

MEHARI: An IP/ATM Traffic Analysis Platform based on Configurable Patterns[†]

M. Álvarez-Campana¹, A. Azcorra², J. Berrocal¹,
A. B. García¹, J. R. Pérez¹

1 Departamento de Ingeniería de Sistemas Telemáticos
Universidad Politécnica de Madrid.
ETSI Telecomunicación, Ciudad Universitaria, 28040 Madrid
www.dit.upm.es

2 Area de Ingeniería Telemática
Universidad Carlos III de Madrid,
28911 Leganés (Madrid)
www.it.uc3m.es

Abstract

The MEHARI system provides a low-cost programmable and scalable tool for the analysis of IP/ATM traffic. The system runs on one or several PCs, thus allowing increasing the processing capacity as required. The tool can be configured to capture either IP headers or whole IP packets, being possible to concentrate the analysis on specific traffic flows by using VPI/VCI or IP address filters. The captured traffic is analyzed on the fly so that the samples are discarded as soon as the different application modules process them.

The application modules perform the analysis of the IP packets/headers. The system has been designed so that the output of one module may feed the input of the following. This approach allows the addition of new applications at different processing levels.

An especially interesting application developed for the MEHARI system is the Traffic Classification Module. This application performs a classification of IP flows (packets sharing the same IP origin/destination addresses and TCP/UDP ports) into traffic categories (e.g. academic, commercial, leisure, ...) The particular categories and classification criteria can be tailored in a quite flexible way by simply adding the appropriate patterns and heuristic rules in the configuration data base.

The MEHARI system has been successfully tested on a real scenario: the RedIris ATM backbone (the Spanish Academic & Research Internet branch). The traffic classification application has proved to be highly effective, even though the programmed patterns and heuristics are rather simple.

Some of the potential applications for which the MEHARI system could be used in the future include: accounting and billing, characterization of user profiles, identification of most visited hosts and servers, detection of security threats and attacks. In particular, we think that one of the most promising areas on which the MEHARI system could be used is the supervision and enforcement of Acceptable Usage Policies (AuPs) on Internet.

[†] This work has been partially supported by the CICYT project MEHARI-M within the Spanish R&D Program.

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. Furthermore, the lack of usage policy restrictions poses almost no limit to the amount of traffic that a single user can draw from the network. As a result, Internet backbones are constantly subject to strong traffic load fluctuations, highly unpredictable.

The combination of all the above problems extraordinarily complicates the dimensioning, management, and operation of the Internet backbones. Such tasks cannot be successfully performed without a precise and up-to-date knowledge about what is going on in the network. From a *traffic-engineering* point of view, this knowledge can be obtained by using conventional traffic analysis methodologies and equipment. In this sense, 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. 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.

In the case of academic and research networks, there is an additional dimension of the problem where traditional traffic analysis fails: the *network usage* perspective. The fact that these type of networks are usually financed with public funds, combined with the lack of usage policy restrictions, is leading to claims from some sectors about inappropriate usage of public resources. In this sense, there is an increasing interest in the development of Internet traffic analysis tools for monitoring and auditing the network usage. These types of tools require the application of new methodological approaches for traffic analysis, as the one followed by the MEHARI system [3]. This paper shows how the combination of content analysis, heuristics, and subjective techniques, can be successfully used to obtain network usage measurements.

The rest of the paper is structured as follows. Next section presents the functional architecture of the MEHARI system. Section 3 describes the field trial of the MEHARI system in a real network scenario: the Spanish Academic & Research Network (RedIris). Section 4 provides a summary of the traffic classification results obtained by the MEHARI prototype on the RedIris backbone. Finally, section 5 summarizes the main conclusions of our work.

2 MEHARI Functional Architecture

Figure 1 shows the functional architecture of the MEHARI system, which consists of the Traffic Capture Subsystem and the Traffic Analysis Subsystem.

The Traffic Capture Subsystem (TCSS) is responsible for capturing the traffic samples, which are subsequently analyzed 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).

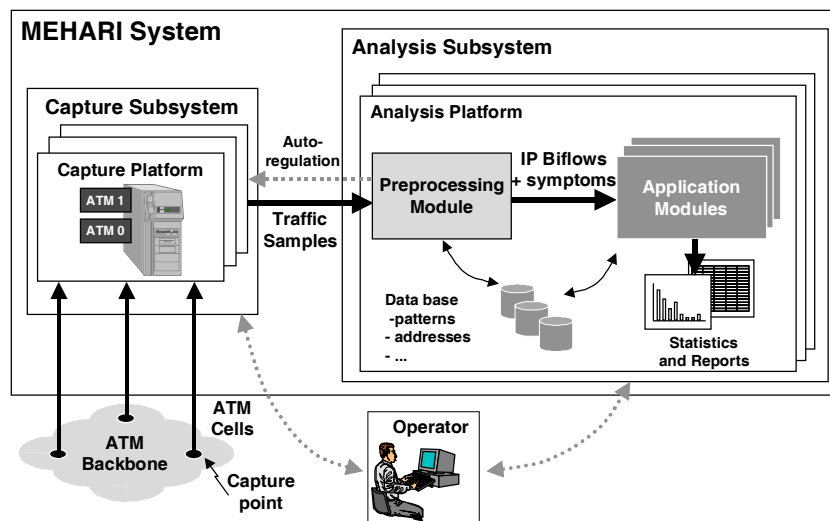


Figure 1. Functional Architecture of the MEHARI System

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 analyze 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) with a special firmware (OC3MON [4]). Because of the use of standard hardware components, the total cost of the system is quite low (approximately 6,000 Euro per capture platform).

The Traffic Analysis Subsystem (TASS) is responsible for the analysis of the traffic samples generated by the TCSS subsystem. The TASS contains the Preprocessing Module (PPM) and the Application Modules (APMs). Because of the high traffic volume that 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 preprocesses 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 preprocessed is crucial for the overall performance of the system. In the MEHARI project, we concentrated in the analysis on IP flows, which led us to develop a somehow CPU intensive preprocessing of the traffic samples. Other applications would probably require less processing capacity. In any case, the TASS subsystem has been designed so that it can be physically realized 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 that adapts the capture ratio to the TASS processing capacity.

The traffic analysis in the TASS is actually performed by the Application Modules (APMs). Different types of APMs, corresponding to different types of analysis can be present in the TASS. 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.

Section 4 describes the traffic classification APM used in the MEHARI field trial on the Spanish R&D Internet IP/ATM backbone (RedIris). Next section provides an overview of the RedIris network architecture and test scenario used to validate the MEHARI system.

3 MEHARI Field Trial Configuration on the Spanish R&D Network

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 3 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 [5] frames according to the RFC 1483 [6].



Figure 3. Topology of the RedIris IP/ATM backbone

Taking into account the star topology of the RedIris backbone, and in order to take simultaneously measurements from the 17 regional links, we decided to locate the monitor point at the RedIris central node in Madrid. Figure 4 shows the configuration of the measurement equipment. A capture platform consisting of a PC with two ATM NICs (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 PC. In order to achieve a high traffic capture ratio, it was decided to use a separate machine (another standard PC) for data analysis. This allows the capture PC to concentrate on the acquisition of traffic samples.

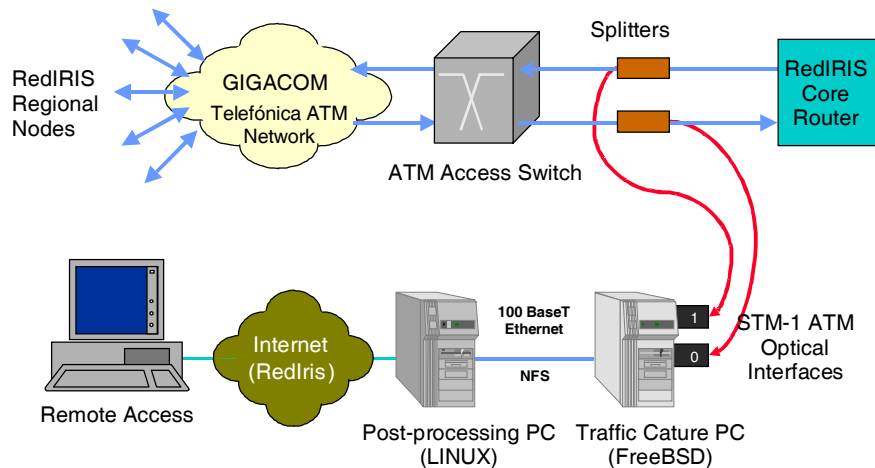


Figure 4: MEHARI Field Trial Scenario on the RedIris Backbone

The data collected by the capture PC 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 capture machine. 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.

4 Traffic Usage Classification Sample Results

As an example of the application areas for which the MEHARI system can be used, this section describes the Application Module (APM) used for traffic usage classification in the MEHARI field trial on the RedIris. Figure 2 shows the structure of the MEHARI Traffic Classification Module (TCM).

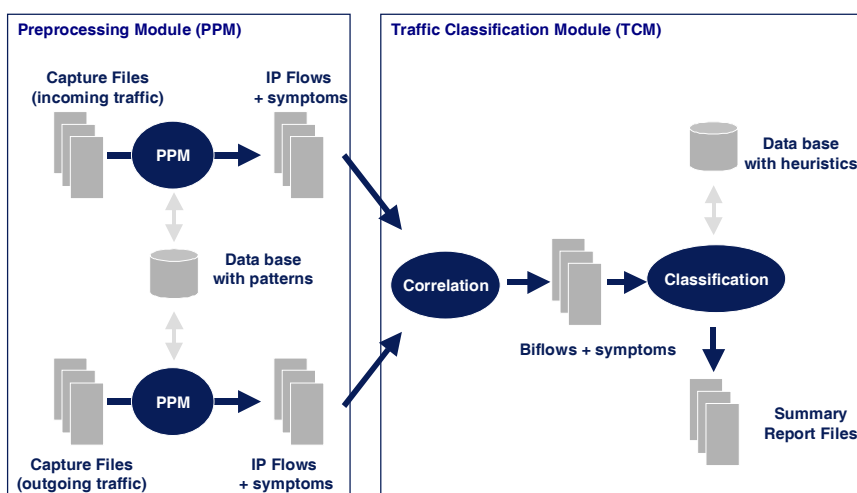


Figure 2. Traffic Classification Module (TCM)

The TCM analyzes the data coming from the PPM (see section 2), 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 tuple 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.

The IP flows are further analyzed by the TCM, which performs the correlation of the data corresponding to incoming and outgoing traffic. The result is a collection of bidirectional IP flows (biflows) with a weighted 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. Finally, the TCM generates a summary file reporting the traffic volume under each traffic category considered.

Figures 5 and 6 show some sample results obtained by applying the TCM on the RedIris backbone. The data correspond to the average values obtained during a capture period of approximately four and a half months (15/9/98 to 3/2/99). The graph in figure 5 represents the percentages of incoming bytes on the RedIris backbone for each of the four traffic categories considered in the MEHARI project (academic, leisure, commercial and undetermined). Figure 6 shows traffic classification results for each of the 17 regional links.

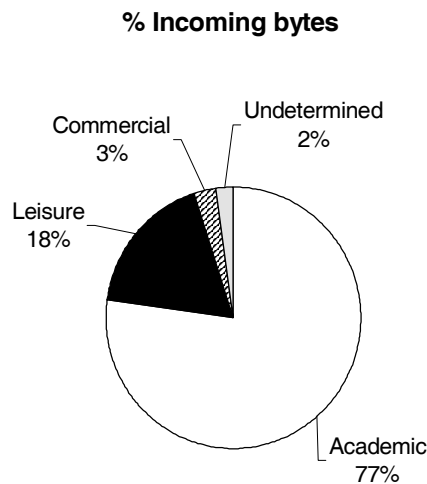


Figure 5: Traffic classification by usage (global results)

It must be pointed out that although the heuristics used in the MEHARI project are rather simple, they have been proved highly effective. We have checked that the individual verdicts generated by the TCM for each bidirectional IP are correctly emitted in a quite high proportion (approximately a 95% in our experiments). Obviously, the effectiveness of the TCM can be improved by using more elaborated rules and heuristics. These can be easily tailored according to particular traffic classification criteria desired. However, 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.

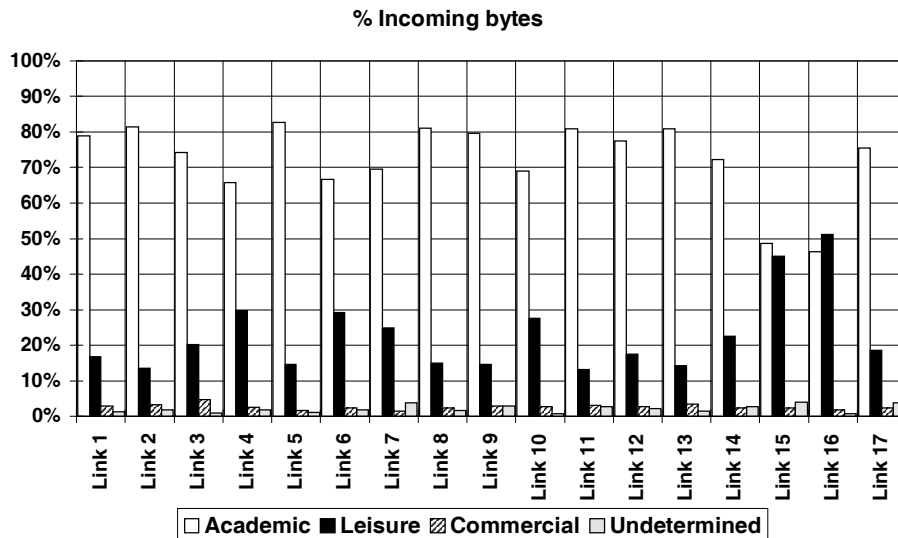


Figure 6: Traffic classification by usage for each RedIris regional node

5 Concluding remarks

In this paper we have described the main characteristics of the MEHARI system, a low-cost programmable and scalable Internet traffic analysis tool. The system has been designed so that it can be easily expanded both in terms of processing capacity and traffic analysis capabilities. The former can be done by increasing the number of hardware platforms (based on conventional low-cost PC components). The MEHARI analysis functionality can be easily extended or modified both by configuration and by the addition of new analysis software modules.

As an example of the potential capabilities of the MEHARI system, we have developed a traffic classification application based on configurable patterns and heuristics. The application has been successfully tested on a real network scenario (the Spanish Academic IP/ATM Internet backbone). The application can be easily tailored for other traffic classification criteria, by simply configuring the appropriate patterns and heuristics. As it has been proved in the MEHARI project, the effectiveness of the traffic classification tool may be surprisingly high even using rather simple masks and rules.

Among the different applications for which the MEHARI system could be used in the future, we can mention the following: network dimensioning, accounting and billing, characterization of user profiles, identification of most visited hosts and servers, detection of security threats and attacks. Finally, we think that one of the most promising areas on which the MEHARI system can be used is the supervision and enforcement of Acceptable Usage Policies (AuPs) on Internet.

References

- [1] K. Thompson, G. Miller, R. Wilder, "Wide-Area Internet Traffic Patterns and Characteristics", IEEE Network Magazine, November 1997.
- [2] 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.
- [3] 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.
- [4] J. Aspirdof et al., "OC3MON: Flexible, Affordable, High-Performance Statistics Collection", INET'97, Malasya, June 1997.
- [5] ITU-T, "Recommendation I.363.5 - B-ISDN ATM Adaptation Layer specification: Type 5 AAL", August 1996.
- [6] J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", RFC 1483, IETF, July 1993.