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

Pedro J. Lizcano, Comisión Interministerial de Ciencia y Tecnología (CICYT),
Rosario Pino, 14-16, 28020 Madrid, Spain

Arturo Azcorra, Universidad Carlos III de Madrid (UC3M),
Butarque 15, 28911 Leganés (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 for monitoring and analysing the traffic on National R&D Networks (NRN). The main design requirement of the MEHARI system was to cope with the main challenges currently faced by most of these networks: Service usage, charging schemes, network dimensioning and auditing procedures according to a given Acceptable Use Policy (AUP). 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 can report information on the usage of the Internet services, the main traffic origin and destinations, etc. The MEHARI system has been tested by running a field trial on the Spanish NRN (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, Internet traffic analysis, National R&D Networks, Acceptable Use Policy auditing.

1. Introduction

The full development of the *Information Society* paradigm has led to the development of new IP infrastructures worldwide, driven by commercial interest and by different government initiatives because of strategic reasons. Even the EU has play a leading roll in this matter with the *TEN-34* and *TEN-155* initiatives to foster the trans-national collaboration among the different National NRN researches.

But these deployments are facing with two major problems:

- In most cases, NRNs are funded by public money. Nevertheless, budgets can not follow the traffic increasing rates that users demand. Therefore, a new funding model which leads to an use-responsible approach is needed.

- Market forces, in current full competitive scenarios, demand from these public funded networks a transparent and acceptable use policy (AUP) to assure no distortions in the emerging Internet-based services business.

To cope with these two open issues, tools specifically tailored for auditing AUP are required for monitoring the actual use of the public funded networks. The possible collaborations of European or NRNs networking initiatives with their leading counterparts in North America (i.e., Internet 2 and Canarie) depend on having these tools available to ensure the coherence of the Americans and Europeans AUPs concerning the traffic to be exchanged among them.

Furthermore, the above mentioned full development of the Information Society also complicates the dimensioning, management, and operation of Internet backbones. These tasks cannot be successfully performed without precise and up-to-date knowledge about what is going on in the network. Thus, there

are two different dimensions to this problem: one is from the *traffic* perspective, and the other is from the *purpose usage* angle. The former may need contents analysis in cases in which application identification by port is not possible, and its results are reasonably objective. The latter definitely needs contents analysis, and its results are based on heuristics and subjective techniques.

On the traffic analysis aspect, it is worth mentioning some recent studies on the characterisation 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 the results of traffic analysis are therefore only valid in the short-term.

On the purpose usage aspect, different heuristic techniques may be combined according to 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 the processing power required. It is recommended that heuristics be manually checked from time to time in order to make the analysis more tight or relaxed, depending on the deviations between real traffic and the diagnosis performed by the tool.

This paper describes the design of the MEHARI system, a high performance traffic processor for analysing 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 performed 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 modules 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 adding on new analysis software modules. Three examples of the type of information that can be obtained from the current implemented traffic analysis modules are the following: 1) a heuristic classification of the traffic in academic, commercial, leisure and undetermined categories; 2) a classification of the traffic according its origen and destination; and 3) the verification of port assignment of applications. The system including these three has been field tested and the results of the trials are also described in the article, both in terms of functionality and performance

Data provided by MEHARI may serve different purposes, some of which are:

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

The rest of the paper is structured as follows. The 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 development. Section 5 provides some sample results of the field trial of the MEHARI system in a real network scenario: the Spanish NRN (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 be performed, the user can select either dumping the whole AAL5 frame or just part of it (e.g. the first 48 bytes).

It should be noted 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 the capture ratio to be increased as required. The capture platforms are standard PCs running Free BSD. 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 euros per capture platform).

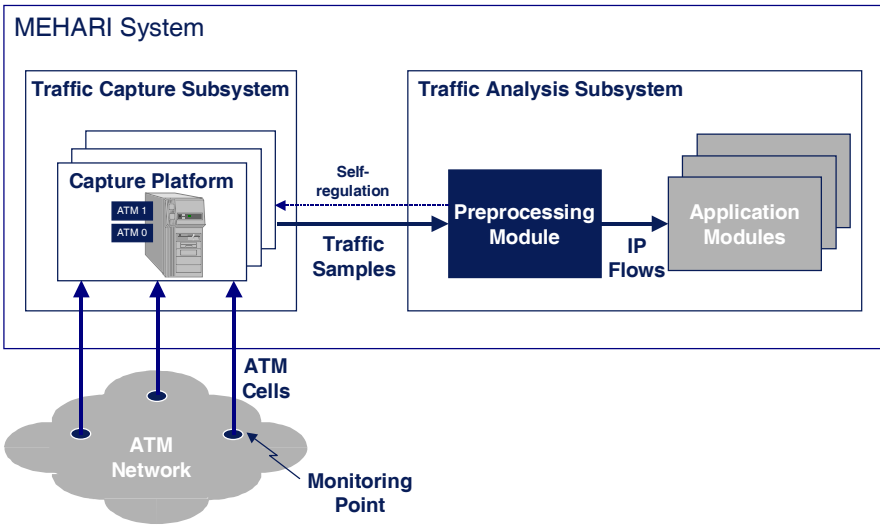

Figure 1. Functional Architecture of the MEHARI System.

- *Traffic Analysis Subsystem (TASS):* This functional block is responsible for analysing 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 that can be transported on the ATM links that are monitored, the TASS must discard the traffic samples as soon as possible. This is precisely the reason for 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 on analysing IP flows (see Section 3.1), which led us to develop a somewhat 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 was designed so that it can be physically implemented on one or several machines. This approach allows the processing capacity to be increased as required by adding more PCs or Workstations. 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 easily be 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 are 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 NRN (RedIRIS) in the CASTBA project. The results of CASTBA revealed 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 sufficient accuracy [5]. The reason is the existence of applications making use of unregistered ports and registered ports used for purposes other than those asigned. This resulted in a large percentage of traffic (10-20%) whose nature could not be determined or, what it is worse, 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 modest attempt to overcome the limitations of TCP/UDP port traffic analysis turned out to form the *Traffic Classification Module*, which has been proved to be highly effective for characterising 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 is based on the concept of IP flows and pattern recognition. The PPM performs the statistical aggregation of all the packets belonging to a same IP flow during a configurable period of time. 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 in it.

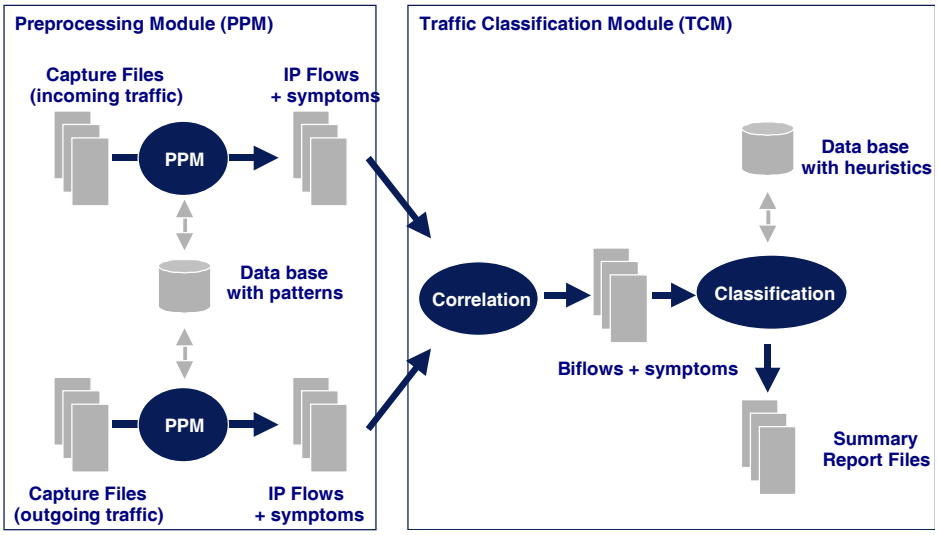

Figure 2. Traffic Classification Module (TCM).

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 weighted list of symptoms for each direction. The next step consists in applying a set of heuristics, which can easily be 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, leisure 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 Subsection 3.1). The functional diagram of the TODM module is represented in Figure 3.

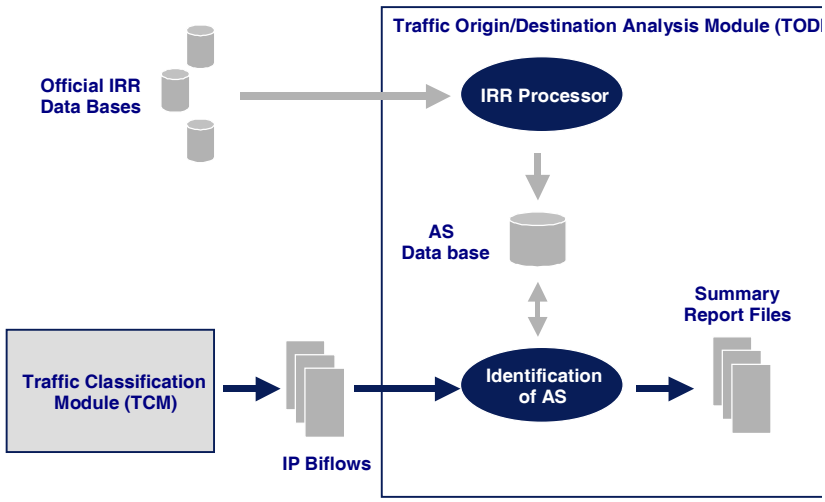

Figure 3. Traffic Origin/Destination Analysis Module (TODM).

To determine 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 the 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 subnetwork 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:

- *Assessment of the external links usage*: The system provides a set of reports showing the use of resources per sub-network of the monitored academic network. This information is essential for future network re-design, billing schemes, etc.

- *Billing and charging for the use of network resources:* The distribution of traffic according to 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 many traffic flows (packet sequences with the same source and destination) as possible, and then to check the coherence between the TCP/UDP port and the application within each flow. It differentiates between the traffic that performs according to the standards (the TCP or UDP port corresponds to standard assignment) and the traffic that does not. It provides

statistics on the most visited local (inside the academic network) and remote (outside the academic network) servers.

The classification is a result of the verification of the services contained in the routed traffic. This verification is carried out by checking 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 better performance and less consumption of computing resources. Some possible applications of the IHM module could be billing or auditing network usage, verification of the services routed by the network, recording unknown/unusual traffic, etc.

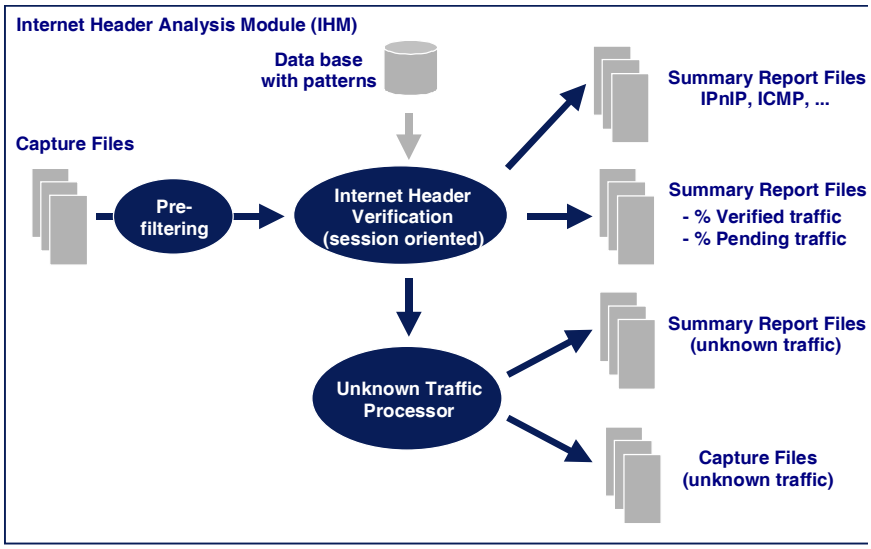

Figure 4. Internet Header Analysis Module (IHM).

Some possible applications of the IHM module could be the following:

- *Verification ratio*: Traffic distribution classified as: 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 according to the local or remote server addresses. This classification allows the main remote and local servers to be highlighted, either according to the service they provide or not.

- Classification of unknown traffic: Some of the following heuristics can be useful for classifying 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. 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 NRN (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 on Catalonia's RedIRIS access network, with a network scenario very similar to the one used in the central node.

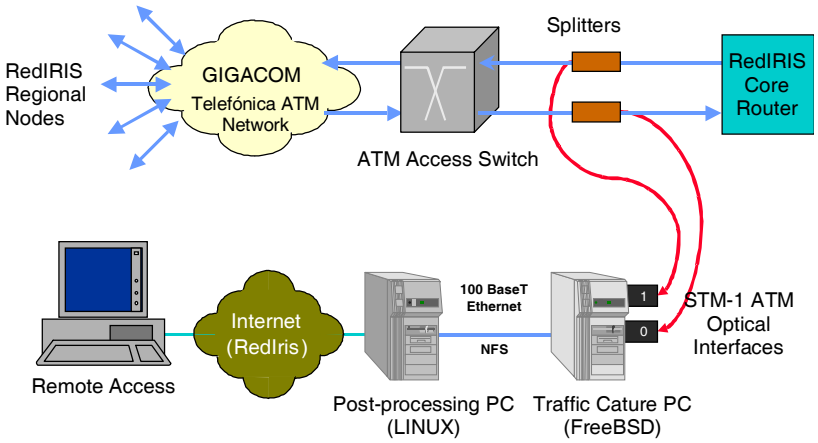

Figure 5. Trial network configuration.

It must be stressed that the results presented hereinafter aims at providing some examples of the MEHARI reporting capabilities. The figures included in these reports do not necessary correspond to the actual value of the involved traffic parameters in the Spanish NRN. However, these values may correspond to a monitoring process in a given and limited time period, and therefore are not full representative.

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

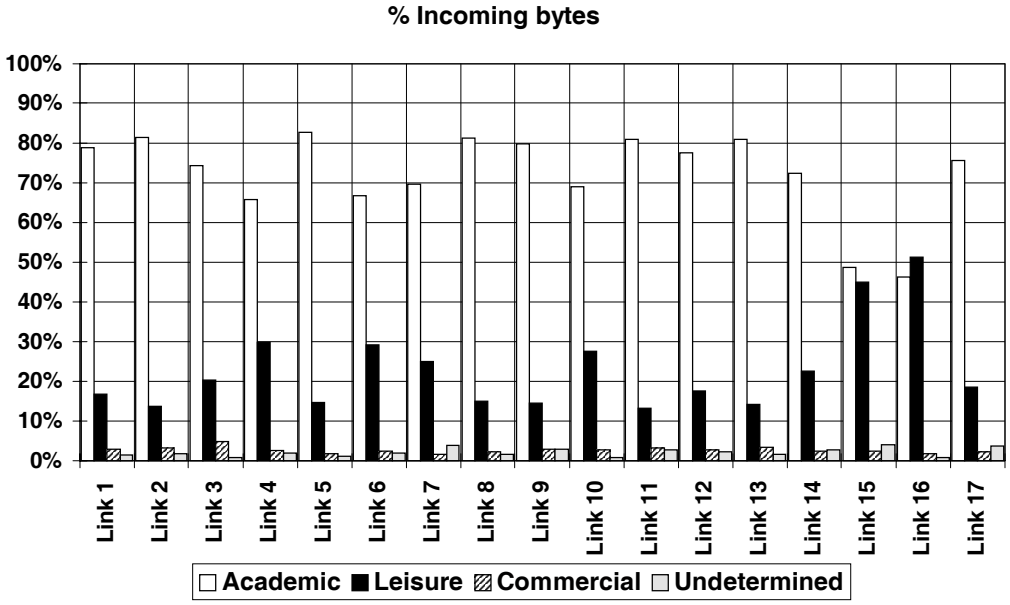

Figure 6. Traffic classification by usage (per RedIRIS regional nodes).

Figure 7 shows the overall input traffic distribution for the 17 regional nodes of RedIRIS as whole for the same period of capture than in Figure 6.

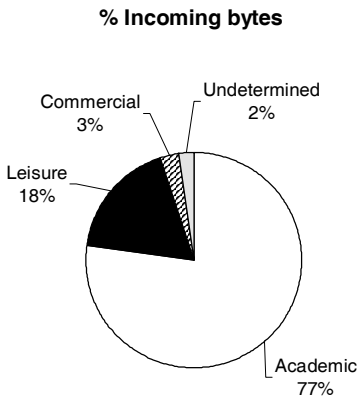

Figure 7. Traffic classification by usage (global distribution).

For an example of the results provided by the Traffic Origin/Destination Analysis Module (TODM), see Figures 8 to 12. Figure 8 shows the main traffic origins of the RedIRIS traffic obtained from an input traffic capture made during a period of approximately four and a half months (from 15 September 1998 to 2 February 1999). The graphics show the percentage of input traffic to each of the 17 RedIRIS regional nodes and the following origins: other RedIRIS sites, the Spanish commercial Internet, the European research networks (TEN-34/155), and the rest of Internet (through the RedIRIS to USA link).

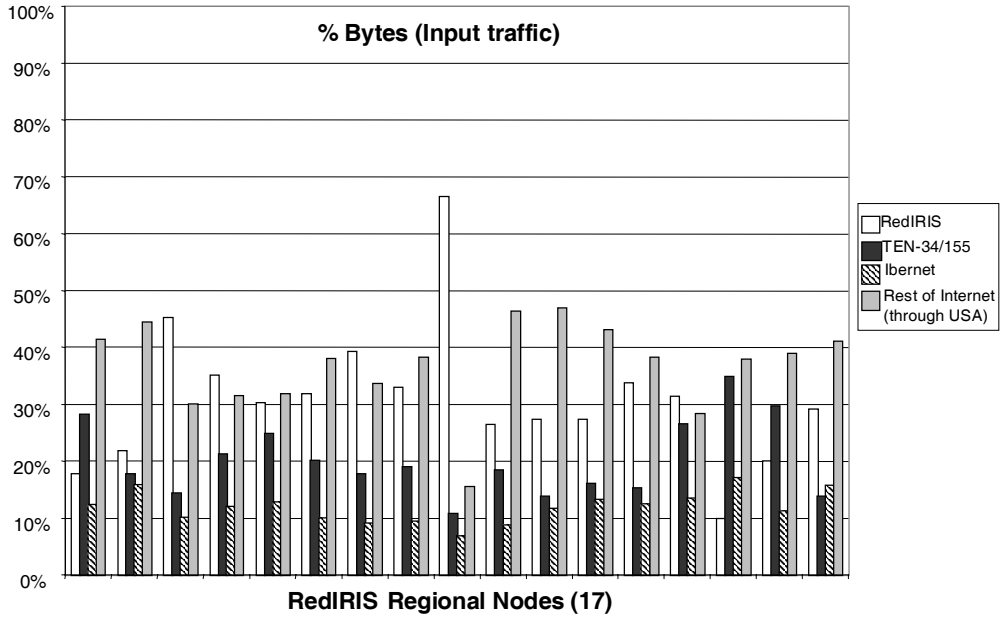

Figure 8. Main origins of the RedIRIS traffic (per RedIRIS regional nodes).

Note than, in both cases Figures 6 and 8, the fact that the distribution is given in percentage of captured traffic hides the differences in amount of traffic handled by the different regional nodes of RedIRIS.

Figure 9 shows the overall input traffic distribution for the 17 RedIRIS regional links for the same capture period as in Figure 8.

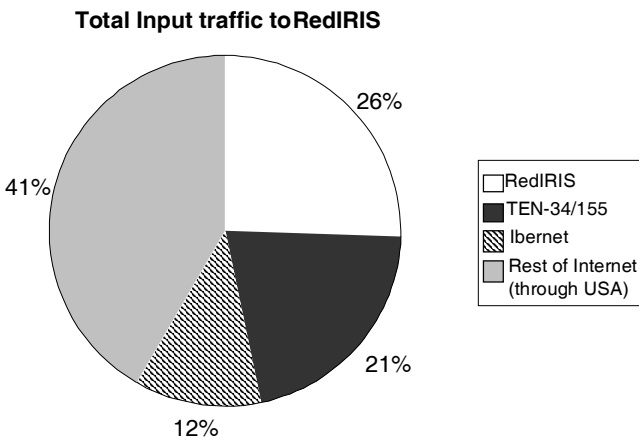

Figure 9. Main origins of the RedIRIS input traffic (for the whole RedIRIS network).

Figure 10 shows an attempt to measure the percentage of academic traffic in the link between RedIRIS and USA. Note that Figure 9 is based on the pessimistic approach, and only the traffic from USA (input traffic to RedIRIS) originating at an institution declared as being an academic or research centre in the Internet Routing Register is considered academic. The rest is considered commodity traffic. This result corresponds to one month's capture (between September and October 1998).

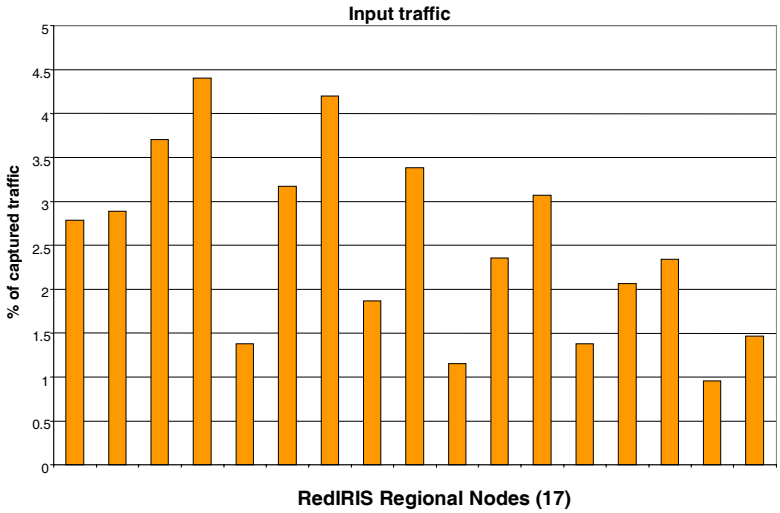

Figure 10. Percentage of academic traffic found in the USA to RedIRIS link (according with the IRR description).

As regards traffic classification by commercial destinations, the TODM module provides separate statistics for each of the Spanish NRN regional nodes. Figure 11 shows the main commercial origins of the input traffic to the Catalan node. The capture period in this case goes from January to February 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 compiled in the last column. Note that the 25 most visited sites collect around 60 % of the captured traffic.

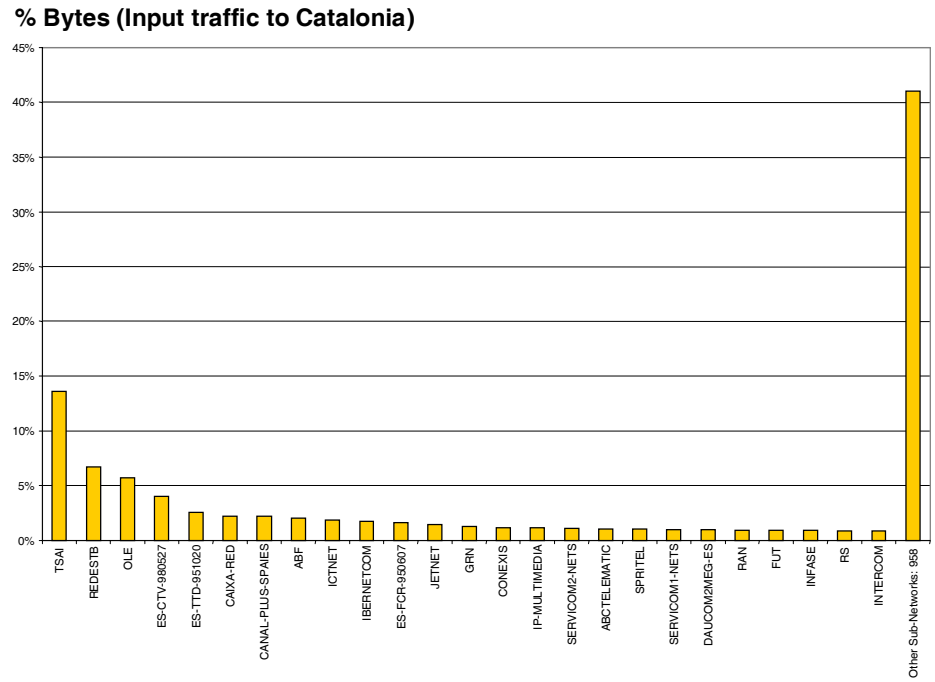

Figure 11. The 25 most visited commercial sites in Catalonia (from input traffic captures).

Another interesting result is the traffic classification by the most visited ASs of the European academic network (TEN-155). As for the main commercial origins, the TODM module provides separate statistics for each of the Spanish NRN regional nodes. Again, as an example, Figure 12 includes the statistics on the main visited ASs of the Catalan node for the same capture period as above (January and February 1999).

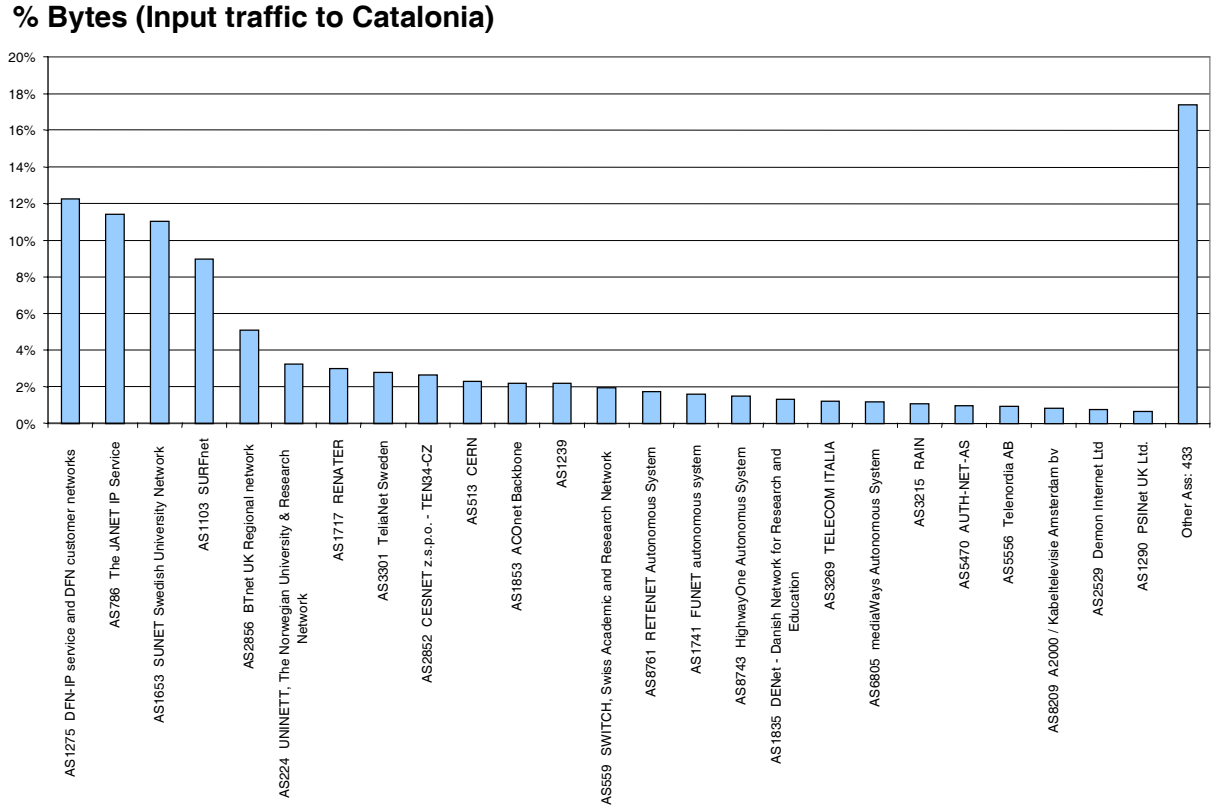

Figure 12. The 25 most visited TEN-155 ASs in Catalonia (from input traffic captures).

Figures from 13 to 16 show some examples of the results provided by the Internet Header Module (IHM). Figure 13 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 in January and February 1999.

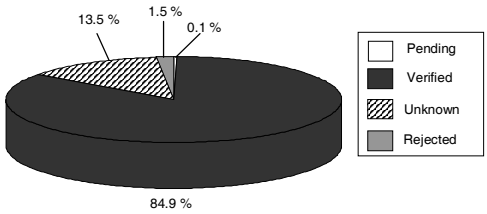

Figure 13. Traffic verification results.

Figures 14 and 15 present 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 1999. In these graphics the input and output traffic are computed together. Note that remote servers collected 74 % of the total amount of captured traffic (see Figure 14). Local servers collected only 26%, the remainder of the captured traffic (see Figure 15). Also note that in both remote and local server graphics the last column represents the traffic collected by the rest of the servers, other than the 25 most visited. The number of these servers is 865,811 and 153,900 respectively.

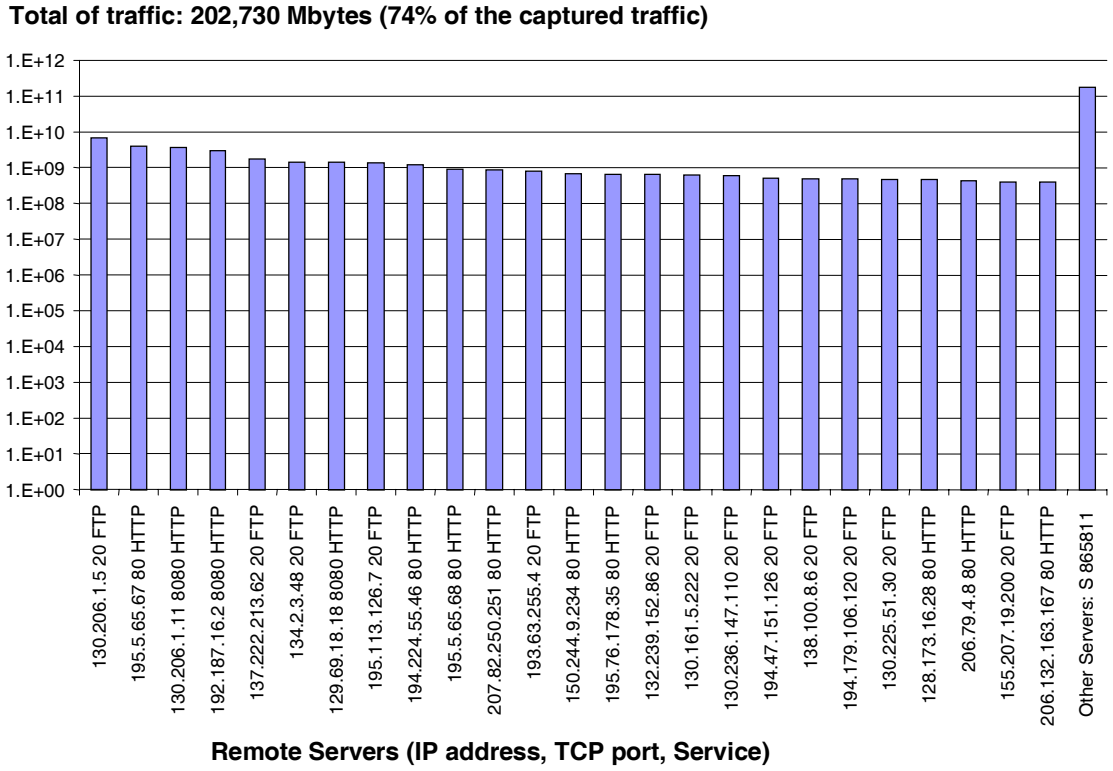

Figure 14. The 25 most visited servers from Catalonia.

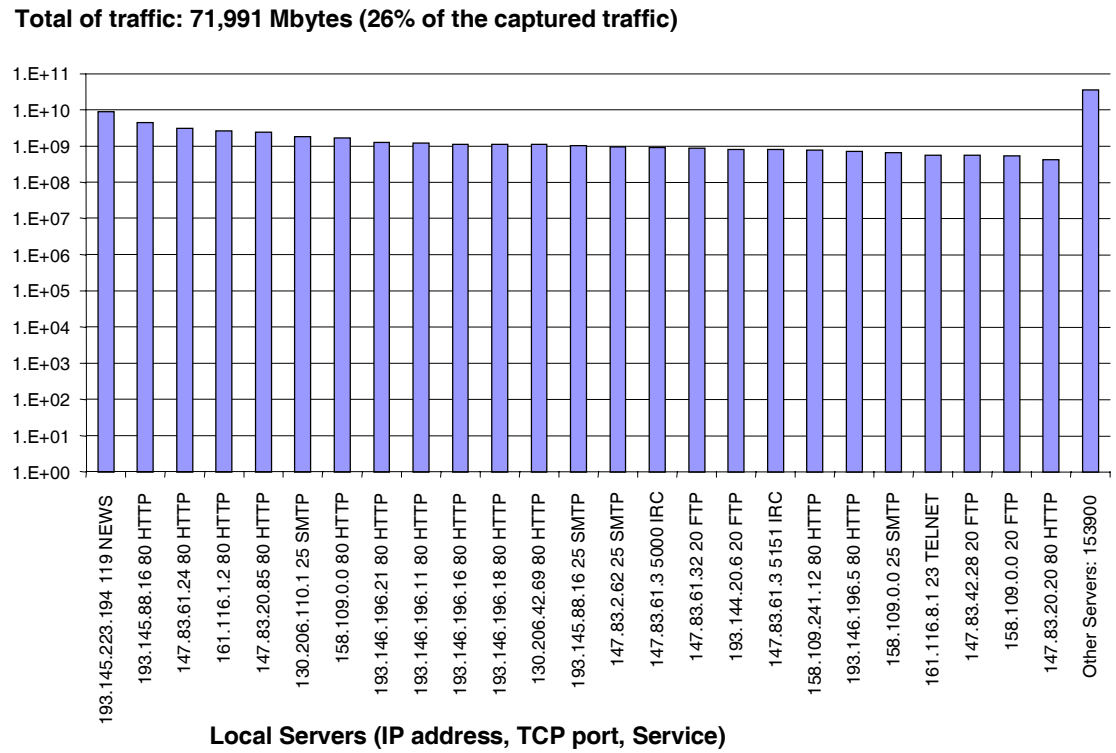

Figure 15. The 25 most visited servers of Catalonia.

Figure 16 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). The graphics of the Figure 16 corresponds to a traffic capture of a single day.

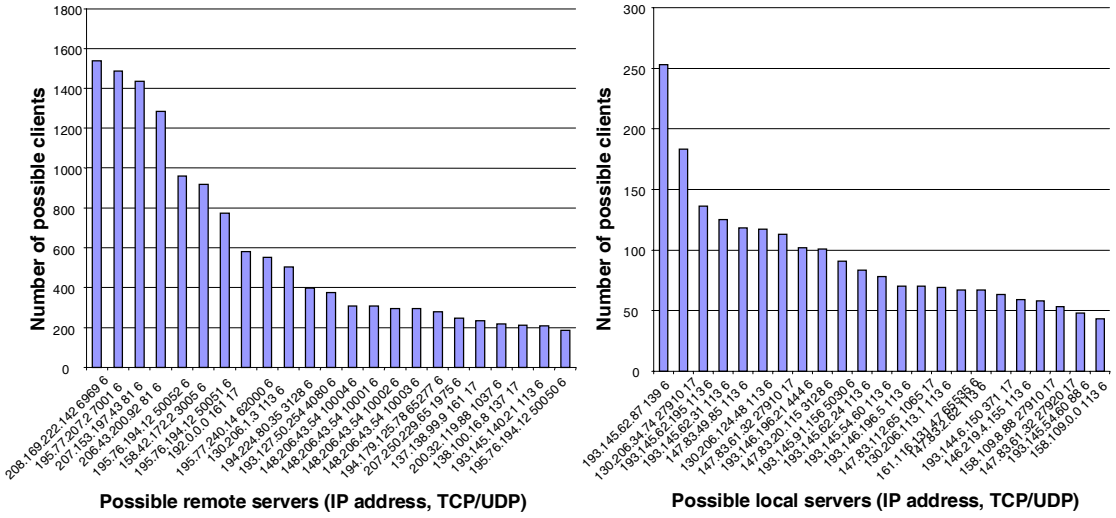

Figure 16. Statistics of the unknown traffic.

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 use 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 analysers, also having a much higher range on the number of VCIs snooped. As capture is at the AAL5 level, it is independent of the higher level protocol, and may be customised to capture only part of the packets, or certain packet types. Furthermore, its modular design allows the traffic capture probes to be distributed throughout the network, consolidating pre-processed flow information to a central site.

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. They can be customised and/or extended to suit the changing traffic information requirements of network administrators.

The system field trial has been running 24 hours a day, 30 days a month for 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 give great 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 ratios. More traffic processing modules are also being developed, mainly targeted at 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 National R&D Networks AUP (Acceptable Use Policy)

Acknowledgements:

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

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 ATM Network", HP-OVUA Workshop, Rennes (France), April 1998.

Author Biographies:

Pedro J. Lizcano (lizcano@cicyt.es) holds a Telecommunication Engineering Master degree from the Universitat Politècnica de Catalunya (UPC, Barcelona, in 1983). From 1983 to 1988 he was with the R&D Department at Telefónica de España, involved in advanced communications systems development such as a multi-access radio system for rural applications (*MARD* fully industrialised by Telettra and deployed world-wide), and a second generation X.25 switching packet system (*TESYS-B*, a major joint development between Telefónica industrial group companies and Fujitsu, deployed nation-wide, in which he led the hardware platform development). In 1988 he joined Telefónica R&D as the Packet Switching Division manager. From 1991 to 1993 he worked with AT&T Bell Laboratories in Denver (Colorado, USA) as Resident Visitor in a forward-looking work project dealing with SONET / SDH high speed

broadband interfaces. In 1993, back at Telefónica R&D, he led the ATM Systems Division, fully involved in the development of a B-ISDN platform (RECIBA) for benchmarking the emerging multimedia services. He also worked on a number of EURESCOM projects (concerning network synchronisation) and in RACE II and ACTS EU-programme projects (dealing with advanced networking and services). Since 1994 he has been working as external expert and adviser with the EU Commission in programme definition, proposal selection and technical audit processes. From 1996 to now he has held the positions of head of the areas dealing with International Development, Basic Technologies and Product Engineering. He is currently leading the Technological Innovation Programme area.

Since 1995 he has been a technology adviser at the Interministerial Commission of Science and Technology, leading a National R&D Programme on Telematic Services and Applications. He is the national representative in the European Networking Policy Group (ENPG, currently as Vice-President). He is also member of the Board of Trustees of the International Computer Science Institute (ICSI) in Berkeley (California, USA). He has been a member of the IEEE for over 20 years and a member of the Planetary Society.

Arturo Azcorra (azcorra@it.uc3m.es) was awarded an M.Sc. degree in Telecommunications from the Technical University of Madrid in 1986 and a Ph.D. from the same university in 1989. In 1990 he became an associate professor at the Department of Telematics Engineering, Technical University of Madrid, Spain. Since 1998 he has been lecturing at U. Carlos III of Madrid, Spain. Current research projects include broadband networks, multicast teleservices, intelligent agents, and IP/ATM convergence.

Josep Solé-Pareta (pareta@ac.upc.es) was awarded 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 UPC. Since 1992 he has been an Associate Professor with this department. He spent the summers of 1993 and 1994 at the Georgia Institute of Technology. He has participated in the Spanish R&D Programme for the development of the Broadband Communications in Spain (PLANBA). He is a member of the Advanced Broadband Communications laboratory of 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 characterisation, traffic management and QoS provisioning. He is a member of the IEEE and the ACM (Sigcomm).

Jordi Domingo-Pascual (jordid@ac.upc.es) is an Associate Professor of Computer Science and Communications at the Universitat Politècnica de Catalunya (UPC) in Barcelona. There, he received the engineering degree in Telecommunication (1982) and the Ph.D. Degree in Computer Science (1987). In 1983 he joined the Computer Architecture Department. Since 1994 he is co-founder and researcher at the Advanced Broadband Communications Center of the University (CABA) which participates in the Spanish National Host and in the PLANBA demonstrator. He was visiting reseracher at the International Computer Science Institute in Berkeley (California) for six months. His research topics are Broadband Communications and Applications. Since 1988 he has participated in RACE projects (Technology for ATD and EXPLOIT), in several Spanish Broadband projects (PLANBA: AFTER, TR1 and IRMEM), ACTS projects (INFOWIN, MICC, IMMP), and spanish research projects (CASTBA, MEHARI, SABA). A more detailed information may be found in: http://www.ccaba.upc.es.

Manuel Alvarez-Campana (mac@dit.upm.es) received his M.S. and Ph.D. degrees in Telecommunication Engineering from the Universidad Politécnica de Madrid. In 1989 he joined the Departamento de Ingeniería Telemática at the Universidad Politécnica de Madrid (DIT-UPM), where he is an associate professor since 1996. He has participated in several projects within the framework of European and Spanish R&D programs. His current research interests include design and analysis of broadband networks, integration of services, traffic characterisation, and network performance evaluation.