Some Multimedia Traffic Characterisation and Measurement Results

Jon Crowcroft, Steve Hailes, Mark Handley, Ajit Jena, David Lewis, Ian Wakeman

Department of Computer Science
University College London,Gower Street, London WC1E 6BT,United Kingdom.

ABSTRACT

Videoconferencing is becoming a commonplace facility. In the past this has been over dedicated lines, but increasingly, communities are experimenting with video and audio over today's packet switched digital networks without using very expensive quantities of bandwidth, in attempt to economise by using the existing communications infra-structure.

This paper is about traffic characterisation and measurements made of experimental multimedia traffic. This includes conventional data, shared windowing applications, uncompressed audio and slow scan video. This work is intended to guide future developments in communications and operating systems.

This work has been carried out under a number of programs, including the SuperJANET initiative and the RACE program projects CAR and PREPARE.

1. Introduction

Videoconferencing is becoming a commonplace facility. In the past this has been over dedicated lines, but increasingly, communities are experimenting with *multi-media*, including conventional and shared applications (either "conference aware", or using a "shared window" paradigm), video and audio over todays packet switched digital networks without using very expensive quantities of bandwidth.

This paper* is about traffic characterisation and measurements made of experimental multimedia (including audio and video) traffic on the CAR and PREPARE experimental networks, and the IP pilot service on the UK JANET Network [Crow91a, Crow91b] (see figure 1).

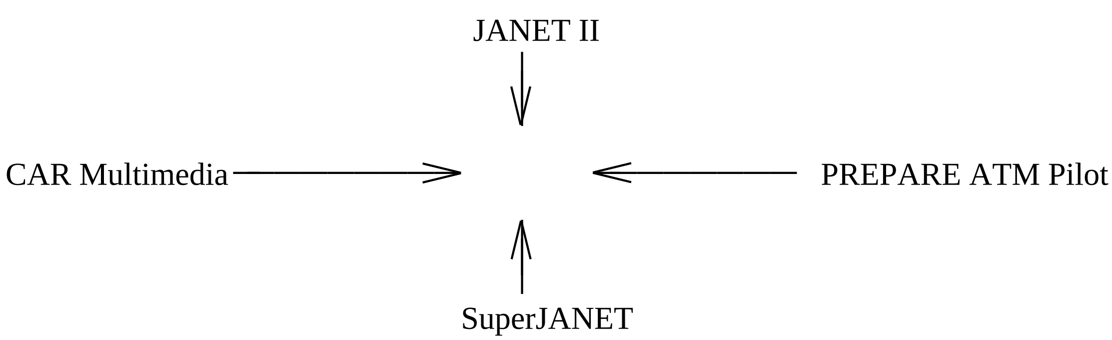

Figure 1. Inputs to multimedia work at UCL

In the next section we describe the PREPARE experimental architecture and the proposed SuperJANET [Wilb92a] architecture which it is envisaged will be very similar.

In the subsequent section, we summarise the characteristics of the various media and some of the architectural issues concerning multimedia conferencing. Following that, we present detailed results for measurements of shared files, shared windows, audio and video on the current JANET II Network.

Finally, we draw some conclusions about the future of multimedia services.

---

* Available as a research note, reference RN/92/29, Computer Science Dept, University College London

2. Network Support for Multimedia Conferencing

2.1 The PREPARE Network

One of the primary aims of the PREPARE project is to develop an integrated broadband communication testbed for the demonstration of the network management ideas addressed by the project. This will involve the interworking of LANs (Token Ring and Ethernet), B-ISPBXs, a DQDB MAN [IEEE86a] and an ATM WAN[Pryc91a] as depicted in figure 2. To demonstrate the realistic use of such a network multimedia teleconferencing has been selected as the application to run at the end user sites. This application is based on a multimedia conferencing system developed at UCL for the RACE CAR project combined with the transmission of packetised video and audio also being studied at UCL.

The network will connect several sites in Denmark with a possible further site in Germany. The full network is planned to be integrated and operating multimedia conferencing by early 1994.

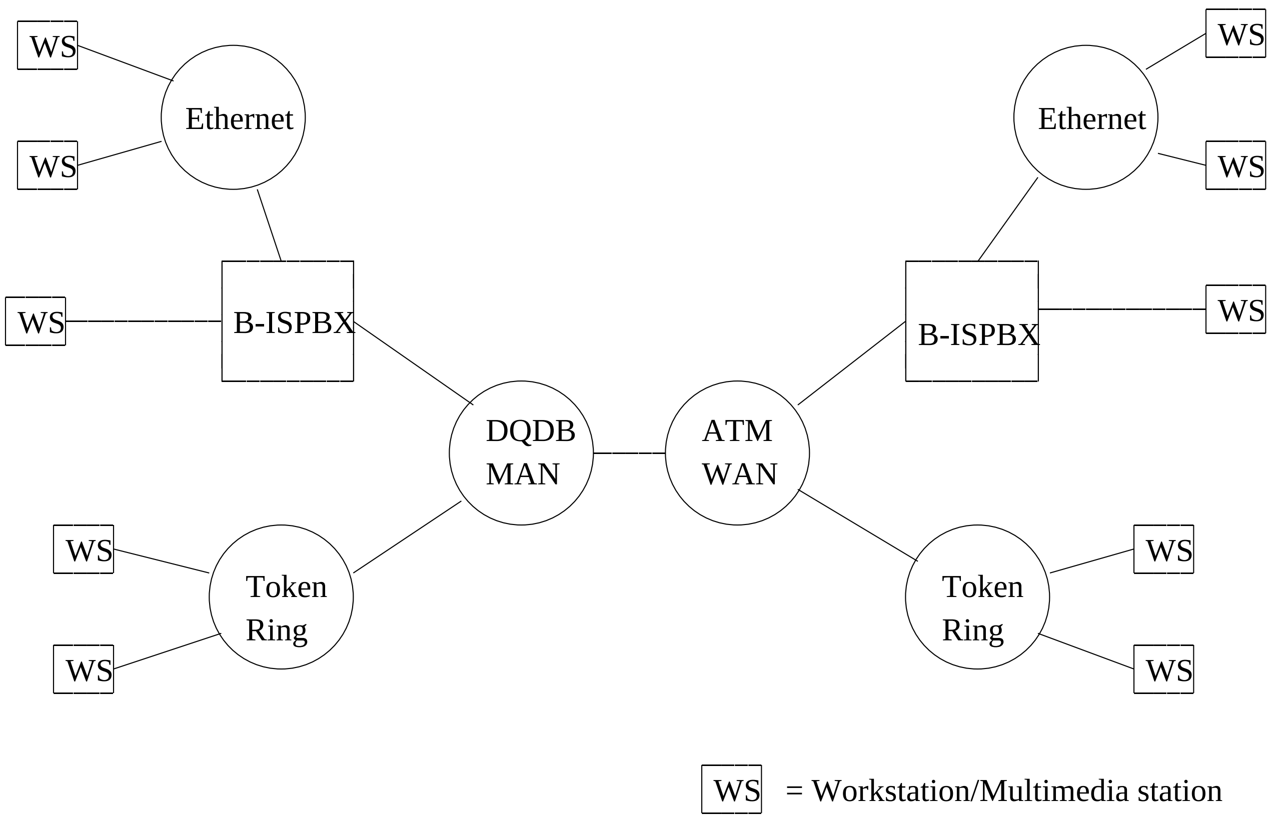

Figure 2. PREPARE network scenario

2.2 Super JANET

Plans are being laid for a new UK academic network to provide a basis for high bandwidth and novel applications including multimedia teaching and conferencing. The original JANET X.25 network was one of the first European national academic and research networks. Over the years it has grown in size and performance to the point where some 200 UK sites are connected, with 25% of sites connected at 2 Mbps.

With the growth of workstation based systems, it was decided to provide an IP overlay on the JANET network, i.e. IP packets were tunnelled over JANET. Approximately one third of the 2 Mps backbone now is used for this JIPS (Janet IP Service) traffic. Since it became available in October 1991, JIPS traffic has seen explosive growth.

Projecting the existing demands forward, data services alone will probably need 34 Mbps trunks before the end of the decade. In addition there will be potentially exponential growth in demand for multimedia services as increased processing power drastically reduces the specialised hardware required. Clearly a new network is required if this potential is to be exploited.

To this end, the SuperJanet Project was founded. Its aims are:

- from an early stage, to provide higher capacity conventional data communications on a service basis allowing existing use to grow naturally;
- to provide a testbed for users to develop new applications which only become feasible with significant bandwidth;
- to allow the communications and systems research community to use part of the network for experiments in high performance communication;
- to allow collaborating companies from European industry access to a large-scale testbed for novel communications and applications.

The current thinking is to use SDH (Synchronous Data Hierarchy) technology, with perhaps three rings covering roughly the south, middle and north of the UK. Each ring might operate at 622 Mb/s initially with cross connects between the rings. Connections to user sites would probably be at 155 Mb/s initially. Although in many ways ATM (Asynchronous Transfer Mode) standards would offer an ideally flexible way of sharing bandwidth, it seems likely that products and standards will be too immature for the initial network. However, an ATM pilot service is planned for early in 1993 with multimedia applications.

2.2.1 Multimedia Services and SuperJanet

Eventually isochronous multimedia services such as voice and video will be provided over ATM, but initially it is likely that 2 Mbps circuits will be set up between H.261 CODECs, or packetised and sent over the packet networks.

The first uses of these services will be to provide cost effective connections between the existing local video networks and studios such as the University of London Livenet, and similar networks at the Universities of Ulster and Wales (C5C). These sites already have video teaching rooms in place. The cost of suitable CODECs to connect these sites is currently about £50,000, which is not expensive when compared with the cost of equipping a studio.

3. Traffic Characterisation and Conference Architecture

3.1 Summary of Traffic QoS Characterisation

This section characterises the quality of service required in communicating the different media types employed in multimedia teleconferencing.

The four primary information flows involved in a multimedia teleconference are:

- video,
- audio,
- shared windows/workspace,
- file transfer.

Figure 3 represents the basic functional units required by a participant in a multimedia conference.

In the following subsections the four information flows are studied separately but it is assumed that they can be activated simultaneously by the participants.

3.1.1 Video Traffic Characteristics

The key item of equipment that needs to be examined when considering the transmission of digital video is the CODEC. The basic options currently available (based on equipment used at UCL) are:

- Uncompressed video:

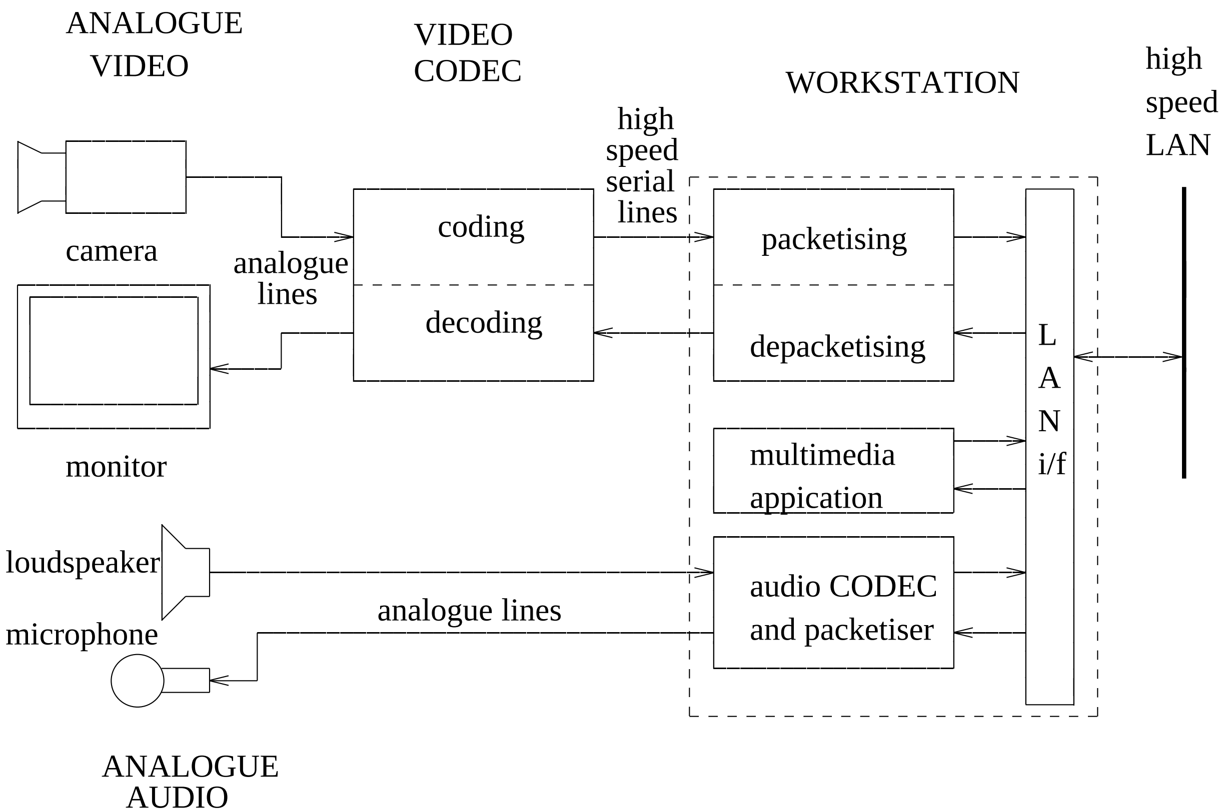

Figure 3. Multimedia Conference Station

This involves the straightforward digitisation of an analogue video signal with no intra-frame or inter-frame compression. The values presented below for this provide a base line for comparison of the other techniques.

- Slow scan:

  This is similar to the above in that no compression is undertaken but the values given are specifically for Sun Microsystems' VidoePix card. This is designed to operate in a single S-Bus slot of a SPARC workstation and is basically a frame grabber that is capable of producing a motion video signal at about 5 frames per second.

- DVI CODEC:

  Values for this are based on the operation of an experimental card again designed to operate over the S-Bus interface of a SPARC workstation. This CODEC utilises Intel's DVI (Digital Video Interactive) [Tink89a] compression technique.

- H.261 Compressed Video CODEC

  The values here refer to CODECs that use the CCITT H.261 standard compression technique. In particular we have examined a standalone CODEC manufactured by GEC/Plessey Communications that allows for connection to communication equipment via a 2Mbps line.

The following table attempts to present the requirements of these different video digitisation techniques (direct comparisons should not be made without careful reference to the accompanying notes).

Notes:

a. This figure is based on the transmission of video frames in Common Intermediate Format (CIF) at a rate of 25 frames per second.

technique	bandwidth	end-to-end delay	jitter	bit error rate
uncompressed video	**36.5 Mbps note a**	**200ms note d**	**note e**	**note f**
slow scan	**7.3 Mbps note b**	"	"	"
DVI CODEC	**1.8 Mbps**	"	"	"
H.261 CODEC	**0.384 - 1.9 Mbps note c**	"	"	"

TABLE 1. Digital video characteristics

b. This figure is based on a CIF frame rate of 5 per second.

c. The H.261 standard allows for bit rates of p x 64kbps where p is varied in integer values to allow achieve sufficient bandwidth for a particular application and to conform to bandwidth restrictions. The figures quoted are those found suitable for video conferencing applications

d. In general delay is application dependant and based on the perceptual tolerances of the user. In conversational interaction delays greater than the given figure start to become noticable though higher delays may still be quite tolerable. The actual transmision delay will consist of this figure minus the coding and decoding delays which for some techniques can be considerable.

   A further restriction on acceptable delay is "lip-sync", i.e. the synchronisation of the image of a participant talking and the accompanying sound of his or her voice. A lip-sync delay of more than half a frame period (approximately 20ms for full speed video) becomes disconcerting for the viewer. Usually this problem is solved by buffering the audio data at the receiver (due to the video’s longer coding delay the audio will tend to arrive first) and playing out using synchronising timestamps present in both data streams.

e. The amount of tolerable jitter is dependent on how much, if any, buffering is used in the receiver though obviously a bound is placed by the maximum tolerable delay.

f. This again may be very application dependent especially when the output is simply shown to a human since the eye is very tolerant to losses. Note, however, that some compression techniques result in a higher picture loss for a given transmission error rate. An important point to consider if specifying an error rate is that using retransmission to meet a specified error rate will produce delays that will probably render the video signal less acceptable than the original error. For this reason Forward Error Correcting codes are attractive for video transmission protocols.

3.1.2 Audio Traffic Characteristics

For telephone quality audio, 64kbps is adequate while 325kbps gives a high quality signal ideal for teleconferencing requirements. The acceptable delay is again around 200ms for the same reasons of conversational tolerances.

In order to avoid too much delay due to retransmissions a "best attempt" transport protocol may well be used for audio (as for video). However in general participants find that high quality sound accompanied by low quality video is more acceptable than any great loss in quality of the sound thus giving some guide to acceptable audio data loss rates.

3.1.3 Shared Windows/Workspace Traffic Characteristics

The solution used on the CAR project is Shared-X developed by DEC. This implements a shared-X bridge (see figure 4) which is a process that sits between an application and its X-server. This not only passes the user interface information to the console of the machine that it is being run on but also to any other selected machines. It can then be instructed to take user input actions from any one of these machines and pass it back to the application process. The corresponding output is replicated identically on each machine. In CAR a special nesting function is used to control the Shared-X bridge thus allowing the sharing of the

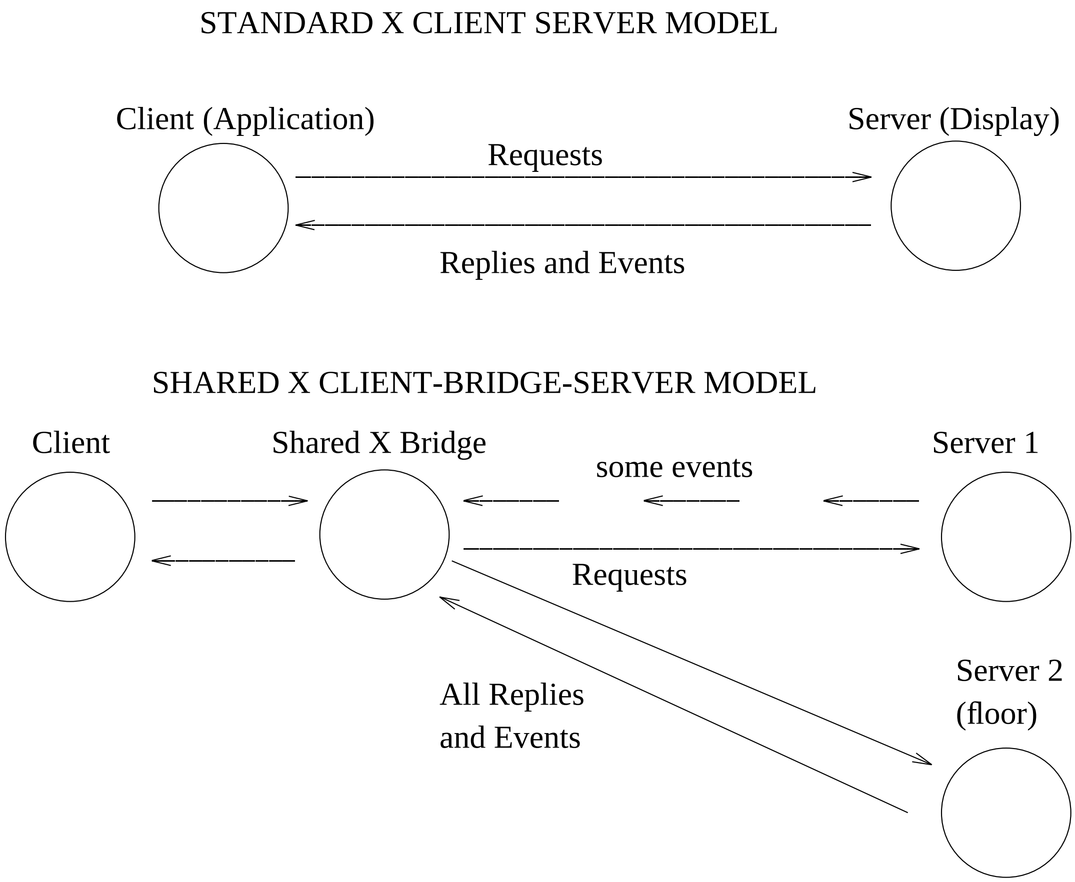

Figure 4. Shared-X model

application to be under the control of the conference control mechanisms.

The throughput during usage is very much dependent in the speed that the user works at and the amount of change to the display as a result of the users actions. See the later section on Shared-X measurements for some throughput values.

A simple drawing package as used in the measurements would suffice in many teleconferencing scenarios of a general nature where it would operate as a shared sketch pad for graphically expressing ideas quickly in much the same way as a Whiteboard does in a normal meeting. CAR was envisaged with more specific conference scenarios where the sharing of complex CAD packages was required.

The acceptable delay (typically 500ms) is again bounded by what the user find tolerable.

A reliable transport protocol (e.g. TCP) would be used so bit error rate is only important in its effect on the overall delay*.

3.1.4 File Transfer Traffic Characteristics
The most obvious solution is to use FTP over TCP/IP, although other alternatives exist such as NFS over UDP/IP.

FTP will take as much bandwidth as is available, bounded by the top transmission rate of the workstation (8Mbps for a SPARC workstation). Delay would be noncritical with this media and since a reliable transport protocol (e.g. TCP) would be employed neither is the bit error rate.

*** Note that TCP is conservative in offering load in the face of congestive loss**

3.2 Conference Network Architecture

This section discusses the various issues that have been considered in the design of the PREPARE teleconference network architecture for the various traffic types.

3.2.1 Video

For video traffic the primary consideration is how the video signals are distributed among participants.

The most basic option is for all participants to receive video signals from all other participants continuously as shown in figure 5.

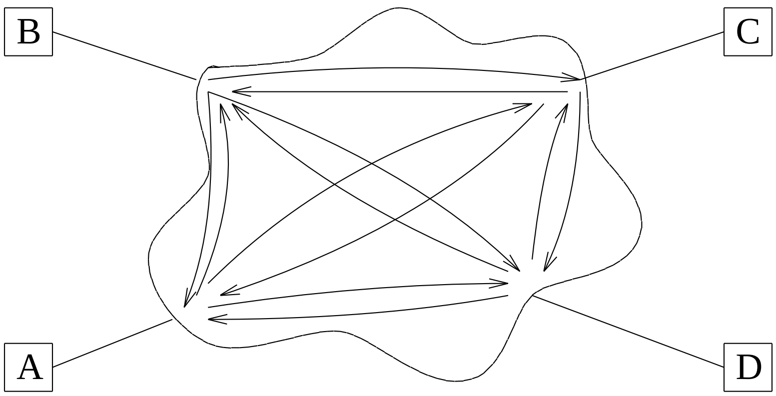

Figure 5. All-to-all transmission architecture

This is very unappealing for conferences with more than two participating sites due to the number of network channels required as well as the increase in the number of CODEC channels and monitor equipment required at each site.

Most conferences tend to operate with only one participant active at a time, usually under the control of a chairperson. A teleconferencing system can take advantage of this by requiring each participant to see only the currently active participant. This introduces the concept of floor control into the requirements for video teleconferencing. The CAR conferencing application operates a simple method of floor control where any participant may request and gain the floor at any time. It is this participant whose video image is viewed by all the other participants as well as having the current control over any shared applications that are involved in the conference. A small delay is required between one participant gaining the floor and a request from another being granted to ensure simultaneous requests do not result in a confusingly rapid switching of the floor.

The introduction of floor control next raises the issue of how the switching between various video traffic sources is accomplished. The two solutions to this considered in PREPARE were centralised switching and distributed switching.

In the centralised switching scheme one participant site would have the task of receiving video signals from all the other participants and then, depending on who has the floor, transmitting one signal back to all the other participants as in figure 6. In PREPARE this option held some attraction since video distribution under the existing CAR conferencing system was based on local area analogue video cabling being centred on a central analogue video switch and thus presented a least effort evolution for implementation in the wide area. However implementing in the wide area using digital transmission techniques required that the central switching site be in possession of a large number of expensive CODEC channels in addition to the analogue switch. Also this scheme introduces an extra transmission, coding and decoding delay to each signal as well as an additional degradation in picture quality due to its experiencing the coding/decoding process twice.

In the distributed switching scheme each participant transmits its video signal if and only if it is informed by the conference management software that it currently has the floor, see figure 7. This scheme is much

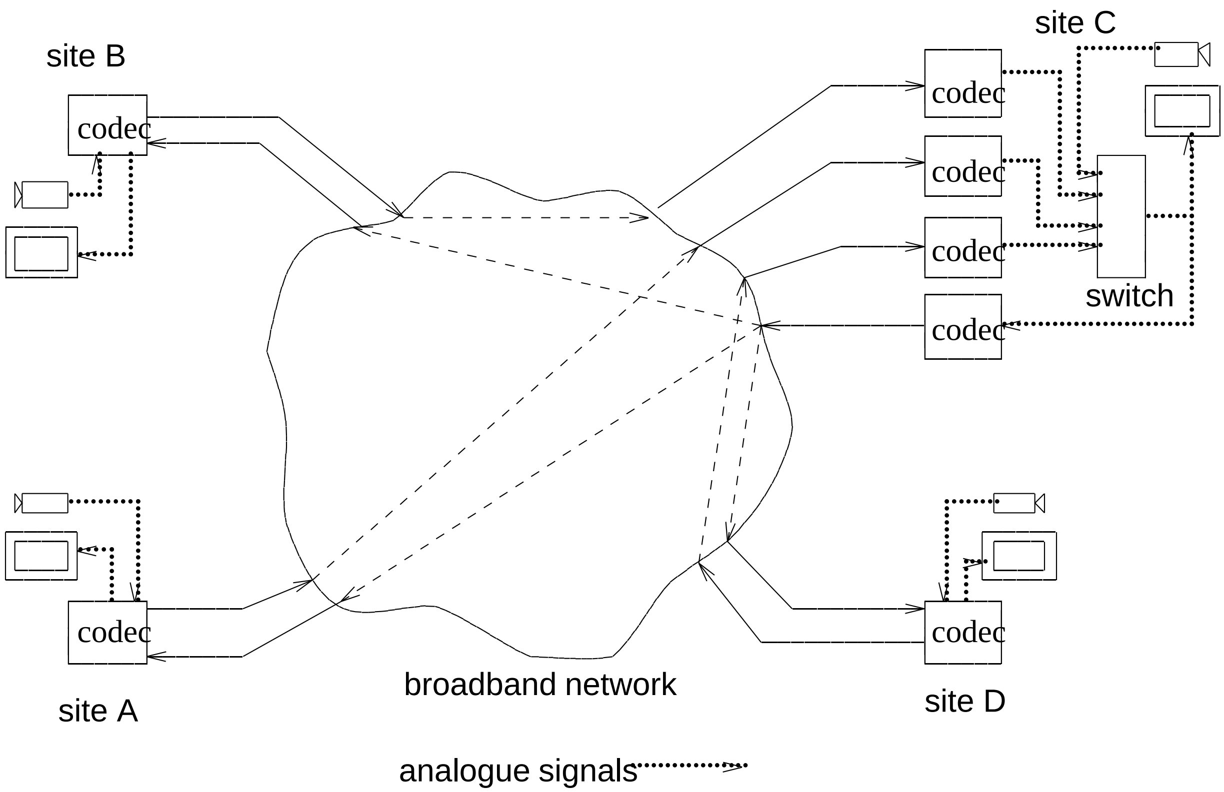

Figure 6. Centralised videoconferencing architecture

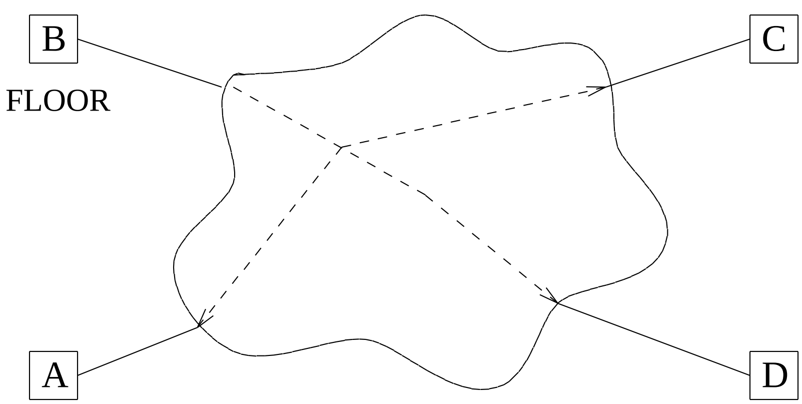

Figure 7. Distributed videoconferencing architecture

less demanding on bandwidth since only one channel per participant is required at any one time. Use of multicast capabilities will improve this further still and is an essential requirement for realistic teleconferencing over public networks. It avoids the delay and picture quality problems experienced by centralised switching as well as allowing the number of participants to be increased without the extra hardware provision required at the central site. A slight problem may arise in using such a scheme over a public wide area network in that switching between participants would probably involve some delay as a new set of calls are established. This switching delay, however, would probably be more acceptable to the users than the potential absolute delay and picture quality loss incurred by centralised switching.

A distributed switching scheme is currently envisaged for use in the PREPARE testbed.

3.2.2 Audio

Audio differs from video in that humans are able to make sense of several audio signals heard at once while they can only deal with one visual signal at a time. This introduces the requirement on the conferencing system to mix all the audio signals before presenting them to the user. Again this can be done in a centralised or distributed manner. It is however well within the capabilities of modern workstations to digitally mix several incoming PCM coded audio channels as demonstrated by the VAT audio conferencing application developed at Lawrence Berkeley Laboratories [Jaco92a]. This therefore favours distributed mixing for the same advantageous reasons as distributed video switching. In addition, mixing at end-systems facilitates the addition of side-conferencing capabilities to the system.

3.2.3 Shared Window/Workspace and File Transfer

These would use a multicast or multiple call scheme. Control of the shared applications would be decided by floor control mechanisms.

3.2.4 Protocol Stack

For PREPARE, the transport mechanism used is IP based. Above this we may use UDP [Post80a] or PVP [Cole81a] and NVP [Cohe81a] over ST [Topo90a]. However, this will not affect aspects like call control and bandwidth reservation which will be performed out of band by the network management functions. For file transfer and shared applications, TCP would probably be employed [Post80b].

3.2.5 Partial Teleconference

One feature that may be central to realistic teleconferencing is the ability to operate only a partial conference, with some or all partners able to use audio or shared applications but not video. This could be due to bandwidth restrictions or equipment availability. If these restrictions are temporary the conference should be able to dynamically introduce and retract certain conference facilites (primarily video) for individual participants, in order to provide the best service under the given conditions.

Another important feature for large teleconferences is the ability for two or more participants to have an aside separate from the main conference.

4. Measurement Configuration and Results

We have carried out a number of measurements of multimedia traffic in support of our traffic characterisation. Most of these were carried out over the JANET II IP Service. We must stress that we were **not** interested particularly in the effect of the traffic on the network. Nor were we seeking predominantly to measure the effect of the network control algorithms such as IP router’s queueing and drop algorithms, or X.25 flow control strategies. The measurements are primarily to elicit traffic from real users rather than simulations, or local emulations, so that we can gain an accurate picture of the *offered* traffic parameters. The experiments used the configurations in Figure 8.

Figure 8. JIPS/JANET-II Measurement Path

The JANET II protocol parameters were as follows:

1. JIPS = JANET IP Service, the pilot Shoestring network.
2. The X.25 level 3 window was 7; The packet size 1024 bytes.
3. Voice Traffic is typically generated with 180 byte samples (20 byte IP header + 8 byte UDP header)
4. The CISCO router output queue on the X.25 side is 40 packets at MTU size.

4.1 Video

4.1.1 Cambridge Video Pix Results

A simple videophone program (VP) was written to read the frame buffer as fast as possible on a SPARC workstationstation equipped with a SUN *VideoPix* card. The frames are read in XImage format, so that they can then be displayed on a remote machine running X windows. There are two approaches to remote display:

1. Simply use the X Protocol (given the size of the image, this is equivalent almost to a pure TCP connection, so the X overhead is negligible).
2. Stripe the XImage into a number of UDP packets, send these to a server, and reconstruct the XImage, and then locally send to the X Display Server.

The first approach was tried locally for basic performance measurements and to see how fast the SPARC workstation could generate frames. The most acceptable rate was achieved for 128x<128 monochrome pixel frames, and was 7 to 8 frames per second. This is equivalent to a data rate of around 132000 bps, or 132kbps. This should be achievable with little loss over the JIPS path. However, any loss will incur a TCP retransmit, which will hold up the display of the next frame. Since the UCL-Cambridge delays are of order 100 msec, this would limit us severely (if each frame required a retransmit, it would take at least a round trip time per frame to detect this, which would limit frame rate to about 5 per second at most).

It was decided to split the VP program into client and server side, and use the UDP striping protocol as a transport. No retransmission would be involved, but at a suitable frame rate, any loss of packets would now show as a missing stripe in an XImage. Since the server (receiver) does not zero the XImage between frames, this would show as a stripe staying the same between one frame and the next, which will often be the case anyway. Obviously, a random striping algorithm could avoid systematic correlation between natural stripes in a scene, and artifacts caused by loss.*

4.1.2 Video Emulation

Below is a table for a constant 64Kbps data stream, sent in different size packets with correspondingly different output rates. For each of the packet sizes, a total of 15Mbyte of data was transferred, using UDP, in 6 sets of 960kbyte and 1 set of 9600kbyte in order to spread the measurements over time.

---

* **Note horizontal versus vertical striping can be chosen based on human vision/perception balance being in favour of vertical detail.**

Pkt Size (bytes)	No. sent	Freq. (Hz)	No. lost	% loss
512	30720	16.00	480	1.56%
576	27306	14.22	231	0.85%
640	24576	12.80	145	0.59%
704	22340	11.64	111	0.50%
768	20480	10.67	82	0.40%
832	18904	9.85	27	0.14%
896	17553	9.14	17	0.10%
960	16381	8.53	11	0.07%
1024	15360	8.00	14	0.09%

It is clear that we can get a reasonable frame rate out of the JANET II, with minimal packet loss, by introducing a little buffering (but not very much, since average H.261 frame sizes in H.221 are of the order of 80 bytes anyway).

4.1.3 Transmission Delay

We found some anomalous delay results where there is a bi-modal distribution of packet delays - from around 100msecs and 1 second- at an offered steady rate of 64kbps, 1 kbyte packets, there are periodic (of order every 10 secs) spikes of delay at 1 sec, whilsts the vast majority of packets get through at approximately 100 msecs, and overall there is almost 0% loss. If we send at a very low data rate, the 1 second glitches happen more occasionally. We suspected 3 possible causes:

1. The call is dropping out despite our traffic - if the call setup is of order 1 sec, and since the CISCO has around a 40 packet queue between IP and X.25, it won't lose any packets in a second (we are sending at 8 pps), but we would see the extra delay for short bursts of packets that got stuck behind the call setup.
2. The CISCO is opening a second call but queueing all packets for a burst on that new call.
3. Some scheduled event internal to the router such as creating or processing a routing or management packet.

After consultation with the router vendor, it eventually transpired that the problem was caused by the latter cause, due to a misuse of the simple real time scheduler in the T20 router.

4.2 Audio

4.2.1 UCL to Cambridge:

Drop rates on JANET II links such as between UCL to Cambridge are low. We ran a couple of tests with data at 64kbps transmitted in 180 byte packets at a constant rate (45.51 Hz = 1/(21.97 ms)).

Test 1:

- Sent 10,000 packets in 220 sseconds,
- Loss rate = 0.08%,
- Statistics for 9992 packets:

Mean	22.027		
Variance	525.071	Coefficient of Variance	1.040
Average Deviation	11.705	Standard Deviation	22.914
Skew	7.558	Kurtosis	99.296

TABLE 2. Interarrival time stats (ms)

n	Data Above Mean + (n x stand. dev.)	Mean delay (ms)
1	5.154%	44
2	2.332%	67
3	0.911%	90

TABLE 3. Interarrival distribution

t	Min.	Mean	Max.
44	49	425	5909
67	69	940	9609
90	120	2405	23719

TABLE 4. Statistics for packet interarrival times for all interarrival times greater tha t (ms)

Test 2:

- Sent 60,000 packets in 22 mins,
- Loss rate = 0.82%,
- Statistics for 59509 packets:

Mean	22.211		
Variance	1567.838	Coefficient of Variance	1.783
Average Deviation	14.978	Standar Deviation	39.596
Skew	20.462	Kurtosis	798.300

TABLE 5. Interarrival time stats (ms)

n	Data above mean + (n x stand. dev.)	Mean delay (ms)
1	3.065%	61
2	1.244%	101
3	0.817%	140

TABLE 6. Interarrival Distribution

t	Min.	Mean	Max.
61	69	723	8420
101	109	1768	20070
140	149	2672	32030

TABLE 7. Statistics for packet interarrival times for all interarrival times greater tha t (ms)

Since the accepted delay for interactive audio is around 200ms, these results indicate that the JANET-II network can support limited audio at a reasonable quality level.

4.3 Shared Windows Appliations

4.3.1 Shared X Measurements

The following throughput measurements were taken to try and get a feel for the amount of traffic involved in operating a shared application using Shared-X and were performed on an Ethernet LAN. The applications tested were *idraw*, a simple drawing package, and *DesignView*, a CAD package.

The measurements are split into three phases:

- Replication:

  This is the initial phase when the local application is replicated on a remote machine using Shared-X.

- Local operation:

  The application is operated from the workstation from where it was originally started with the all event and display changes being replicated of the remote machine.

- Remote operation:

  The application is operated from the remote machine.

	App to Remote, bytes/s	Remote to App,bytes/s
replication	614.8	2784.3
local operation	6.79	0.94
remote operation	2.35	3.13

TABLE 8. Measurements for *idraw*

	App to Remote, bytes/s	Remote to App, bytes/s
replication	5196.1	4173.7
local operation	70.7	1.95
remote operation	1.3	13.7

TABLE 9. Measurements for *DesignView*

As would be expected the intial replication of the application requires very much higher bandwidth than that required by operation of the application which is basically driven by the human speed of operation. It is quite possible however for a single operator action to cause a large change in the appearance of an application thus providing a high bandwidth burst similar to that seen at replication.

4.3.2 X on VMS Measurements

The tests below were all carried out under VMS 5.3-1, with a MicroVAX 3600 running the clients, and a VAXStation 3100 running the server. The server wasn't doing anything else of any signicant load whilst the tests were running, but the client system is used for timesharing*.

In the tests we compare the communication costs when the server and the client reside on the same VAXStation 3100. The local transport is used where client and server reside on the same machine. On VMS, it is implemented using shared-memory queues, so there is no buffer copying between client and server, the packets are just linked and unlinked from different queues. In the table, D stands for Decnet, T for TCP/IP, L for local, C for CONS and O for OSI.

We performed three tests on these transports. The first used a modified *muncher* (an X intensive pattern displayer) which had the randomness removed, and involved the transmission of about 8.08 Megabytes of data from client to server. The second test used the synchronization time reported by *x11perf*. This is a measure of the time taken to do a getimage request which *x11perf* uses to ensure that the server hardware has actually performed the tasks requested of it. It is a crude measure of round-trip time. The final test was *x11perf's putimage500* test, intended as a raw throughput test.

*** These experiments were carried out by the JNT's ISO X SIG.**

Each test was performed 5 times, and the average results are shown below. Variability was very small (all figures were within 0.5% of the average) for all transports except TCP/IP, which exhibited results which ranged within 2% of the average figures on all tests. Figures are in seconds for *muncher* and *putimage*, and in milliseconds for *x11perf's* synch. Figures in brackets below are relative performance figures, with Decnet as 1.

Application/Transport: O	T	C	D	L
Muncher	79 (1.32)	77 (1.28)	71 (1.18)	60 (1.0)
X11perf	33.9 (1.83)	27.0 (1.46)	33.9 (1.83)	18.5 (1.0)
X11perf putimage500	1.67 (1.36) r (1.15)	1.42 (1.31)	1.61 (1.0)	1.23 (0.60)

TABLE 10. X over VMS measurements

4.4 File Transfer and Shared File Access

4.4.1 NFS results

The most straightforward way to measure NFS read or write performance is to time the copying of a file between server and client. A suitably large file (e.g. 1 Mbyte or larger) should be used to minimise the effect of local NFS caching.*

To measure the NFS retransmission rates and round-trip times the *nfsstat* command can be used on the client (Unix) system. This can be invoked with the -c option before and after a test copy of a large file, and will give the number of NFS requests and retries involved in the copy. (If it is used before the copy with the -cz option, the counts will be zeroed, making interpretation easier.) Thus, for example, the following output from *nfsstat -c* shows that there were 27 retransmissions from a total of 319 calls, and that all were due to timeouts on the requests.

Client rpc:

calls	badcalls	retrans	badxid	timeout	wait	newcred	timers
319	0	27	0	27	0	0	109

TABLE 11. NFS measurements

Invoking *nfsstat* with the *-m* option gives the average round-trip times for requests, on a per-mount basis. For example, the following shows that the current average round-trip time (srtt) for NFS reads is 322 ms, and for lookups (i.e. resolving filenames) 185 ms:

```
/usr/Ucl.cd-rom from 128.16.6.12:/cdrom (Addr 128.16.6.12)
Flags:    dynamic  read size=4096, write size=512,  count = 5
Lookups:  srtt=74  (185ms), dev=31 (155ms), cur=24 (480ms)
Reads:    srtt=129 (322ms), dev=4  (20ms),  cur=18 (360ms)
Writes:   srtt=0   (0ms),   dev=0  (0ms),   cur=0  (0ms)
All:      srtt=108 (270ms), dev=21 (105ms), cur=24 (480ms)
```

* **Results courtesy Bob Day, the JIPS Manager.**

The *dev* value is the standard deviation and the *cur* value is the current backed-off retransmission value [suitable reference here]

5. Discussion

In this paper, we have discussed multimedia conference traffic characteristics, and presented some results for network configurations including:

1. LANs
2. IP and X.25

At the moment, Interworking between these three technologies been achieved using multi-protocol routers. Under normal traffic conditions, these are sufficiently performant to carry a reasonable mixture of traffic.

As we can see from the figures presented (and was evidenced from human perception) 64kbps voice/audio across a 2Mbps IP on JANET X.25 service is perfectly feasible. There is occasional drop out, when competing for switch/CISCO buffers under extreme load. With compression (e.g. ADPCM) this should be very attractive. Also, video at slow frame rates without compression, or at reasonable frame rates with compression is perfectly feasible as well.

The future will bring us even more heterogeneous networks together with interworking problems, including:

1. Basic Rate ISDN
2. Primary Rate ISDN
3. Broadband ISDN
4. Metropolitan Area Networks

Interworking units for these newer technologies have yet to be devised, although we envisage a two fold evolution, which will continue for some time:

1. Multi-protocol routers continue to support data traffic between LAN and MAN islands over primary and broadband ISDN backbones, with limited multimedia being carried in slightly altered version of the data communications protocols.
2. ATM starts to move into the LAN and local exchange/customer premises network/subscriber loop, and true integrated services and networks will arrive.

The results presented here are extremely encouraging, since if an underengineered 2 Mbps mixed protocol network is capable of carrying a single multiway conference, then the envisaged SuperJANET at 155 to 620 Mbps should be capable of carrying between 75 and 300 simultaneous high quality conferences.

References.

Crow91a. J Crowcroft and R Bradshaw, , " ‘‘DoD Internet Protocols and JANET", ’’, *EARN/RARE IAB Networking Conference/Blois,* , pp. Elsevier (May 1991).

Crow91b. J Crowcroft, " ‘‘JANET IP Network Service"’’, *Connexions Magazine,* , pp. Published by Advanced Computing Environments (June 1991).

Wilb92a. S Wilbur, ‘‘The UK National Research Network: From JANET to SuperJanet’’, *Proc IEEE Eurocom 1992 Conference on Computer & Telecommunication Networks,* (May 1992).

IEEE86a. IEEEProject802, ‘‘Local and Metropolitan Area Network Standard: Overview, Interworking and Systems Management ’’, *IEEE 802.1, Draft D,* (August 1986).

Pryc91a. Martin de Prycker, *Asynchronous Transfer Mode, solution for broadband ISDN*, Ellis Horwood (1991). /.

Tink89a. Michael Tinker, ‘‘DVI Parallel Image Compression’’, *Communications of the ACM* **32,7**pp. 844-851, ACM (July 1989).

Jaco92a. Van Jacobson and Steven McManne, ‘‘VAT(1)’’, *Unix Manual Page*, (February 1992).

Post80a. J. Postel, ‘‘User Datagram Protocol’’, RFC 768, DARPA (August 1980).

Cole81a. Cole, E, ‘‘PVP - A Packet Video Protocol’’, *USC-ISI Tech Report*, , USC-ISI (Aug 1981).

Cohe81a. Cohen, D, ‘‘Network Voice Protocol NVP-II’’, *USC-ISI report*, , USC-ISI (April 1981).

Topo90a. Topolicic, C, ‘‘Experimental Internet Stream Protocol Version II’’, *RFC 1190*, , SRI-NIC (Oct 1990).

Post80b. J. Postel, ‘‘Transmission Control Protocol’’, RFC 793, DARPA (September 1980).

CONTENTS

LIST OF FIGURES

LIST OF TABLES