UCL Experiences with High-Speed Pilots

J. Crowcroft, M. Handley, A. Hutton, N. Ismail, P. Kirstein, D. Lewis, P. Su

University College London

Abstract

Many of the papers in the Broadband Islands conference treat the experience of using various networks for particular projects. University College London (UCL) has a set of such projects, mostly involving the Department of Computer Science (UCL-CS). These projects have involved the testing of emerging wide area high speed infrastructure for various application traffic types and the piloting of multimedia applications over this infrastructure. This paper describes some of the activities of such projects which have been put together at UCL both using SuperJanet and the European ATM Pilot (EUATMP), and the experience we have been gaining in their use.

1 INTRODUCTION

Over recent months much of the world-wide efforts in the development and testing of high speed services has moved out of the laboratory and onto wide-area pilots. These pilots have involved the national high speed infrastructure now available in many countries in Europe, there have also been trans-national projects, under initiatives such as the EU-funded RACE and ESPRIT programmes that have been enabled by the EUATMP. These activities range from establishing and characterising basic high speed connectivity over heterogeneous equipment through to establishing the performance possible from existing multimedia applications using various high speed network configurations available.

The Computer Science Department of University College London has been involved in a number of these piloting activities through different research projects. These projects cover those specifically aimed at high speed network piloting activities as well as those concentrating on specific client application areas, e.g. service engineering, network management and multimedia conferencing. Within the UK these projects have relied on the high speed infrastructure offered by SuperJanet which interconnects several academic and research sites as well as providing access to the EUATMP. A brief overview of some of the activities in which the various projects are involved in is given below:

- Detailed measurements of basic workstation and application performance over SuperJanet
- Piloting of audio/video conferencing over SuperJanet (these include the EU Telematics-funded DIMAGIO project demonstrations of surgical operations and tele-lectures)
- High-speed visualisation over SuperJanet and the European ATM Pilot
- Access to digital library information over SuperJanet (the BLRDD-funded CODA project as well as UKERNA funding) [5]
- Piloting of desktop multimedia conferencing over SuperJanet and the European ATM Pilot (the ESPRIT-funded MICE project) [6]
- Measurement of the characteristics of high speed video data transmission (the UKERNA funded TITAN project) [1]

- Integrated management of multi-domain broadband networks and multimedia services over the European ATM Pilot (the RACE-funded PREPARE project [3]).

- Application of Open Services Architecture [11] to high speed multimedia application management (the RACE funded DRAGON project)

These piloting activities have largely dealt with applications that use IP-based protocols. These protocols form the basis of much of the distributed application piloting and research work at UCL-CS and there is therefore considerable interest in its performance over the emerging high- speed infrastructure as well as the limitations this infrastructure might impose as is it colonised by the Internet. The use of IP protocols has had a major impact both on the topologies adopted for the various activities and the type of problems we have had to face compared to activities concerned with applications using pure ATM protocol stacks.

This paper attempts to relate some of the experiences in high speed networking and in the piloting of applications over these networks garnered through these activities at UCL-CS. The paper first outlines the general infrastructure available locally at UCL, nationally in the UK and internationally using the European ATM Pilot-covering network components, management facilities, protocol stacks and applications. It then describes some of the major issues that were encountered during the piloting activities. Next some examples of the actual work performed are presented providing a view of the specific types of problems that were encountered and some of the results obtained. Finally the further work being performed in this area at UCL is considered and what our assessment of the high speed network piloting activities are given. It is hoped that this paper will provide people planning future high speed network pilots and services some insight into the problems and characteristics of today's high speed infrastructure.

2 Infrastructure

The success of any high speed piloting activity consists of the suitable functionality being supplied by, and the successful integration of the following components:

- computer applications with requirements for high speed data transmission,
- application hosts with suitable high speed network adapters,
- local area high speed infrastructure,
- high speed wide area networks and access to them,
- network management facilities,
- traffic generation and performance measurement tools,

This section details the infrastructure that is used in the piloting activities, covering the pan-European infrastructure, the UK infrastructure, the local UCL infrastructure, the applications and hosts used and the test and measurement facilities available. For each of these, the level of management functionality available is also outlined.

2.1 Pan-European Infrastructure

The Pan-European high speed infrastructure used for the work described here consisted of various high speed national pilot networks inter-connected by the EUATMP. This EUATMP is a European-wide ATM network sponsored by the European Public Network Operators (PNOs) to promote the introduction of broadband services in Europe. The memorandum of understanding that defined the agreement to implement this test network states that the use of international links between the access points in individual countries could not be charged for. Each national access point is able to provide at least 34Mbps ATM access channels to the international network;

some countries have access to several such channels and some even to 155 Mbps international links. The access to this international resource is allocated on the basis of requests from individual projects. These requests have to be considered in concert with the allocation of resources on national networks connecting the sites from where individual projects were operating to the EU-ATMP access sites in each country. The requests for international links had to be requested from project representatives in each of the participating countries using a standardised form (TFD) - though there was some variation in the way it was handled in the different countries. In some of the participating, countries this representative is also responsible for co-ordinating the national part of the projects' connection to the EUATMP, often this still requires some interaction with other organisations whose networks are part of a project's end-to-end communication path. Topology and resource allocation can be modified over the lifetime of the EUATMP, through requests to national representatives, though this has to be closely co-ordinated by the projects with corresponding changes in the national access and local area networks. No automated management visibility or control of the EUATMP is available to the Pilot users, so diagnosis of problems usually has to be resolved by personal contacts between the organisations involved.

2.2 UK Infrastructure

Within the UK access between different testing sites and access to the PNO Pilot is provided by UKERNA via SuperJanet. This currently provides the following types of access mechanisms:

- An SMDS service operated by British Telecommunications plc (BT) connecting routers at 30 UK academic sites and used solely for IP services. Access is currently available at 10Mbps.
- A PDH network connecting 12 sites using 34Mbps lines. These lines are used to connect DV2 ATM cross-connects from GDC at each sites. These cross-connects are used both for connectivity of further ATM equipment at each site as well as legacy H.261 video-conferencing equipment.
- An SDH network connecting six sites via 34Mbps lines connecting Fore Systems ATM ASX-200 switches in a star configuration to a central SDH cross-connect operated by BT. The Fore switches provide access to further on-site ATM equipment.

The routers, GDC switches and Fore switches are all remotely manageable (either by remote terminal access or via SNMP MIBs); this function is currently performed by the Network Operation Centre at Edinburgh University.

2.3 Local Infrastructure

The high speed infrastructure at UCL was split between the Computer Science Department (UCL-CS) and Information Systems Division (UCL-ISD), with the latter providing access via SuperJanet nodes to the high speed WAN. The infrastructure in UCL-CS consists of:

- several SPARC workstations fitted with Fore Systems SBA-200 ATM host adapters providing 100Mbps multi-mode. The operating systems used were SunOS 4.1.3_U1 rev b, Solaris 2.3 and 2.4. (TAXI) interfaces.
- two Fore Systems ATM switches (one ASX-100 and one ASX- 200).

The supporting infrastructure at UCL-ISD consists of:

- a SuperJanet-managed GDC DV2 ATM switch connected to the SuperJanet PDH infrastructure.
- a SuperJanet-managed Fore Systems ASX-200 ATM switch connecting to the SuperJanet ATM SDH infrastructure.

- a CISCO 7000 router with an interface to the SuperJanet SDH service via a connection to the ASX-200 switch and a direct interface to the SuperJanet SMDS service.

The connection between the UCL-CS infrastructure, which acted as traffic sources for most of the activities, and the WAN infrastructure was either directly via ATM between UCL-CS and UCL-ISD or via the departmental Ethernet and the campus FDDI ring to the router at UCL-ISD.

Additionally, for the G7 summit demo and the connection to Inria and Telesytemes, a dedicated Ethernet was installed between a workstation at UCL-CS and the Cisco 7000 router at UCL-ISD, to ensure isolation from other traffic sources.

Figure 1 shows the UCL and UK High Speed Infrastructure

2.4 Applications, Protocol Stacks and the Interconnecting Technology

As mentioned in the introduction, the piloting and measurement work described in this paper was predominantly conducted using IP protocols, as the use and effects of high speed networks on the Internet were our prime concern. The different high speed IP activities that could be undertaken were obviously restricted by the availability of equipment and infrastructure (both local and international). However configurations were possible that have enabled us to gain experience in two important areas: the operation of IP over end-to-end ATM networks, i.e. ATM to the desktop, and the use of high speed network as the backbone connecting Internet routers

The emphasis on these areas drive requirements on the underlying architecture of the pilots towards certain areas that might be of less concern if pure (native) ATM applications were used. For example, the connectionless/routed nature of IP tends to pressure us to seek solutions to topology problems through IP addressing and routing, where a native ATM stack would only be concerned with signalling. Furthermore, the "layer isolation" that IP causes between the application's offered traffic and ATM's policing (and possibly even shaping) functionality leads to traffic management concerns that could be avoided in a set of native ATM (bespoke) applications.

This also helps explain why the effort in making the pilot work for applications in the different projects has been highly variable. For example, if we look at the protocol architecture and topology, there are at least four different "inter networks" supporting these high speed IP activities which are illustrated in Figure 2:

1. IP over an end to end ATM service. See Figure 2a.
2. IP over Ethernet LAN interconnected by routers which are long-haul interconnected by ATM. See Figure 2b.
3. IP over Ethernet LAN interconnected by routers which are long-haul interconnected by using SMDS - which itself is supported by the European ATM Pilot for inter-national links. See Figure 2c.
4. IP over ATM, but with different ATM management domains (e.g. SuperJanet, the European ATM Pilot, Renater etc.) separated by IP routers. See Figure 2d.

Because of the problems mentioned above, the first of these has been the hardest to make work, whilst the second was the easiest, and the third relatively straightforward. We are just starting to look at the fourth. An additional activity using native ATM video coders and decoders is also being performed and is also proving to be relatively straightforward.

The number of different organisations involved in the piloting work, i.e. those running the end applications, the national pilot operators and the various PNOs involved in the European ATM Pilot, meant that any co-ordination on the purchase of equipment was nearly impossible; to many of the organisations only became involved in the piloting work after making their initial investment in high speed technology. Though the variety of interface options offered with many of these products ensure that end to end network connectivity is possible, the state of signalling and network management standards, and their implementations in the products, precluded any automation of connection set-up or quality of service management over all network domains.

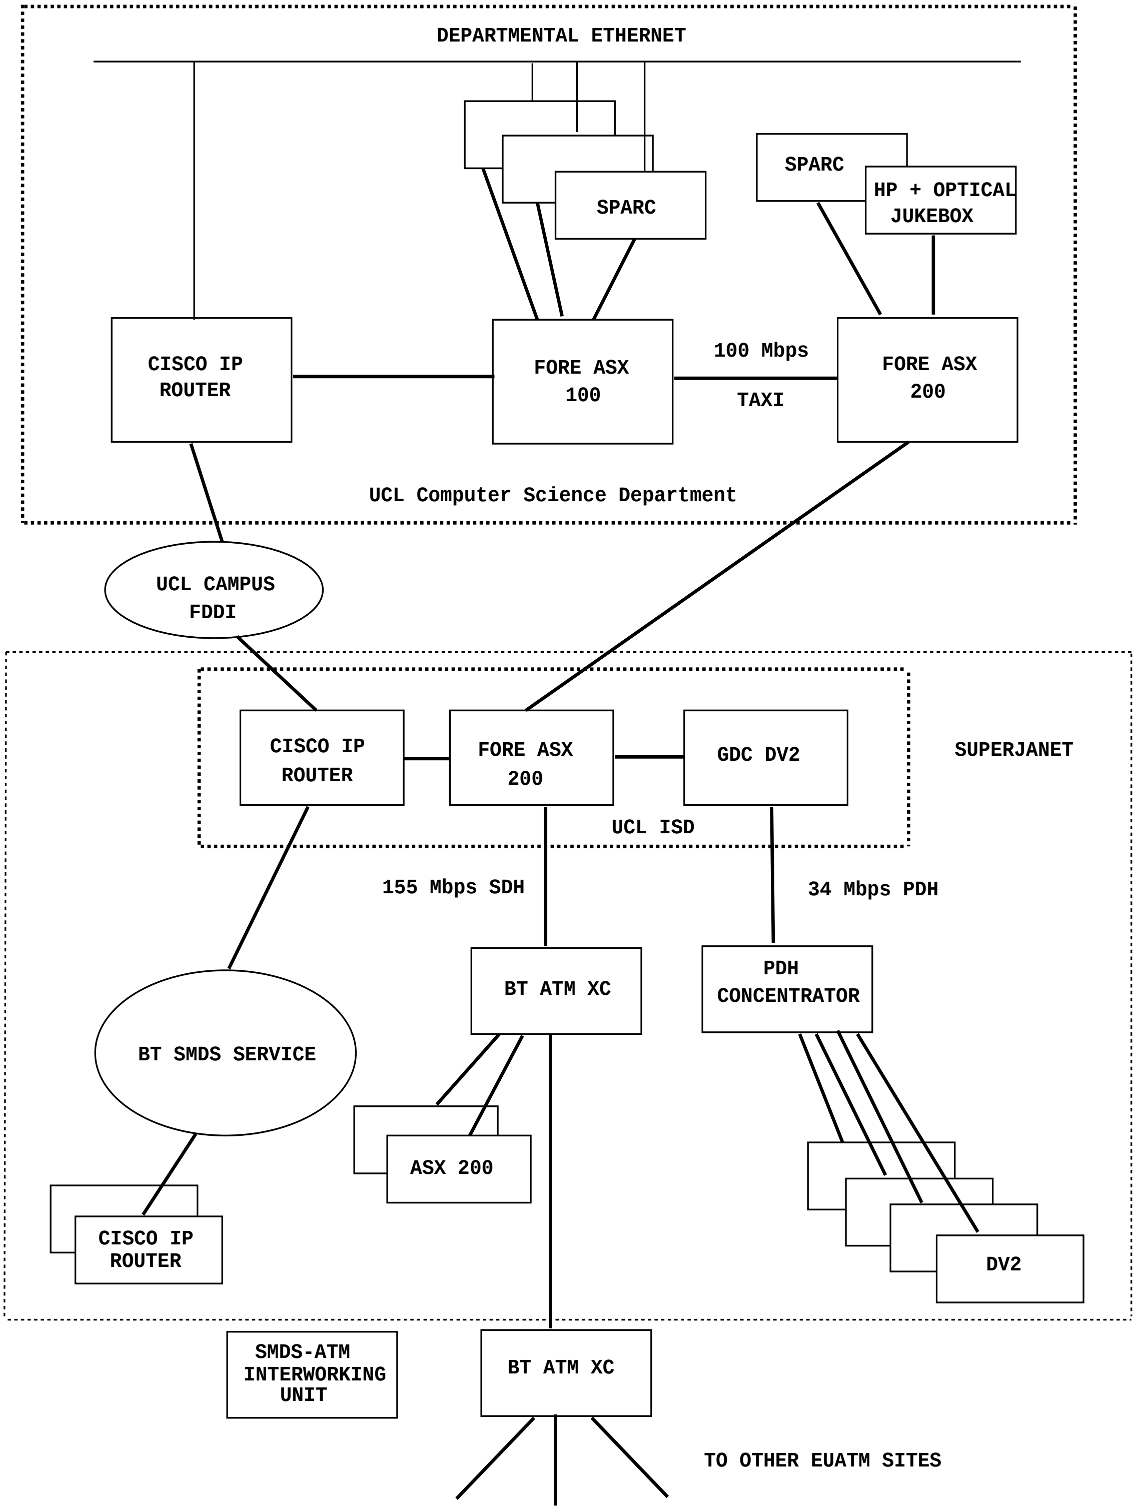

Figure 1: UCL and UK High Speed Infrastructure

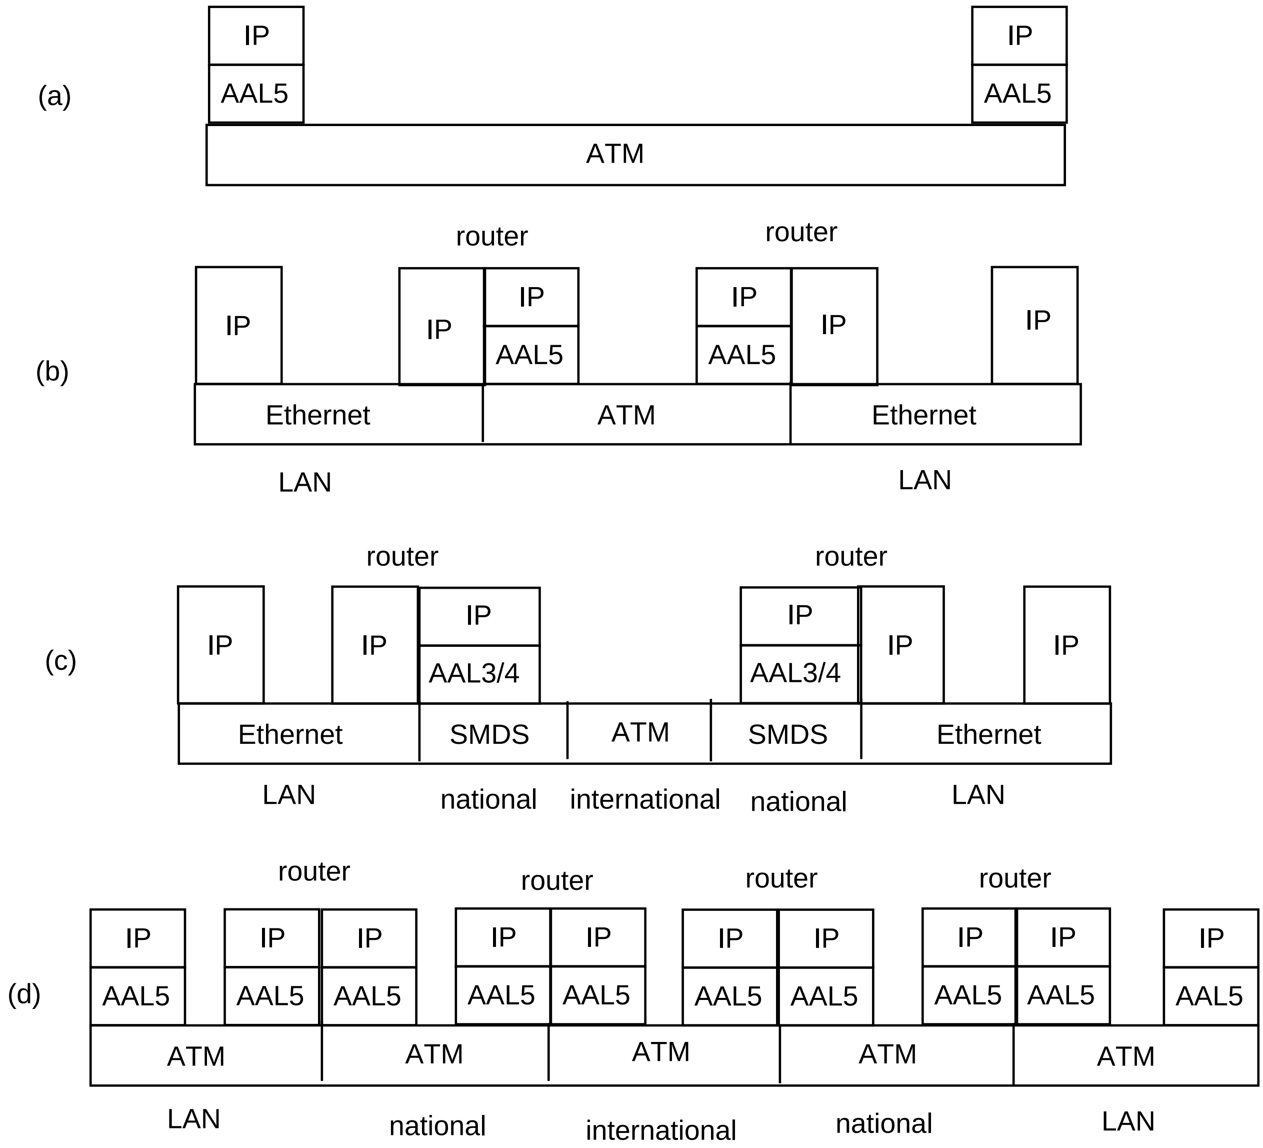

Figure 2: Network Configurations Used in Piloting Activities

3 Key Issues

3.1 Addressing and Routing

Due to the lack of end-to-end signalling and the profusion of administrative domains involved, addressing and routing at the ATM level was performed separately in each domain. In the EUATMP requests to the PNO for resources were responded to with individual VP values being allocated at the interface point between each network domain. End sites then agreed on a VC for a particular VP, and set up appropriate PVCs, specifying the AAL layer.

While many of our applications are point-point, others, e.g. the MICE conferencing one, are essentially multicast - a facility not provided currently in SuperJanet or the EUATMP. In view of the current six participating MICE countries it is already awkward and wasteful to set up the point-point routes needed even for the international connectivity. For this reason, we have constructed an international tree, in which each country has direct connection to at most three other countries at the application level; national connections all coming through one international MICE node as shown in figure 3. Further expansion in the international nodes, to ensure that the participating countries can be connected without exceeding the above fan-out with a similar scheme for national distribution, needs very careful design. This is being pursued in the context of the Terena activities of section 4.3. In a similar activity in BAGNET in North California, the participants have chosen to adopt a fully connected network of VPs. With even 20 nodes, this leads to 380 VPs - an administrative nightmare which does not scale.

3.2 Policing, Shaping and Performance

As many of the ATM switch products used in the EUATMP were only able to support Constant Bit Rate (CBR) traffic, this is the only type service that can currently be offered over the Pilot. The network boundary policing used to implement this service can normally be made only on a per VP basis; as a result, transmission capacity has to be allocated based on the maximum demand - which is a very inefficient way of using limited bandwidth. In some of our projects (e.g. PREPARE), we allocate a number of Virtual Paths (VPs) between each pair of sites; we must then control carefully the bandwidth required for each VP. In others (e.g. MICE), we have deliberately put all the traffic on a single end-to-end VP - and have then to control the maximum traffic on that VP.

The use of CBR obviously has serious implication for the typically bursty traffic generated by IP applications. Even multimedia applications that generate IP traffic at a constant rate are not able to influence the rate at which this data is transmitted as cells. Since the shaping of traffic to fit specific bit rate service is not a function currently available on the switches used before the Pilot boundary, efforts had to be made to ensure traffic shaping was performed correctly at the IP traffic source, i.e. workstation ATM adapters or router ATM interfaces. Some problems with shaping on end host adapter cards became apparent in our first wide area EUATMP link up, and prompted a series of tests carried out between UCL-ISD and UCL-CS. These tests used a UDP traffic generator at UCL-CS and a Hewlett Packard ATM analyser at UCL-ISD. This revealed significant unevenness in the host adapters CBR traffic shaping. We had noticed during the first wide area tests that we could only achieve a continuous ping of 108 bytes using some equipment. The cause of this turned out to be near zero of delay between the last 2 cells in a packet, this "breached" the policing on the EUATMP switch. This phenomenon had not been noticed in the local area or indeed UK ATM tests because CBR policing was not enforced in the same way. This traffic shaping problem was subsequently corrected. This phenomenon was first encountered in packet switches; where the specific buffer sizes initially limited the size of packets that could be used. With high-speed networks it is more efficient if the workstation uses long packets; now even with adequate spacing between packets at the workstation level, a single packet can produce enough cells to violate the traffic policing criteria, if there is inadequate inter cell pacing. [12] [13] Problems were also encountered with switches that could not deal well with both constant bit rate (CBR) and variable bit rate (VBR) traffic. A CBR video subnet has been set up on SJ which is disturbed by even a small amount of VBR traffic through the same switches The rectification of this fault delayed the start of some of the EUATMP trials.

4 ACTIVITIES

This section aims to give a glimpse of some of the activities described above and their results in a bit more detail.

4.1 Video traffic testing over SuperJanet

These tests were conducted as part of the UKERNA funded project TITAN (Transmission of Isochronous Traffic over Academic Networks). The involved testing video traffic characteristics between:

- workstations connected via Ethernet to local routers which in turn are connected by the SuperJanet SMDS
- workstations connected directly to ATM switches that are in turn connected across the SuperJanet PDH network.

Though SMDS is primarily intended for data traffic, there are many reasons why it should be suitable for supporting real time multimedia traffic as well. Firstly it provides high bandwidth access which is a basic requirement for multimedia traffic especially video. SMDS specifies that 95% of all traffic should be delivered in less than 30 milliseconds. This bounded transit delay is quite important to real time multimedia traffic, which is stream-oriented, and is quite sensitive to delays, provided that the delayed traffic is evenly distributed throughout time. SMDS sets the upper limit of the error ratio to be 1 in 10^{10} bits, the misdelivered ratio to be 5 in 10^8 data units and the undelivered ratio to be 1 in 10^4 data units. These upper limits are quite attractive for compressed multimedia traffic and hence it is important to ensure that they are adhered to.

The main objective of the tests was to investigate the effects of the ATM (PDH) and SMDS networks on multimedia IP-based traffic regarding loss rate and delay variance.

In general, video and audio data streams can tolerate losses more than delay variance, however, the effect of losses on compressed streams can be disastrous. Most of the multimedia compression standards used in IP-based video conferencing software can use inter-framing encoding techniques. These techniques create dependency within the compressed data streams which can cause the effects of a single error to be felt over a significant period of time. With inter-frame coding, a 2% packet loss is already quite noticeable for multimedia communication. The burstiness of the losses was also investigated. Burstiness is defined as the loss of more than one packet in a row.

Multimedia data is a continuous stream of traffic that is very sensitive to delay variance. Packet switched networks will usually impose delay variance upon transmitted traffic. As a result of this, applications dealing with multimedia traffic transmitted over packet switched networks must compensate for the effect of the network and ensure a continuous delivery of data to the output display devices. Usually the receiving end of an application will use a playout buffer to smooth the delay variance effects of the network. The higher the delay variance is, the bigger the playout buffers are and hence the longer the end to end delay between data transmission and data display will be. The recommended upper limit bound for end to end delay in video conferencing is 100-200 milliseconds.

Within the tests, both real as well as simulated multimedia traffic were used. H.261 [7] video streams were chosen to represent multimedia traffic in both the real and the simulated traffic tests. The reasons for this were :

- Compared with other types of multimedia traffic such as audio, the video traffic imposes the highest bandwidth requirements on the network.
- The H.261 international standard is implemented by many commonly used video applications such as IVS from INRIA [8] and VIC from LBL [9] .

- H.261 compressed video streams consist of a hierarchy of logical units such as pictures, groups of blocks, macroblocks and blocks. These logical units can easily be used as testing units.

The tests consisted of two phases. The initial phase involved a general testing of the SMDS and ATM connections which were later used in the second phase. During this phase, constant packet length streams were transmitted across the ATM network. The main objective was to ensure that the results of phase two represent the actual behaviour of the ATM network with neither interference nor odd interaction coming from the testing machines or from the ATM switches. The second phase involved transmitting both real and simulated multimedia traffic over the SMDS and ATM networks.

For each phase, two types of tests were carried out. The first type was direct tests where traffic was transmitted between a machine in UCL-CS and a machine in UKERNA connected via the SuperJanet ATM network. The second type was loop back tests where traffic was transmitted between two machines in UCL-CS. The traffic was consecutively looped back through one of the four ATM switches used in the direct tests.

The results of the real and simulated traffic tests were encouraging, when switches present no policing problems, at transmission rates of 2Mbps on a VP of10Mbps. For ATM error rates were low; in the range of 0.04-0.004%. All traffic losses were bursty in nature. The delay variance caused by the network was not significant with more than 98% of all traffic had jitter values less than 2 milliseconds during the simulated traffic tests.

For SMDS error rates were low; in the range of 0.12-0.001%. 95% of the traffic lost was bursty in nature. The delay variance caused by the network was not significant with more than 96% of all traffic had jitter values less than 2 milliseconds during the simulated traffic tests.

4.2 IP traffic measurements over the EUATMP

Traffic measurements have been carried out on the networks see figures 2a and 2b in order to compare traffic throughput using ATM capable Cisco routers and Fore switches with Fore ATM end host adapters. Results vary depending on hardware and software version used.

The traffic generator/simulator used is implemented in software as a client and server. Traffic transmitted by the simulator can be of either fixed or variable size. The packets are UDP, to simulate packet based video traffic. The rates per second at which the packets are sent determines the initial delay between them. It is possible to set the length of time a test runs and whether packets are fragmented. As well as using fixed packet sizes the simulator can also read from a file of varying packet sizes.

Fixed packet size tests were run on both of the above network configurations to identify the point at which there was *no loss or virtually no loss*. The packet size used was 1000 bytes. All tests were run repeatedly for 120 seconds over a time span of several days between UCL-TeleDenmark in Copenhagen over the network configuration in figure 2a and between UCL-Telesystemes over the network configuration in figure 2b.

The tests were run over a 2Mbps VP using PVCs carrying IP over AAL5, with a packet size of 1000 bytes, and an IP protocol overhead of 28 bytes (20 bytes for the IP header and 8 bytes for UDP). If LLC/SNAP encapsulation were used an additional 8 bytes of overhead would be added.

There are 3 sources of protocol overhead in ATM:

- the 5 byte header of the 53 byte cell
- the overhead of dividing the initial payload into an even amount of 48 byte loads padding out the last cell where necessary.
- Finally, there is the overhead of the particular adaptation layer used. AAL5 has a lower protocol overhead than AAL3/4. 8 bytes of protocol overhead are added by AAL5 - however, this can incorporated into the payload of the last cell, bytes permitting. [14]

Size of IP Datagram (bytes)	1028
No. of 48 byte payloads	21.4
No. of 48 byte payloads + padding	22.0
AAL 5 header included	Yes
Total size in bytes	22 * 53 = 1166

Table 1: Protocol Overhead

So given an initial packet size of 1000 bytes and an UDP/IP datagram size of 102 8 bytes the total size in bytes affter the addition of ATM headers becomes 1166 bytes as shown in the Table 1.

Given a bandwidth of 2Mbps and 1000 byte packets (1166 bytes including protocol overhead), if it were possible to use all of this bandwidth then we should be able to achieve around 214 packets per second. This was not achieved in any of the tests. Table 2 shows the typical results. Please note the versions of Fore drivers used were:

SunOS 3.0.1-1.114 patched to deal with the problem of traffic shaping.

Solaris 3.0.1_1.28

Rate/Mbps	**Cisco %**	**Fore/SunOS %**	**Fore/Solaris %**
1.8	6	15	1.80
1.72	2	11.5	0.4
1.6	0.025	5	0.3
1.59	0.07	4.5	0.26

Table 2: Loss related to Hardware/Software

Loss would be expected at 1.8Mbps, since at this rate packets including overhead should exceed the 2Mbps peak rate by around 98Kbps.

As can be seen from the table, we have found very variable end-to-end losses with the EUATMP. These are probably due to performance or defensive programming in the ATM switch interfaces or the interworking units. It would take another full paper to desseminate the measurements made - let alone explain them all.

4.3 IP based multimedia applications over the EUATMP

This involved local LANs and routers being interconnected using the EUATMP for activities by both the MICE and DRAGON projects.

In the case of MICE, the network is used for multimedia conferencing using multicast IP based applications such as IVS, VIC and VAT [10]. The first wide area link up happened in February 1995 using this technology for a demonstration at the G7 summit in Brussels. For this multicast has been used between 3 sites - Brussels (the summit), Nice (Inria), London (UCL-CS) using tunnels set up across the IP over ATM links, effectively using the ATM network as a higher bandwidth backbone.

The above setup was a precursor for the recent MICE/TERENA linkup for the IETF (Internet Engineering task Force) in July in Stockholm.

See figure 3

Dragon is a RACE project on open service management. [11] A Service Machine is built using Integrated Engineering Approaches and Open Service Architecture. [4] It has the functions of service provision, control, deployment and accounting management. The Dragon project uses an Open Service Architecture based Service Machine for the provision, control, deployment and

AN EUROPEAN MULTICAST-IP over ATM NETWORK
provided by MICE-project and TERENA ATM-TF
used by Mbone

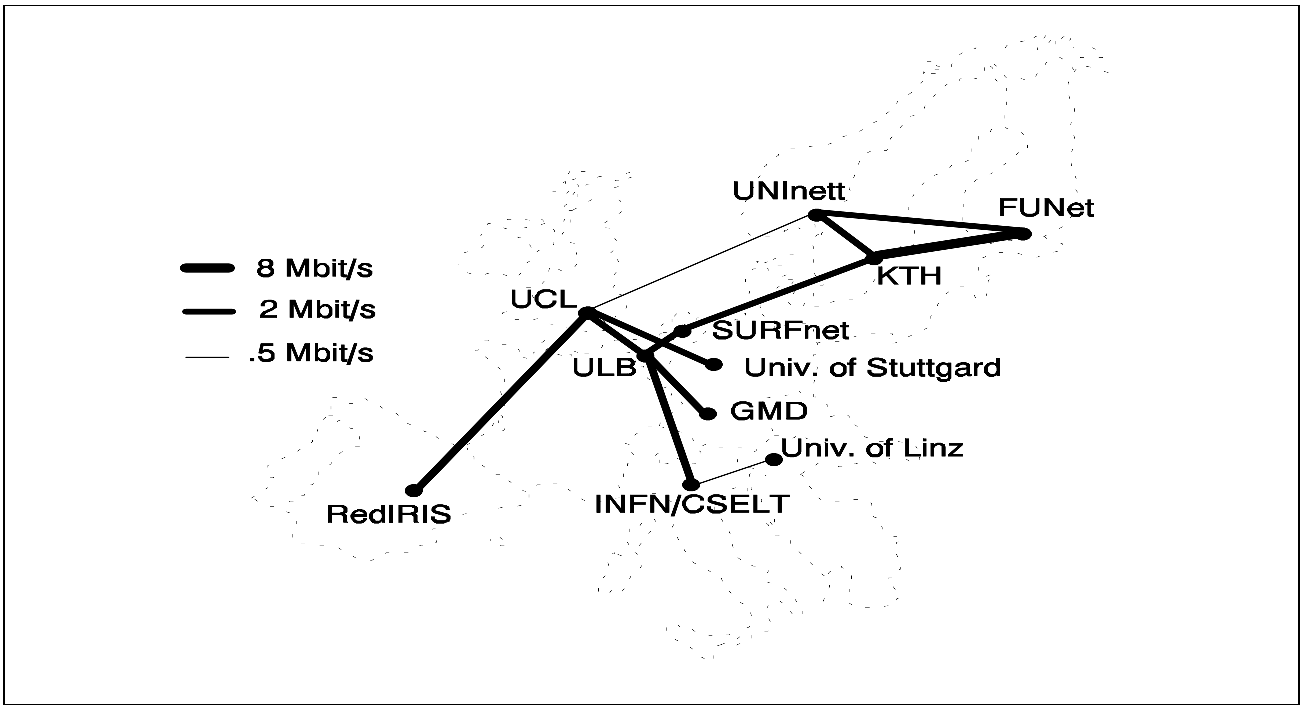

MICE-project:
M.Handley@cs.ucl.ac.uk
http://www.cs.ucl.ac.uk/mice/mice.html
TERENA ATM-TF:
Victor.Reijs@SURFnet.nl
http://www.terena.nl/terena/technical-programme/task-forces/tf-atm/
Mbone:
Erik-Jan.Bos@SURFnet.nl
http://www.nic.surfnet.nl/surfnet/persons/bos/mbone/

IETF/VR-1, July 14, 1995

Figure 3: MICE/TERENA ATM MBONE

management of multimedia conferencing services over high speed networks. This conferencing service is based on the VIC for video, VAT for audio and WB for shared work-space conferencing. A 2 Mbits/s ATM link was established for test and technical integration purposes between UCL-CS and Telesystems in Paris. We plan to supplement this with an SMDS link to Alcatel in Stuttgart; a full demonstration of the DRAGON scenario is planned for Telecom '95 in Geneva. The DRAGON project also plans other demonstrations of applications in remote visualisation and editing.

When using VIC in uni-cast mode to transmit video, the maximum rate obtained were 300kbits/s for H.261 encoding and 1.4 Mbits/s for JPEG encoding. The different results here being due mostly to the processing load on the encoding workstation imposed by the different techniques. In these trials, it was also found that when one end transmit less than 40 Kbits/s video, the other end can transmit almost as much data as the video encoder could generate, i.e. about 300kbits/s for H.261 and about 600kbits/s for CellB and 1.4 Mbits/s for JPEG. However, when both ends transmit at a higher rate, losses occurred. Some of the experimental results are shown in table 1 while h.261 encoder is used in both ends. The same phenomena is also observed when JPEG and CellB are used. The reason is under investigation.

Transmis. Rate (Kbits/s)		*Loss Rate (%)*	
UCL-CS	TS	UCL-CS	TS
300	30	0	0
30	300	0	0
150	60	up to 70	up to 45
150	150	up to 80	up to 80
90	90	0	0

Table 3: Video over ATM — Transmission rate v. loss rate

5 ASSESSMENT

During these trials the IP throughput performance has often been disappointing and less than the expected levels. Issues of traffic shaping, policing and intelligent cell discard need urgent attention if ATM is to be widely used in the Internet.

In the EUATMP, the user-defined switching is minimal; VPs are set up manually. This is awkward both from an experimenter's and an operator's viewpoint. It has led UCL-ISD, for example, to desire fixed routing with few daily changes in the establishment of VPs.

Access to cell count and cell loss statistics across all domains is essential for deployment of such high speed network and lack of remote access to this information can significantly slow down this process. The current effort on switch and network management is much needed by the operators.

While EUATMP has not been charging for international traffic, the national charges have often been too high for commercial partners in EU projects. This has resulted in only limited use by commercial partners - unless they were PNOs.

6 Further Work

Further work at UCL-CS in this area will be aimed both at continued piloting activities in different application areas and an assessment of improved product functionality as it becomes available.

On the application front the PREPARE project is implementing an integrated, multi-domain network and service management architecture that couples the management of multimedia teleservices with active management of high speed network elements in the private and public domain [2]. This addresses many of the inter-domain management problems encountered during trials with the EUATMP and it is expected that results of this work will be considered EURESCOM in planning further international pilots.

The MICE/TERENA activity will continue to develop and assess the use of wide area international high speed link for use in the IP multicast backbone. By September this is intended to embrace at least a 10 country multicast MBONE over EUATMP; there will be several hosts attached in each country. This will strain our managerial capability; we hope to succeed, however, in establishing a viable and rugged experimental MBONE over the EUATMP.

Imminent work will also address:

- the connection of signalling aware ATM LANs over signalling blind public network
- comparisons of host based and switch based traffic shaping for CBR data
- assessment of new products providing early packet discard and more sophisticated traffic shaping and policing functionality, and the impact of this on IP traffic

In the longer term we also expect to be involved in assessing the introduction of signalling and a wider rage of data rate service (e.g. ABR, VBR and UBR) as they become available in the public network. We also hope to use our expertise in IP implementations to examine broader issues of data transport over ATM and the possibility of dual ATM/IP stacks.

7 CONCLUSION

UCL-CS has been developing a wide range of expertise in the deployment and testing of high speed networks and the piloting of IP applications over them. It is clear, however, that though the operation of multimedia applications is feasible over today's Pan-European high speed infrastructure progress must be made in several areas before this can be offered as a service. Workstation interface vendors must provide better control over their traffic shaping; switch vendors more graceful overload protection; all must provide more usuable configuration management. They also need to work with operators to ensure the rapid deployment of signalling standards and techniques for end-to-end management visibility and control is multi-domain services are ever to be a reality.

8 Acknowledgements

The PREPARE, DRAGON and MICE projects are funded by the Commission of the European Union under contracts RACE 2004, RACE 2114 and ESPRIT 7606 respectively. Access to SuperJanet has been assisted by BT under the University Research Initiative, and was only possible thanks to the close co-operation of UKERNA, British Telecom and UCL-ISD.

References

[1] Ismail, N. Traffic Measurement Over SuperJanet SMDS Network 1994, http://www.cs.ucl.ac.uk/staff/nismail/JNT-WP1/wp1.html

[2] Lewis D, End to End Multimedia Services using the Telecommunication Managment Network, these proceedings

[3] Lewis D., Kirstein P., A Broadband Testbed for the Investigation of Multimedia Services and Teleservice Management, Proceedings of the 3rd international Conference on Broadband Islands, 1994

[4] Tag M., et al, Design and Implementation of a minimal service machine, these proceedings

[5] Kirstein P., Montasser-Kohsari G., Experiences with On-line Access to Chemical Journals, Proceedings INET 1995

[6] Kirstein P., Clayman S., Handley M., Sasse A., Recent Activities in the MICE Conferencing Project, Proceedings INET 1995

[7] Recommendation H.261: Video codec for audiovisual services a p x 64kbits/s, CCITT, 1990

[8] Huitema C., Turletti, T.: H.261 Software Codec and Workstation Videoconference, Proc. INET'92, Kobe, Japan, Aug. 1992, pp. 501-508

[9] VIC - video conference, on-line manual page, S. McCanne, Lawrence Berkley Laboratory, Berkeley, CA, 1994

[10] VAT - X11-based audio conferencing tool, on-line manual page, V. Jacobson, S. McCanne, Lawrence Berkley Laboratory, Berkeley, CA, 1992

[11] RACE R2049 CASSIOPEIA. Open Service Architectural framework for integrated service engineering Draft deliverable Nov. 1994, RACE Central Office Brussels.

[12] Armitage G. J., Adams K. M., "Packet reassembly during cell loss," IEEE Network, vol. 7, pp. 26–34, Sept. 1993.

[13] Bagwell C. E., Daigle J. N., "ATM cell delay and loss for best-effort TCP in the presence of isochronous traffic," in Proceedings of the Conference on Computer Communications (IEEE Infocom), (Toronto, Canada), June 1994.

[14] Cavanaugh J. D., Protocol overhead in IP/ATM networks, tech. rep., Minnesota Supercomputer Center, Aug. 1994, ftp://ftp.magic.net/pub/magic/overhead.ps