

Working Paper Series  
ISSN 1170-487X

## **Measuring ATM Traffic**

### **Final Report for New Zealand Telecom**

**by Prof J Cleary, Prof I Graham,  
Dr M Pearson and Dr T McGregor**

Working Paper 98/14  
October 1998

© 1998 Prof J Cleary, Prof I Graham,  
Dr M Pearson & Dr T McGregor  
Department of Computer Science  
The University of Waikato  
Private Bag 3105  
Hamilton, New Zealand

# **Measuring ATM Traffic**

## **Final Report for New Zealand Telecom**

Prof. J. Cleary,

Prof. I. Graham,

Dr. M. Pearson,

Dr. T. M<sup>c</sup>Gregor

Department of Computer Science,

University of Waikato.

October 22, 1998

### **Abstract**

This report describes the development of a low-cost ATM monitoring system, hosted by a standard PC. The monitor can be used remotely returning information on ATM traffic flows to a central site. The monitor is interfaced to a GPS timing receiver, which provides an absolute time accuracy of better than 1  $\mu$ sec. By monitoring the same traffic flow at different points in a network it is possible to measure cell delay and delay variation in real time, and with existing traffic. The monitoring system characterises cells by a CRC calculated over the cell payload, thus special measurement cells are not required. Delays in both local area and wide-area networks have been measured using this system. It is possible to measure delay in a network that is not end-to-end ATM, as long as some cells remain identical at the entry and exit points. Examples are given of traffic and delay measurements in both wide and local area network systems, including delays measured over the Internet from Canada to New Zealand.

## Table of Contents

1	Introduction.....	1
2	Initial Measurement System.....	2
3	Dag 1 .....	3
4	CRC Generation.....	5
5	Synchronisation .....	6
6	Remote Monitor.....	8
7	Dag 2 .....	8
8	Traffic Generation .....	9
9	Results.....	10
10	Conclusions.....	15
11	References.....	15

## LIST OF FIGURES

Figure 1 ATM Traffic Measurement System.....	3
Figure 2 Block Diagram of the Dag 1 Board .....	4
Figure 3 DAG, outline Xilinx internal logic .....	4
Figure 4 Photo of the Dag 1 Board.....	5
Figure 5 Example set-up using two measurement devices.....	7
Figure 6 Synchronisation using a common oscillator.....	7
Figure 7 NIC Synchronisation using GPS.....	7
Figure 8 Clock Difference between two NIC cards over time .....	8
Figure 9 Block Diagram for the Dag 2.....	9
Figure 10 Traffic measured on a NFS server to workstation link (40 sec bins).....	10
Figure 11 Traffic measured on a NFS server to workstation link (1 sec bins).....	11
Figure 12 Cell delay through a lightly loaded switch.....	11
Figure 13 Cell delay under heavier loading .....	12
Figure 14 Experimental setup for Hamilton -> Wellington delay measurements .....	13
Figure 15 Histogram of Hamilton to Wellington cell delays over OPERA.....	13
Figure 16 Histogram of delays over the Internet from Calgary to Waikato .....	14

# Measuring ATM Traffic

## Final Report

Prof. J. Cleary, Prof. I. Graham, Dr. M. Pearson, Dr. T. M<sup>c</sup>Gregor

Department of Computer Science,

University of Waikato.

October 22, 1998

This is the final report for the project titled "Measuring ATM Traffic" being carried out by a research group in the department of Computer Science at the University of Waikato New Zealand for New Zealand Telecom. In the original proposal the following objectives were defined:

1. Measuring traffic over ATM networks;
2. Timing and synchronising of measurements at geographically remote sites.
3. Generating synthetic traffic over ATM networks;

### 1 Introduction

ATM (Asynchronous Transfer Mode) is becoming widely accepted as one of the technologies used in the implementation and expansion of large telecommunications networks. It is very attractive because of its ability to utilise statistical multiplexing techniques, provide good expansion paths and its ability to carry a wide variety of traffic types concurrently. However, if traffic on these networks is not managed carefully side effects of these very attractions have the ability to severely impact performance. For example statistical techniques assume that the traffic is Poisson distributed. However if the traffic is bursty in nature then the delays and cell loss on a link can be much higher than for a smoother Poisson Distribution.

Also the ability of these networks to carry different traffic types concurrently has the potential to adversely affect the performance of networks because of the differences in their characteristics. These can be categorised into two main groups: real time and non-real time traffic (Habib and Saadawi, 1994). Real time traffic, such as voice and video, has strict time delay and delay variation requirements that must be met, but can tolerate a small cell loss rate. Non real time traffic, on the other hand, can tolerate delay but cannot tolerate cell loss. These differing and contradictory requirements of different traffic types combined with the estimated sizes of future ATM networks promise to make the task of network design and management an extremely difficult one. It will be important for network designers and engineers to be able to assess and predict traffic performance of current and future networks to ensure that they are used and implemented in the most economic way. Bad decisions may lead to networks that are excessively expensive and/or perform badly, adversely affecting users of the network.

The effective design and maintenance of these ATM networks will require the development of efficient and accurate network planning tools, which in turn require accurate models of ATM traffic. To construct these models usually involves collecting and analysing traffic measured on actual networks. To date most traffic measurements made to characterise different traffic sources have been obtained using existing LAN and WAN networks. While this data can provide great insight into the characterisation of a particular traffic source some care needs to be taken, as results obtained by measuring traffic on real ATM networks may vary greatly. For example, differences in available bandwidths and latencies on ATM networks are different from existing technologies, which may alter the way people use them and consequently alter

parameters of the model. Hence it is becoming increasingly important to construct models based on traffic collected from actual ATM networks. The types of measurements that can be used to help analyse the behaviour of ATM networks include:

1. **Throughput** which records the number of cells that pass the measurement point in a specified time so that a link's utilisation can be determined.
2. **Cell arrival time** which records the time that each cell passes the measurement device and can be useful for characterising traffic sources. To perform this type of measurement an accurate clock is required so that a timestamp can be generated for each cell that is measured.
3. **Cell Delay** which records the length of time that it takes cells to travel between two points in a network. This requires two measurement points where data is collected at each, and is later merged so the delay times can be calculated. This requires a mechanism to allow cells to be matched when the two sets of measurements are merged. Also a mechanism is required to synchronise the clocks which can be difficult when the measurement points are geographically distant.
4. **Delay variation** which records the variation in the length of times that it takes for cells to travel between two points in a network. This is important in real time applications such as video conferencing where it is important that cells arrive at regular intervals
5. **Protocol analysis** where the full content of each of the cell is recorded so that it can be used to perform analysis on the higher level protocols such as TCP/IP.

To date ATM test equipment available on the market is either very expensive or incapable of collecting the measurement data necessary to develop accurate traffic models. Previous work at the University of Waikato developed a measurement system consisting of a PC and a modified ATM Network Interface Card (NIC) (Graham, 1996). This system allows traffic on either a 25.6 Mbps or a 155Mbps ATM link to be measured and is described in the next section. Section 3 describes the design of a board that has extended the range of line rates that can be measured by the system. Section 4 then details the use of GPS systems to allow accurate timings to be recorded for sets of measurements being performed at geographically remote sites concurrently.

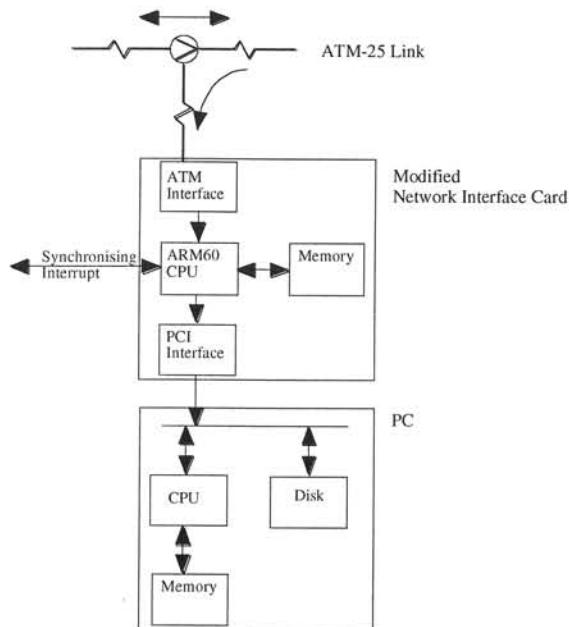
## 2 Initial Measurement System

Figure 1 shows the layout of the initial measurement system that is based on an ATM Ltd Virata Link network interface card (NIC)<sup>1</sup>. This card interfaces ATM25 (25 Mbps) over unscreened twisted pair (UTP) cable to the PCI bus of an IBM-compatible PC. The card has been reprogrammed by us to record accurate cell arrival times along with either the entire cell, or the cell header and an optional CRC calculated on the cell payload.

In our application the ATM link level code that normally runs in the NIC is replaced by a simple program that waits until an incoming cell enters the receiver interface. It then reads a 4 MHz timer, and copies the required portion of the cell to the PC memory via the PCI bus. The process is repeated for the next cell. A data collection program running in the PC copies the output from the PC memory via the PCI bus onto disk, using main memory as a buffer.

---

<sup>1</sup> See <http://www.escalate.net>



**Figure 1** ATM Traffic Measurement System

To extend the measurement capability of the system two ATM 155Mbps prototype cards were obtained and modified so that they could be used for measurements. While testing the ATM 155 cards it was found that while the cards were fast enough to capture a timestamp and cell header for each cell, they were not fast enough to capture the entire cell contents. This is because the section of code responsible for transferring the contents of the cell from the ATM interface to the PC interface takes longer to execute than 2.8  $\mu$ sec (the length of time to transmit a cell on an ATM 155 link). Another problem with these prototype cards is that only two were ever produced, and the manufacturer had no plans to produce any more. To solve both of the above problems and a requirement to measure traffic at E3 (34Mbps) and DS3 (45Mbps) it was decided to develop a special purpose card (called the DAG) that could be used in the measurement system.

### 3 Dag 1

The Dag Board has been designed to attach to an ATM 25 NIC card using an expansion bus on the NIC card. It provides additional ATM interfaces for 155Mbps fibre and twisted pair, DS3 (45Mbps) and E3 (34Mbps). These interfaces are connected to the ARM bus on the NIC card via a full 32-bit path. Figure 2 shows a block diagram of the board design. It comprises two independent physical layer sections (one to capture ATM 155 traffic and the other to capture either DS3 or E3 traffic). It is intended that the board will only collect one data stream at a time so both the designs share a common CPU interface. The purpose of this CPU interface is to:

- generate a timestamp for each cell arrival;
- demultiplexes bytes arriving from an ATM framer device into 32-bit words for transfer over the ARM bus;
- buffer timestamps and cell data for transfer over the ARM bus;
- generate a CRC on the cell payload if necessary (described in Section 4).
- provide a GPS interface to allow timestamp clocks on different Dag cards to be synchronised (this is described in section 5).

The design of the ATM-155 and DS3/E3 ATM interfaces were strongly based on reference designs (PMC,1995 and PMC,1995a) to improve the chances of success of the board.

The CPU interface is implemented using a XILINX<sup>tm</sup> field programmable gate array chip<sup>2</sup>. These XILINX devices are based on a SRAM technology, meaning that they require reprogramming after power-up or a reset operation. On this board the image used to program the XILINX device can be downloaded directly from the host PC via the NIC card as required. The advantage of this approach is that it is possible to download different images dependent on the functionality required to perform a particular set of measurements. Another great advantage of this approach is that it greatly simplifies the development cycle as a new ROM device containing the FPGA image does not have to be burnt each time a design is to be tested (which is the normal approach).




Figure 2 Block Diagram of the Dag 1 Board

Figure 3 shows the outline design of the internal logic used in the Xilinx device. The cell multiplexer and filter read each cell received by one of the physical layer interfaces. If the cell passes the filter parameters it is stored in the next free block of the cell buffer. As each cell arrives a time stamp is generated, and a cell payload CRC is computed as the payload bytes of the cell are read. These values are stored in the buffer along with the cell header and payload. Time signals are received by a GPS receiver at a 1 Hz rate. These also generate timestamps, which can be post processed to correct internal clock timestamps to UTC. The operation of the GPS interface is more fully described in section 5.

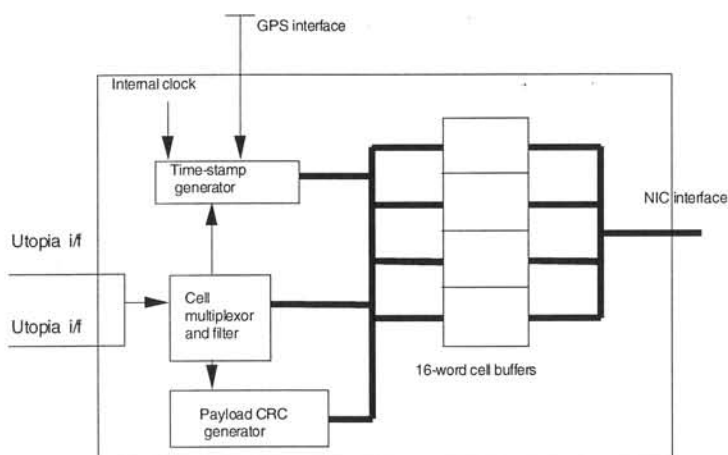


Figure 3 DAG, outline Xilinx internal logic

<sup>2</sup> see <http://www.xilinx.com>




The ARM 610 processor on the NIC card has direct access to the cell buffers in the XILINX. The processor can read any word of the cell, the header, CRC or time stamp. With the 32 MHz ARM processor the complete contents of a cell buffer can be read into its memory in less than 1.25  $\mu$ S

In future it is intended that the various gate array designs are going to be performed using the VHDL hardware description language and automated synthesis tools. VHDL allows designs to be performed at a level of abstraction similar to that provided by high level languages in the software development world. These designs can then be simulated to ensure their correctness, and then synthesised to generate an implementation suitable to program the field programmable gate array. A photograph of the completed board is shown in Figure 4.

The A7M155 side of the board has been fully tested and has been used to conduct a number of measurements have been carried out on the departments Atm Network, the Opera<sup>3</sup> Network and between New Zealand and Canada. The results of these measurements can be seen in Section 6 below. These measurements have demonstrated the feasibility of building dedicated measurement cards.

This has led to the design of a single board solution called the Dag 2 that combines the functionality of the Dag 1 card and the parent ATML NIC card that it plugs into.



**Figure 4** Photo of the Dag 1 Board

## **4 CRC Generation**

In order to measure cell delay we need to be able to identify cells at the entry and exit points of the network we are testing. Conventionally this is done by injecting special test and measurement cells into the network, these cells have some characteristic that distinguishes them from cells in the normal traffic. However, this technique has two disadvantages. Firstly, an active device is required to inject new cells into the network,

---

<sup>3</sup> A trial being run jointly by Telecom New Zealand and a consortium of New Zealand universities and Crown Research Institutes – see <http://www.opera.net.nz>

secondly the very injection of cells into a network may change the parameters that we are trying to measure. Therefore it would be better if we were able to use the existing traffic in the network for measurement. To do this we need a way of characterising cells in the existing traffic flow, and this must be based on the cell payload, as the cell header will change as the cell moves through switches. It is not practical to use the entire cell payload; there would be too high a volume of traffic back from the monitor to the matching station, matching would be slow as all twelve words from each cell would have to be matched, and the network user might rightly be concerned about security if their traffic was being copied and re-transmitted. Thus we have investigated the use of a 32-bit CRC calculated over the cell payload to characterise the cell. This has the advantages of being easy to calculate in hardware, of reducing the data to be transmitted to the analysis site from 48 bytes to 4, and of being an irreversible operation, in that the cell payload cannot be recovered from the CRC which reduces security problems.

We have chosen to implement the AAL5 CRC32 in our hardware. This CRC is calculated in parallel with cell storage in the buffer as each 32-bit word is received, and the CRC is transferred to the cell buffer at the end of the cell time. The CRC transfer adds only one cycle of the FPGA clock (24 MHz in this implementation) to the time taken to receive a cell.

In order to match one cell stream against another in the analysis software, and thus measure delay through a network, each payload CRC from the cell stream at the entry to the network is compared with CRCs recorded at the exit point. This comparison is carried out for exit cells within a time window relative to the timestamp of the cell recorded at the entry point. The size and position of this window depends on the circumstances of the measurement.


For example, if we are measuring delay through a switch that has a minimum delay of 20  $\mu$ sec, the window will initially be set to start at say time 0, and may run forward to 1 msec. If we are measuring the delay between two points on the Internet separated by a light-speed travel time of 40 msec the window will start at 40 msec, and go forward to about one second. The width of the window is a compromise between increasing the probability of correctly matching cells delayed for a very long time, and reducing the probability of a false match.

False matches can occur if two cells with the same payload CRC are transmitted during the width of the time window. To study the distribution of CRC values we have recorded several million cells from NFS traffic between a workstation and its file server. We have found that the overall number of duplicate CRCs is very small, less than one per cent of the total. However, the important statistic for judging whether or not cell matching will be accurate is the number of duplicates that occur within the time window used for matching. Investigations into this statistic are still proceeding, using recordings of different types of network traffic.

However, at present we conclude that the AAL5 CRC is suitable for cell matching, and that there is a very low probability of false matches for delays in the region of 0 - 1 second.


## 5 Synchronisation

To perform certain types of analysis of ATM traffic requires measurements to be performed at several parts of the network simultaneously. For example to calculate cell delays through a switch, it is necessary to record the times each cell enters and leaves the switch. Using the Waikato measurement system this can be achieved by using two measurement systems as shown in Figure 5. The data resulting from the measurements can then be merged for analysis. If this analysis is to be meaningful the timestamps generated for each cell on each of the NIC cards must be accurate with respect to each other. As each of the NIC cards uses its own local clock for generating timestamps it is important to compensate for any clock drift that might occur in order to maximise accuracy.




**Figure 5** Example set-up using two measurement devices

One of the approaches used in the Waikato measurement system to compensate for clock drift is to connect a common oscillator circuit to the NIC cards as shown in Figure 6. The cards are then programmed so that on every pulse from the oscillator, a special record containing the current local time added to the measurement data. These records can be used when merging two or more sets of data to model clock drift and make appropriate compensations to the timestamps in the measurement data. A number of tests including connecting two measurements systems to the same ATM link have showed that this is an accurate method of compensating for clock drift.



**Figure 6** Synchronisation using a common oscillator

While this approach works well for measurement systems which are physically very close together, it is impractical if the measurement systems are geographically remote. The approach that is used in this situation is to use GPS receivers attached to each measurement system. GPS receivers provide accurate global timing information in addition to global positioning information. For example the Trimble™ Palisade models<sup>4</sup> provide a pulse every second that is globally accurate to 100ns. Figure 7 shows an example network with two geographically remote measurement systems attached.

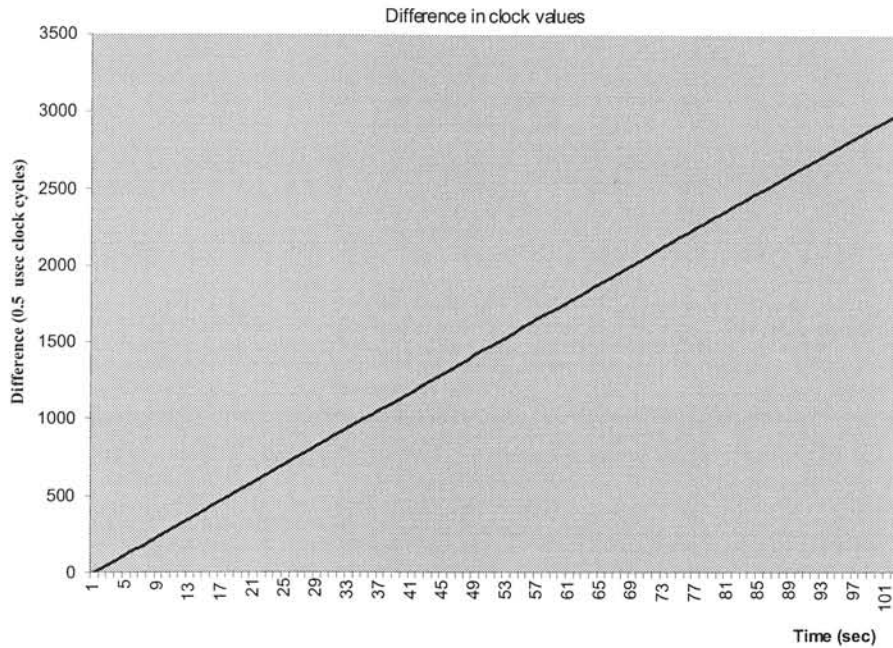


**Figure 7** NIC Synchronisation using GPS

An initial concern with using the Trimble™ Palisade GPS receivers to provide the synchronisation pulse was that the frequency of the timing signal (once a second) would be too low to accurately model clock drift. To determine if this was the case an experiment that recorded the differences in the clocks of two NIC cards over a period of time. Figure 8 is a plot of the difference between the timers on two NIC cards over time. As can be seen the relative rate of drift is constant over a long period of time (29 clock cycles per

<sup>4</sup> See [http://www.trimble.com/cgi/products.cgi/pd\\_om008.htm](http://www.trimble.com/cgi/products.cgi/pd_om008.htm)

second in this experiment). As the clock drift closely approximates a linear function it was decided a 1Hz synchronisation signal was sufficient.



**Figure 8** Clock Difference between two NIC cards over time

## 6 Remote Monitor

The present version of the analysis software runs on a LINUX system, and is written in C. Its main functions are to set up communications with the remote monitors, to set up filter parameters, and to request the return of data on filtered cell streams. The software can also display graphs of traffic and delay histograms in real time.

The cell timing information in the received data stream is corrected to UTC, using calibration information also carried in the data stream. The corrected data can be written to disk, or may be matched with a cell stream from another monitor to calculate delay and delay variation statistics. Development of analysis software is continuing, with a web-based version as a high priority.

## 7 Dag 2

The Dag 1 board design was used to prove that it is possible to build an ATM measurement device. This Dag 1 board however must be connected to an ATML VL2000 NIC card. This poses two main problems. First the manufacture of this model of card has now ceased and so it will not be possible to obtain further boards in the future. Second, the ATML board and the DAG 1 card when connected together take up the space of two PCI slots limiting the number of additional cards that can be installed in a PC. This is particularly problematic if two measurement systems are to be installed in the same PCs as most of them only have 3 PCI slots. This has led to the design of the Dag 2 board.

The Dag 2 design is a single board solution that contains the functionality of both boards in the previous system. The design contains the A7M155 interface and CPU interface of the Dag 1 board plus an ARM 710 processor, memory and a PCI interface. Figure 9 shows a block diagram for the Dag 2 board. The A7M155 line interface and framer and the CPU interface components have been taken from the Dag 1 design.

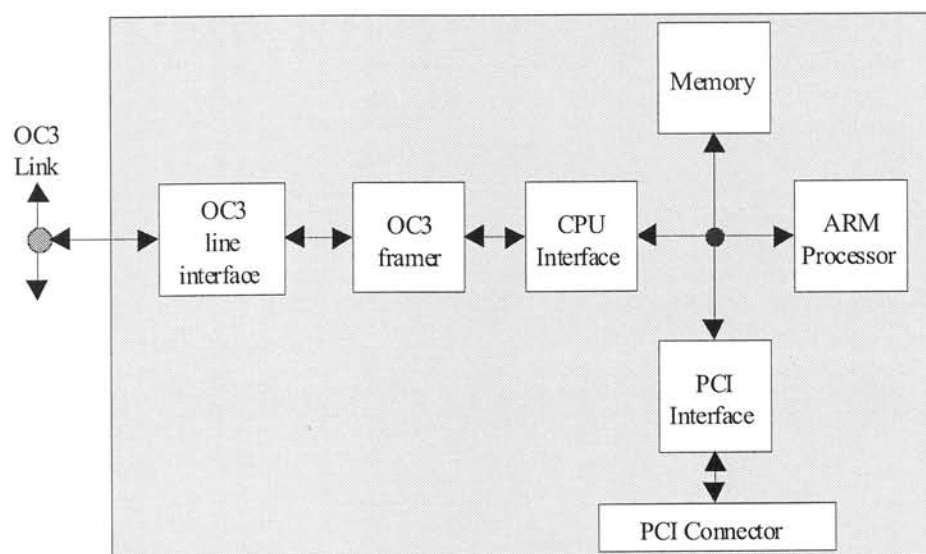


Figure 9 Block Diagram for the Dag 2

The final design is a six layer PCI card sized board. An initial run of three prototype boards has been manufactured, assembled and fully tested. We are currently in the process of building a batch of 15 production boards.

## 8 Traffic Generation

To allow synthetic traffic to be generated the original ATM25 NIC cards can be programmed to send cells at predetermined times. Initially the system is only capable of generating constant bit rate (CBR) or burst streams and the burst size, number of bursts, inter-burst time and inter-cell gap are parameters to the traffic generation program.

This traffic generator has been used in a number of experiments on the OPERA network. A number of these experiments involved saturating an E3 (34Mbps) link of OPERA. As the ATM25 (25Mbps) cards are incapable of generating enough traffic to saturate a E3 link two of the cards were paired up (by switching the traffic from both of these cards onto the same PVC on one of the output ports of the groups ATM switch) to generate cells rates in excess of that of an E3 link. Using this configuration it was possible to conduct the necessary experiments on the OPERA network.

To extend the traffic generation capability of the system work has been done to develop XILINX images that will enable the Dag 2 cards to generate synthetic traffic up to A7M155 rates. The traffic generation programs that drive the Dag 2 cards in this mode will use similar parameters to the ATM25 based traffic generation programs. In addition a mode to send cells at the times specified in a file of timestamps is planned. This will allow the results generated from the ATM-TM simulator to be fed into an actual system so the actual and simulated behaviour can be compared.

## 9 Results

An initial requirement of the system was that it should be capable of collecting ATM cell arrival times over long, continuous periods, to enable us to carry out statistical analysis of ATM traffic patterns. An example is shown in figures 10 and 11, where NFS traffic between a SUN server and a workstation has been recorded. Figure 10 shows a record over a period of about two hours, displayed as counts within 40 second bins. Figure 11 shows the same traffic, but with 1 second binning, for a 200 second period starting at about 3900 seconds into the record.




Figure 10 Traffic measured on a NFS server to workstation link (40 sec bins)






Figure 11 Traffic measured on a NFS server to workstation link (1 sec bins)

#### Switch delay measurements

To demonstrate the timing accuracy and resolution of the system, measurements were made of time delay for cells passing through an ATML Virata switch. The cell stream used for these measurements was NFS traffic from a server to a workstation. Figure 12 shows the delay histogram for the switch under the relatively lightly loaded condition produced by typical NFS traffic.

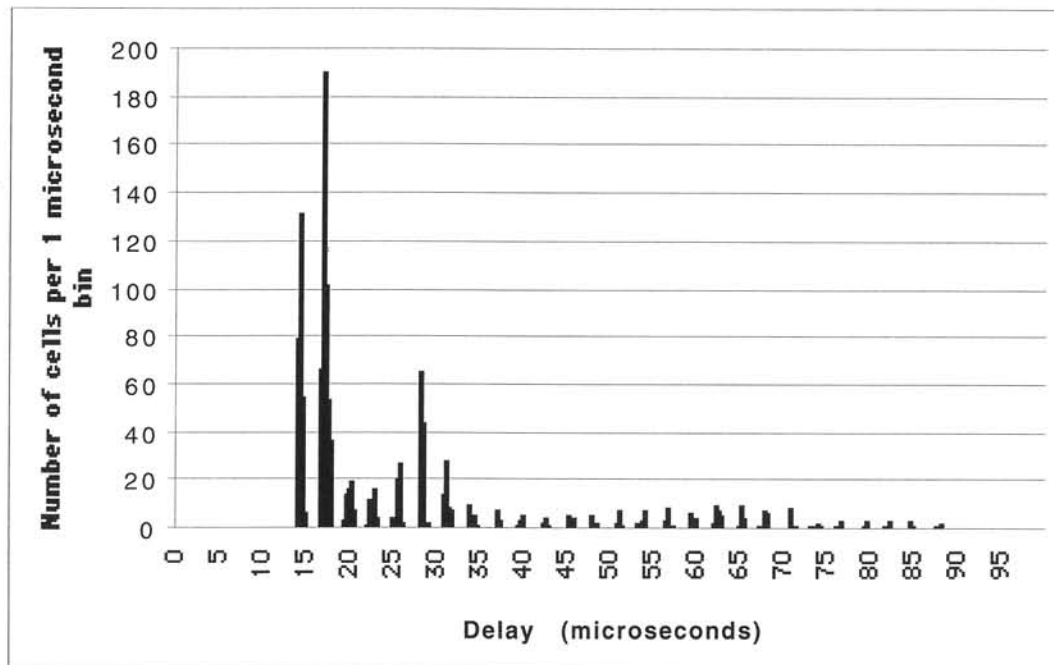



Figure 12 Cell delay through a lightly loaded switch



**Figure 13 Cell delay under heavier loading**

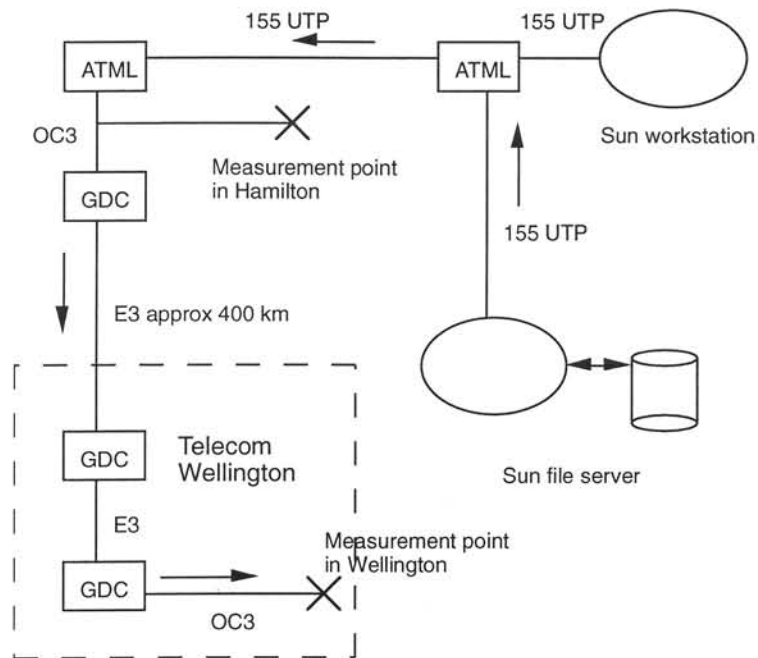
In figure 13 a heavier switch loading has been simulated by sending the same data twice through the switch. These diagrams illustrate that time resolution of the system, the 2.75 msec quantisation of delays due to slotting of cells in the A7M155 frame is clearly visible.

#### ***Delay measurements over OPERA.***

Measurements were made over an approximately 400 km path from the University of Waikato, in Hamilton, to a Telecom New Zealand site in Wellington, New Zealand. The main part of the circuit was carried over E3 links (34.5 Mbps) which are part of the OPERA experimental trial.

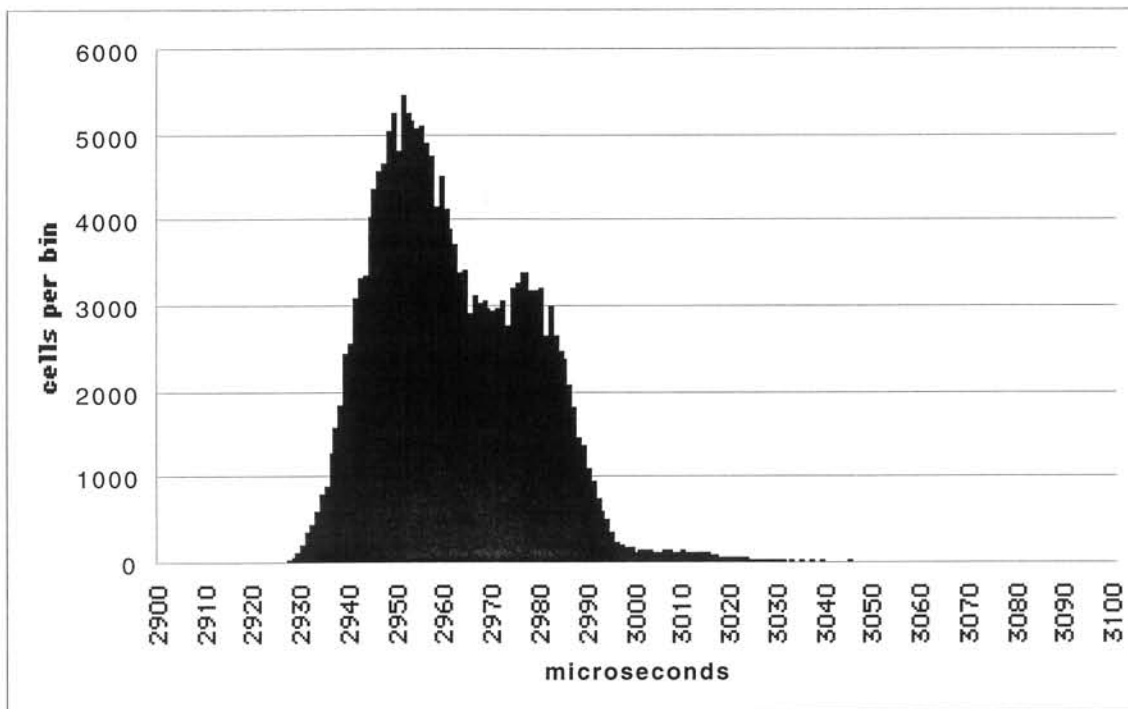
The traffic flow for the experiment was a copy of live NFS traffic flowing between a Sun server and a Sun workstation. The traffic was copied by creating a multicast circuit within the ATM Ltd Virata switch that connects the workstation to its server (figure 14).





**Figure 14** Experimental setup for Hamilton -> Wellington delay measurements

The histogram in figure 15 shows a typical result, where 196,346 cells were matched in an experiment lasting several minutes.




**Figure 15** Histogram of Hamilton to Wellington cell delays over OPERA

### *Intercontinental Internet delay measurements*

The remote monitor system has recently been used to measure delay over the Internet between the universities of Calgary (Alberta, Canada) and Waikato (Hamilton, New Zealand) a great circle distance of 12,000 km, corresponding to a time of flight of 40 msec at the speed of light.

As only a unidirectional route could be established, from Calgary to Waikato, the measurement were made on UDP traffic. The test traffic was generated by a workstation in the computer science department at Calgary, and routed via WNet<sup>5</sup> to another workstation on campus. This workstation acted as a gateway to the Internet, and the traffic then followed a path across the Pacific to Australia and onwards to New Zealand. Within Waikato University the traffic was routed into our ATM LAN and re-encapsulated as classical IP over an ATM PVC.

Most measurements were made on UDP packets containing 400 bytes of data, which were encapsulated as 10 ATM cells per packet. It was observed that the CRCs of the first and last cells in each UDP packet could not be matched, although all intermediate cells were matched. When UDP/IP is encapsulated by AAL5 and split up into the cells the first cell contains the UDP and IP headers. Most of the fields in these headers remain constant as it passes through the IP network, however the time to live (TTL) field is updated as it passes through each router. The change in the TTL field affects the payload CRC of the first cell, and also the AAL5 CRC in the last cell, and thus the payload CRC of that cell. Figure 16 shows the measured results for approximately 680 UDP packets, for which 5412 cells were matched.



**Figure 16** Histogram of delays over the Internet from Calgary to Waikato

<sup>5</sup> see <http://www.wnet.ca>

## 10 Conclusions

In this paper we have described the design of a simple and low-cost system for making cell arrival time measurements on ATM systems. Despite its low cost the system achieves very high time resolution and absolute accuracy, and is capable of a range of measurements that would otherwise require expensive test gear. By using two or more monitor system it is possible to make accurate measurements of cell delay and loss, and hence measure the Quality of Service (QoS) parameters of an ATM link.

Work is continuing at Waikato on higher performance versions of the monitor hardware and software, capable of making measurements on 622 Mbps ATM links. We have also developed a similar Ethernet-based system for monitoring uni-directional Internet delays in the Pacific Rim/Asia region. This was described at the recent INET'98 conference (Graham, Donnelly, Martin, Martens and Cleary, 1998), and further results of this work are shown on our web site<sup>6</sup>.

## 11 References

Graham, I. D., Cleary, J.G. (1996) *Cell Level Measurements of ATM Traffic*, Proceedings of the Australian Telecommunication Networks and Applications Conference, Monash University, Melbourne, pp 495 – 500

Graham, I.D., Donnelly, S.F., Martin, S., Martens, J. and Cleary, J.G. (1998) *Non Intrusive and Accurate Measurement of Unidirectional Delay and Delay Variation on the Internet*. Proceedings of INET'98, Internet Society, CDROM and Web Publication only – [http://www.isoc.org/inet/proceedings/6g/6g\\_2.htm](http://www.isoc.org/inet/proceedings/6g/6g_2.htm).

Habib I. B. and Saadawi, T.N. (1994) "Multimedia Traffic Characteristics in Broadband Networks", IEEE Communications Magazine, pp4 – 9.

PMC, (1995) "S/UNI-Lite Optical Reference Design", PMC-Serria Inc, 8501 Commerce Court, Burnaby, BC Canada V5A 4N3, September, 1995

PMC, (1995a) "Reference Design for S/UNI-PDH", PMC-Serria Inc, 8501 Commerce Court, Burnaby, BC Canada V5A 4N3, September, 1995

---

<sup>6</sup> See <http://atm.cs.waikato.ac.nz/atm/>