TRANSITION TO HIGH SPEED NETWORKS
SuperJANET EXPERIENCE

Kicheon Kim, Steven Simpson, David Hutchison, Andrew Scott

Computing Department
Lancaster University
Lancaster LA1 4YR UK

{kk, ss, dh, acs}@comp.lancs.ac.uk

Abstract - For the time being, trials to establish the Information Superhighway are booming. In Britain, JANET has provided wide-area computer communication, and has recently been upgraded to SuperJANET, increasing the throughput by a factor of five to 10 Mb/s, with some sites having PDH access at n x 34 Mb/s. In this paper, the technological changes seen from a user perspective are addressed. A multimedia communication-based distance learning project on SuperJANET is introduced and the network performance measurements for this project are presented. These measurements suggest the employment of reservation protocol and packet scheduling. We also provide a mechanism for on-the-fly playback of continuous media.

1. Introduction

Recent changes in the technology of the British academic network, SuperJANET, have had implications for multimedia communications, particularly continuous media (CM), and in our case, for a multimedia project called MUNSTER. We start by tracing the history of the national academic network.

1.1 JANET

JANET (Joint Academic Network) was started in the early 1970's to allow computer communication between academic establishments across Britain. Initially providing 1 Mb/s for X.25 traffic, its usage grew steadily in the 1980's and was upgraded to 2 Mb/s.

The amount of IP traffic over the X.25-based JANET rapidly grew, and by the end of 1992, the amount of IP-carrying X.25 traffic had reached 90% of the total traffic. This justified a introduction of native IP networks into the JANET backbone network, and further encouraged JANET users to switch to IP, with services from the JANET administrators, UKERNA, to assist in the transition for X.25-reliant services. The X.25 backbone has been downgraded, and is intended to be withdrawn in the near future. [1]

1.2 SuperJANET

The main contributing factor to the increase of IP traffic on JANET has been the widespread deployment of local-area networks. Switched Multimegabit Data Service (SMDS) technology was developed to meet the increased demand for the high-performance LAN interconnection services. [2] The successor to JANET, SuperJANET, adopted SMDS technology to accommodate increased IP traffic on JANET.

The SuperJANET SMDS backbone network is built on a 34 Mb/s transmission link, and is expected to be upgraded to 155 Mb/s in the second phase of development. A range of protocols can be routed over this network, including IP as specified in RFC 1209. [3] SuperJANET has been available to 60 sites since April 1994.

Another component of SuperJANET is the PDH backbone, which is available to 16 sites, and will be replaced with SDH technology in the second phase. The main purpose of this network is to support the development of an ATM network. [4]

1.3 Applications

There is no doubt that the increased bandwidth of SuperJANET is used by multimedia applications. There is already multimedia traffic, mostly generated by World Wide Web (WWW) applications. In addition, distance-learning applications based on multimedia are under

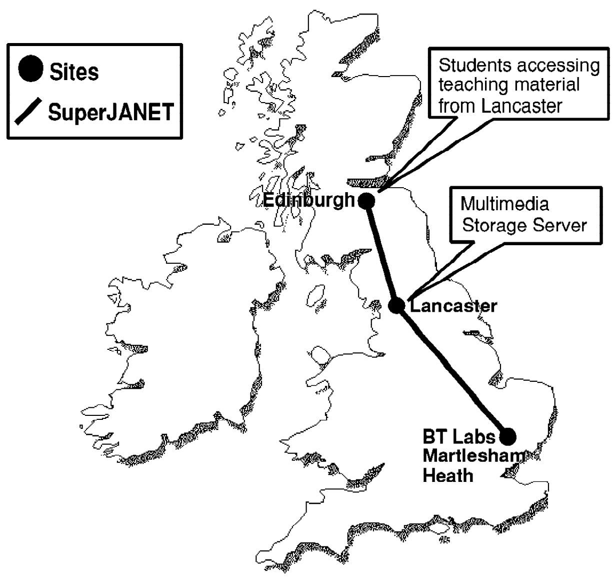

Fig 1: MUNSTER Network Configuration

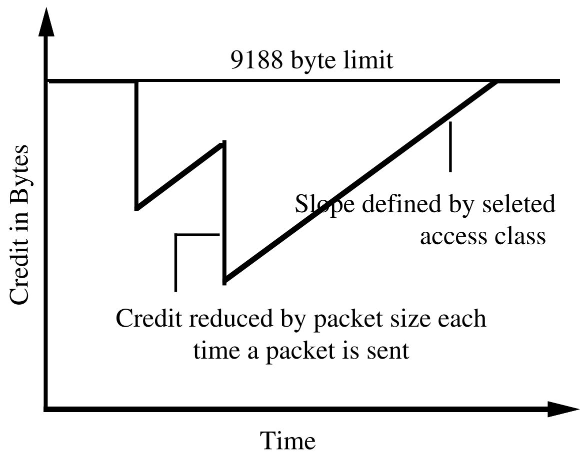

Fig 2: The Access Class Credit Manager

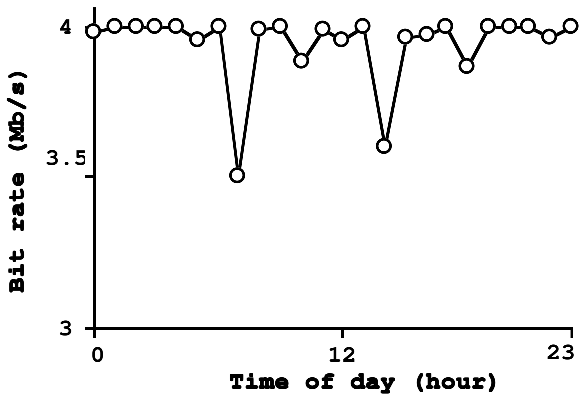

Fig 3: Achieved bit rate throughout the day

development, on SuperJANET. As an example, Fig 1 shows the sites involved in a distance-learning project where Lancaster University is engaged.

This project is called MUNSTER, and is funded by JISC and BT to investigate long-distance learning with multimedia. Using SuperJANET as the long-distance communication medium, multimedia teaching material can be stored at one site, while students can access it from another. Compressed video and audio is stored on a server at Lancaster University, while students at Edinburgh University can playback the information on their terminals.

The requirements of the MUNSTER project (to transfer real-time data across SuperJANET) partly motivate this research.

2. Traffic Analysis

2.1 Achieved Throughput Measurements

At Lancaster University, we wanted to find out if it was possible to use our SuperJANET SMDS link to send continuous media (MPEG video and sound) between a multimedia file server and a client.

The traffic to the SuperJANET SMDS network is controlled by a credit manager in the router and limited to the sustained rate at 10 Mb/s. There are 4, 10, 16 and 25 Mb/s access classes available on the 34 Mb/s SuperJANET network. Fig 2 shows the mechanism of the credit manager. The credit manager has a credit counter which indicates the maximum amount of data transmitted to the network. When a packet is transmitted, the packet length is subtracted from the credit counter. According to the rate of the access class, the credit counter is increased. If the packet size is larger that the credit counter, the packet is discarded. When there occurs instant congestion on the router to the SMDS network, packet loss is expected.

A simple server (not multimedia) was established at Edinburgh, and a machine at Lancaster executed a client program. The client issues a request to the server to produce n frames of size k KB at r Mb/s, and transmit them to the client. As the frames arrive, the client time-stamps them to produce raw data which can then be analysed to obtain statistics about the bit rate and packet loss rate of the connection. Frames that arrived after a suitable time-out, or did not arrive, were considered lost. Twenty-four measurements ($n = 1000$, $k = 8$, $r = 4$) were made throughout one day, at hourly intervals, starting at 11:00am. From each measurement, bit and achieved bit rates were evaluated, to produce Fig 3.

Fig 3 shows the achieved bit rate varying with the time of day. The measurements used UDP, because TCP implementations would be unable to make full use of the bandwidth available due to restricted TCP window size. Also, TCP provides facilities that are unnecessary for transmission of video, such as re-transmission, and so are undesirable if retained but not used. The measurements were obtained from a connection between Lancaster and Edinburgh 250 km apart.

4 Mb/s was chosen to match an earlier experiment [5] in which the bit rate was not specified directly (it was inferred from other arguments), and did not exceed 4.31 Mb/s. The error rate caused by packet loss depends on the time of day. Compared to the traffic level, when the error rate is high, the traffic level is also high. This is because the bandwidth allowed to the router to SMDS is limited to 10 Mb/s, and packets above the bandwidth limit are discarded. Local access networks may also contribute directly to packet loss, but it is not clear to what degree, Jitter caused by the local networks at the source could cause occasional bursts which exceed the rate allowed by the router, causing indirect packet loss. The error rate during working hours is higher than at other times. This also shows the working patterns in Britain. There is also a peak in loss rate at 7:00am, whose origin is not clear, but may be due to network maintenance.

So, if we wish to transmit continuous media, like a compressed video stream, through this network, the QoS of the connection cannot be guaranteed, especially during working hours, although the network is not congested all the time. This also means that for the multimedia communication we need a new protocol that is not influenced by the traffic level. This kind of reservation protocol is now under study. [6]

2.2 Buffer Size Measurements

Although the server attempts to transmit packets at regular intervals (every 15.625 ms), the packets are received with a wider distribution of inter-packet delays. In part, this is due to the limited capabilities of the server's scheduler, which controls the timing of packet submission (in fact, the server transmits almost entirely at 10 ms and 20 ms, with a mean of 15.625). Network behaviour further distributes the inter-packet delays to be mainly within the range 0-35 ms.

In order to study the effects that this distribution has on the receiver and intermediate nodes, we have simulated packet buffers into which packets are injected according to arrival times obtained from the throughput measurements. Packets are consumed from the buffer at the rate that they were originally transmitted, i.e. every 15.625 ms.

Measurements at two times have been chosen for use in the simulations. The 2:00am measurement was chosen for being a 'quiet' time, i.e. there was no packet loss, while 2:00pm was chosen for being a 'busy' time, i.e. there was a large packet loss of 9.8%. Figs 4 and 6 show the distributions of inter-packet arrivals at these times.

Fig 8 shows an example of the changing state of a 5-packet buffer during a simulation. The trough in the 2nd second demonstrates buffer starvation, or underflow. A

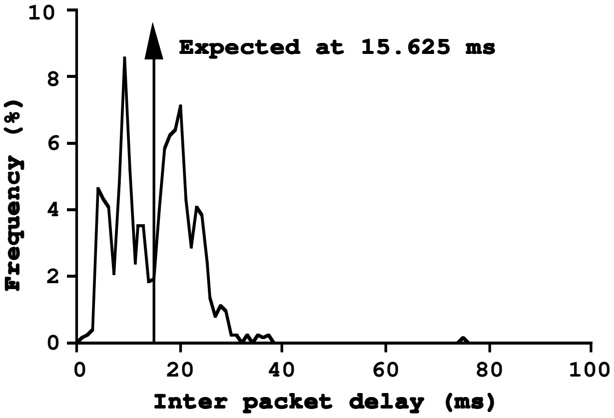

Fig 4: Inter packet delay distribution at receiver (2am)

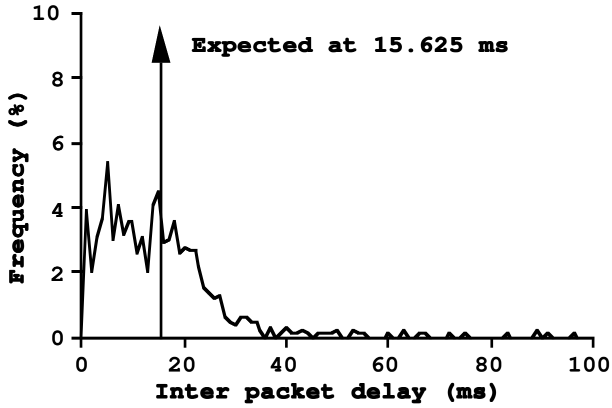

Fig 6: Inter packet delay distribution at receiver (2pm)

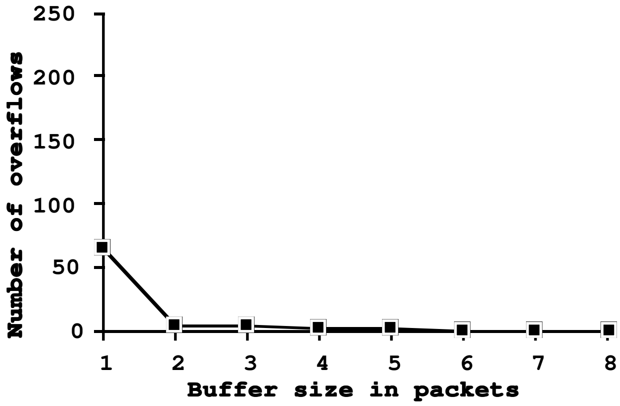

Fig 5: Overflow frequency vs. buffer size (2am)

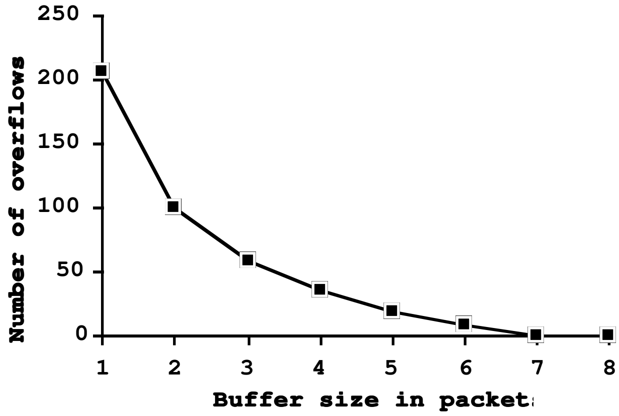

Fig 7: Overflow frequency vs. buffer size (2pm)

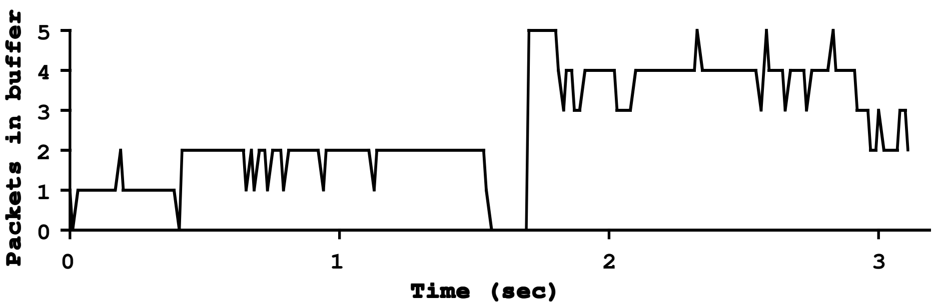

Fig 8: Packets in the buffer at the receiver

peak immediately follows the underflow, and any packet that arrives during that time will be lost as the buffer overflows.

Overflows and underflows are counted, and overflows are displayed against buffer size (in packets) for the 2:00am and 2:00pm measurements in figs 5 and 7 respectively. The difference between underflows and overflows was found to be constant for each particular set of arrival times, as each overflow causes an underflow later on, and this adds to a constant number of underflows that are present regardless of buffer size. (For 2:00am and 2:00pm, that constant is 6 and 106 respectively.)

The figures show that during quiet periods, small buffers are adequate, while at other times, larger buffers are needed, or packets will overflow and be lost. This loss is due to the jitter of the packets as they arrive, and this can occur anywhere in the transmission. Any buffering after such spreading out is more likely to experience over- and under-flows. On the other hand, buffering throughout the link may help to regulate inter-packet arrival times, before the effects become too great.

2.3 CM-Web Integration

The advent of World Wide Web (WWW) servers and browsers has proven to be the major cause of the increase in bandwidth requirements. [7] Web applications enable users to retrieve a limited form of multimedia data including hypertext and still images. The Web browsers can support sound and moving pictures through external 'player' software, but do not support direct playback of these continuous media. Audio and video objects are treated as files, and as such they must be transferred completely before playback can begin. This means users often wait a long time while whole files are transferred. If the continuous media can be played while they are being transferred, it can save time and storage at the player. We call this scheme *on-the-fly* playback.

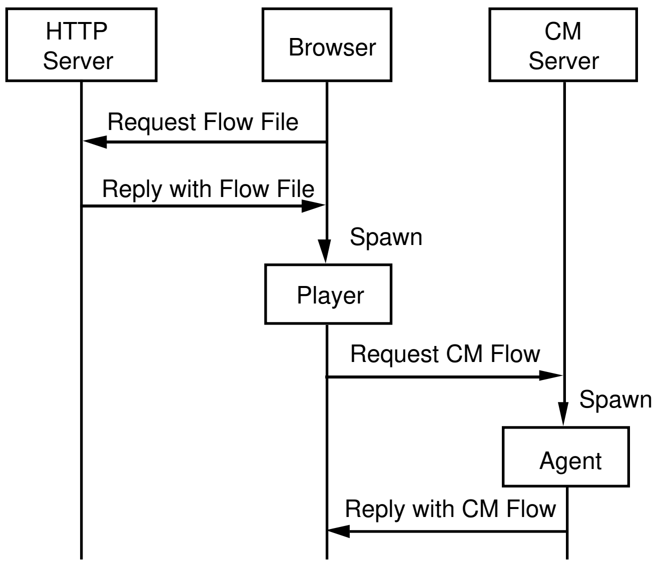

Fig 9: Procedure to retrieve CM data on the Web

This scheme is not feasible on a large scale over the current Internet because, at the moment, the network cannot support the reservation of the network resources in terms of bandwidth and guaranteed QoS. Under the restricted condition where there is abundant bandwidth available, this on-the-fly scheme can work. Usually a local area network environment provides enough bandwidth for a number of small screen MPEG-1 streams and the measurements in the former section suggest that the SuperJANET SMDS service could also handle a restricted number of CM streams at a lower QoS due to the less predictable characteristics of the WAN service.

However, the Web server implementations do not support the protocols for the on-the-fly scheme and the browser implementations do not have capability to play back CM streams on-the-fly, although the HTTP protocol can be extended to accommodate new media types, when server support is implemented. [8] So, basically the Web provides a good environment to search and retrieve data around the world, but current server implementations do not support the continuous media. Here we suggest intermediate solutions to integrate the Web and the continuous media using an external on-the-fly player and dedicated CM server. There are a number of implementations, at Lancaster and elsewhere, of CM-Web systems based on MBONE applications, but these have no way of reserving network bandwidth and offering a guaranteed QoS.

Fig 9 shows the procedure to retrieve *a flow file* and playback CM stream. The flow file specifies the resource identifier and the flow spec of the CM flow. When the browser retrieves the flow file from the HTTP server, it spawns the external player. The player requests the CM server for the CM flow with the flow file specification. The CM server spawns an agent dedicated to this connection and the agent generates the CM flow according to the requested flow spec. The player which has been waiting works on the flow on-the fly. Usually, CM servers are specially designed for CM data. So separation between HTTP server and CM server has benefits.

Our on-the-fly scheme works well on the local area network where there is abundant bandwidth and across SuperJANET occasional picture freezing is unavoidable. Ultimately, the request to the CM server will be supplemented with the resource reservation protocol. When the reservation protocol is employed in network equipment such as routers, the playback of CM flow across networks would be popular, require excessive bandwidth, and eventually cause the employment of new technology such as ATM.

5. Discussion

Usually, when the design of a high speed network is considered, the technology for the network comes into question. Although the advantages of ATM technology seems clear, it is unlikely that everybody who uses computers will want ATM immediately. The ATM technology usually means higher bandwidth (which can easily exceed 100 Mb/s), faster switching, and full motion video without compression (which decreases the transmission delay). In order to access the ATM network however, new cabling, or at least pieces of expensive network equipment, are required.

Usually, network establishment must be done within a limited budget, but the number of computer users that want network connections is increasing over and over again, although their individual bandwidth requirements are not increasing as rapidly. It may be more cost-effective to expand existing conventional packet networks rather than increasing bandwidth through use of ATM.

In the case of SuperJANET, the number of users is still increasing faster than the increase in per-user bandwidth. So the expansion of current packet networks is unavoidable. In addition, the current bandwidth of the individual packet network user is increasing due to the use of new applications, such as Web browsers. A primary factor influencing the change of the network is the type and behaviour of user traffic, as we saw in the case of JANET in 1992.

From traffic analysis, we can infer that a reservation protocol is required. The study of reservation protocols on the conventional packet network is making it feasible to transmit continuous media over these networks. This will also make multimedia communication more widely available without excessive cost.

The ATM network is expected to be widely deployed when novel applications, such as advanced distance-learning and remote sensor imaging, experience greater demand. So, the use of high-speed ATM technology for end users will be limited to specific research areas, for the time being.

In addition, computer communication networks in the future are hardly expected to be unified into ATM networks. There is a huge installation base of packet networks, especially under local area network environments including ethernet, token ring and FDDI subnetworks. Even mobile communications networks are expanding rapidly and should be provided with the multimedia communications capability, but the use of mobile ATM technology is not expected to be employed because of inefficiency and lack of bandwidth to wireless user. So, a heterogeneous network environment is unavoidable and higher layer network abstractions are needed still for global accessibility.

An IP network on top of the heterogeneous network environment is called a logical IP network and a good candidate for the abstracted network. The current IP network has huge number of subnetworks and uses, but does not provide multimedia communication services such as guaranteed QoS. The proposed IPv6 has a packet format which can specify traffic type and enable intermediate nodes to schedule packets according to their traffic types. [9] RSVP is proposed for the signalling protocol on the IP network. Further because different subnetworks use different QoS parameters, a QoS mapping framework on the heterogeneous network environment has been proposed. [10]

Acknowledgements

Kicheon Kim from Korea is supported by the British Council. Steven Simpson is a research student supported by an EPSRC-CASE studentship award (ref. 9456391X) with sponsoring from BT Laboratories. This research is carried out in the context of the JISC-BT funded MUNSTER project. The authors would also like to thank their colleagues on the MUNSTER project: Barry Forde, Phillip Lougher and Doug Shepherd.

References

[1] Network News, Number 37, November 1992, Joint Network Team

[2] Bellcore Technical Advisory, TA-TSY-000772, "Generic System Requirement in Support of SMDS," Issue 3, October 1989.

[3] D. Piscitello and J. Lawrence, “The Transmission of IP Datagrams over the SMDS Service,” RFC 1209, Bell Communications Research, March 1991

[4] Network News, Number 40, November 1993, Joint Network Team

[5] S. Simpson and D. Hutchison, “Experiments with SuperJANET/SMDS,” Second Communication Networks Symposium, Manchester, July 1995.

[6] L. Zhang, R. Braden, D. Estrin, S. Herzog, and S. Jamin, “Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification,” Internet Draft, July 1994

[7] Tim Berners-Lee, Robert Cailliau, Ari Luotonen, Henrik Frystyk Nielsen, and Arthur Secret, “The World-Wide Web,” Communications of the ACM, Volume 37, Number 8, August 1994

[8] Tim Berners-Lee, Roy T. Fielding, and Henrik F. Nielsen, “Hypertext Transfer Protocol - HTTP/1.0,” Internet Draft, March 1995

[9] S. Deering and R. Hinden, “Internet Protocol, Version 6 (IPv6) Specification,” Internet Draft, March 1995

[10] Andrew Campbell, Geoff Coulson, and David Hutchison, “A Quality of Service Architecture,” ACM Computer Communications Review, April 1994