MEHARI: A System for Analysing the Use of the Internet Services

Pedro J. Lizcano, Comisión Interministerial de Ciencia y Tecnología (CICYT),
José Abascal 4, 08003 Madrd, Spain

Arturo Azcorra, Universidad Carlos III de Madrid (UC3M),
Butarque 15, 28911 Leganes (Madrid), Spain

Josep Solé-Pareta and Jordi Domingo-Pascual, Universitat Politècnica de Catalunya (UPC),
Jordi Girona 1-3, 08034 Barcelona, Spain

Manuel Alvarez-Campana, Universidad Politécnica de Madrid (UPM),
ETSI Telecomunicación, Ciudad Universitaria, 28040 Madrid, Spain

Abstract. This paper describes the MEHARI system, a system for monitoring and analysing the traffic of Academic Networks. The major design requirement of the MEHARI system was to cope with the main challenges that currently have most of these networks (service usage, charging schemes, network dimensioning, improper usage, etc.). The resulting MEHARI system consists of a low cost hardware traffic capture platform, and several traffic analysis software modules running on top of this platform. These modules may report information on the usage of the Internet services, the main traffic origin and destinations, etc. The MEHARI system has been tested running a field trial on the Spanish academic network (RedIRIS). Some relevant data on the performance and efficiency of the system configuration used in the field trial is presented in this paper.

Keywords: Internet Services Monitoring, Academic Networks Auditing

1. Introduction

The growth in number of users and traffic levels has led to the deployment of high-speed ATM-based Internet backbones in the main academic and research networks. Bandwidth availability, far from solving all the problems, is introducing some new ones. To begin with, it is giving rise to the birth of new applications creating the need for even higher capacities. The proliferation of applications is, in turn, altering continuously the Internet traffic patterns. Besides, the networks lack the technical means to enforce acceptable usage policies, leading to claims from some sectors about inappropriate usage of public resources.

The combination of all the above problems complicates the dimensioning, management, and operation of Internet backbones. These tasks cannot be successfully performed without a precise and up-to-date knowledge about what is going on in the network. There are two different dimensions on this problem: one is under the *traffic* perspective, and the other is on the *purpose usage* perspective. The former may need to do contents analysis in those cases in which application identification by port is not possible, and the its results are reasonably objective. The latter surely needs content analysis, and its results are based on heuristics and subjective techniques.

On the traffic analysis aspect, it is worth to mention some recent studies about the characterization of Internet traffic [1,2]. The traffic measurements obtained in these studies have revealed some interesting results about the Internet traffic patterns on IP/ATM backbones [3]. The main conclusion is, however, that Internet traffic is highly variable and unpredictable and, therefore, the results of traffic analysis are only valid in the short-term.

On the purpose usage aspect, different heuristic techniques may be combined depending on the nature of the traffic classification desired. As content analysis is costly in terms of CPU, a balance has to be established between confidence in the objectivity of the results and processing power required. It is

recommended that heuristics be manually checked from time to time in order to make their analysis more tight or relaxed, depending on the deviations between real traffic and the diagnosis performed by the tool.

The applications of the data obtained may serve different purposes, to cite some:

- Network engineering to adjust capacity to traffic matrix and hourly profiles.
- Characterisation and analysis of users (or user groups) traffic.
- Identification of the information services, information servers and client sites available in the network.
- Classification of the network traffic by application and furthermore by purpose usage.
- Validation the appropriate usage policy.
- Billing and charging.
- Detection and identification of security threats and attacks [4].

This paper describes the design of the MEHARI system, a high performance traffic processor to analyse the headers and contents of Internet traffic. The system characteristics are modularity, scalability, high-performance, low-cost, flexibility and adaptability. The MEHARI hardware is based on low-cost PC components, and the application processing may be done on any UNIX machine. The MEHARI platform consists of PCs with STM-1 cards for traffic capture and pre-processing. The pre-processed traffic is fed from the platform to the MEHARI traffic analysis software modules. The MEHARI capture and processing capabilities can be expanded in a modular way by increasing the number of hardware elements in the system configuration. The MEHARI analysis functionality can be extended or modified both by configuration and by the addition of new analysis software modules. Three examples of the capabilities of the current implemented traffic analysis modules are: 1) a heuristic classification of the traffic in academic, commercial, fun and undetermined categories, 2) a classification of the traffic according its destination, and the verification of port assignment of applications. The system including these three has been field tested, and 3) the results of the trials performed, both in terms of functionality and performance, are also described in the article.

The rest of the paper is structured as follows. Next section presents the functional architecture of the MEHARI system. Section 4 describes the application modules currently supported by the system and the reasoning behind its developing. Section 5 provides some sample results of the field trial of the MEHARI system in a real network scenario: the Spanish Academic & Research Network (RedIRIS). Finally, section 6 summarises the main conclusions of our work.

2. MEHARI Functional Architecture

Figure 1 shows the functional architecture of the MEHARI system, which consists of the following subsystems:

- *Traffic Capture Subsystem (TCSS):* This functional block is responsible for capturing the traffic samples, which are subsequently analysed by other blocks of the MEHARI system. The TCSS can be configured to capture ATM cells either in a given list of VPI/VCI pairs or in promiscuous mode (all the VPI/VCI pairs). Cells are reassembled into AAL5 frames, which are periodically dumped to disk for further processing. Depending on the type of analysis to perform, the user can select either to dump the whole AAL5 frame or just part of it (e.g. the first 48 bytes).

  It should be noticed that, although the MEHARI system was initially conceived for monitoring IP/ATM traffic, the TCSS does not make any assumption about the content of the AAL5 frame. This way, the system could be used in the future to analyse other protocols encapsulated over ATM.

  The TCSS has been designed so that it can be physically implemented on one or several capture machines, thus allowing increasing the capture ratio as required. The capture platforms are standard PCs running UNIX. The capture itself is performed by two ATM Fore network interface cards (one for each transmission direction) using a special firmware (OC3MON [3]). Because of the use of standard hardware components, the total cost of the system is quite low (approximately 6,000 Euro per capture platform).

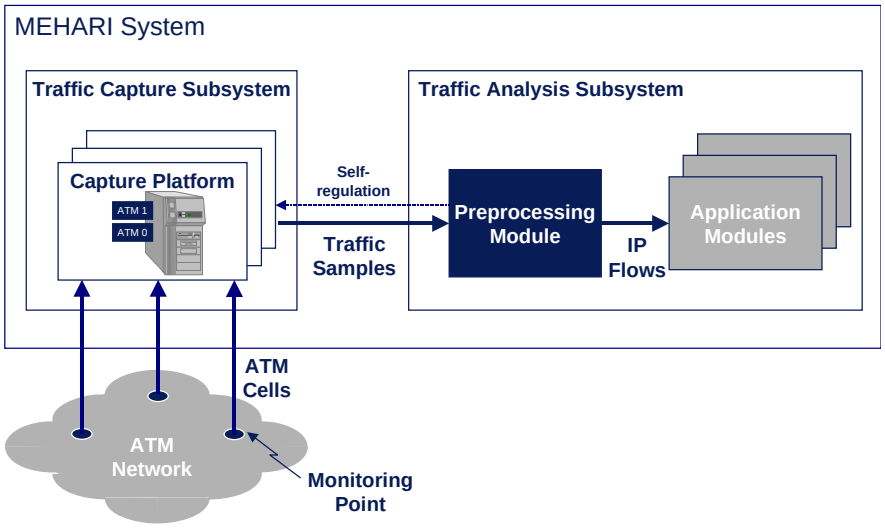

Figure 1. Functional Architecture of the MEHARI System

- *Traffic Analysis Subsystem (TASS):* This functional block is responsible for the analysis of the traffic samples generated by the TCSS subsystem. The TASS contains the Pre-processing Module (PPM) and the Application Modules (APMs).

  Because of the high traffic volume than can be transported on the ATM links monitored, the TASS must discard the traffic samples as soon as possible. This is precisely the reason of introducing the PPM, which pre-processes the traffic samples on the fly so that capture files are deleted once the parameters of interest have been extracted.

  Note that the rate at which the traffic samples are pre-processed is crucial for the overall performance of the system. In the MEHARI project, we concentrated in the analysis on IP flows (see section 3.1), which led us to develop a somehow CPU intensive pre-processing of the traffic samples. Other applications of the MEHARI system will probably require less processing capacity. In any case, the TASS subsystem has been designed so that it can be physically realised on one or several machines. This approach allows increasing the processing capacity as required by adding more PCs. In any case, note the existence of a self-regulation mechanism, which adapts the capture ratio to the TASS processing capacity.

  The traffic analysis is actually performed by the APMs. Different types of APMs, corresponding to different types of analysis can be present in the TASS. Section 3 describes some of the APMs developed for the MEHARI system. It should be pointed out that the MEHARI system has been designed so that new APMs can be easily added. Besides, there is the possibility of using several PCs to perform a distributed analysis, so that the processing capacity can be matched to the requirements of the different APMs considered.

3. Application Modules

This section describes the Application Modules (APMs) specifically developed for the MEHARI project. As mentioned in section 2, these APMs only an example of the different types of analysis that can be carried out by the MEHARI system. Our main purpose in describing the following APMs is to show the capabilities of the system and to provide a sample of the wide range of possibilities for which it could be considered.

3.1 Traffic Classification Module

The MEHARI project was conceived as the continuation of the traffic measurements and network usage characterisation studies carried out on the Spanish Academic & Research Network (RedIRIS) in the

CASTBA project. The results of CASTBA demonstrated a number of difficulties in characterising Internet traffic by using conventional traffic analysers. We found out that a simple analysis based on protocol and TCP/UDP port numbers did not provide the sufficient accuracy [5]. The reason is the existence of applications making use of not registered ports and registered ports used for purposes other that the intended. This resulted in an important fraction of traffic (10-20%) whose nature could not be determined or, what it is worst, an undetermined amount of traffic classified under the wrong categories. One of the main motivations of the MEHARI project was precisely to develop an Internet traffic monitoring and analysis tool to solve these uncertainties. What started as a timid attempt to overcome the limitations of TCP/UDP port traffic analysis, turned out in the *Traffic Classification Module*, which has been proved to be highly effective for characterising the Internet traffic and services.

Figure 2 shows the structure of the MEHARI Traffic Classification Module (TCM). This APM analyses the data coming from the PPM, which are based on the concept of IP flows and pattern recognition. The PPM performs the statistical aggregation, during a configurable period of time, of all the packets belonging to a same IP flow. An IP flow is defined as the aggregation of packets sharing the quadruple IP source address, TCP/UDP source port, IP destination address, and TCP/UDP port. Furthermore, the PPM explores the contents of the packets trying to determine the presence of specific patterns, which can be programmed by the user. As a result, the PPM periodically generates a file containing a summary of the IP flows detected and the number of times that a given pattern has been detected within it.

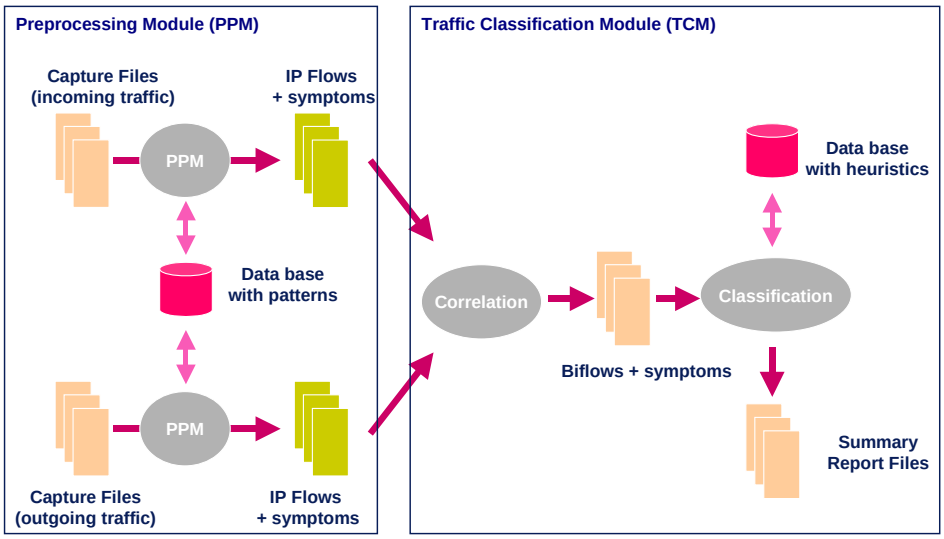

Figure 2. Traffic Classification Module (TCM)

The IP flows are further analysed by the TCM, which performs the correlation of the data corresponding to incoming and outgoing traffic. The result is a collection of bi-directional IP flows (biflows) with a pondered list of symptoms for each direction. The next step consists in applying a set of heuristics, which can be easily configured by the user. The application of the heuristics causes a biflow to be classified under a particular traffic category, which has been previously defined by the user. For example, in the field trial described in section 4 of this paper, we considered four traffic categories, namely academic, commercial, fun and undetermined. Finally, the TCM generates a summary file reporting the traffic volume under each traffic category considered. Section 4 shows an example of the traffic classification summary reports that can be obtained with the TCM.

3.2 Traffic Origin/Destination Analysis Module

The Traffic Origin/Destination Analysis Module (TODM) performs a traffic classification based on the origin/destination Autonomous System (AS) of the bi-directional IP flows identified by the TCM module (see section 3.1). The functional diagram of the TODM module is represented in Figure 3.

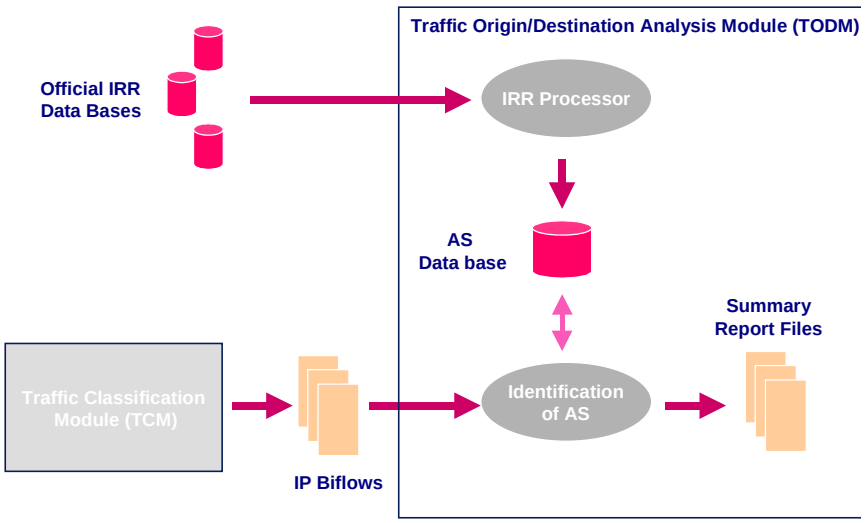

Figure 3. Traffic Origin/Destination Analysis Module (TODM)

To determinate the AS to which an IP source/destination address belongs, the TODM makes use of the information stored in the Internet Routing Register (IIR) databases. This information is periodically downloaded by the *IIR Processor Block*, which generates a local AS database that is used by the *AS Identification Block*. The AS database includes the name of organisation responsible for the different ASs as well as the identification of the subnetworks belonging to each AS.

The AS Identification Block analyses the IP addresses by looking them up in the local AS database. As a result, it periodically generates a statistics report file with the traffic volume exchanged between each subnetworks belonging to the network under study (e.g. RedIRIS) and the different ASs registered in the IIR databases. The TODM application module can be used, for example, to determine the fraction of traffic exchanged with a specific academic or commercial network, or to obtain a list of the most visited ASs. Other possible applications of the TODM module could be the following:

- *Accounting the external links usage*: The system provides a set of reports showing the use of resources in function of the sub-networks monitored academic network. This information is essential for future network re-design, billing schemes, etc.
- *Billing and charging the use of network resources:* The distribution of traffic in function of sub-network or AS allows supporting billing and charging models based on destination where tariffs rely on source and destination IP address or mask, autonomous system number or chain (route).

3.3 Internet Headers Analysis Module

The Internet Headers Analysis Module (IHM) provides a traffic classification based on the analysis of the three basic headers of the Internet protocol architecture: the IP packet header, the TCP/UDP segment headers, and the service PDU header. The structure of the IHM module is shown in figure 4.

The main objective of the IHM is to differentiate as much traffic flows (packet sequences with the same source and destination) as possible to then check the coherence between the TCP/UDP port and the application within each flow. It differentiates between the traffic that performs according the standards (the TCP or UDP port corresponds to standard assignment) and that that does not. It provides statistics on the local (inside of the academic network) and remote (outside of the academic network) most visited servers.

The classification is a result of the verification of the services contained in the routed traffic. This verification is carried out through the check of the coherence between the TCP/UDP ports and the service PDU header. In order to achieve this verification, we produced a dictionary of the main well-known services containing a set of patterns extracted from each service specification. The verification is server-oriented, which provides a better performance and less computing resources consumption. Some possible applications of the IHM module could be billing or auditing of network usage, verification of the services routed by the network, accounting of unknown/unusual traffic, etc.

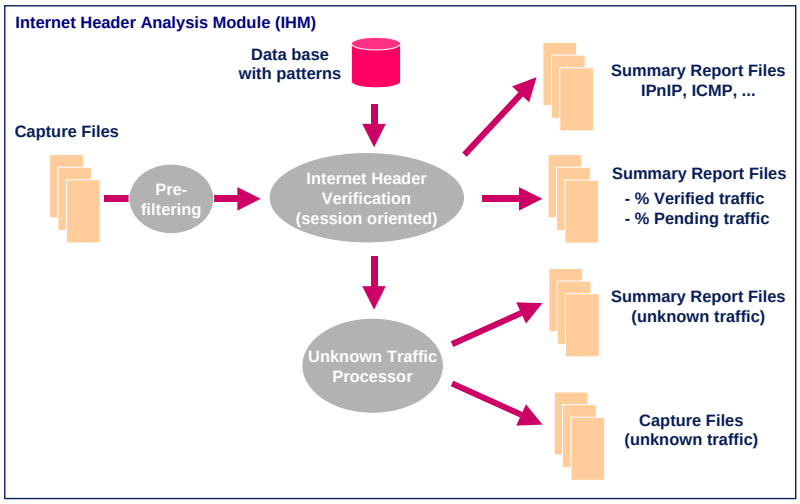

Figure 4. Internet Header Analysis Module (IHM)

Possible applications of the IHM module can be the following:

- *Verification ratio*: Traffic distribution in function of verified (checked ok), pending (patterns not found in the dictionary), unknown (ports not included in the dictionary) and rejected (transport protocol different from TCP/UDP).
- *Identification of the main local and remote servers*: The verified traffic is classified in function of the local or remote server addresses. This classification allows highlighting the main remote and local servers either dependent or independent of the service they provide.
- *Classification of the unknown traffic:* Some of the following heuristics can be useful in order to classify the main servers and services that remain unknown for the network managers: Reports on remote and local addresses sorted by bytes served. Report of possible servers detected sorted by number of possible clients connected to. Report of possible hid. All these reports highlight addresses, volumes, transport protocols and application ports that enable system auditors to inspect traffic and guess applications.

4. MEHARI System Tests

The MEHARI system has been successfully tested on a real network scenario: the IP/ATM backbone of the Spanish Academic and Research Network (RedIRIS). Two MEHARI system units were deployed at two different monitoring points of the RedIRIS backbone. One of the MEHARI units was installed in the RedIRIS central node in Madrid, according to the configuration shown in figure 5. The other MEHARI unit was placed at the RedIRIS access network of Catalonia, with a network scenario very similar to the one used in the central node.

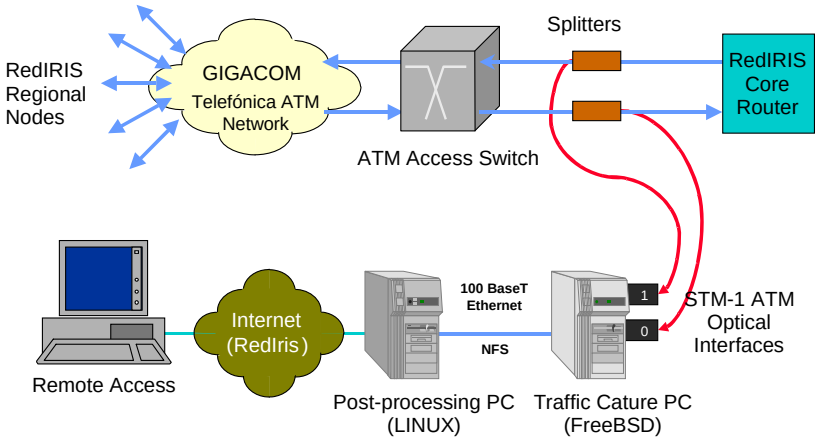

Figure 5. Trial network configuration

Figure 6 shows some example of the results obtained by the Traffic Classification Module (TCM) at the central node site in Madrid. The histogram represents, for each of the 17 regional links, the percentage of traffic corresponding to the four categories of traffic considered in the MEHARI project (academic, fun, commercial and undetermined). The results correspond to the average obtained during a capture period of approximately four and a half months (from September 15, 1998 to February the 2nd. of 1999).

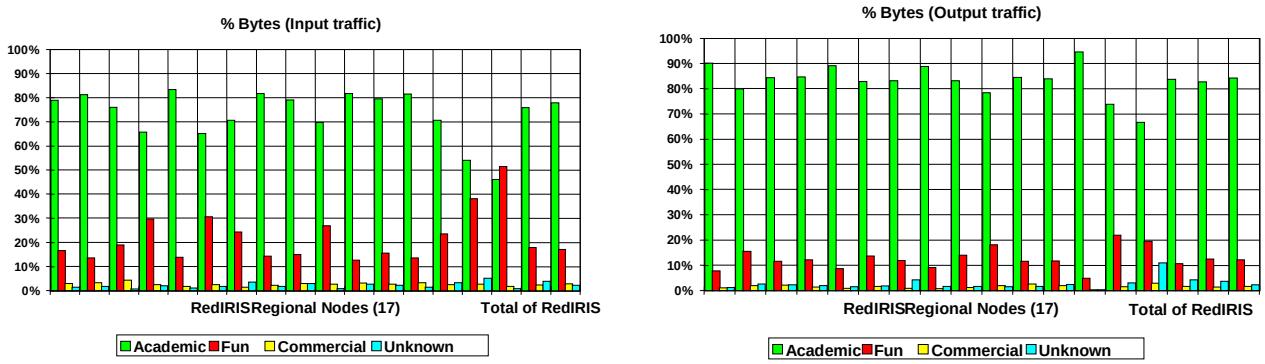

Figure 6. Traffic classification by usage (per RedIRIS regional nodes).

As an example of the results provided by the Traffic Origin/Destination Analysis Module (TODM), figures from 7 to 11. Figure 7 shows the main traffic destinations of the RedIRIS network during a capture period of approximately four and a half months (from September 15, 1998 to February the 2nd. of 1999). The graphics show the distribution of incoming/outcoming traffic exchanged between each of the 17 RedIRIS regional nodes and the following destinations: the rest of RedIRIS users, the Spanish commercial Internet, the European research networks (TEN-34/155), and the rest of Internet (through the RedIRIS to USA link).

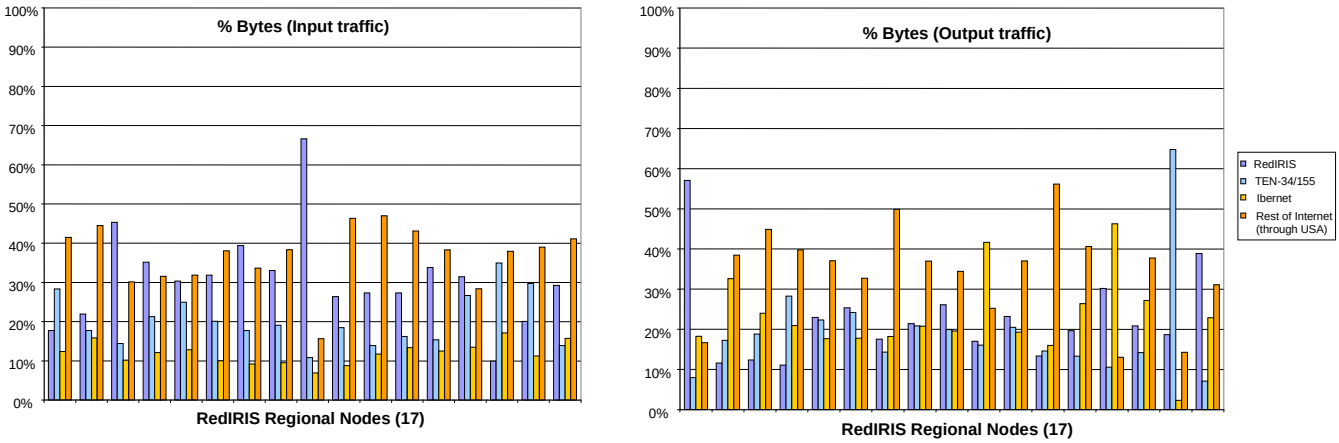

Figure 7. Main destinations of the RedIRIS traffic (per RedIRIS regional nodes)

Figure 8 shows the overall traffic distribution for the 17 RedIRIS regional links for the same period of capture than in Figure 7. Note than in both cases (Figures 7 and 8) the distribution is given in percentage of captured traffic, which hides the differences in amount of traffic handled by the different regional nodes of RedIRIS.

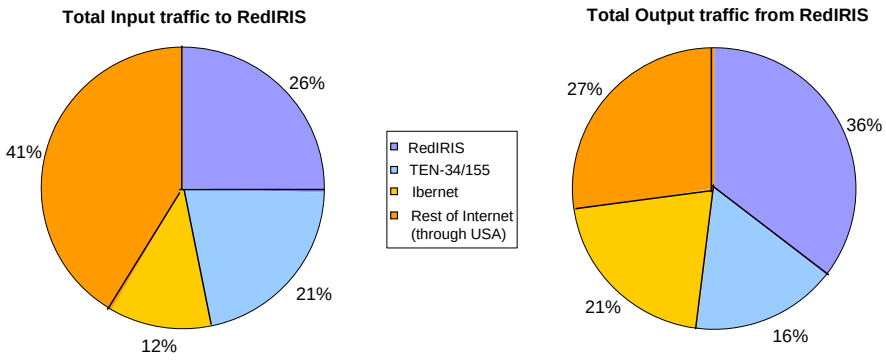

Figure 8. Main destinations of the RedIRIS traffic (for the whole RedIRIS network).

Figure 9 shows an attempt measure the percentage of academic traffic in the link between RedIRIS and USA. Note that Figure 9 is based on the pessimistic approach, only the traffic from USA (input traffic to RedIRIS) which was originated in an institution declared as academic or research centre in the Internet Routing Register is considered as academic. The rest is considered as commodity traffic. . This result corresponds to one month of capture (between September and October of 1998)

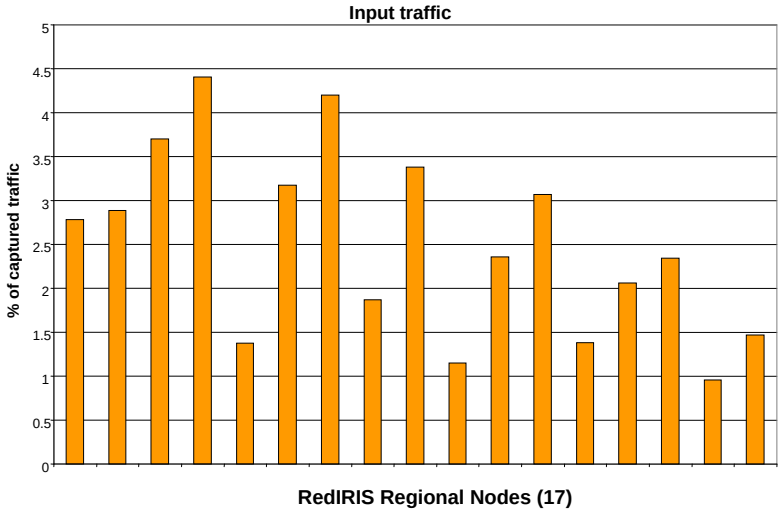

Figure 9. Percentage of academic traffic from USA to RedIRIS (according with the IRR description).

Regarding the traffic classification by commercial destinations the TODM module provides separate statistics for each of the Spanish academic network regional nodes. Figure 10 shows the main commercial destinations from the different subnetworks under the Catalonian regional node. The period of capture in this case goes from January to February of 1999. In this Figure only the 25 most visited sites are depicted, the rest (958 sites for the input traffic graphic and 990 for the output traffic graphic) are cumulated in the last column. Note that, in both graphics (input and output traffic) the 25 most visited sites collect around the 60 % of the traffic.

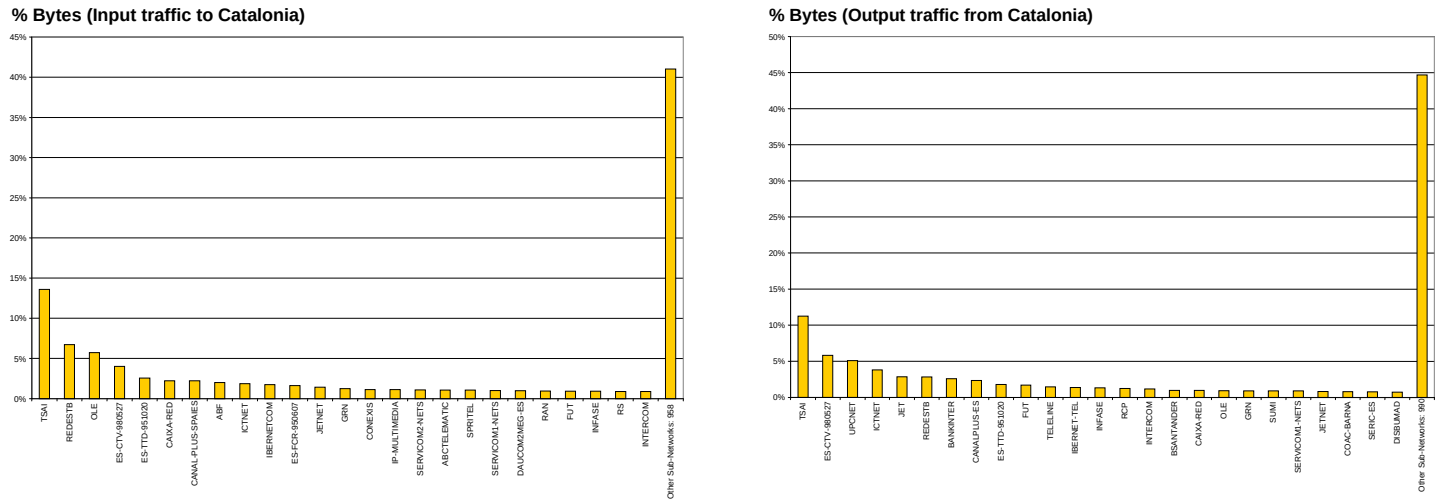

Figure 10. The 25 main commercial traffic destinations in Catalonia (January-February´99).

Another interesting result is the traffic classification by the most visited ASs of the European academic network (TEN-155). Like for the main commercial destinations, the TODM module provides separate statistics for each of the Spanish academic network regional nodes. Again, as an example Figure 11 includes the statistics on the main visited ASs of the Catalan node for the same capture period than above (January and February of 1999).

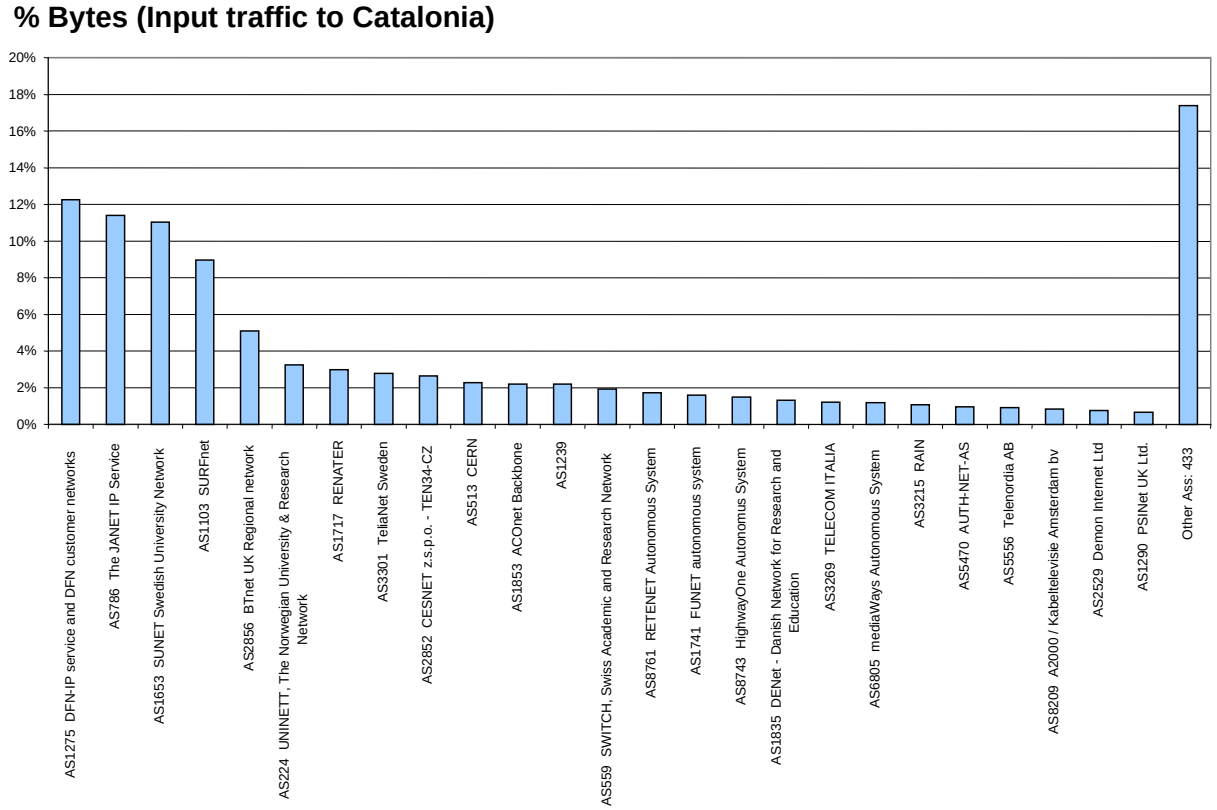

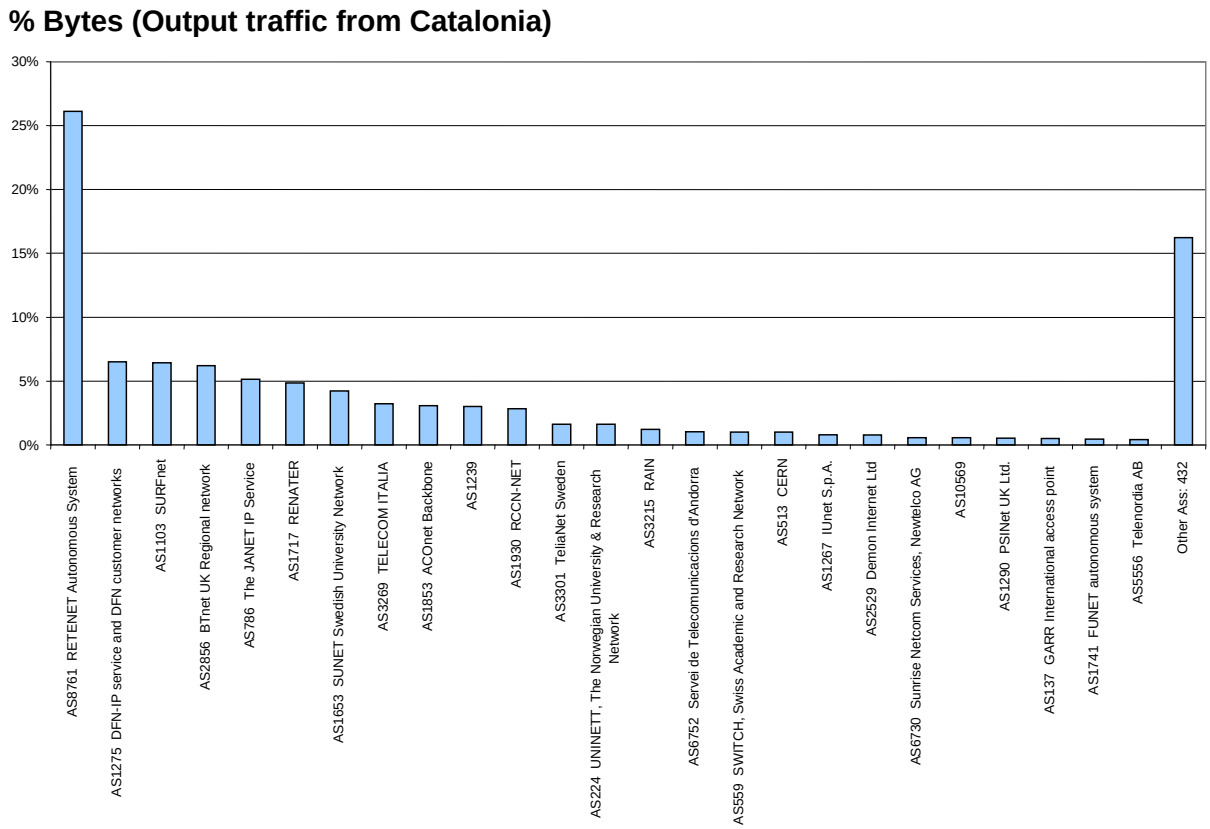

Figure 11. The 25 most visited TEN-155 ASs in Catalonia (January-February´99).

Figures 12, 13 and 14 show some examples of the results provided by the Internet Header Module (IHM). Figure 12 shows the percentage of verified traffic, pending traffic (no pattern found), unknown traffic (ports not registered) and rejected traffic (other protocols than TCP/UDP), for the traffic captured in the Catalan sub-network during January and February of 1999.

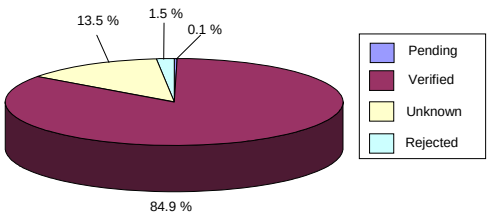

Figure 12. Traffic verification results (January-February´99).

Figures 13 presents one application of the tool: finding the most relevant servers. Host and application are determined for the 25 most visited remote and local servers respectively from the traffic captured during January and February of 1999. Note that remote servers collected 74 % of the total amount of captured traffic (in these graphics the input and output traffic are computed together). Local servers collected only the 26% (the remainder of the captured traffic). Also note that in both remote and local servers graphics the last column represents the traffic collected by the rest of servers, others than the 25 most visited. The number of these servers is 865,811 and 153,900 respectively.

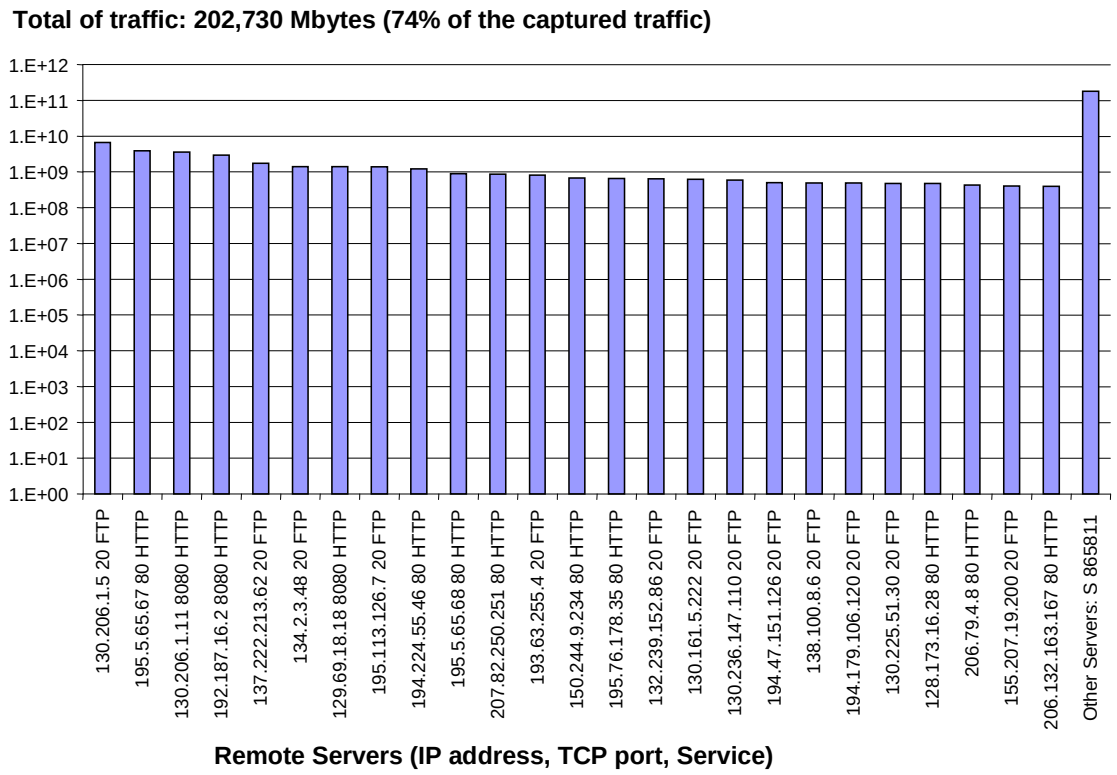

Figure 13.a. The 25 most visited servers from Catalonia (January-February´99).

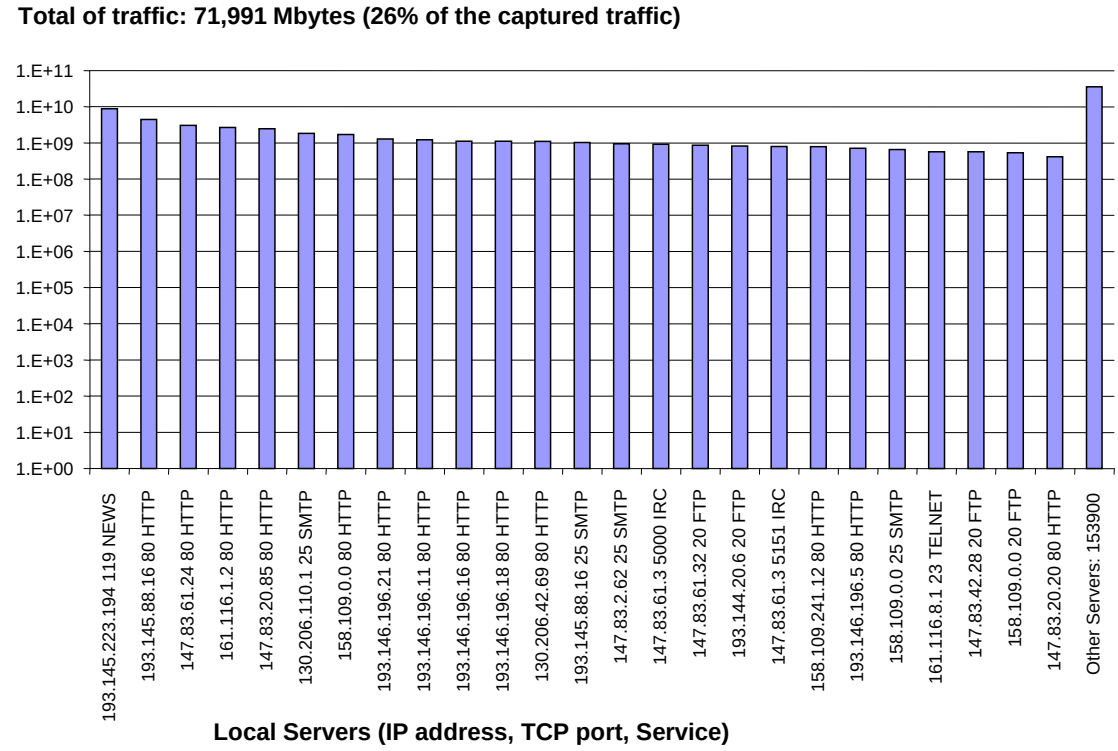

Figure 13.b. The 25 most visited servers of Catalonia (January-February´99).

Figure 14 shows an example of the possible reports on the unknown traffic that can be provided by the IHM. It is an attempt to find possible servers by the number of their possible clients (number of sources sending traffic to that destination).

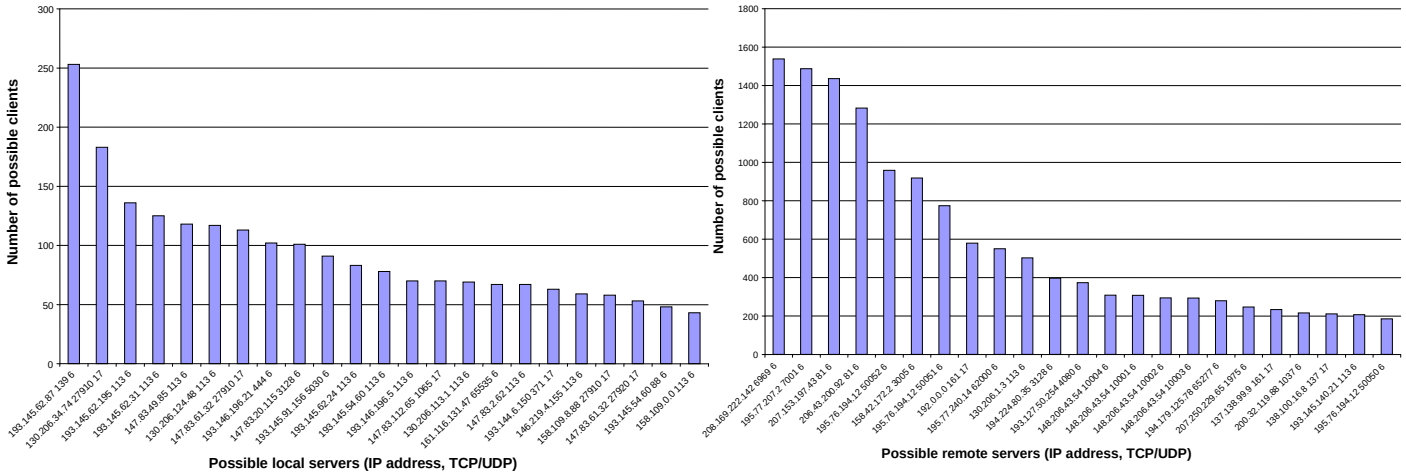

Figure 14. Unknown traffic statistics (one day traffic capture).

7. Conclusions

The flexibility of the IP protocol combined with broadband networks creates different difficulties for the dimensioning, management, and operation of this type of infrastructures. The permanent availability of detailed and updated information about network traffic is critical for network tuning, accounting, enforcement of acceptable user policy and security purposes.

The MEHARI platform allows data capture rates on STM-1 lines with price/performance ratios ten times better than commercial protocol analyzers, having also a much higher range on the number of VCIs snooped. As capture is done at the AAL5 level, it is independent of the higher level protocol, and may be customized to capture only part of the packets, or certain packet types. Furthermore, its modular design allows to distributing the traffic capture probes throughout the network, consolidating pre-processed flow information to a central site.

The MEHARI traffic processing modules currently provide information about conventional traffic statistics, the autonomous systems traffic matrix by user group, validation of the port-to-application assignment and heuristic detection of unacceptable traffic flows. It can be customised and/or extended to suite the changing traffic information requirements of network administrators.

The system field trial has been running 24 hours/day, 30 days a month during several months. The field trial results have been cross-checked, to verify the correctness of objective measurements, and the statistical validity of heuristic results on traffic types. These results provide high confidence in the resilience and accuracy of the MEHARI system.

Current work is expected to improve the capabilities of the capture and pre-processing platform, in order to obtain better price/performance capture rations. More traffic processing modules are also being developed, mainly targeted to security and charging applications.

Two practical applications of the MEHARI system are: monitoring the traffic to support billing and charging mechanisms for Academic and Corporate Networks and auditing the academic networks AUP (Acceptable Use Policy)

References

[1] Kevin Thompson et al., *"Wide-Area Internet Traffic Patterns and Characteristics"*, IEEE Network, November/December 1997.

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

[3] J. Aspirdof et al., "OC3MON: Flexible, Affordable, High-Performance Statistics Collection", INET'97, Malasya, june 1997.

[4] David Newman et al., "Intrusion Detection *Systems, Suspicious Finds"*, Data Communications, August 1998.

[5] M. Alvarez-Campana, A. Azcorra, J. Berrocal, D. Larrabeiti, J.I. Moreno, J.R. Pérez, "CASTBA: Internet Traffic Measurements over the Spanish R&D ATM Network", HP-OVUA Workshop, Rennes (France), April 1998.

Acknowledgements:

The MEHARI project has been funded by the CICYT (Comisión Interministerial de Ciencia y Tecnología) under contracts TEL97-1897-E and TEL97-TEL97-1893-E. The authors thank the people who, together with the authors, have been participating in the MEHARI project their contribution to this work. These people are: Xavier Martínez, Carles Veciana, Albert Renom and Sergi Sales of the UPC, David Larrabeiti of the UC3M and Julio Berrocal and Ana B. García of the UPM. The authors also thank the people from RedIRIS, Telefónica I+D and C^4 (Centre de Computació i Comunicacions de Catalunya), which have been following the work done within this project, for their valuable comments and suggestions. Finally, many thanks to the OC3-MON development team for the useful they provided us in order to adapt the modified FORE firmware to the MEHARI capture modules.

Author Biographies:

Pedro J. Lizcano

Arturo Azcorra

Josep Solé-Pareta received his Master's degree in Telecommunication Engineering in 1984, and his Ph.D. in Computer Science in 1991, both from the Universitat Politècnica de Catalunya (UPC). In 1984 he joined the Computer Architecture Department of the UPC. Since 1992 he is an Associate Professor with this department. Josep Solé-Pareta spent the summers of 1993 and 1994 at the Georgia Institute of Technology. He has participated in the R&D Spanish Program for the development of the Broadband Communications in Spain (PLANBA). Josep Solé-Pareta is member of the Advanced Broadband Communications laboratory of the UPC (http://www.ac.upc.es/CCABA). His current research interests are in Broadband Internet, ATM Networks and Optical Packet Networks, with emphasis on traffic engineering, traffic characterization, traffic management and QoS provisioning. He is member of the IEEE and the ACM (Sigcomm).

Jordi Domingo-Pascual

Manuel Alvarez-Campana