REAL ATM TRAFFIC PERFORMANCE ANALYSIS IN BETEUS

F. Robles-Rojí[†], A. Serhouchni[†], S. Tohmé[†], M. Besson[‡], K. Traore[‡], and Ph. Dubois[*]

[†] Ecole Nationale Supérieure des Télécommunications, 46, rue Barrault 75634, Paris cedex 13, France.
{roji, ahmed, tohme}@res.enst.fr

[‡] ASCOM Sophia Research Lab., 2229, route des Cretes, 06560 Sophia-Antipolis, France.
{besson, traore}@ascom.cica.fr

[*] Institut EURECOM, 2229 Route des Cretes, BP.193, 06904 Sophia-Antipolis Cedex, France.
dubois@eurecom.fr

ABSTRACT

This paper analyses the performance statistics gathered by the TMN-based Network Management platform of the BETEUS project. BETEUS performs multimedia application trials in a testbed based on an ATM WAN that interconnects several local ATM LANs. The local ATM LANs, located in Sophia-Antipolis (F), Lausanne (CH), Zurich (CH), Berlin (D), and Geneva (CH), are fully meshed interconnected by 3 Mbps and 4 Mbps Semi-Permanent VPs over the European ATM Pilot. The BETEUS multimedia platform allows geographically distant users to participate in collaborative projects from their personal workstations equipped with audio and video facilities. BETEUS multimedia platform includes all the performance constraints of the existing and future multimedia applications on the multimedia hosts, on the networking interfaces, and on the networks. The QoS provided to the end users is directly affected by the performances of all the networking components along the communication path.

1 INTRODUCTION

BETEUS, which stands for Broadband Exchange for Trans European Usage, is a European project that has developed a multimedia networking infrastructure to provide a real usage of tele-collaboration services with the objectives to:

- give broadband access to tele-teaching organizations,
- develop and spread out the usage of remote education facilities,
- evaluate suitable multimedia applications.

BETEUS participants include universities, research institutes, and telecommunication companies. BETEUS provides a unique experience in the utilisation of applications with multiple and various kinds of multimedia streams (images, audio, text) over a Wide Area ATM Network. The BETEUS platform provides to geographically distant researchers, engineers, students and teachers a virtual community environment in which they can talk to each other, debate, edit common documents, teach, learn and design projects with a feeling of real presence. BETEUS focuses on distributed classroom, informal meeting, multimedia document archival and retrieval scenarii.

In general, multimedia applications place severe and various performance constraints on the networking systems in terms of bandwidth and QoS requirements that the conventional telecommunication technologies fail to fulfil [1]. It should be noted that these conventional technologies have been designed for a few types of traffic with a fixed QoS. ATM is being recognized as the basis of the next generation telecommunication technology in both LANs and WANs. One attractive aspect of ATM is its scalability in terms of bandwidth and QoS provided to the applications. It is therefore the candidate technology for carrying the data, voice and video traffics that are generated by multimedia applications.

BETEUS performs over the European ATM pilot which spans most of western Europe and involves 18 Public Network Operators. The BETEUS local ATM networks, located in Switzerland, Germany, and France, are directly interconnected through the ATM pilot. This real usage of multimedia applications over an ATM Network is a good opportunity to investigate whether the early ATM networks and the existing multimedia hosts satisfy the stringent performance requirements of multimedia applications.

This paper addresses the performance issues of the ATM networking environment provided to BETEUS multimedia applications. It also presents the network management platform that has been deployed for gathering performance statistical data within BETEUS ATM LANs.

This paper is organized as follows: section 2 presents the multimedia platform, section 3 describes the network management platform that has been used for the performance monitoring, section 4 analyses the statistics gathered, and section 5 presents our conclusion.

2 BETEUS Multimedia Platform

The BETEUS applications use multimedia workstations equipped with camera, microphones and loudspeakers to allow users to communicate with the other participants of the multimedia session. Each workstation can transmit, receive and process multiple video, audio and data streams. The applications infrastructure includes the components needed for running collaborative work sessions at the end points of a virtual community. Functionalities such as audio, video, data communication, telepointer facilities, shared work space, and a session management are provided. More generally, the multimedia services provided by BETEUS applications have been influenced by the following scenarii:

- distributed classroom and informal meeting,
- multimedia document retrieval and archiving.

In the distributed classroom scenario a high quality multipoint videoconferencing service and a shared workspace functionality are provided to allow geographically distant students to participate in lectures. Each remote site is shown in a window in addition of the lecturer in another view. Every student at a remote site can ask questions to the lecturer at any moment.

The Informal meeting scenario permits users to participate actively in group discussions from their multimedia hosts. This scenario focuses on the functionality to dynamically join-in and drop out sessions.

The archiving and retrieval multimedia documents functionality enables to compress and store multimedia sessions such as lectures and seminars. They can be replayed through World Wide Web servers. This scenario is mainly characterized by the transfer of large amount of information.

The data network used by BETEUS consists of ATM LANs and the European ATM pilot. The ATM Pilot is an experimental network that aims to confirm the interoperability in a multi-vendor and multi operator environment, evaluate the capacity of ATM as a technology to support broadband service and test applications in conjunction with pilot users. Its users include:

- Europe wide advanced research projects such as BETEUS,
- the European public research community,
- PNO's research labs,
- multinational companies.

The sites running BETEUS applications that are interconnected over the European ATM Pilot include: EPFL (Lausanne), ETHZ (Zurich), and CERN (Geneva) in Switzerland, TUB (Berlin) in Germany, and Eurecom (Sophia-Antipolis) in France (see figure 1). The sites are interconnected through Semi-Permanent bidirectional Virtual Paths (VPs) of 3 Mbps and 4 Mbps. Two kinds of physical interfaces are used to access the ATM Pilot: 34 Mbps (E3 interface) and 155 Mbps (STM-1 interface). Within each VP, three VCs are created for the video traffic, the audio traffic, and the data traffic.

The ATM configuration at each local site consists of a single or two interconnected private ATM switches (FORE ASX 200). The multimedia applications run on Sun Sparc 10 Workstation(s) directly attached to the switch through an ATM adapter card (FORE SBA 200).

Several network protocol stacks are used for communication among BETEUS applications (see figure 2). UDP/IP over AAL5 is used by video and audio streams. TCP/IP over AAL5 is used by the session manager application (contro. The end-to -end protocol may be influenced by these protocol designs and implementations and their performance should therefore be monitored.

3 NETWORK MANAGEMENT

According to ITU-T, performance management provides functions to evaluate and report upon the behaviour of telecommunication equipment and the effectiveness of the network or network element (NE). Its role is to collect statistical data for the purpose of monitoring and correcting the behaviour and effectiveness of the network, NE, and to aid in planning and analysis.

In short, performance management is performed mainly by monitoring the network and ensuring that the quality of service is satisfactory. It is also important to diagnose any chang-

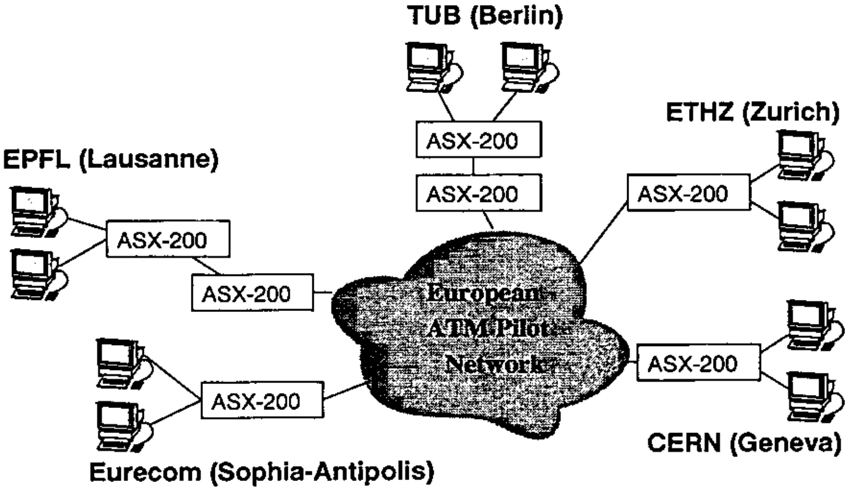

Figure 1: BETEUS Data Network

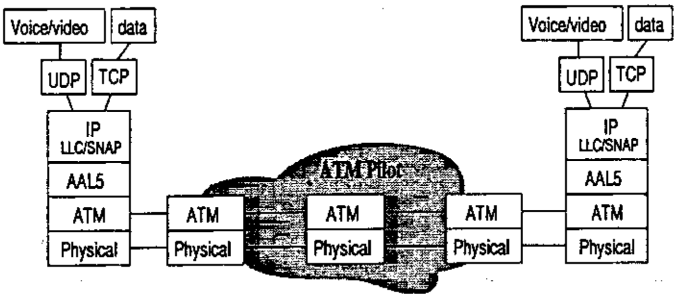

Figure 2: Protocols architecture used in Beteus

es that may modify the QoS. ITU-T [2] identifies the following tasks for performance management:

- performance monitoring,
- performance control,
- performance analysis.

A powerful and flexible capabilities are needed to monitor the performance of ATM networking devices. Measurements should be taken in short time scale and this necessitates a decentralized solution with more intelligence in the managed nodes. Quick management response to the performance degradation is mandatory in ATM networks since the occurrence of events is in the order of the microsecond. The TMN model offers suitable facilities to design and configure a flexible management system that can cope with the complexity of ATM networks. Therefore, a TMN network management system has been used to measure the performance parameters of the individual component of each local multimedia networking environment, including the local ATM switches, the multimedia host and its associated ATM interface.

The International Telecommunication Union (ITU-T) is contributing to the network management standardisation effort by introducing the Telecommunications Management Network (TMN) concept. The TMN specifications [3] aim to provide a framework for telecommunications management. The TMN is a logically separated network that interfaces at some points to individual Data Communication Networks (DCNs). At these points it controls the DCNs operations and can exchange information with them. TMN standards make use of OSI system management functions, protocols (i.e.,CMIP), and management information (X700 Recommendations). The TMN architecture is described in three views:

- the Functional Architecture,
- the Information Architecture,
- the Physical Architecture.

TMN functionality consists of the following components:

- the Network Element Function (NEF) communicates with the TMN in order to be monitored and controlled,
- the Workstation Function (WSF) block that allows human user to interpret TMN information,
- the Operations System Function (OSF) allows the NEF management. It processes information relating to management,
- the Mediation Device (MD) stores, adapts, filters, and condenses threshold information passing from network elements to operations systems,
- the Q Adapter Function (QAF) acts on the content of information passing between TMN function blocs and non-TMN entities.

Basically, the idea behind a TMN is to define an architecture that allows OSs and NEs to exchange management information specified in terms of OSI messages through standardized interfaces (Q interfaces).

The local ATM Switches and the adapter cards can be managed using SNMP only since only SNMP agents are provided by the manufacturer of these equipment. Other SNMP agents have been developed for the management of the multimedia host systems (UNIX Workstation) and the multimedia applications.The SNMP proprietary MIB implemented by the Fore ATM Switch contains information on the switch board, the buffers, the hardware ports and the different VPs and VCs. The ATM switches use buffering strategies at input/output ports or within the switch fabric for solving the possible resource contentions. Buffer overflows, and consequently cell losses, could occur due to the undeterministic variations of the traffic being queued.

The Fore Adapter cards provide also a SNMP proprietary MIB that maintains aggregate information for the different layers of the ATM protocol (e.g. Physical, ATM, AAL). The ATM layer performs no processing on the ATM payload field. Therefore there is no error control at the ATM layer. It carries cells without checking their contents. Errors in the payload field may appear only at the Adaptation Layer. These errors may result in CS layer errors in the case of AAL5.

The System MIB, designed for the UNIX hosts running the multimedia applications, includes information on the system activity (CPU utilization, I/O activity, context switching, etc.). Recall that UNIX, which is not a real time system, provides an unpredictable performance. The standard MIB II [4] is used to gather the performance statistics of the IP, UDP, and TCP layers on the UNIX hosts.The most common problem that occurs between the UDP layer and the IP layer is lack of resources (buffers) on the data transfer path. It may cause IP segments and therefore UDP packets to be dropped.

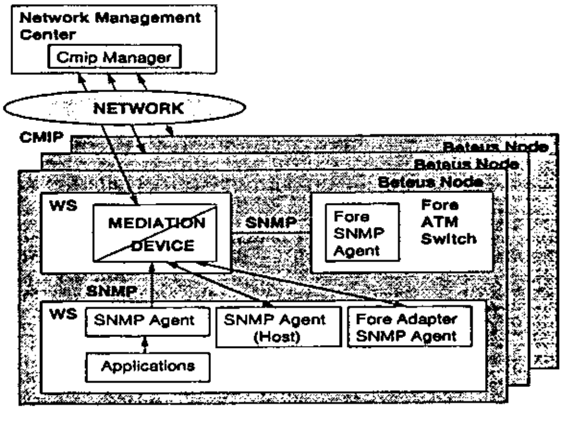

Figure 3: Local Network Management Configuration

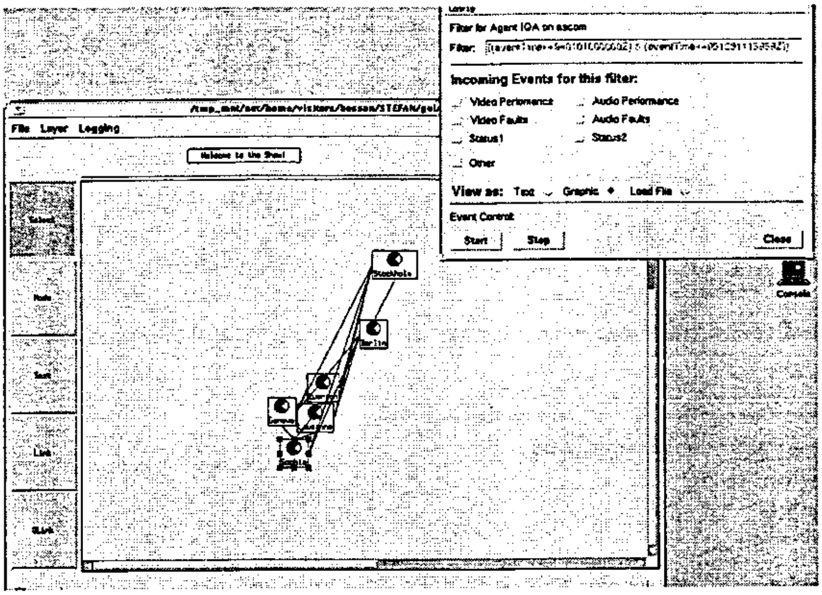

Figure 4: The BETEUS NMC

An application MIB has been designed to provide end-to-end performance statistics directly connected to the QoS perceived by users. The information of this MIB is related to the performances of the image and audio packets.

Since SNMP is considered as a non-TMN entity, a QAF has been used in each local sites to translate the CMIP messages into SNMP and conversely.

A Mediation functionality has also been used at the local sites to perform, filtering, polling, threshold checking, data summarization and statistical analysis. The network management configuration used in Beteus is shown in figure 3.

The WAN management traffic consists mainly to CMIP notifications containing performance data sent regularly by the NEs to the Network Management Centre (NMC) that is located in Sophia-Antipolis (see figure 4). The NMC is an WSF block. It should be noted that this traffic would have doubled with the SNMPv1 used locally and that allows to get performance data only with a polling mechanism. The polling is very dependent on the number of managed resources and may not allow the NMC to get a quick view of the performance of each individual component at the local sites.

The NMC includes the following components:

- A Presentation Service implements the Graphical User Interface (GUI) organized in maps and symbols representing the network devices in a topological view,
- A Data Management Service permits the storage of the incoming events,
- A Notification Management Service specifies the reception of specific events by registering Event Forwarding Discriminators,
- Communication Protocols (CMIP).

4 TRAFFIC ANALYSIS

4.1 Introduction

A global performance analysis of the BETEUS platform [5] was derived concerning the end terminal (protocol stack, operating system) and the ATM transport network (end to end analysis). The information collected during the performance monitoring period has been stored in flat text files and analysed off-line.

4.2 End Terminal Protocol Stack Performance Analysis

The performance analysis of the full protocol stack is done at a SUN SPARC station located at the ETHZ site. The multimedia application runned over the workstations generates variable size JPEG image frames sent at a constant rate that must be more than one packet sent every `200ms` (for an acceptable grade of QoS for the user). The multimedia application is set to send one JPEG image every `120 ms`. A bottom-up performance analysis covering the ATM layer, the AAL5 layer, the IP layer, and the application layer is presented to derive performance correlations between protocol layers. The host performance statistics are studied as well as the information concerning the virtual path connection between the workstation and the local ATM switch.

Audio/video application data follows an encapsulation process across the protocol stack where header information is added at each protocol layer (see figure 5.). The transport layer (UDP layer) adds 8 bytes of control information. Then, IP adds 20 bytes of header information. The IP datagram follows an LLC/SNAP encapsulation as proposed by the IETF (see RFC 1483 [6]) which adds another 8 bytes of overhead. A trailer is then added to the LLC/SNAP packet (e.g. the AAL5 SDU) to form an AAL5 PDU (8 bytes plus eventually, additional pad bytes). The AAL5 PDUs are then segmented

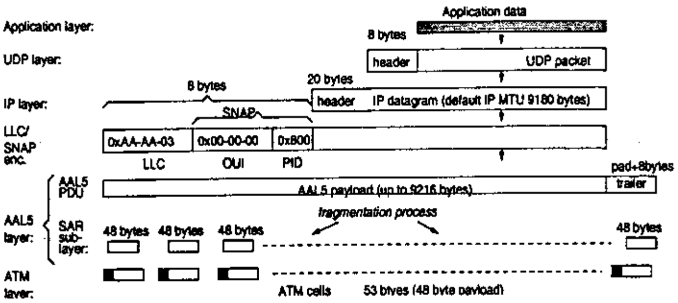

Figure 5: Encapsulation (disencapsulation) process of the user data

in 48 byte packets by the SAR sub-layer which are passed down to the ATM layer. The ATM layer adds the ATM cell header (5 bytes) and cells are emitted to the ATM network.

4.2.1 Physical Layer Performance Analysis

Physical layer errors include the indication of the number of ATM cells received with incorrect physical layer framing and the number of ATM cells received with bad CRC values. During the time duration of the test that lasted 4 hours 58 minutes and 35 seconds, the FORE ATM adapter card did not experienced any physical layer errors.

4.2.2 ATM Layer Performance Analysis

The ATM performance monitoring agent collects all the ATM cell information for the whole end station's ATM adapter card. The ATM layer objects collected by the agent at the end station included the *ATM emitted* and *received cells*, and the *unconnected VCIs cells*. The average time interval between two SNMP agent queries is 15 seconds.

Figure 6. shows the variation of the effective ATM emission and reception cell rate over the 4 hours and 18 minutes test period duration. The time axis is given in seconds and represents the actual time where the monitoring tests were performed.

In figure 6 (a), we identify seven main time periods of the ATM layer effective emission cell rate variation:

- The first time period represents ATM traffic due to signalling messages sent in bursts with an effective average cell rate of `1.8 cells/sec`.
- The second time period represents a single JPEG image emission flow of the local multimedia application with an average effective cell rate of `1053.8 cells/sec` (having a fairly constant traffic behavior with a standard deviation of `83 cells/sec`). With this average cell rate, an average size JPEG image is sent every `142 ms` (instead of every `120 ms` as set at the application level). The emission time variations are induced by the overall protocol layers and by the operating system but these are still acceptable for the multimedia application grade of service requirements.
- The third time period shows the combination of two JPEG image emission flows of the multimedia application (the average effective cell rate is 2492.8 cells/sec). Over this period the system does not alter the instantaneous observable cell rate of two superposed emitted image flows.
- The fourth and fifth time periods show first, three JPEG image flows being emitted and secondly, a single JPEG image flow continuing to transmit. Over both periods of time, the system does not alter the instantaneous observable cell rate.
- The sixth period of time shows a `90` second inactivity period where only signaling and initialization messages are emitted (as in the first period).
- Finally, the seventh period is characterized by the emission of two JPEG image flows with an average cell rate of 2475.2 cells/sec and a standard deviation of 161.62 cells/sec.

In general terms, ATM emission cell rate follows a fairly constant behavior during multimedia application activity pe-

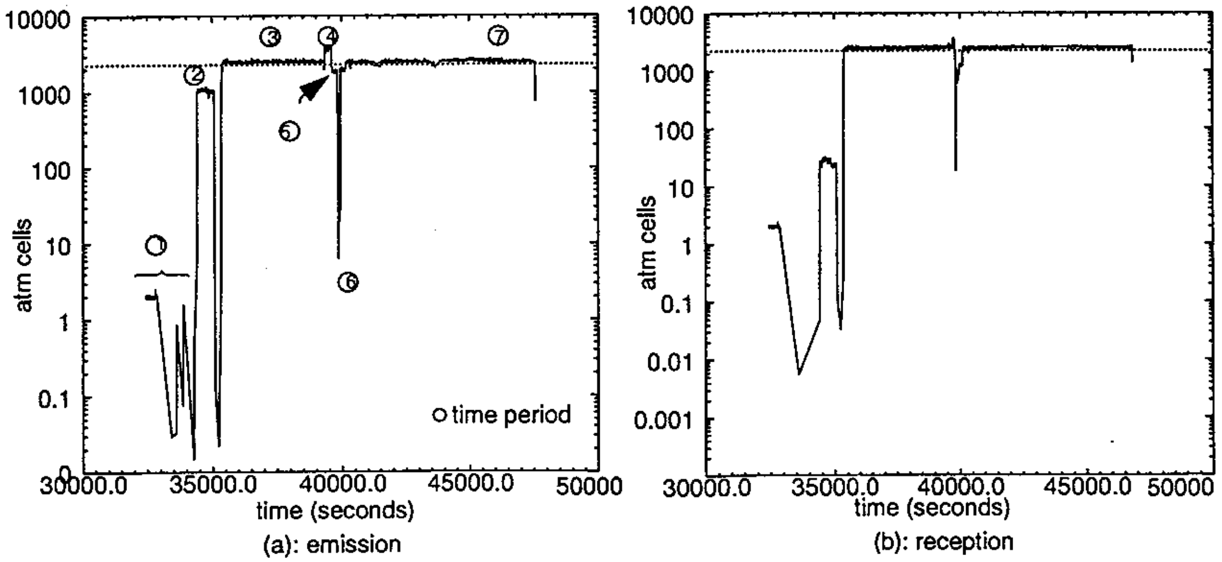

(a): emission

(b): reception

Figure 6: ATM layer effective (a) emission / (b) reception cell rate variation

riods. The superposition of multiple JPEG image flows does not disturb this constant cell rate behavior. However, the queuing of the multimedia traffic at intermediate ATM switches and at the end stations is one of the main factors that could modify the cell rate variation characteristics of the JPEG image traffic to be received by end stations.

The reception side (see figure 6(b)) closely follows the behavior of the ATM transmission cell rate variation (figure 6(a)). This similarity follows the fact that the local multimedia application establishes communications having symmetrical characteristics with the other multimedia applications located at the other sites of BETEUS. It is then important to notice that the network has a low impact on the JPEG image data traffic time variation between source and destination applications. There is a linear delay introduced by the network (e.g. propagation and processing delays) but waiting times in the ATM switch queues are negligeable. Cells entering the switches do not generally wait in the queues. This is particularly important in order to preserve the time variations of the JPEG traffic across the network. Virtual channels capacities are also well dimensioned to support the aggregated traffic. The ATM switch output traffic behavior and output buffer queue length will be analyzed in the ATM switch performance analysis section.

ATM Unconnected VCI's are the number of ATM cells received with VCI values in range but that did not correspond to active connections. There were 30 962 ATM Unconnected VCI cells during the test with no observable effect in the JPEG flow.

4.2.3 ATM adaptation layer (AAL) performance analysis

The ATM adaptation layer performance monitoring information is collected by the SNMP agent for the whole adapter instead of being by AAL connection. The average time interval between two SNMP agent queries is 15 seconds. AAL5 adaptation layer is used by the multimedia application data flows. AAL5 layer objects collected by the SNMP agent include the *aal5TransmittedCells*, the *aal5ReceivedCells*, the *aal5TransmittedPDUs*, the *aal5ReceivedPDUs*, the *aal5-CRCErrors*, the *aal5CellsDiscards*, and the *aal5PDUsDiscards*. Figure 7. shows the effective emission(a) and reception (b) rate variation of AAL5 type ATM cells and AAL5 PDUs over the test period time (4 hours, 12 minutes and 32 seconds

The AAL5 layer performance analysis allows us to verify the average number of cells per AAL5 PDU, AAL5 PDU measured loss rate, ATM layer cell loss rate estimation and to observe traffic behavior correlated to the ATM layer traffic behavior.

ATM layer and AAL5 emission/reception layer performance analysis comparison

- Globally the number of AAL5 emitted/received cells ***equals*** the total ATM cells emitted/received by the SBA200 ATM adapter. This confirms that all emitted ATM traffic is supported by the AAL5 layer.
- Both the ATM layer and the AAL5 layer effective emission/reception cell rates follow the same dynamic variations over the test period (see figure 6. and figure 7. respectively). No significant mismatches between these layers were detected. This leads to a minimum overhead introduced by the SAR sub-layer and by the ATM layer in the process of adapting `48 bytes` SAR PDUs into `53 bytes` ATM cells (and in the reception process of adapting 53 bytes ATM cells into 48 bytes AAL5-SAR SDUs). The system has a low load and AAL5 cells are crossing the ATM layer with almost no queueing time that could modify the traffic dynamic variation between layers.

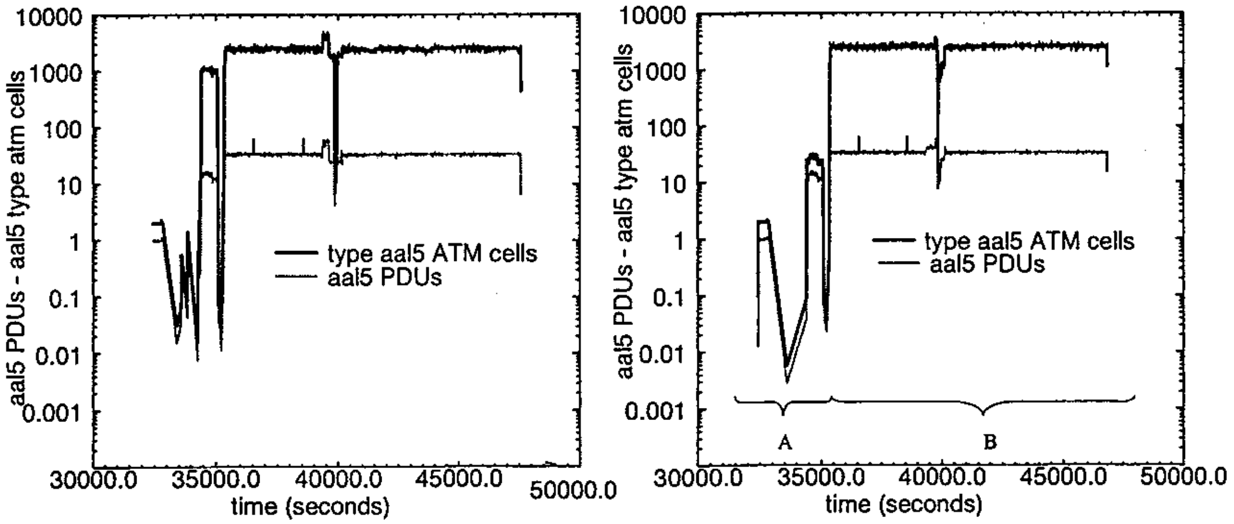

(a): emission

(b): reception

Figure 7: AAL5 layer effective variation of AAL5 type ATM cells and AAL5PDUs. (a) emission side / (b) reception side.

- No significant mismatches between the AAL5 PDU and the AAL5 ATM Cell effective rate variations were detected (see Figure 7. (a) and (b)). The AAL5 segmentation and reassembly process has no significant impact on the instantaneous traffic rate variations.
- AAL5 emission cell rate variations over smaller time periods and the multimedia data flow evolution correspond to the same variations as described in the ATM layer performance analysis section.

AAL5 emission PDU Size performance analysis

The AAL5 cell rate and AAL5 PDU rate variations (as shown in figure 7 (a) and (b)) allow us to consider that in average, the number of cells per AAL5 PDU corresponds to a simple calculation of the number of AAL5 cells divided by the number of AAL5 PDUs over a determined period of time. This consideration is made because the instantaneous traffic variations of AAL5 PDUs and AAL5 cells are almost similar in corresponding periods of time.

At the beginning of the test, 1 AAL5 PDU generates 2 ATM cells on the average. Over the 96 bytes of payload contained in 2 ATM cells, only 52 bytes are made available to the application data (the throughput is only a 54.1% of the usable data in the payload field).

In the periods where the multimedia applications were active we find an average value of 74 ATM cells per one AAL5 emitted PDU and 73 cells for one AAL5 received PDU. The average AAL5 PDU size is then equal to 74 x 48 bytes = 3552 bytes (emitting side) and equal to 73 x 48 bytes = 3504 bytes (receiving side). The emitting and receiving applications have a stable behavior generating AAL5 PDUs with similar performance characteristics. The variable size compressed JPEG images induce the AAL5 size variations but it keeps their values near the average value (the standard deviation for the number of cells per AAL5 emitted PDU is 16.46 cells/AAL5 PDU and the standard deviation for the number of cells per received AAL5PDU is 22.14 cells/AAL5 PDU). The 3552 bytes size of the emitted AAL5 PDU includes the AAL5 PDU trailer, the PAD bytes, and the payload data (LLC/SNAP header + IP header + UDP header + Application data). The application data represents 98.76% of the total payload data of the 74 ATM cells (only 1.238% is due to the AAL5 trailer, the LLC/SNAP, the IP and the UDP headers) which shows an improvement the emission performance throughput.

The performance throughput can be improved if IP datagrams are set to their maximum size of 9180 bytes (as described by the RFC 1626 [7]). User data represents 99.6% of the AAL5 PDU maximum size. The multimedia application compressed images are segmented in UDP packets of a maximum size of 4 Kilobytes (average size around 3500 bytes). If the application provides higher image segment size up to the MTU (minus the LLC/SNAP, the IP and the UDP header bytes) the emission performance throughput could be slightly improved.

AAL5: CRC errors, Cell discards, aal5PDU discards, and Loss Rate

At the AAL5 reception layer, every AAL5 PDUs discard was caused by a CRC error. The ATM cells belonging to a partially discarded PDU (in the course of being dropped) are also discarded. The information collected at this level allowed to verify the average number of cells created by an AAL5 PDU. For example, we list some information results gathered from the SNMP agent:

```
1.Event  time:  09:30:41  aal5CRCErrors:  20
aal5CellsDiscards: 40 aal5PDUsDiscards: 20

2.Event  time:  09:30:56  aal5CRCErrors:  9
aal5CellsDiscards: 18 aal5PDUsDiscards: 9

 ...                                       ....
                ......               ......

3.Event  time:  09:51:11  aal5CRCErrors:  1
aal5CellsDiscards: 74 aal5PDUsDiscards: 1

4.Event  time:  09:51:42  aal5CRCErrors:  1
aal5CellsDiscards: 73 aal5PDUsDiscards: 1
```

The number of discarded ATM cells and AAL5 PDUs follow the number of cells/AAL5 PDU behavior as described in the previous sub-section. The first two time interval sets confirm that at the beginning of the test period, 20 discarded AAL5 PDUs represent 40 discarded ATM cells which makes an average value of two cells dropped by a dropped AAL5 PDU. The following time periods confirm that in average, a dropped AAL5 PDU represents 73 discarded cells.

AAL5 layer AAL5 PDU loss rate.

During the whole test period, 37 AAL5 PDUs were discarded among a total of 394 413 received AAL5 PDUs. This leads to a total of 394 376 AAL5 PDUs passed to the IP layer.

During the initial test phase (area A of figure 7 (b)), a total of 29 AAL5 PDUs (58 ATM cells) were discarded (this is, *78.37%* of the total AAL5 PDUs discards during the whole test period). The AAL5 PDU loss rate during this period is equal to the number of lost PDUs (29) divided by the 481 received AAL5 PDUs giving a loss rate of $60x10^{-3}$ which is extremely high. Startup application settings (i.e. fixing the application windows at screen locations of the workstation) and bursts of signaling traffic information induced this high loss rate.

During the multimedia application activity period (area B of figure 7 (b)) there were only a total of 8 AAL5 PDUs lost among 394818 received AAL5 PDUs. This leads to an AAL5 PDU loss rate of (8/394818) $20.26x10^{-6}$ which is lower and which follows the constant and less bursty behavior of the multimedia data transfer.

ATM layer cell loss rate estimation.

An AAL5 PDU is discarded by the loss of a single ATM cell up to the loss of the total of cells forming the AAL5 PDU. For the application, the loss of a single AAL5 PDU induces the loss of a JPEG image (i.e. for the application, it is as if all the cells of the AAL5 PDU were lost). We make an *estima-*

tion of a lower and upper bound of the ATM layer cell loss rate in the initial test phase and in the activity phase depending on the number of cell lost by AAL5 PDU.

- Considering that the occurrence of cell loss is uniformly distributed over the time, a single ATM cell loss causes an AAL5 PDU to be discarded. This leads to a pessimistic approach since errors are normally correlated. In the initial period, 2 ATM cells form in average an AAL5 PDU. The lower bound ***ATM layer cell loss rate***$_{initial\ phase}$ = $60x10^{-3}/2 = 30x10^{-3}$. Then, during the activity phase, an average of 73 ATM cells form an AAL5 PDU. The lower bound ***ATM layer cell loss rate***$_{activity\ phase}$ = $20.26x10^{-6}/73 = 2.77x10^{-7}$.
- Considering that errors are normally correlated, the upper CLR bound considers that *all* the cells in the AAL5 PDU are lost (worst case). The upper bound ***ATM layer cell loss rate***$_{initial\ phase}$ = $60x10^{-3}$ and the upper bound ***ATM layer cell loss rate***$_{activity\ phase}$ = $20.26x10^{-6}$.

These values obtained are just estimations of the possible values that the ATM layer CLR could take. The performance analysis made at the local workstation - local ATM switch give real measured CLR results for the emission side.

4.2.4 IP layer performance analysis

The IP layer information is collected by the SNMP agent for *all* the received IP traffic every 5 seconds (on average) and for a test time period of 4 hours, 14 minutes and 59 seconds.

Information collected at this level includes the *received IP datagrams*, the *delivered IP datagrams to upper layers*, the *global average of IP datagrams delivered to higher layers*, the *IP datagram fragments that need to be reassembled*, and the *successfully reassembled IP datagrams.*

Figure 8. shows the variation of the effective IP datagram reception rate, the variation of the effective rate of IP datagrams to be delivered to higher layers, and the variation of the AAL5 PDU effective reception rate (for comparison purposes). Delivered IP datagram and received IP datagram effective rate variations are superposed in the following figure due to their almost identical instantaneous time variations.

The IP reception rate behavior can be analyzed in the following time periods (see figure 8):

1. *First time period:* the AAL5 PDU reception rate (representing multimedia application related traffic) is low but there is already a significant IP traffic data flow at the reception entity. This traffic will be considered as an external traffic which could be generated by different IP traffic sources. External traffic sources can include the Network Management Center traffic, the IP traffic received by the local Ethernet connection, X11 traffic, etc.

   On the average, 38.61 IP datagrams/second are generated by external sources while only 7.86 IP datagrams/second are generated by the multimedia application related traffic transported by the AAL5 layer.
2. *Second time period:* two incoming JPEG image flows (e.g.from distant multimedia applications) are added to the existing external IP traffic. The average value of the IP traffic flow is 74.34 IP datagrams/second while the AAL5 traffic (JPEG flow) rate has an average value of 33.6 AAL5 PDUs/second. On the average, the external traffic represents a rate of 40.74 IP datagrams/second.

   Over this period of time, IP traffic dynamic variations are mostly due to the incoming IP traffic generated by the AAL5 PDUs (e.g. by the incoming JPEG images). The external IP traffic has a stable behavior over this period of time.

   As traffic dynamic variations between the AAL5 CPCS sub-layer and the IP layer are closely preserved, queueing delay between this layers has almost no influence and the systems does not shape the received traffic. There is almost no waiting times in the process of receiving an IP datagram from the AAL5 layer and delivering it up to the IP layer.
3. *Third time period:* JPEG image flow activity has stopped. The receiving IP datagrams correspond to the external traffic generated by other IP traffic sources. The effective receiving IP datagram rate has a value of 40.1 IP datagrams/second. When the JPEG image applications stopped, the external IP traffic sources continued to generate IP datagrams without modifying their traffic characteristics.

<u>IP datagram reassemblies</u>

These are the IP fragments received which needed to be reassembled at this layer. During the test period, a total of 5658 IP fragments needed to be reassembled. All IP fragments were successfully reassembled.

The workstation at the ETHZ site has an Ethernet network interface and an ATM adapter interface. When using the ATM adapter interface, the default IP MTU value for use over ATM AAL5 is 9180 bytes. From this, IP will fragment packets issued from UDP, that are bigger than 9160 byte (e.g. user data + UDP header). Since the workstation maximum UDP packet size is 9000 bytes, fragmentation will never occur at the IP layer for data using the ATM adapter. Furthermore, the multimedia application sends fragmented JPEG image packets of around 4Kilobytes to the UDP layer. This is an important factor that reduces additional processing at the source and destination hosts IP layer that could cause QOS performance degradation of the multimedia application. Moreover, fragmentation in the BETEUS ATM network does not occur and all distant end stations have their IP MTU set to the default value of 9180 bytes.

Reassembled IP datagrams do not correspond to the multimedia application data flow but to external IP applications (e.g. applications using the Ethernet network interface).

No IP datagrams were discarded due to IP header errors or due to resource limitations.

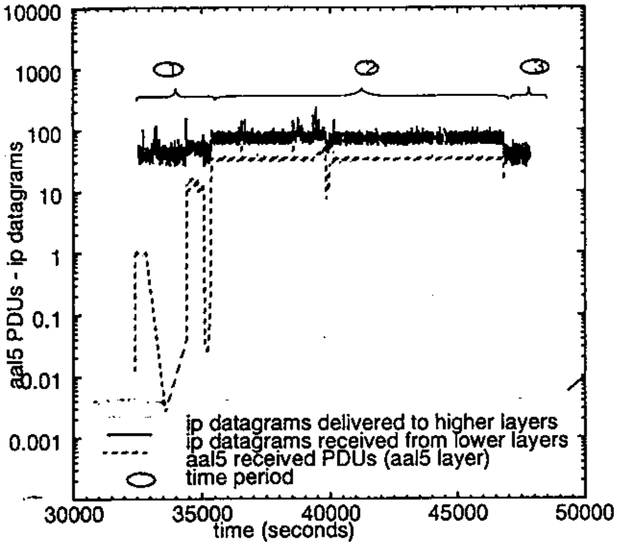

Figure 8: IP layer effective reception rate variations

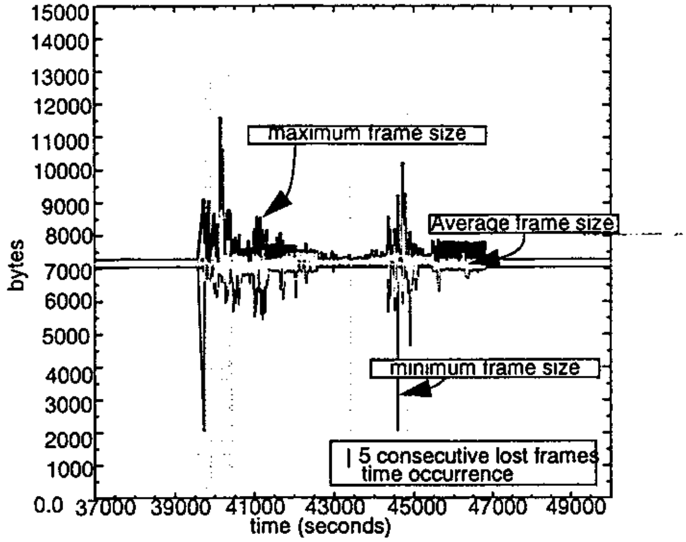

Figure 9: JPEG image frame variations& loss frame distribution

4.2.5 Application layer performance analysis

The application layer performance analysis covers both the video performance analysis and the fault performance analysis of the multimedia application. The performance analysis is made on the reception data flow. The ETHZ workstation receives video frames sent in from the EPFL and from the CERN sites.

4.2.6 Application layer performance analysis

The application layer performance analysis covers both the video performance analysis and the fault performance analysis of the multimedia application. The performance analysis is made on the reception data flow. The ETHZ workstation receives video frames sent in from the EPFL and from the CERN sites. Application event information is sent in UDP messages to a relay workstation station that contains the SNMP agent. The SNMP agent sends a SNMP message to its mediation device which converts the SNMP trap into a CMIP notification that is sent through the ATM network to the NMC. Application event information contains video performance messages sent on a periodic basis (on average every 5 seconds) and video performance degradation messages that are sent when 5 consecutive lost images are detected.

Video Performance Analysis

During the video performance analysis the frame interval rate was effectively found to be (frmInt) 120 ms. At the ATM layer emission performance analysis we found that an average size JPEG image was sent on average every 142 ms. Transmission time variations are induced by the protocol layer's processing and by the operating system but they do not shape the instantaneous rate variation of the emitted traffic (the induce a linear delay).

The application layer periodic messages included 3 values concerning the frame size variations over a 5 seconds average interval times. These values include the minimum detected frame size of 2050 bytes, the average frame size of 7139.5 bytes and the maximum detected frame size of 11598 bytes. Figure 9. shows the video frame size variations (minimum, mean, and maximum variations) over the last 7240 seconds of the test. The frame size variations belong to the traffic received from the EPFL site and the CERN site.

Reception rate performance analysis between the application and lower layers

Since the IP layer reception rate information includes the combination of the video traffic and of external IP applications, we assume that the AAL5 PDU effective reception rate is a good approximation to verify the application layer and lower layer performance correlations.

For a single video data flow, the application reads two incoming packets issued from the UDP layer (with an average packet size of 3500 bytes each) to form a JPEG image of around 7 kilobytes. The application will read this 3.5 kilobyte packets at a rate of 16.66 packets/second to maintain the rate of 8.33 JPEG images/second.

Since the application is reading **two** incoming video flows (over this activity period), the average reception rate is doubled to 33 packets/second (with an average packet size of 3500 bytes each). This ensures that the application will read a video image from each incoming video flow every 120 ms.

A bottom-up reception rate analysis is presented to verify layer performance correlations. This analysis is made over the period starting at time 10:59:30 hrs (39569 seconds) and ending at time 13:00:15 hrs (46815 seconds) were two JPEG image flows are received.

- The ATM layer presented a reception cell rate of 2490.5 ATM cells/second. All ATM cells are passed to the AAL5 SAR sub-layer.
- At the AAL5 layer, the AAL5 ATM cell reception rate equals 2490.5 AAL5 ATM cells/second. The measured AAL5 PDU reception rate is 33.249 AAL5 PDUs/second. This represents 74.9 ATM cells per AAL5 PDU which makes an AAL5 PDU size of 3595.41 bytes (74x48bytes).
- The application layer receives 33.249 packets/second with an average size of 3595.41 bytes. This gives an application reception rate of 119 543.78 bytes/second. We then consider the measured average value of a video frame size (7139.5 bytes) to calculate the average video frame reception rate value of 16.74 video frames/second.

This confirms that the application receives two JPEG image flows with an average rate of 8.37 JPEG images/second each. On the average, JPEG images are read every 119.5 ms from the Parallax image card buffer.

Error Performance Analysis

At the application layer, two kind of lost frame event report indications are issued towards the NMC. The first kind is sent in a periodic basis while the second is sent upon performance degradation (5 consecutive detected lost frames).

On a periodic basis, 4 lost frames were detected on UDP port 20001 of the EPFL - ETHZ multimedia session.

When 5 consecutive images are lost, a fault event report is issued to the NMC. Lost frames occur in bursts of at least 5 consecutive lost images. There were 60 fault event reports of this kind issued to the NMC. During the 7246 seconds test time period, a minimum of 300 images were lost among the 120718 total received images. This gives a lower bound frame loss rate of $2.4x10^{-3}$.

Figure 9. shows the loss frame distribution over the time period along with the frame size variations. Vertical lines represent the time occurrence of 5 consecutive detected lost frames.

Frame sizes around the average frame size value did not represent a consecutive lost frame problem. The system is well dimensioned to operate with frames sizes around the average size of 7139.5 bytes. Globally, consecutive frame loss occurred when frame sizes were over a value greater than 8.5 kilobytes due to system resource limitations.

Consecutive frame loss is not uniformly distributed over the test time period. We find two main periods of time where the application performance degradation was concentrated. For example, over a time period of 250 seconds (starting at time 11:09:00 hrs (40140 seconds) and ending at time 11:13:10 hrs (40390 seconds), 40 fault event reports representing 66.6% of the total frame loss were issued by the application to the NMC. Outside this small performance degradation periods, the multimedia application behaved with almost no consecutive frame loss.

4.2.7 Local ATM switch performance analysis

In this section, we analyze some performance aspects of the ATM local switch. We focus this analysis on the ATM connection between the switch and the workstation at the ETHZ site.

ATM switch hardware port queue length

The local ATM switch hardware port queue length (output buffers) of the connections with the local workstation and with the ATM Pilot switch were monitored. Both connection output buffers presented a maximum value of 1 ATM cell during the test. That event occurred only two times for the workstation connection, six times for the ATM pilot connection and the queue length remained empty during the rest of the time.

On the average, no waiting times were introduced by the hardware port queue of the ATM switch. The switch introduces a linear delay while maintaining the dynamic time variations of the ATM traffic (as will be confirmed in the *path* and *output path* cell analysis).

ATM switch Path Cells

The ATM switch path cells are the ATM cells received by the switch from a specific ATM connection (i.e. the ETHZ workstation connection). We observed that the received path cells variations at the ATM switch closely follow the ATM emitted cell variations of the local workstation. Measured average values of emitted ATM cells (from the local workstation) and the received ATM cells at the ATM switch are very similar. Cell rate dynamic variations are closely preserved. The input buffers of the ATM switch does not alter the instantaneous observable ATM cell rate.

The total number of path rejected cells (i.e. cells over this path that were rejected by the hardware due to traffic violation) is 14 cells among a total of 31 014 625 received path cells.

During the initial phase (i.e. no JPEG image reception activity), a total of path rejected cells were detected among 714 path cells received (this is 57.14% of the total rejected path cells detected). This gives a cell loss rate value, $\mathbf{CLR_{initial\ phase}= 11.204x10^{-3}}$.

During the activity phase, the JPEG image traffic is less bursty since policing was only performed over 6 cells (a total of 6 path rejected cells were detected among 31 013 911 path cells received). The cell loss value for the activity phase is very low, $\mathbf{CLR_{activity\ phase} = 1.93461x10^{-7}}$.

The measured cell loss rates have lower (better) values than the estimated cell loss rate values calculated in section 4.2.3. No direct comparison can be made since path cells are the emitted cells from the ETHZ work station and that CLR values obtained in section 4.2.3. were estimated for traffic entering the station. Nevertheless, we can observe that both initial phases show high CLR while both JPEG image data transfers show a low CLR preserving the QoS required by the end user.

ATM switch Output Path Cells

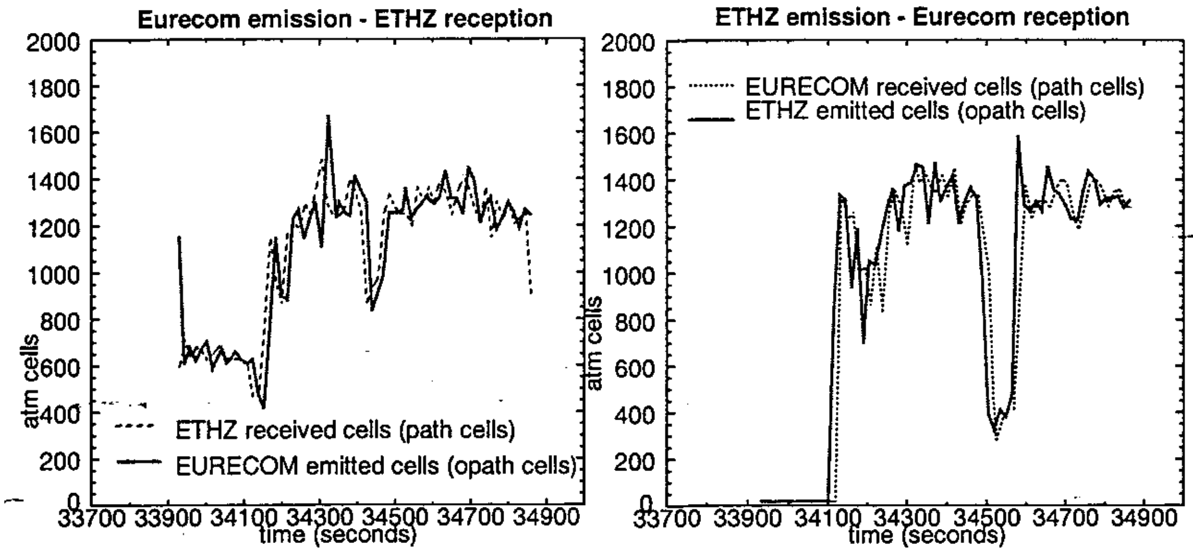

Figure 10: Distant private ATM switch emission and reception cell rate variations

ATM switch output path cells are the ATM cells emitted by the switch into a particular connection (i.e. the ETHZ workstation connection). This analysis showed in general terms that the output buffers of the ATM switch does not alter the instantaneous observable ATM cell rate. Furthermore the ATM cell rate dynamic variations are closely preserved between the output of the switch and the reception of the adapter card.

4.2.8 Host Performance Analysis

The CPU performance analysis is hardware and software dependent and must be integrated within the global performance analysis of the system [8]. The Host performance monitoring information is collected for the whole SUN SPARC 10 operating system SUN OS version 4.1.3. We use a simple technique to determine OS performance related to the emitting and receiving of multimedia application data. For example, we measure global OS performance information (with emitting and receiving multimedia flows and system internal flows). Then we measure the OS performance when no multimedia flow is being received and determine the OS performance due to the receiving multimedia information. Lower layer performance analysis (AAL5 layer for instance) allow us to determine correlations due to the multimedia application with the operating system performance.

We analysed the Host performance of Context Switches, Device Interrupts, Pages-In, Pages-Out, the CPU idle time and the user CPU time.

Interesting results showed that the number of context switches per second have a linear relation with the dynamic variations of the emitted and received multimedia traffic (compared with the AAL5 layer). On the average the emission and reception of an AAL5 PDU generate 4 context switches. The emission or reception of a JPEG image then generate (on the average) 8 context switches.

Device interrupts are generated when a host device (e.g. the ATM adapter) indicates to the host processor that he needs his attention. On the average, the process of receiving an AAL5 PDU and transferring the user data up to the application, generates 3 host processor interrupts.

4.3 End-to-end performance analysis: ETHZ - EURECOM

The multimedia traffic between the EURECOM site and the ETHZ site was analysed over a period of 15 minutes (in both directions). External multimedia traffic (from other sites of BETEUS) was added at the VP level between the private ATM switch and the local workstation (of both sites). In order to isolate the EURECOM-ETHZ ATM traffic, the performance analysis was done at the VP connection linking both distant private ATM switches. This analysis allows to observe the ATM pilot network impact over the dynamic rate variations of the emitted and received ATM traffic as well as the local ATM output buffer and hardware port queues influences.

Figure 10. shows the ATM cell rate dynamic variations measured at the output and input ports of distant private ATM switches. The EURECOM and the ETHZ graphics are out of phase of an average time value of 16 seconds because the clock between workstations are not synchronized in time and the SNMP time stamps depend on their internal system clock.

The average value of the emitted and received atm cell rate corresponds to a single multimedia JPEG image flow (around 1100 cells/second for the Eurecom-ETHZ flow and around 983 cells/sec for the inverse flow). Instantaneous cell rate variations correspond to the emission of variable size compressed video images generated by the multimedia application.

The previous figure shows that the dynamic time variations of the received ATM cell traffic follows very closely the behaviour of the dynamic time variations of the emitted ATM cell traffic. The ATM Pilot network does not have an impact on the dynamic time variations of the emitted and received ATM traffic. There is a linear delay introduced by the network but the received ATM traffic is not severely modified by varying delays. This shows that cells are not buffered (e.g. do not wait) in the ATM pilot network intermediate switches. The following results may mean that the ATM network is lightly loaded thus, not modifying the dynamic time variations of ATM traffic between the EURECOM and ETHZ sites.

Output Buffer queue length and Hardware Port queue length measurements performed on both distant private ATM switches showed that there is generally no queuing.

In section 4.2.2 (ATM layer performance analysis) we observed that the received ATM traffic faithfully followed the dynamic rate variations of the emitted ATM traffic as symmetrical multimedia applications were used between distant sites. It was then concluded that the ATM network has a very small impact on the emitted and received ATM traffic rate variation. The end-to-end performance analysis has shown that over two different test periods, the multimedia application ATM traffic runs over a lightly loaded network. Cells almost do not wait at the intermediate and the private ATM switches. Output buffers and hardware port queues at the local ATM switches showed to be generally empty. The dynamic rate variations of the ATM traffic is preserved along the virtual connections up to the ATM adapters at the workstations.

Globally the ATM network is not a factor that could cause performance degradation of the multimedia application. Round trip times are then an important factor that could cause performance degradations of the multimedia application.

5 CONCLUSIONS

This paper presented the BETEUS Multimedia and Management platforms followed by the ATM traffic analysis.

End Terminal Protocol Stack performance analysis conclusions

ATM traffic performance analysis was done at the ETHZ site at a local workstation and at the ASX200 FORE systems ATM switch. Multimedia applications emitting and receiving JPEG image data flows were used between the tested sites of BETEUS.

A bottom-up performance analysis was performed at the local workstation and performance correlations were found between protocol layers.

ATM traffic behavior showed to have fairly constant delay between protocol layers at the emission and reception side. The instantaneous traffic rate is not altered while data information transit between adjacent protocol layers. The multimedia application runs over a lightly loaded system because the transit delay shows to be fairly linear.

The increase of the IP MTU length, used over ATM, avoids datagram fragmentation of multimedia information at the IP layer thus decreasing processing times. On the average, the multimedia data represents 98.7% of the entire AAL5 PDU length. The throughput is actually very close to the maximum value of 99.6% of user data contained in a maximum size padded AAL5 PDU.

At the application layer, we noticed that image frame loss was due to the reception of large frame sizes. The application assigned buffers could not support image frames bigger than 8.5 kilobytes on the average. Frame sizes that remained around the measured average frame size did not induce a considerable frame loss. ATM layer loss rate is small. Its effect is not very much perceived at the application layer. The test involved symmetrical systems in the network sites. Applications established multimedia sessions with similar characteristics. At the local workstation, we found that the input and the output traffic dynamic variations are similar. The network showed to have a low impact on the JPEG image data traffic time variations between distant sites. Output buffers queue length at the local switch confirmed that no variable delay was introduced at this point. There is a linear delay introduced by the network with a very low waiting time in the ATM switches.

End to end performance analysis conclusions

The multimedia traffic between the EURECOM site and the ETHZ site was analysed in both directions over the VP connection linking both the private ATM switches. This analysis permitted to observe that the ATM pilot network had no significant impact over the dynamic rate variations of the emitted and received ATM traffic.

Output buffers and hardware port queues at the local ATM switches showed to be generally empty. The dynamic rate variations of the ATM traffic is preserved along the virtual connections up to the ATM adapters at the workstations. Propagation delays and round trip times are then important factors that could cause performance degradations of the multimedia application.

End to end performance analysis showed that the BETEUS ATM platform is able to fairly support isochronous traffic (e.g. voice and video).

REFERENCES

[1] «CCITT Recommendation I.350, General Aspects of Quality of Service and Network Performance in digital networks, including ISDNs», 1988.

[2] "Recommendation M.3400: TMN Management functions", ITU-T.

[3] "Recommendation M.3010: Principles for a telecommunications management network", ITU-T.

[4] "Management Information Base for Network Management:of TCP/IP-based internets: MIB II RFC1213", K. McCloghrie, M. Rose, May 1991"

[5] "BETEUS Workpackage WP4, project number M1010. Broadband Exchange for Trans-European Usage" September 1995

[6] «Multiprotocol Encapsulation Over ATM Adaptation Layer 5, RFC 1483», J. Heinanen. IETF July 1993.

[7] «Default IP MTU for use over ATM AAL5, RFC 1626», R. Atkinson. IETF, May 1994.

[8] «The Design of the Unix Operating System», M. Bach. Prentice Hall, Englewood Cliffs, NJ, 1986.