Packet Delay and Loss at the Auckland Internet Access Path

Klaus Mochalski[1], Jörg Micheel[2], Stephen Donnelly[1]
{klaus,sfd}@cs.waikato.ac.nz
joerg@nlanr.net

1 Waikato Applied Network Dynamics
Computer Science, G Block
The University of Waikato
Private Bag 3105
Hamilton, New Zealand

2 NLANR Measurement and Network Analysis Group, SDSC, UCSD
10100 John Hopkins Dr
92093–0505 La Jolla
California, USA

Abstract–– **In this paper we study packet delay and loss for IP data traversing the University of Auckland Internet access path. The ISP uplink of the university has been the subject of previous measurements and studies [1]. The focus of this paper is to evaluate the information derived from a multipoint measurement, Auckland–VI [2], which was collected for a duration of four days during June 2001 at the ISP uplink and Ethernet hubs outside and inside relative to the Internet firewall host. The specific value of this data set lies in the hop–by–hop instrumentation of all devices operating at the packet level, combined with the duration of observation and the fine–grained timestamping process of packet arrivals. We provide measures for the normal day–to–day operation of the access link system. We show the impact of operating the same links at 10MBit/sec Ethernet vs. 100MBit/sec Ethernet hubs. The data set includes a number of areas where service quality (delay and packet loss) is extreme; we examine the causes and impacts on network users.**

Index terms–– **packet delay, packet loss, multipoint measurement, passive measurement**

A. Introduction

Network capacities are being deliberately overengineered in today's commercialised Internet. Any network provider, be it a commercial Internet Service Provider (ISP) or an Information Technology Services department at a company or university site, will design network bandwidth resources in such a way that there be virtually no data loss, even during the worst possible network utilisation scenario. Thus, the service delivered by today's end–to–end wide area Internet would be perfect – if it wasn't for the inter–domain connections, such as Internet access links to the (next higher level) ISP or peering points between ISPs. Those points are carefully rate limited and their capacity is controlled by Service Level Agreements and Peering Arrangements. It is at these points that data packets will experience delays and loss due to deliberate reductions in network bandwidth.

The specific value of the analysis presented in this paper lies in the hop–by–hop instrumentation of all devices operating at the packet level. Unlike single point measurements, this scenario allows the study of *changes* in traffic patterns of the same data stream relative to different points of observation and to investigate the *individual contributing factors* to the changes observed. There is little doubt that this approach will lead to a better understanding of general Internet traffic patterns.

Our approach to understanding network behaviour is based on the properties of the devices present in the network under observation as well as the offered traffic load. We consider such properties as:

- peak link data rates
- typical packet sizes and their distributions
- link layer overheads, interpacket gaps, padding and other properties such as collisions which modulate packet timing behaviour on a given link
- rate shaping and packet buffering at routers

The rest of the paper is organised as follows. In section B we describe the measurement environment and technology used to collect the data sets and discuss their accuracy and limitations. Section C describes our notion of packet delay specific to the network under observation. In section D we discuss delay and loss across the access router. Section E focuses on packet delay across the firewall device. In section F we look at the impact of changing the Ethernet hubs at the access path from 10MBit/sec

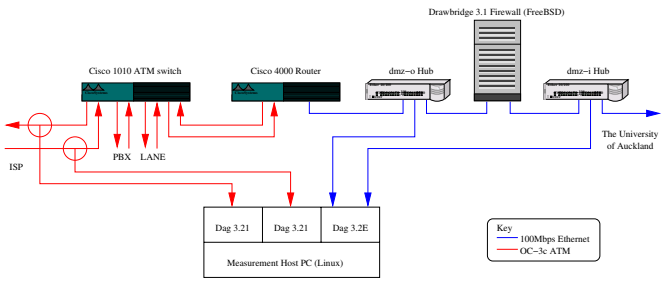

Figure 1: University of Auckland Internet infrastructure

to 100MBit/sec. In section G we simulate the packet delay and loss figures across the access router with a simple leaky bucket scheme and show whether this allows us to appropriately model the delay and loss effects observed. Section H discusses the analysis; section I concludes the paper.

B. Measurements

In this paper we study the Auckland−VI data set, which is a four day continuous collection of three−point IP header trace files with GPS−driven precision timestamping [4], obtained at the University of Auckland Internet access path (Figure 1). This path has been instrumented at the ATM OC3c access link to the university's ISP, Clear Communications [1], which is rate limited to 4.048MBit/sec, although the traffic shaping upstream and downstream is implemented in different ways (Figure 2).

The second and third measurement points monitor the 10/100MBit/sec Ethernet hubs before and after the Internet firewall host, a Drawbridge 3.1 running the FreeBSD operating system. It should be noted that a dual−speed Ethernet hub does not implement one single broadcast domain operating at either speeds, but, at a minimum, has two distinct domains bridged to each other; one operating at 10MBit/sec, the other at 100MBit/sec. In this measurement the Ethernet Dag cards have been connected to the 100BaseTX collision domain.

The disparity between local network capacity and available bandwidth at the upstream ISP link enables the detailed study of network bandwidth demand at the university and the supply or lack of network capacity at the access link, which results in delays caused by queuing at the router, as well as packet loss, once buffer resources have been exhausted.

The data collection system consists of a pair of Dag3.2 OC3c/OC12c measurement cards [3] (as used in previous studies [1]), as well as a Dag3.2E, which is a dualport 10/100MBit/sec Ethernet card. All of the Dag3.2 series of

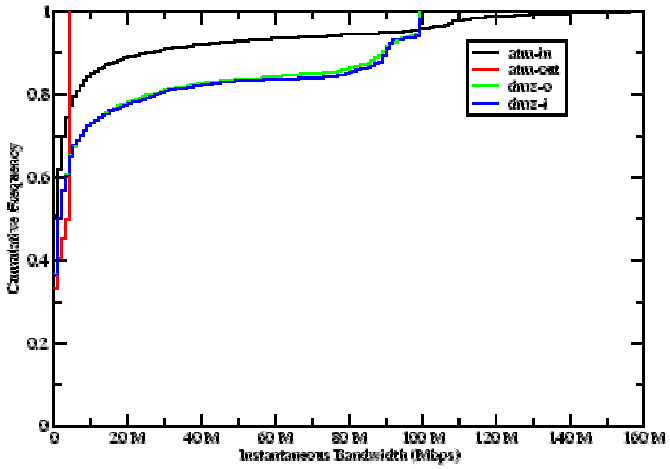

Figure 2: Instantaneous bandwidth graph for all monitored links

network measurement cards implement a conditioned clock operating at 16.66 MHz, coined the DUCK [4], which provides for a 60−nanosecond clock resolution. The actual accuracy of the clock is controlled via synchronisation to a Trimble Palisade GPS system and is within about 600 nanoseconds to UTC at any point in time during the measurement session. Log files of the timekeeping process have been recorded for verification. In addition, the instantaneous bandwidth (Figure 2) as well as the firewall delay analysis (Figure 12) do not reveal inconsistencies in the timestamping process, which reassures our confidence in the data set.

To put the precision of the timestamping process into perspective, the smallest distinguishable events on a 100MBit/sec Ethernet are back−to−back 64−byte packets with a minimum interarrival time of 5.1 microseconds; ATM cells on the ATM OC3c links are a minimum of 2.7 microseconds apart. Thus, the timestamping system provides for about an order of magnitude better precision relative to the object of observation. All of the cards deployed in this scenario deliver the first 40 bytes of the IP packet and generally contain most, if not all, of the IP/TCP, IP/UDP, and IP/ICMP header information. The Ethernet records also provide the Ethernet MAC header and the true size of the Ethernet packet as it appears on the network, including padding.

The Auckland−VI data set consists of two subsets. The first trace was taken for one hour in May 2001 with the Ethernet hubs operating the DeMilitarised Zone (DMZ) at 10MBit/sec. A month later, the 3Com hubs had been replaced by Allied Telesyn CentreCom hubs (Figure 1). The second subset covers a total of four days, including a weekend and two busy working days. The size of the entire data set is about 20GB compressed and contains 312 million IP packet headers from each of the three collec−

tion points.

The data is available to the public via request from the WAND research group in the form of anonymised compressed IP header files on DDS4 tape or via anonymous ftp from the NLANR/MNA Group PMA data server [6]. A CD–ROM with sample data is also available (as of late March 2002).

C. Definition and understanding of packet delay and loss

We define *delay* as the difference in time between the arrival of the same packet at two different points of observation.

The arrival of a packet at a network link is not an atomic event, but due to bit deserialisation, it is a function of the packet's size. The Dag cards implement packet header timestamping [4]. To account for the packet wire time, we consider two different notions of delay. We use the term *delay* for the simple difference in timestamps as seen at two points within the network. We use the term *processing delay* to express time between the completed arrival of a packet at one point relative to the first appearance at another. The difference between *delay* and *processing delay* is the packet deserialisation time at the incoming link.

The algorithm used to identify the same packet at two distinct measurement points matches the first 40 bytes of the IP packet, however masking the TTL and IP checksum fields, since these are modified at a router. A match must occur within a two second window from the arrival at one link. If no match occurs, the packet is considered unmatched. If more than one matching packet originates from the same link, the entire match is considered ambiguous and discarded from delay analysis. We investigate duplicate packets in another paper [7].

Delay calculation can yield positive and negative values (which indicate the direction in which the packets travel). The delay values for each direction are analysed separately.

Unmatched packets can be attributed to several factors:
a) local cross traffic at either of the Ethernet hubs,
b) for the IP router: packets destined for, or originating from, one of the router's interface IP addresses,
c) for the firewall: packets discarded according to filtering rules, and
d) true packet loss due to overload (offered packet load exceeding link capacity).

Since the firewall rules are unknown, we do not provide loss figures for traffic crossing this device. To distinguish cross traffic at the outside DMZ hub from packets destined for the access router, we use the Ethernet MAC address as a filter. In addition, packets with source or destination IP addresses of the access router are dropped from loss analysis.

D. Packet delay and loss across the access router

From the 96 hours of observation, we have chosen a representative six–hour window to illustrate and discuss the details we have observed within the trace.

Figures 3and 4 show delay values of packets passing the router over a period of six hours. The most obvious feature is the huge disparity of the distribution ranges: 99% of all incoming packets experience a delay below 0.53 ms, however for outgoing traffic, the 99th percentile is at 82ms – a difference of two orders of magnitude.

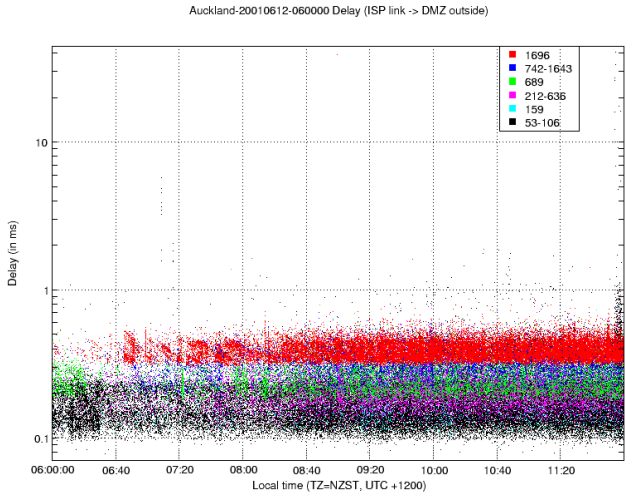

Figure 3: Scatter plot of delay values for incoming traffic

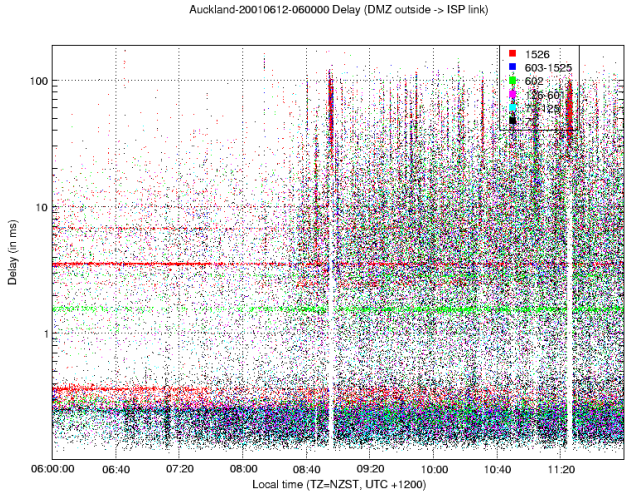

Figure 4: Scatter plot of delay values for outgoing traffic

The reason for this disparity is the huge link speed difference (100 vs. 4.048 MBit/s). Most of the time inbound traffic passes the router unqueued (only one packet at a time is stored in the router's buffer), whereas bursty outbound traffic has to be queued, thus introducing a delay depending on buffer occupancy. The strong bottom band in Figure 4 represents the unqueued portion of outbound traffic and looks similar to Figure 3.

Both figures display a horizontal banding at certain delay levels. Analysis of packet size distributions (Figures 5 and 6) reveals a strong prevalence of typical IP packet sizes (i.e., 40, 576, and 1500 bytes). We apply a colouring scheme to the delay distribution graphs which reflects these IP packet sizes and intermediate values. For Ethernet the groups are: 72, 602, 1526, 73–125, 126–601, and 603–1525; for ATM they are 53–106, 689, 1696, 159, 212–636, and 742–1643. Note that a specific link layer group may contain a range of IP packet sizes due to padding [2].

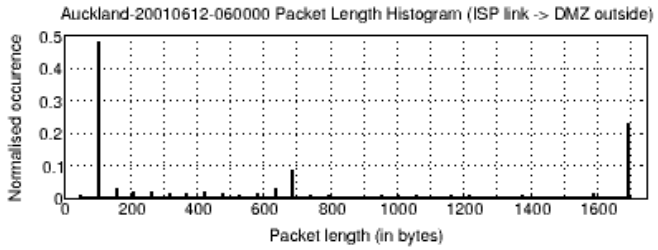

Figure 5: Link layer packet size histogram for inbound traffic

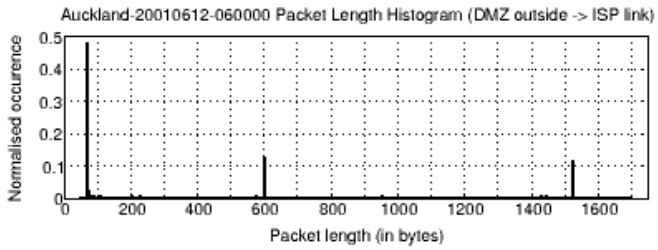

Figure 6: Link layer packet size histogram for outbound traffic

The colourisation illustrates the dependency between delay and packet sizes. In the graph for incoming traffic, the packet size groups are neatly stacked above each other in increasing order. The bottom band for outgoing traffic looks similar, if less pronounced (Figure 4).

The spread of delay values for unqueued packets can be solely attributed to packet deserialisation. This can be seen by looking at the bottom band of the corresponding processing delay graph (Figure 7) where all colours become intermingled red dots (1500–byte packets) showing up at the same low levels as black dots (40–byte packets). Unfortunately, the calculation of processing delays

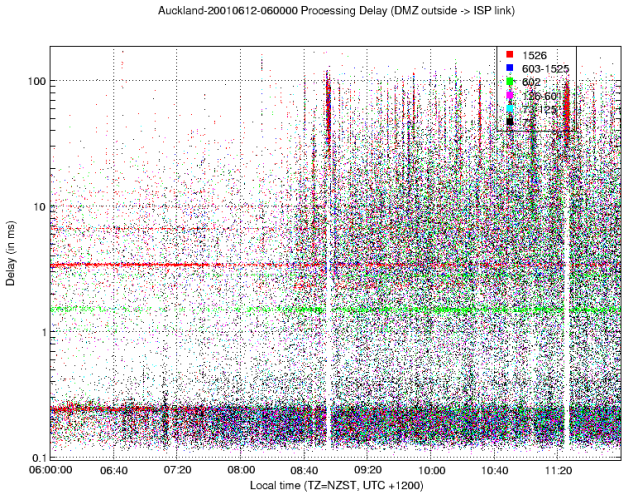

Figure 7: Processing delay for outgoing traffic

is only possible for outgoing traffic with packets arriving at a fixed packet rate of 100MBit/s. As shown in Figure 2, packets arrive at the University with unpredictable cell rates, which makes it impossible to determine the arrival of the last cell of a packet and thus the packet deserialisation time.

In addition to the strong bottom band there are at least four discernible bands at much higher levels – two green ones representing 576–byte IP packets, at about 1.4 and 2.8ms, and two red ones, representing 1500–byte packet at about 3.3 and 6.8ms. Calculating serialisation times for these packets with the transmission rate of 4.048MBit/s and link layer sizes of 689 and 1696 bytes yields 1.361ms and 3.351ms respectively. These values and their multiples match the bands very closely. The bands represent a router buffer state where two or more packets are queued behind each other. We further investigate the router's queuing behaviour in section G.

The most disturbing features in Figure 4 are two periods of about 5 minutes (around 8:50 and 11:30am) during which the minimum packet delay appears to be no less than 20ms. Apart from these two very strong events there are several smaller artefacts, notably a black streak (representing many 40–byte packets) at about 7:10 and several shorter spikes during the high volume part of the measurement starting at about 8:40.

Zooming into the region around 8:50 of Figure 4 and comparing it with a graph showing the protocol mix during this period reveals a clear coincidence with a sudden rise in SMTP traffic (Figures 8 and 9).

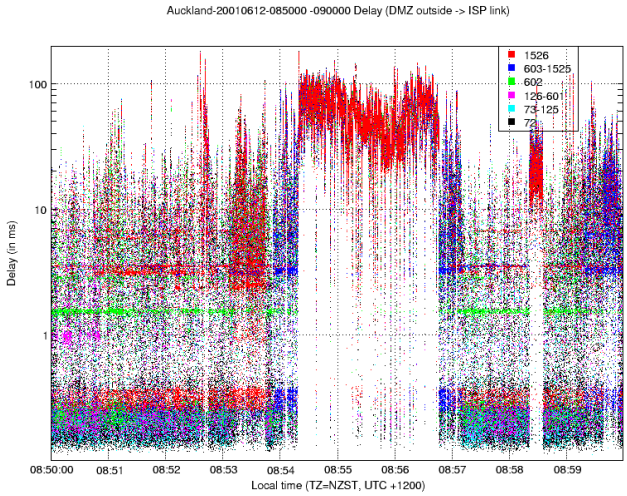

Figure 8: Zoomed into 8:50–9:00 of Figure 4

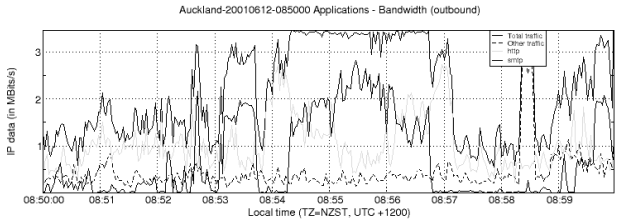

Figure 9: Outbound SMTP and HTTP traffic 8:50–9:00

The period from 11:20 to 11:30 exhibits similar correlation (Figures 10 and 11).

It also reveals a shorter period of generally high delays between 11:23 and 11:24 which again coincides with a SMTP spike whereas spikes of HTTP traffic (around 11:29) do not have such an adverse effect on delay.

Most of the time, packet loss across the access router varies around a few dozen packets per second, or, typically less than one percent of the packet link load. However, several spots with high loss do exist. Our intention is to find a correlation of periods of high delays with those of packet loss; we study those in section G. Another interesting anomaly is the loss of packets upon entering Auckland from the outside. Since the rate limiting in this direction is not explicit, bursts of packets may arrive back–to–back at full link speed and will get lost eventually due to the disparity of the peak link speeds at the 155MBit/sec link and the first 100MBit/sec hub.

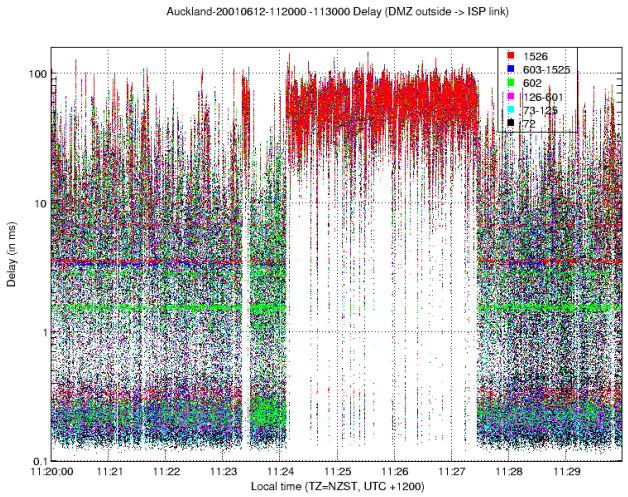

Figure 10: Zoomed into 11:20–11:30 of Figure 4

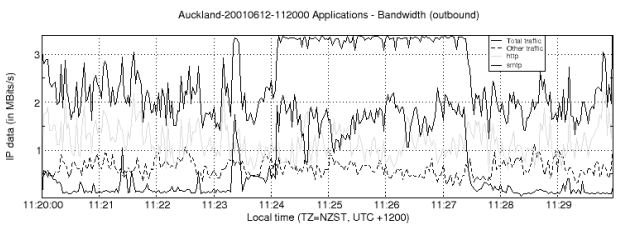

Figure 11: Outbound SMTP and HTTP traffic 11:20–11:30

E. Packet delay across the firewall

Applying the same technique used in section D shows a high consistency of delay values at a very low level. 99% of all packets for both directions show delays of less than 0.3ms. Again, a packet size dependent banding appears.

Figure 12 shows the inner portion of a delay–packet size graph. It includes approximately 99% of all delay values.

The inclined inner borders quite consistently show the minimum delay values for each packet size. Their slope represents the firewall's forwarding bandwidth. Figure 13 shows the corresponding processing delays. For incoming traffic the minimum delay line becomes nearly vertical implying a forwarding rate close to 100MBit/s independent of packet size, while the line for outgoing traffic maintains a certain slope which is equivalent to 90MBit/s for packets larger than 512 bytes and only 80MBit/s for smaller packets.

The three common IP packet sizes of 40, 576, and 1500 bytes show up as sharp horizontal lines. More interestingly, there are two vertical lines dropping from the areas where the inclined minimum delay and common packet sizes lines meet. Packet level analysis reveals that these lines stem from smaller packets arriving immediately af–

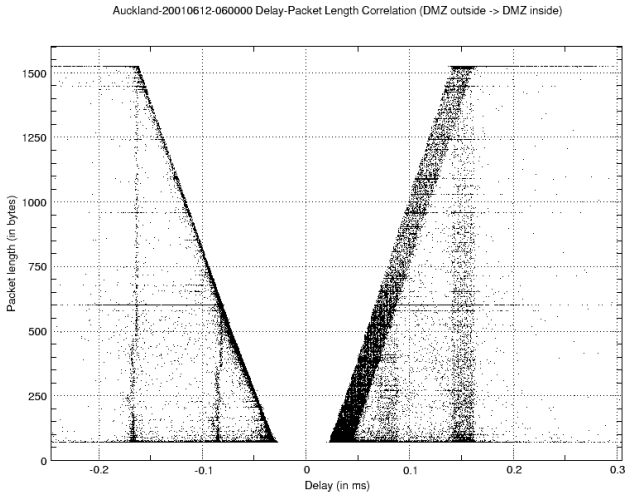

Figure 12: Delay−packet size correlation of packets passing the firewall

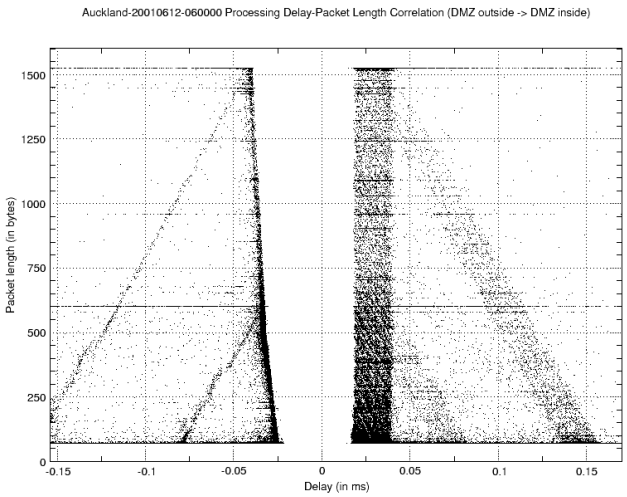

Figure 13: Same as Figure 12 normalised by subtracting deserialisation time

ter a 576− or 1500−byte packet thus experiencing the same delay as the long packet.

Another notable feature is the spread of delay values for incoming traffic inside a clearly defined band which does not appear for the opposite direction. Knowledge about the machine configuration suggests firewall rule processing, host memory allocation, and NIC interrupt latency as potential causes. However, since that band is only 0.025ms wide, this effect is negligible.

One of the authors observed a periodic component of increased delay values which occur at one−second intervals [5].

F. The impact of bandwidth on delay and loss − 10 vs. 100MBit/sec

Comparison of the trace captured in May in a 10MBit/s DMZ environment and a similar period of the June trace captured at 100MBit/s shows a clear bandwidth impact. While outgoing traffic is nearly unaffected, the incoming traffic sees a large decrease of delay and delay variation. The 99th percentile of delays decreases from nearly 28ms to about 0.5ms.

We attribute high delays in the 10Mbit/s DMZ environment to congestion. With 4MBit/s configured for each direction of the Internet access link, spikes of inbound traffic of up to 155MBit/s and cross traffic, the Ethernet hub is operating at its capacity limit, forcing queuing delays at transmitting Ethernet stations.

G. Simple leaky bucket simulation describing access router behaviour

The purpose of this simulation study is to understand whether it is possible to predict the egress traffic pattern of the router solely on knowledge of its configuration and the ingress traffic pattern.

If the simulation succeeds we might consider packet delay and loss independent of specific router implementations, making it possible to predict the behaviour of larger portions of the network such as end−to−end Internet paths.

We use a simplistic model to simulate queuing behaviour. A single point trace is passed through a FIFO memory structure for which input and output rate, link media, and buffer size, determine when a queued packet is due for departure. The algorithm is clocked by the timestamps of arriving packets. The simulated delay for each packet is directly comparable to the delay of the same packet in the two−point measurement, providing a measure for the quality of the model.

Figure 14 shows the resulting delay value distribution corresponding to that seen in Figure 4. At a first glance, the graphs show a striking similarity for such a basic simulation. Not only does it reproduce the horizontal banding thus confirming our suspicion from Section D that they are due to packet queuing, but it also captures most of the vertical artefacts.

Despite the promising visual similarities of both distributions, a detailed analysis is necessary to truly evaluate the simulation accuracy.

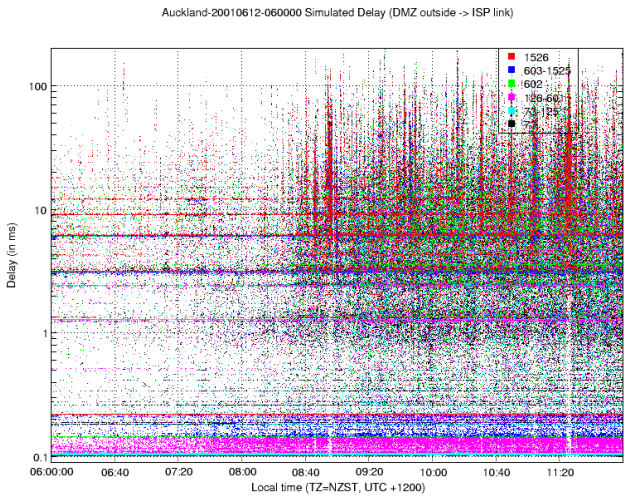

Figure 14: Simulated delay distribution

Figure 15 shows the relative error of the simulated delay over time. Although there is a strong cluster of points around zero, there are many values with huge errors. 45% of simulated delays have an error of less than 10%; 58% have an error of less than 20%.

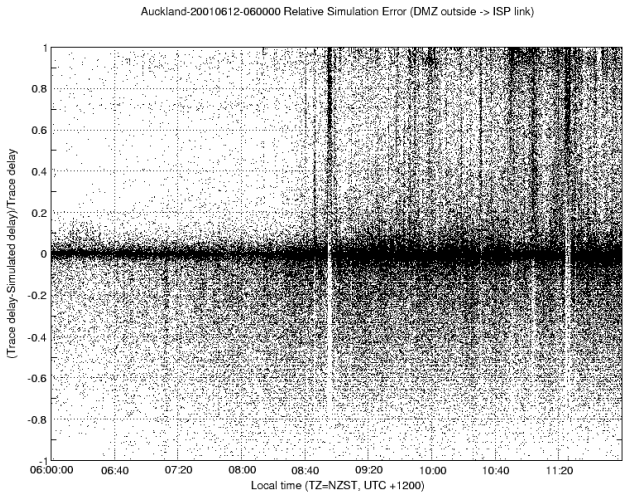

Figure 15: Relative error of simulated delay

The simulation only reproduces queuing behaviour. It yields a constant delay for unqueued packets of equal size. These delays are small compared to those of queued packets. Thus, any additional delay jitter introduced by the router has a strong effect on relative error. A first indication for this higher simulation inaccuracy for unqueued packets is the different appearance of the lower bands representing these packets in Figure 4 in comparison to Figure 14. Further confirmation gives a recalculation of the proportion of simulated delay values – considering only queued packets – which provide a good match for the observed delays. 52% of these delay values are within a 10% threshold and 63% are within 20%.

The visual difference between the earlier and later part in Figure 15 suggests a higher simulation accuracy during periods of lower traffic volume. To confirm this, we plot the observed against simulated delay of the first (black) and the last (grey) 1 million packets out of about 9 million (Figure 16). The error spread from the ideal x=y line is much higher for the high volume part. Given the considerations in the last paragraph, this is counterintuitive. With a queuing simulation one would expect increased accuracy if the traffic volume and thus the proportion of queued traffic increases. However, this does not happen.

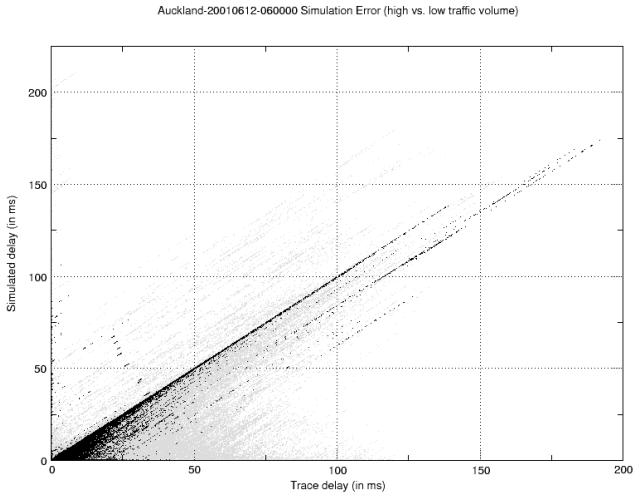

Figure 16: Comparison of absolute simulation errors during periods of high (grey) and low (black) traffic volume

Figure 15 shows two periods of particularly high errors at about 8:50 and 11:20. They perfectly coincide with the high delay artefacts described in Section D. This provides further proof that high load situations adversely affect the simulation accuracy. Zooming into the simulated delay distribution at these times (Figures 17 and 18) – corresponding to Figures 8and 10 respectively – shows that although the simulation is well able to qualitatively catch these events, it largely fails to reproduce their actual severity. It highly underestimates most of the delays during these periods.

Depending upon the definition of the accuracy threshold, the simulation is able to reproduce between 45% and 63% of the observed delays. This still leaves us with about half of all delays unmatched by the simulation. During the process of improving the simulation accuracy we investigated several potential error sources. A discussion follows.

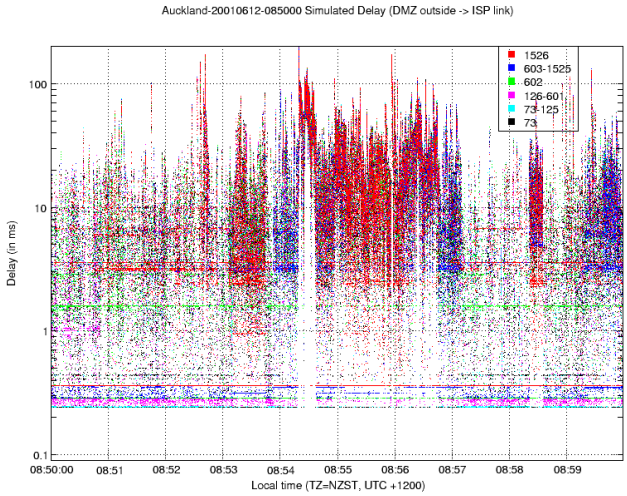

Figure 17: Zoomed segment of simulated delay distribution between 8:50 and 9:00

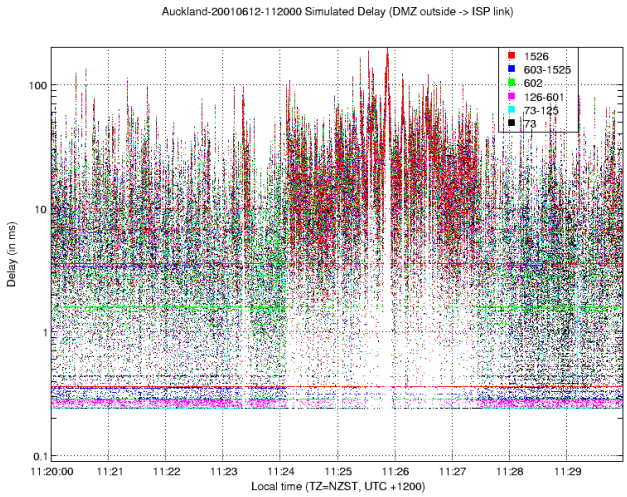

Figure 18: Zoomed segment of simulated delay distribution between 11:20 and 11:30

Our first finding was a delay offset of 0.24ms to deseri–alisation time and buffering delay for a large portion of packets. Compensating for this offset across all data greatly improves the simulation accuracy. However, it also results in overestimation of delay values for many packets.

We the investigated the dependency of simulation error on other variables, notably packet size and buffer occu–pancy. In addition, we tried to improve the simulation by implementing the buffer as a shared resource for inbound and outbound traffic.

Neither the correlation of absolute error and packet size (Figure 19) nor that of relative error and packet size re–vealed any suspicious patterns. Both show a clear cluster around 0 and a far spread of errors over the whole packet size range. The concentration of high relative errors for small packets results from their higher susceptibility to relative errors and a disproportionally high number of them.

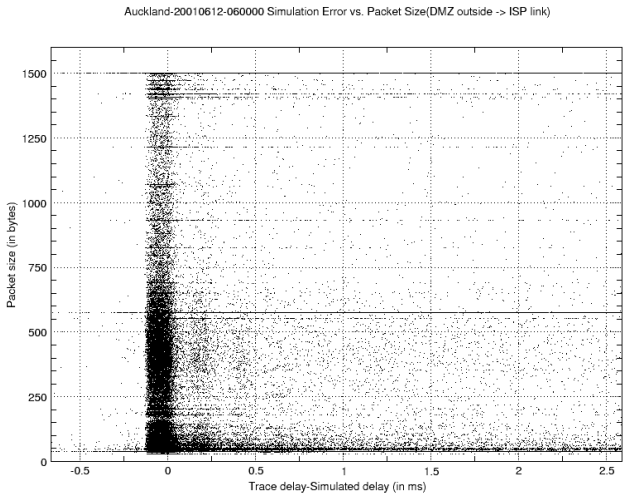

Figure 19: Absolute simulation error vs. packet size

The correlation of simulation error and buffer occupancy does not show any anomalies (Figure 20). For most packets the errors were closely clustered around 0. The stronger spread at the bottom is due to the prevalence of packets encountering a sparsely occupied buffer.

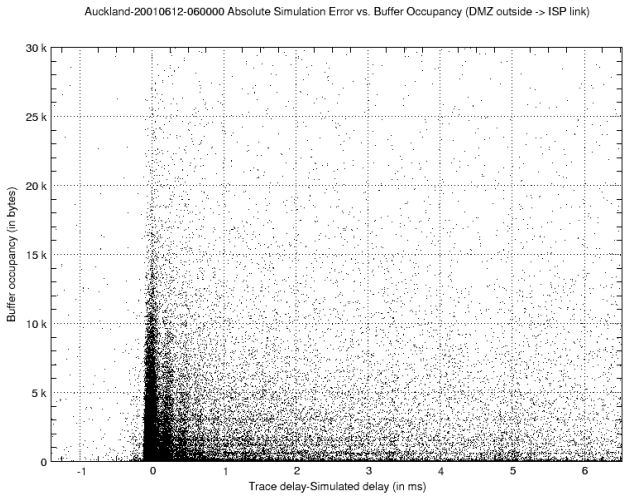

Figure 20: Absolute simulation error vs. buffer occupancy

A strong indication for the weakness of the simulation is its failure to predict the most extreme delay values – packet losses. They usually occur if the FIFO buffer is exhausted. The leaky bucket simulation should be well suited to correctly reproduce a high proportion of ob–served packet losses. However, the simulation predicts only two isolated occasions of about 2,300 lost packets

while we observed about 50,000 lost packets distributed over the whole six−hour trace.

There are several potential sources for this problem. We assume a buffer size of 128kB, derived from the observed maximum delay value. However, all six−hour traces of the whole data set display slightly varying maximum de−lay values. Thus, the usable buffer space may vary from our estimate. We also do not know the router's memory allocation and alignment scheme, which further decreases the confidence in our 128kB estimate.

Perhaps the strongest mismatch of the simulation and the router's actual behaviour is our disregard of traffic pass−ing the router in the opposite direction that might use the same buffer as a shared resource. We enhanced the simu−lation by taking into account all incoming packets, stor−ing them along with all outgoing packets in the same FIFO memory structure. This resulted in a slight increase in the number of predicted losses to 5,400 − still far from matching reality.

Evaluating absolute and relative error distributions as in the previous paragraphs for this bidirectional simulation does not show any notable improvement. We attribute this to the absence of any significant queuing of inbound traffic as can be seen in Figure 3.

H. Discussion

We used a simplistic model in order to understand whether it is possible to describe router behaviour based on the traffic pattern present at the input. We deliberately did not fetch detailed information about its architecture. The simulation captures well the overall behaviour of the router, but fails to accurately describe delay and loss pat−terns on a packet−by−packet basis. We sought to improve the simplistic model by looking for further dependencies, such as those based on packet sizes, queuing depth as well as others and were unable to reveal any patterns. We must conclude that the router under observation acts like an ordinary computer system, which implies non−deter−ministic behaviour when it comes to real−time properties.

A limitation of this study is the observation at only one single measurement point with one particular router as the object. It would be interesting to see if more modern routers (cell switched) behave similarly to the device studied in this paper, or show a more predictable pattern. We are planning a number of future projects to investi−gate modern backbone and enterprise routers.

With this paper we question the utility of simulation studies. While it is possible to observe, interpret and un−derstand general traffic behaviour on the Internet using simulations, we have doubts that there is a good chance to further advance the understanding of network behav−iour with theoretical studies, specifically, when it comes to service quality and guarantees. We found as much as 100% (or more) error between theory and practice for a reasonably large proportion of packets, and we have to attribute those to the specifics of the implementation of the router. With routers each behaving differently, it ap−pears impossible to accurately predict network traffic patterns and behaviour.

Our analysis focused on the behaviour of two typical net−work devices: a router and a firewall. The firewall has very little impact on the overall delay and delay variation behaviour. It is very unlikely to adversely affect all but the most delay sensitive data.

The vast amount of data calls for a highly efficient proc−essing technique. Therefore we purposely developed tools written in C, most of which extract only one specific fea−ture of the data. The lack of standardised tools remains an obstacle to fully exploring data sets of this volume.

We intend to study the impact of delays and losses on in−dividual TCP connections, which will be the subject of another paper.

I. Conclusion

We find that the shortage of bandwidth is a major reason for increased delays. First, insufficient supply of band−width causes queuing delays at network devices. The worst case delay observed at the Auckland access router is equivalent to typical one way delays observed between New Zealand and the US West Coast. About 10% of all packets experience delays unacceptable for any real−time communication, such as interactive audio and video. Second, limited peak data rates add to the per hop delay due to packet deserialisation times.

At several points within this paper we have identified typical packet sizes and their distributions as an impor−tant factor for the delay patterns observed. However, we find the traffic patterns by themselves insufficient to fully describe the observed packet delay and loss figures and we conclude that there is a router specific component which cannot be accurately predicted.

The analysis carried out in this paper is open to third party verification. We publish all the tools used to pro−duce the data plots. We make the complete data set avail−able to the public, either by requesting a tape copy or via download from the NLANR PMA web server [6].

J. Acknowledgements

We would like to thank Nevil Brownlee and the University of Auckland ITSS department for the continuing support of the measurement point and our research.

A portion of Jörg's time is funded via the National Science Foundation's NLANR MOAT Cooperative Agreement No ANI–9807479.

K. References

[1] Jörg Micheel, Ian Graham and Nevil Brownlee,
The Auckland data set: an access link observed
Proceedings of the 14th ITC Specialists Seminar on Access Networks and Systems Barcelona/Gerona, Catalonia, Spain, April, 25–27th, 2001

[2] Auckland–VI trace data – illustrated
http://wand.cs.waikato.ac.nz/wand/wits/auck/6/

[3] The Dag project, http://dag.cs.waikato.ac.nz/

[4] Micheel, J., Donnelly, S. and I. Graham. Precision Time–stamping for Network Packets. ACM SIGCOMM Internet Measurement Workshop, San Francisco, California, Nov. 2001

[5] Donnelly, S., High Precision Timing in Passive Measurements of Data Networks, PhD thesis, The University of Waikato, November 2001

[6] Auckland–VI trace data download
http://pma.nlanr.net/Traces/long/auck6.html

[7] Graham, I., Micheel, J., Chung, S. J., and Sharp, G., Duplicated Packets in an IP Trace. TERENA Networking Conference, Limerick, Ireland, June 3–6 2002, short paper, work in progress