project are being continued in the MEHARI project [18]. This new project is aimed to development of an advanced Internet traffic monitoring and analysis platform. We hope that this new tool will solve some of the limitations of the measurement methodology used in CASTBA, allowing the realization of more sophisticated traffic analysis.

Acknowledgments

The CASTBA project has been funded by the *Comisión Interministerial de Ciencia y Tecnología (CICYT)* within the Spanish R&D Program (contract TEL-2608-E). The authors would like to thank this organism for the interest and motivation shown during the realization of the project.

We also want to thank our colleges at the Centre de Comunicacions Avançades de Banda Ampla (CCABA) of the Universitat Politecnica de Catalunya for sharing with us their ideas.

We are also grateful to RedIris and Telefónica I+D, which acted as promoter and observer entities of the CASTBA project. Finally, we would like to thank Tekelec España and Tekelec Inc. for their assistance in tuning up the Chameleon ATM Analyzer.

References

[1] W. Leland, M. Taqqu, W. Willinger, D. Wilson, “On the self-similar nature of Ethernet traffic”, IEEE/ACM Transactions on Networking, February1994.

[2] V. Paxson, "Growth Trends in Wide-Area TCP Connections", IEE Network, July 1994.

[3] ITU-T, “Recommendation I.363.5 - B-ISDN ATM Adaptation Layer specification: Type 5 AAL”, August 1996.

[4] J. Heinanen, “Multiprotocol Encapsulation over ATM Adaptation Layer 5”, RFC 1483, IETF, July 1993.

[5] A. Adas, "Traffic Models in Broadband Networks", IEEE Communications Magazine, July 1997.

[6] V. Paxson, S. Floyd. "Wide-Area Traffic: The Failure of Poisson Modeling". IEEE/ACM Transactions on Networking, June 1995.

[7] A. Romanow. S. Floyd, "Dynamics of TCP Traffic over ATM networks", IEEE Journal on Selected Areas in Communications, May 1995.

[8] E. Corvella, A. Bestavros: "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," IEEE Transactions on Networking, December 1997.

[9] V. Paxson, "End-to-End Internet Packet Dynamics", Proceedings of ACM Sigcomm '97.

[10] D. Endicott, V. Frost, "Performance Experiences with Wide Area High-Speed Networks", IEEE Communications Magazine, August 1997.

[11] JB.L. Tierny, W.E. Johnnston, J.R. Lee, G. Hoo, "Performance Analysis in High-Speed Wide Area IP-over-ATM Networks: Top-to-Bottom End-to-End Monitoring", IEEE Networks, June 1996.

[12] K. Thompson, G. Miller, R Wilder, “Wide-Area Internet Traffic Patterns and Characteristics”, IEEE Network Magazine, November 1997.

[13] RedIris, "Description of Network Services and Infrastructure", http://www.rediris.es, January 1997

[14] DIT-UPM, “Supervisión de Calidad de Servicios Telemáticos sobre la Red Académica de Banda Ancha", CASTBA-M Project, Spanish R&D Program, Final report, November 1997.

[15] "Chameleon Open/Open-X, ATM Application Module User's Guide", Tekelec Inc., November 1996.

[16] "Chameleon Open/Open-X, ATM C Development Libraries", Tekelec Inc., November 1996.

[17] M. Alvarez-Campana et al., "CASTBA: Medidas de Tráfico sobre la Red Académica Española de Banda Ancha", TELECOM I+D 98, Madrid, October 1998.

[18] DIT-UPM. “Mecanismos y Herramientas para el Análisis del uso de servicios Internet aplicado a RedIris", MEHARI-M Project, Spanish R&D Program, October 1998.