Discovering the Secrets of ATM Networks

Mauro Cociglio*, Robert Cole
Telecom Management Department
HP Laboratories Bristol
HPL-97-64
April, 1997

ATM,
telecoms
management

In the last three years, ATM networks have gone from an experimental, pilot phase to a commercial phase. All the world's major telecoms operators have completed or are completing networks which use this technology at various levels. Many different choices have been made: there are those who use ATM essentially at the trunk level and those who intend to take it all the way to the desktop. However, these networks currently lack any system able to provide full, continuous and reliable monitoring of their Quality of Service (QoS). All the more so since pure ATM networks are not at all common; most networks are mixed, comprising multiple technologies and services.

Currently, various kinds of monitoring policies are enacted, and oftentimes they are not fully integrated. Information from a management center, which today is sure to be incomplete, is used; artificial traffic is generated over paths that are parallel to the user's, and their performance is assessed. Essentially, the physical layer is monitored (it certainly is the one where the most settled techniques and the most economical tools are available), assuming that if it is operating properly, then the upper layers should be operating as well. The attempt is made to integrate these methods, all necessarily incomplete, with access-side monitoring, using RMON (Remote MONitoring) for LAN users, X.25 and Frame Relay instrumentation, etc. The problem, however, is then to correlate all these different sources of information to have a full and up to date view of the QoS provided by the network.

This was the starting point for the collaboration between CSELT and Hewlett-Packard Labs Bristol, aimed at specifying and developing a monitoring system for ATM networks. The goal would be to merge QoS information for the many services which in commercial networks, as we have seen, make use of ATM technology. From the CSELT side, we have drawn on the experience derived both from participating in standard-setting bodies and from our daily work supporting Telecom Italia's network in all the stages of development. For its part, Hewlett-Packard has made available its knowledge in the fields of measurement and management of telecoms systems.

This paper presents the first fruits of our collaboration. After discussing the general problem of monitoring ATM networks, we shall present the system realized by Hewlett-Packard and the applications specified and developed in collaboration with CSELT. We will then look at the issues of testing and tuning the measurement system applied first in CSELT laboratories and then to the commercial ATM network connecting two of CSELT's facilities in Turin. Lastly, we shall examine the issues involved with the commercial use of the system and its prospects for evolution in the short and medium term.

CSELT S.p.A., Via Reiss Romoli, 274, 10148 - Torino, Italy

DISCOVERING THE SECRETS OF ATM NETWORKS

Mauro Cociglio
Dr.
CSELT S.p.A.
Via Reiss Romoli, 274
10148 - Torino, Italy
Tel: +39 11 228 5028
Fax: +39 11 228 5069
E-Mail: mauro.cociglio@cselt.stet.it

Robert Cole
Mr.
Hewlett-Packard Laboratories, Bristol
Filton Road, Stoke Gifford
Bristol BS12 6QZ. U.K.
Tel: +44 117 922 8012
Fax: +44 117 922 8003
E-Mail: rhc@hpl.hp.com

1. The problem of monitoring ATM networks

ATM networks currently installed throughout the world have started to provide the first responses to the promises this technology had made at the beginning of the Nineties: huge bandwidth, flexible services, aggressive pricing policies, better bandwidth exploitation thanks to ever more sophisticated policing algorithms. However, the problem of verifying the performance and quality of service (QoS) offered by these networks remains largely unexplored. Operation & Maintenance (OAM) streams, in particular performance monitoring, standardized by ITU [2], has gone essentially unheeded and not implemented by the manufacturers who have won over the largest market share (Newbridge, Cisco/Stratacom, GDC Apex, Northern Telecom and Cascade), while the management systems for such equipment have so far been limited to monitoring performance within individual nodes, without providing an overall vision of an end-to-end connection. The efficiency of the ATM protocol, deriving from its leanness, with only 5 overhead bytes (the header, out of a cell length of 53 bytes), has deprived it of the necessary information other protocols carry with them to allow for verification of their performance or QoS. The choice of having a single Circular Redundancy Code (CRC) on the header has made it impossible for nodes to verify the correctness of the payload of the cells received and if need be to correct them before they are sent. At the same time, the lack of a sequence number does not allow the detection of lost cells from a single observation point accessible to the network management center, such as a switching node's input interface.

Of course, the very premise of the ATM is that it is to be based on excellent transmission lines, in particular SDH at 155 Mbit/s and up, which guarantee the physical layer such a high level of quality as to minimize errors. Unfortunately, when theory is converted into practice, things do not always work as advertised, and at least for the time being most of the physical lines connecting ATM nodes are PDH (34 or 45 Mbit/s), and node-user connections go as low as 2 or 1.5 Mbit/s. In some cases adaptation layers (AALs) can be of help. AAL1 has a sequence number, albeit limited to modulo 8, whereas CRC32 of AAL5 detects erroneous packets due to lost, errored or misinserted cells, and bit errors. Unfortunately, nothing at all can help if measurements are to be made on parameters like transit delay or delay variance.

To summarize, therefore, we can observe that there are currently only two strategies that provide some assessment of the performance, or of the QoS that ATM networks are able to offer their customers.

The first one is to use the data provided by the management centers, which can count cells in transit over each individual node, subdivided by channels and virtual paths, provide the number of cells with errored header (whether or not it can be corrected through the CRC), count the cells that have undergone a congestion (EFCN bit set to 1), or that have violated policing parameters (CLP bit set to 1), if these two capabilities have been implemented.

The second strategy currently available is to utilize parallel channels to the user's traffic with artificial traffic generated by testing instruments. This would allow the preventive measures prescribed by I.356 [1], but with some drawbacks. In this case, three kinds of problems are to be confronted: the network overhead due to the additional traffic, the inaccuracy of the measurements due to the fact that the traffic, for economic reasons, can only employ a limited bandwidth and, lastly, the lack of any guarantee that user traffic actually behaves the same way.

2. The CSELT-Hewlett-Packard collaboration

The considerations expressed in the previous paragraph have led to consider the development of a system to meet the requirements of an operator who, wishing to take advantage of the opportunities offered by ATM technology, would desire to build and operate a large ATM network for business users. CSELT , as Telecom Italia's research and support center, was tasked with responding to the pressing requirements in this field. We started collaborating with a major industrial partner, Hewlett-Packard, a company with the necessary experience in the field of measurement instrumentation for broad band networks and able promptly to meet the early demands of the ATM market. During 1996 a collaboration agreement was signed by CSELT and Hewlett-Packard Labs Bristol (Hewlett-Packard's European research center) with these main goals:

1) To specify and integrate solutions for distributed testing/monitoring of ATM network to support IP services.
2) To complete applications for actual networks.
3) To test the feasibility of these solutions.

Building on the non-intrusive monitoring system developed by Hewlett-Packard, CSELT has specified a series of measurment applications and ways to present the results graphically or numerically which, on the basis of its own experience of the Italian network, could be particularly helpful to the telecommunications operator.

The experimental system was delivered by Hewlett-Packard to CSELT in September 1996, equipped with the first measurement application ("AAL5 QoS"), while the second application ("AAL0 QoS") was delivered at the end of the year.

Let us now look more in detail at the developments Hewlett-Packard has brought about in this sector.

3. Non Intrusive QoS measurement

As a manufacturer of communications test equipment, most of Hewlett-Packard's effort is directed to testing for R&D and to the installation and measurement (I&M) phases of network technology delivery. These phases are characterized by complex tests involving generating and recording traffic, but under well-constrained laboratory situations. In these phases experienced engineers set about the task of building or adjusting the network until it works to the specification. The choice of test and the results are driven by the network specification. The network is what is measured. The information is used by the engineers to justify the statement that "the network is working".

A project in the Laboratories of Hewlett-Packard in Bristol has taken this measurement expertise and applied it to operational networks. Measuring in an operational network is very different from either of the R&D and I&M phases in terms of what is measured and what is done with the measurement information. In an operational network the information from the network is used to support the business of delivering network services. If a network operator is going to pay for a measurement system then it must deliver a commensurate value to the operator in supporting his business goals. This is the starting point for an operational network measurement system — the measurement information supports the business processes of the operator. Given the competitive pressures on operators selling network services, we focused on the measurement of **quality** as a starting point to meet this need.

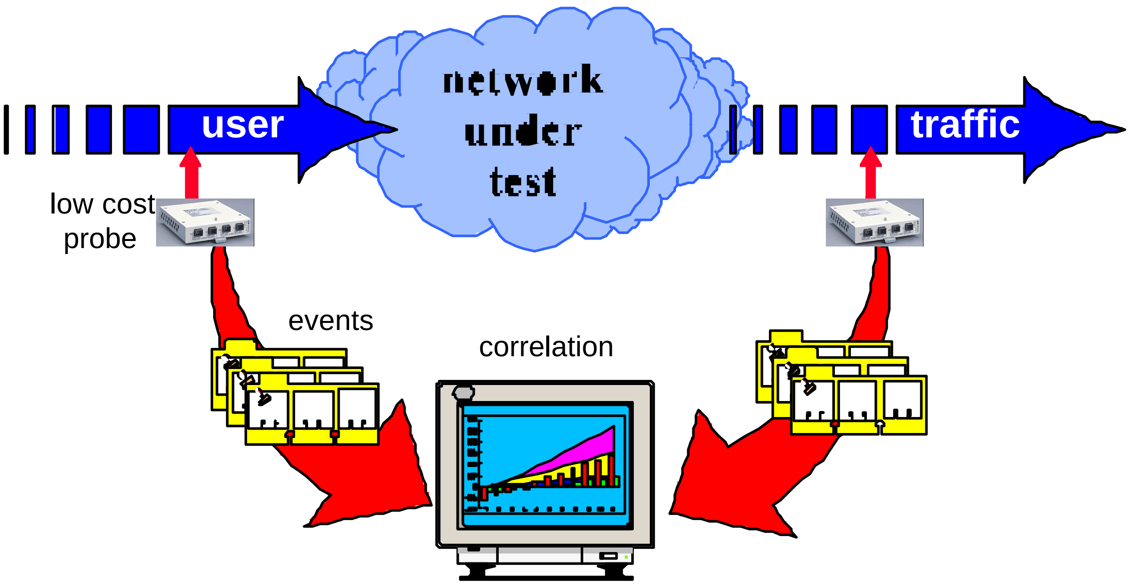

Figure 1. Measuring user traffic

The second difference in the operational phase is the focus of the measurement. In this phase it is the user traffic that needs to be measured and reported on, not the network. Obviously some parts of the network are measured and the performance of some elements can be deduced from measuring the user traffic. However, it is the quality of service the user traffic is experiencing that both the operator and customer want information about. Only the operator is interested in the performance of the network; and he certainly does not want that information to get to customers.

In this trial we have focused on measurements for Quality of Service and used the ITU-T I.356 specification [1] as the basis for ATM measurements. The I.356 specification was obviously not written with a modern systems-based measurement technology for operational networks in mind, but more of the R&D or I&M approach. However, a number of the measurements have been implemented as a first attempt to obtain useful information from the network. Obtaining the measurements for arbitrary user traffic requires placing measurement devices (probes) around the network so that at any time user traffic will pass a probe on entry to the network and another probe on exit from the network. Correlating data from the two probes allows end-to-end measures to be made. Two key traffic characteristics are built up from the measures: the profile of the traffic entering the network, and the profile of the traffic leaving the network. The entry profile enables the operator to see if the user if keeping to the traffic input agreement. Comparing the input and exit profiles enables both parties to see how the network has affected the traffic at any time. Statistical measures can be used to check that service quality agreements have been kept.

Figure 1 shows a scheme of the measurement system applied to a user traffic stream crossing a network. One of the problems with measuring any network traffic is carrying out data reduction so that the measurement traffic does not exceed the traffic being measured! In our scheme some data reduction is carried out in the probes. The probes then send measurement events (abstractions of user traffic) back to a correlator across a data network. Careful selection of events can reduce a 40 Mbit/s user traffic stream to a few hundred kilobits of measurement event data. It is the careful selection of the events that forms a key part of our research.
Figure 1 contains the components used to take a measurement of a user traffic stream and the key features, these are:

- **Multi-point** Measurements are made at two or more geographically distributed points in the network, the data from all of these points is correlated.
- **Correlation** It is the correlation of data, in real time, that provides the most useful data. For instance, detecting cell loss and PDU (Protocol Data Unit) loss (correlating two point measure) and relating them in time.
- **Intelligent probes** Data reduction carried out in the probes enables precise time based measures to be turned into useful information efficiently.

4. A Measurement System for Operational Networks

The measurement system for an operational network needs to support measurement probes geographically distributed around the network. It also needs to be able to deliver valuable information into the business processes of the operator wherever those processes require it. Thus, such a system needs to be distributed and flexible about the distribution. At any point in time the system can be measuring a user's traffic using any of the available probes and delivering information into any business process, as well as displaying visualizations to users and operators at their convenience.
The system we have developed uses distribution technology based on the CORBA standard [4] with some higher level services to support the requirements of application deployment and usage. The system is designed to support a number of simultaneous independent measurement activities, meeting the needs of a number of independent business processes. Figure 2 shows the main components of the architecture.

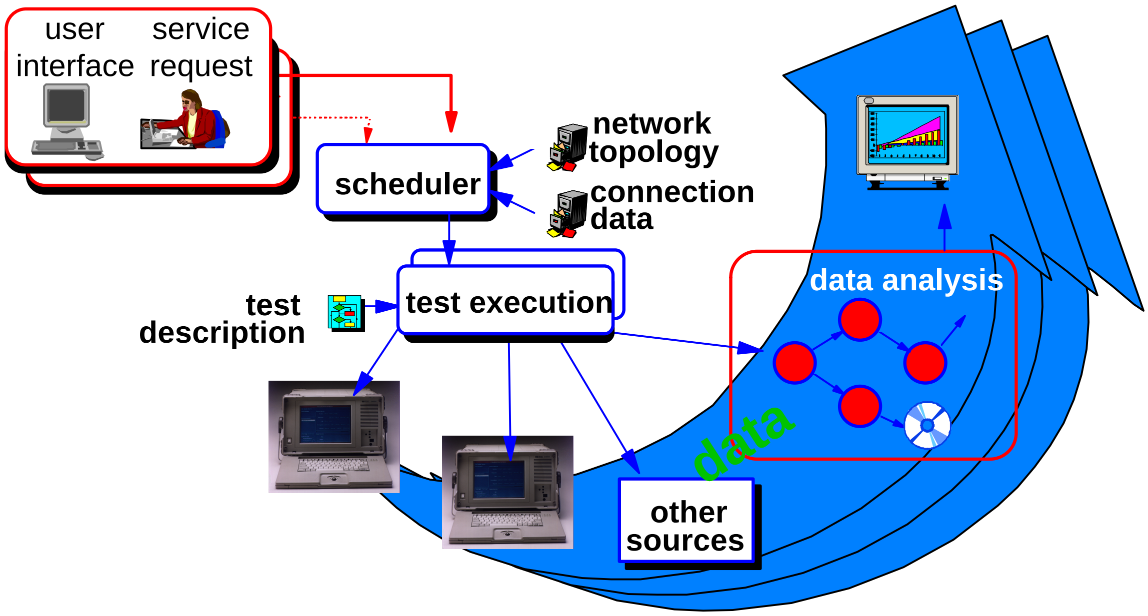

Figure 2. Measurement system architecture

In this figure the measurement process starts at the top left hand corner with a request for measurement data. The request may come from an internal operations source, or it may be external, such as from a customer owned management system. The request will be parameterized with the traffic stream to be measured and the type of measures required. The measurement system looks up the traffic stream in the management databases to determine the route taken through the network topology and therefore the relevant probes to be used. The required measures will determine what events need to be obtained from the probes and the correlation functions needed to process the events into the required information (test description). If a visualization is also required this will be identified and the destination (business process, visualizer, etc.) of the measurement information prepared. When all the resources needed for measurement have been identified the scheduler can decide when the measurement can be run; this may depend on the priority of the request.

When the resources are available the measurement is launched:

1. The probes are set up to identify the user traffic stream (from any other traffic) and to generate the measurement events.
2. The correlation functions (data analysis) are loaded onto appropriate computing resources and connected to the probes. The outputs being connected to the visualizes and or business processes.
3. The system is synchronized and started. Data flows from the probes (and possibly other sources such as network elements) through the analysis modules and out as valuable information into the operator’s business.

A measurement may run for a set time, or until other criteria are met, or it could be controlled from the visualizer. Eventually the measurement will be completed and the resources released for other measurement activities to use.

5. Trial Measurements

For the trial we have implemented two measurement applications (test descriptions) and simplified some scheduling aspects of the system. The two applications focus on traffic and measures at the ATM cell and packet (PDU) levels respectively. The applications are launched from a simple user interface in which the user traffic stream and probes are identified from pre-determined configuration files. Probe and processing resources are locked, and if an application encounters previously locked resources it simply says so and exits.
The measures reported by the two applications are based on those in the I.365 recommendation, with additional measures not conceived of by the standard. The most important aspect of the measurement system is not to report isolated measures but to show correlation of events (such as loss, and delay variance) at cell and PDU levels and with overall traffic service quality. Samples are taken of the user traffic at periodic intervals which produce cumulative measures over short time intervals. Individual cell and PDU activity can be shown over these intervals. The cumulative measures are written to files which can be used to show traffic quality over longer intervals, such as hours or days. The periodic sampling is an effect of the current probes we are using. In the future we will have continuous traffic monitoring capability and the user will be able to select the interval for reporting cumulative measures. The individual cell and PDU activity will then be available through a sliding window.

6. Laboratory testing of the system

Let us now look at how the first tests of the system were conducted in the laboratory, with the dual purpose of testing its functionality, along with the measurement accuracy, and of assessing its effectiveness when subjected to intensive use.
The configuration set up for this test phase is depicted in the following figure.

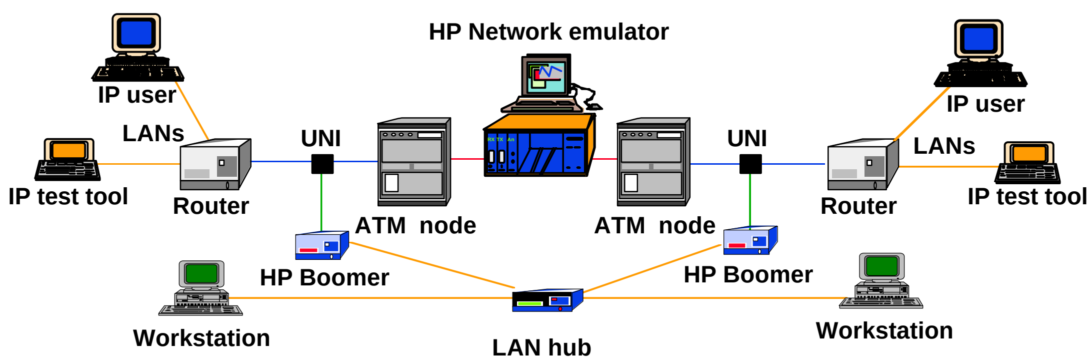

Figure 3. Laboratory configuration

Laboratory testing went through 3 phases. First, generating artificial, known traffic at the ATM, AAL5 and IP layers and checking the consistency of the measurements taken.
The second phase entailed inserting a network emulator that was able to insert impairments in a deterministic manner, in order to verify the accuracy of the measurements made by the system.

The last phase was aimed at recreating an environment that was as close as possible to the actual one, so a two-node ATM network was placed between the two probes and the traffic generated by actual applications was used to carry out long-term measurements.
The essentially positive assessment of the system reached in this period has allowed to make the decision to take the system to the field, albeit in a “friendly” environment like CSELT.

7. Field testing of the system

The monitoring system was moved from the laboratory to the field in January 1997.
The following measurement set-up was used:

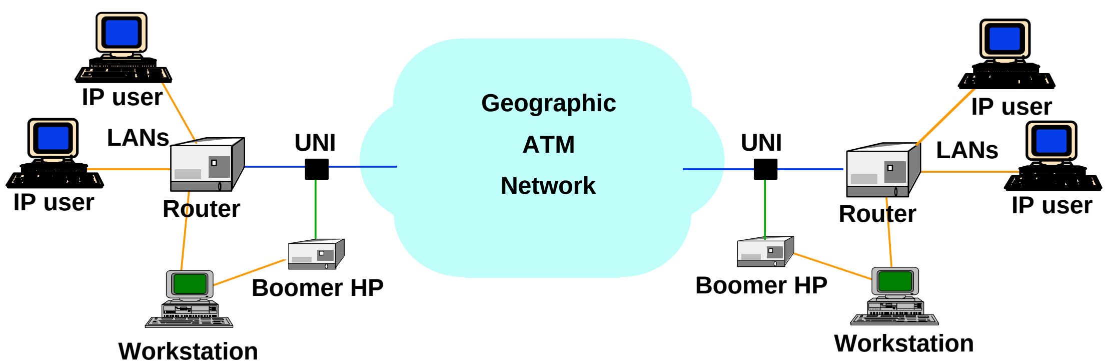

Figure 4. CSELT field trial configuration
The figure shows the following design choices.
We chose to monitor the ATM commercial connection between two routers collecting IP traffic between two CSELT facilities in Turin.
The connection between the Probes (Hewlett-Packard 5200 "boomer") and the control workstations was accomplished in the bandwidth, i.e. through the same CSELT company network using the ATM connection. For a commercial solution, Telecom Italia would be more inclined to separate the two networks (the monitoring network from the commercial network), using ISDN for connecting the probes.
The commercial ATM wide area network (called ATMosfera) is, in this particular case, made up of the single node in Turin connected through a roughly 10 km long, 34 Mbit/s PDH link to the two CSELT facilities.
After completing long term QoS measurements at the ATM and AAL5 layers, we plan to use an application to conduct QoS analysis at the IP layer, currently undergoing specification and development.

8. Some practical examples of utilization

Among the system’s multifold features, we have selected a few significant examples of practical utilization derived from the experience gained in these past few months.
The AAL5 QoS application displays the trend through time of some performance parameters. We have monitored the Cell Loss Ratio (CLR) and CRC32. We noted that, when errors are present

on CRC32 (with subsequent loss of IP packets) but CLR is at wholly acceptable levels, the Fault can be ascribed to the physical layer, essentially bit error. On the other hand, if error peaks on the CLR and on CRC32 were essentially correlated, the likeliest causes were policing violations or network congestions.

Another capability provided by the system is the study of the variation on AAL5,. and thus IP, packets between two observation points. A first assessment can be carried out observing the average length of the packets, expressed in terms of cell gaps in the first and second observation point. If the average length tends to increase, then it can be stated that the network tends to smooth out traffic peaks; if on the other hand it decreases, then peaks are accentuated. The first case can generally be considered to be positive, the second one negative. This study can also be carried out more in detail, by graphically analyzing each packet whose first-to-last cell delay is displayed.

The system's main limitation at this stage of its development, as shown by trials carried out with the network emulator, stems from the fact that not all traffic is analyzed: instead, samples are drawn for monitoring. We have thus been able to observe that errors generated according to constant or Gaussian distributions were assessed correctly, whereas even large discrepancies were possible for measurements on errors with different types of distribution (for instance, exponential). Of course, these issues will be overcome when, as envisaged by Hewlett-Packard's development plans, the probes available will be able to analyze all traffic and not just samples.

9. Using the system commercially

A system of this kind is, of course, ultimately meant for constant monitoring of the QoS offered by commercial ATM networks. During this initial development phase there are a series of practical limitations, due to the experimental nature of the solution, and it would not be commercially feasible to use the system extensively on the actual networks; more focused applications need to be identified.

For a telecoms operator it would be particularly useful, at this stage, to use the system on corporate networks used by particularly strategic clients. Both as a way to tune the offer and to provide the client with traffic reports that could assist him in internally rationalizing his use of the network.

This tool could also be immediately and particularly useful to solve any controversies with the client over the QoS offered by the network and its non-compliance with contract parameters.

Over the longer term, there could be many applications: from integration within network management systems to improve their performance, to help in the testing phase; from a dimensioning and planning tool to a system aiding in the design of new services over existing networks.

10. Conclusions and prospectives

This first phase in the collaboration between CSELT and Hewlett-Packard has allowed us to build, in a short time, a system which may be defined as pre-commercial, already oriented to an operator's requirements, moving from the laboratory to the field in just over a year.

There is surely much still to do; the possibility to analyze upper protocol layers, particularly IP, is currently undergoing development. Another particularly promising research field seems to be the analysis of traffic divided according to the applications employed by the user. The object is to define alarm thresholds, for performance parameters, differentiated among the various applications generating traffic over the ATM network. In this regard, we are currently studying world-wide web and Lotus Notes applications, extensively used over CSELT's company network. An additional field for testing with strategic interest for the near future is the integration with network management systems (particularly, HPOpenview) so as to integrate the performance information provided by network elements with those originated by the monitoring system. This would enable dynamic routing between ATM nodes, not based only on failures or congestions as is currently the case, but also on application QoS measured in real time.

Acknowledgements

The research, development and deployment of this novel measurement system was carried out by teams at CSELT and Hewlett-Packard. At Hewlett-Packard the team consists of Chris Dollin, Francisco Garcia, Patrick Goldsack, Claude Muller, Marc Nijdam, Sam Richards, Dean van Gerrevink.

Bibliography

1) ITU-T Recommendation I.356, "B-ISDN ATM layer cell transfer performance," 1993.
2) ITU-T Recommendation I.610, "OAM principles of the B-ISDN access," 1996.
3) M. Cociglio, M. Concina, A. Forcina, F. Ferrero, "Testing methodology and performance evaluation of the ATM pilot network in Italy," Telecom 95, Geneva, Switzerland, October 1995, proc. vol. 1, pp. 553-557.
4) CORBA 2.0, the Common Object Request Broker: Architecture and specification, Revision 2.0, OMG, July 1995.
5) S. Catanzariti, "ATM Testing: standards pave the way," Data Communication, June 1996, pp. 140-148.
6) E. Guarene, D. Serafini, et. al., "Implementing the Italian Atm Network: Issues and Technical Solutions," ISS '95, Berlin, 23-28 April 1995, proc. vol. 1, pp. 151-155.
7) M. Parker, et. al., "The European ATM Pilot," ISS '95, Berlin, 23-28 April 1995, proc. vol. 1, pp. 146-150.