

Evaluating the Accuracy of Active Measurement of Delay and Loss in Packet Networks

John A. Schormans¹ and Tijana Timotijevic²

¹Department of Electronic Engineering
Queen Mary, University of London, London E1 4NS, UK,
²Technology Development - Network Systems, Vodafone UK.
John.schormans@elec.qmul.ac.uk ,
Tijana.Timotijevic@VF.Vodafone.co.uk

Abstract. Performance measurement of packet (e.g. IP) networks is a vital element in the commercial viability of broadband. Active measurement by injection of probing packets will be very widely used as a method of performance measurement. In this paper we quantify the error when using probing. We discover that, when measuring the mean packet delay across a WAN through a representatively loaded (i.e. 50%-90% utilised) access link, the measurements often have an error of many tens, hundreds or even thousands of milliseconds. Furthermore, we apply results from queueing analysis to show that probing for packet loss will require that the probes be about the same size as the data packets; if small packets are used the measured packet loss probability will be many orders of magnitude smaller than it actually is. When this is accounted for, i.e. by using probing packets the same size as the data packets, then, for constant probing load (i.e. a smaller number of larger packets), the error in the returned delay measurements becomes considerably worse.

1 Introduction

IP and MPLS packet networks are now carrying a heterogeneous mix of traffic, with widely differing Quality of Service (QoS) requirements. The service model of emerging multiservice packet networks, including the packet backbones for 2.5G and 3G mobile networks, is based on the network's ability to guarantee QoS to user applications. In a commercially competitive environment there are two main reasons for performance measurement: a) Traffic Engineering will rely on the accuracy of the performance monitoring capabilities and b) Network Operators and Service Providers must monitor performance to maintain their Service Level Agreements (SLAs) with customers, [1, 2].

Active probing is a technology whereby 'probing' packets are injected into the network, and travel from source to destination along with the packets containing the user information, allowing the delays across the network to be measured. In order to measure in a way that is representative of the users experience during periods when most users are active (i.e. busy hours) we adopt, in this paper, the recent recommendation from Cisco Systems [3] that measurements are carried out over busy

hours. In conjunction with the probing rate this gives us the number of measurement probes. We report results using 2 sizes of probe packet: small (100bytes), in order to reduce the measurement load while probing delay, and large packet probes to ensure that the monitoring probes experience the same loss probability as the data traffic being monitored. For a fuller justification of using representatively ‘large’ packets, see Section 4.

IP networks are evolving in a manner that is essentially heterogeneous – there is no single global network covering the world, rather an interconnected collection of different networks with different owners, and this is itself acting to encourage the use of active network probing. In this paper we address specifically the accuracy that should be expected from using active measurements. Why is this an issue? Any measured probe delay is an experiment whereby the network ‘path’ is tested. In order to know even the mean end to end delay many probes must be used, and the related questions are: 1) how many probes are needed for a desired level of accuracy, and 2) given N probes, what accuracy should be expected (and is it good enough?).

It is known from sampling theory that the greater the variability of the sampled data the more samples are needed for accurate estimation. In the case of actively probing packet networks the variability is dependent on two main factors: the load on the network, and the type of traffic being carried. Highly bursty packet network traffic results in very large variances associated with the number of packets in queues, and hence the packet delays. This is critically important, as using packet networks to carry this sort of traffic implies that a very large number of probes may be used to achieve the required levels of accuracy.

Previous work in assessing the effectiveness of active probing is limited, in [4, 2] to the technological requirements, and in [5] to issues associated with the types of sampling that are available (random, stratified, deterministic etc) and their particular advantages and disadvantages. A study is reported, [6], which addresses the question of whether probing is accurately measuring packet delays. However, this was done by simulation, so the results are necessarily confined to simulations of small numbers of traffic sources in order that the simulations reach steady state. However where there are clear conclusions they tend to support ours.

So this paper is arranged as follows. In Section 2 we combine sampling and queueing theory to find the solution for the variance of the queue depth in a typical buffer at an access link to a WAN. Once the variance is known the number of sampling (probe) packets can be calculated for any utilisation of the access link. In Section 3 we apply these solutions to typical traffic multiplexes at typical access link loads (which will probably be about 70%→80% during peak hours). We use Section 4 to justify why we are concerned to test both short probing packets and longer ones, i.e. probes that are about the same size as the informational packets whose performance they are measuring. In the final section we conclude, attempting to draw together the most significant results for the practice of packet traffic measurement.

2 Analysis for Generic Queueing Scenarios

IP networks are evolving in a manner that is essentially heterogeneous, which is acting to encourage the use of active probing. And network heterogeneity is matched

by traffic heterogeneity: the growth in user applications from VoIP and picture messaging to file transfer – both real time and non-real time – all cause different patterns of traffic to appear in packet networks. However, despite this apparent complexity queueing theory has shown that there is a generic envelope of queueing which has been found to be ubiquitous: that is packet scale and burst scale queueing [7, 8, 9] (at this point we ignore the issues specifically associated with self-similar traffic that strays from this envelope, touching on it in section 2.4).

2.1 Markovian Traffic Buffering: Combining Packet and Burst Scale Queueing

Packet delay is caused by waiting behind other packets in buffers, and packet loss occurs when an arriving packet finds a full buffer. So our focus in this paper is to use mathematical analysis to find the required number of samples (probes) in scenarios in which we know, or can calculate, the probabilities associated with a particular number of packets in a buffer. The known scenario we use is Markovian queueing, i.e. combined packet scale and burst scale queueing [7, 9, 8], which is shown diagrammatically in Figure 1.




Fig. 1. Schematic representation of packet scale and burst scale queueing

In real networks probing will measure delays directly. However, as delays are the result of queueing behind other packets, we must evaluate the number of probes required to estimate with sufficient accuracy the mean of the number of packets in a buffer when a probe arrives.

The previously unexplored aspect of active measurement is whether even the mean of these distributions can be estimated accurately given that they may have very high variance. To quantify this problem in terms of probing accuracy we need an expression for the variance of the queue length in the presence of typical packet network traffic. In section 2.2 we derive the expression we need. In section 2.3 we discuss what is likely to be the result of going beyond probing the mean of the delay

to probing the jitter. In section 2.4 we discuss the potential problems when the traffic is not Markovian but self-similar.

2.2 Variance of the Measured Buffer Queue Level, Markovian Traffic

Combined packet scale and burst scale queueing is a function as follows, [7, 9, 8]:

$$p(k) = (1 - P_{BS}) \cdot P_P(k) + P_{BS} \cdot P_B(k) . \quad (1)$$

where:

$p_P(k)$ = Prob(arriving packet encounters queue length of 'k' packets conditional on packet scale queueing)

$p_B(k)$ = Prob(arriving packet encounters queue length of 'k' packets conditional on burst scale queueing)

P_{BS} = Prob(the buffer is experiencing burst scale queueing)

This can be expressed as:

$$p(k) = (1 - P_{BS}) \cdot (1 - \eta_P) \cdot \eta_P^k + P_{BS} \cdot (1 - \eta_B) \cdot \eta_B^k . \quad (2)$$

where :

ρ = the load on the buffer

η_B = the decay rate in the burst scale part of the queueing distribution;

η_P = decay rate in packet scale queueing part of the distribution

($\eta_P \approx \rho$, see also [8, 9] for more accurate expressions)

Define:

N = the number of probe packets (measurements) available

$t_{N-1, 1-\alpha/2}$ = the Student t-distribution value for $N-1$ degrees of freedom

(i.e. sample size = N), and $(100-\alpha)\%$ confidence interval for the estimate

of

the mean queue size (L)

δ = the standard error in the measurements

L_N = mean number of packets in a queue delaying an arriving probing packet, estimated from N measurements

S_N = standard deviation of the measurements

S_N^2 = variance of the measurements

In order to establish a relationship between the measurement accuracy and the number of probes (samples) needed for that accuracy, we will use the Student t-distribution. This will allow us to relate the standard deviation of the queue length, the accuracy of the measurement and the number of probes. Then, once we establish the relationship between the measurement accuracy, number of probes and the standard deviation, we can use the results from queueing theory to find the queue length variance (and hence standard deviation) for specific queueing scenarios. Queueing scenarios are described by both scheduling discipline and a specific traffic arrival

pattern. We therefore use specific traffic arrival characteristics to find the queue length variance in these well known scenarios, and therefore the number of probes (i.e. measurement load) needed to estimate the variance with the desired accuracy.

For the number of probes needed to estimate the mean L_N , the $100(1-\alpha)\%$ confidence interval for the true mean is given by:

$$L_N \pm t_{N-1, 1-\frac{\alpha}{2}} \cdot \frac{S_N}{\sqrt{N}} . \quad (3)$$

For the sampled mean L_N to be within the error $\pm\delta$ of the true mean, N must be such that:

$$t_{N-1, 1-\frac{\alpha}{2}} \cdot \frac{S_N}{\sqrt{N}} \leq \delta . \quad (4)$$

Equations (3) and (4) follow from sampling theory. Now we must use queueing theory to find a formula for the variance in the measured queue length, from which we can establish the number of probes needed per hour for a given level of accuracy (or the accuracy achieved from a given number of probes per hour).

As noted already, we have to imitate the packet traffic in a realistic way, and this means Markovian models with a packet scale and burst scale queueing envelope (see figure 1). In this case we need equation (4), but must account for the fact that the unconditional mean is the weighted sum of both the packet scale and the burst scale:

L = mean queue length for combined effect of packet and burst scale queueing

$$L = (1 - P_{BS}) \cdot L_p + P_{BS} \cdot L_B . \quad (5)$$

In (5) L_p and L_B are the mean number of packets delaying an arriving packet conditional on packet scale queueing and on burst scale queueing respectively. We now need to move from an expression for the mean to an expression for the variance, and thereby the standard deviation so that we can apply the results from sampling theory.

If 'q' is the random variable representing the size of the packet queue delaying the packet or probe:

$$S_N^2(q) = E(q^2) - (E(q))^2 . \quad (6)$$

$$E(q) = L = (1 - P_{BS}) \cdot L_p + P_{BS} \cdot L_B . \quad (7)$$

From (8) and (9), and the basic definitions of Moments, we find:

$$\begin{aligned} S_N^2 = & (1-P_{BS}) \cdot L_p^2 / \eta_p + P_{BS} \cdot L_B^2 / \eta_B + (1-P_{BS}) \cdot P_{BS} \cdot L_p^2 \\ & + (1-P_{BS}) \cdot P_{BS} \cdot L_B^2 - 2(1-P_{BS}) \cdot P_{BS} \cdot L_p \cdot L_B . \end{aligned} \quad (8)$$

In equation (8) we need to find P_{BS} , L_p , L_B , η_B and η_p in order to find the standard deviation, S_N . The mean number of packets in the queue conditional on burst scale queueing is [9]:

$$E[\text{number of packets} \mid \text{burst scale queueing}] = L_B = \eta_B / (1 - \eta_B) . \quad (9)$$

Convenient expressions for η_B , and P_{BS} are (given homogenous On-Off multiplexing), [7, 9]:

$$\eta_B = \exp[-(N_o / b)] \cdot (1 - \rho)^3 / (4\rho + 1) . \quad (10)$$

where:

b = average number of packets in an ON period of an individual source

N_o = number of active sources needed to have burst scale queueing, [9, 7].

P_{BS} can be found as:

$$P_{BS} = \frac{\exp[-(\rho \cdot N_o)] \cdot (\rho \cdot N_o)^{\lfloor N_o \rfloor}}{(1 - \rho)^2 \cdot N_o \cdot (\lfloor N_o \rfloor !)} . \quad (11)$$

Combining equations (8), (9), (10) and (11) we can find the variance and standard deviation of the number of packets in the queue, and substituting this into (4) we can find how this related to the required number of probes for the desired accuracy. We concentrate on the effect of the packet queue in the access buffer, as this is where the bulk of the packet queueing will take place.

2.3 Error When Measuring the Prob(delay > T) for Delay Variation (Jitter)

Prior work, e.g. [10], has shown the importance of keeping the probing overhead as low as possible. [11] shows that to measure the delay jitter will probably require orders of magnitude more probes per unit time than does measuring the mean. Therefore attempts to get delay jitter estimates from the number of probes assumed here implies errors considerably greater than would be experienced when measuring the mean delay.

2.4 Measuring Networks Carrying Self-Similar Traffic

Recent literature indicates that the presence of self-similar traffic in packet networks is now well appreciated by most of the networking community. Results, both from simulation and analysis, clearly show that where the input traffic has active periods that are power law distributed, the overall traffic 'pattern' is likely to be self-similar, and therefore the distributions associated with packet queueing are likely to be power law distributed too. Power law distributions may feature infinite variance, and this would cast doubt on the ability of active traffic measuring technologies ever to determine mean values within a desired level of accuracy. However, further work is

required here, as results recently reported on passive queue monitoring have shown some promise in resolving queue state probabilities, albeit without the analytical support specifically aimed at guaranteeing the statistical accuracy of the results [12].

3 Results – Application to Realistic Network Examples

Probing across a WAN involves passing the measurement packets through a series of buffers in routers end to end. However, it is now well understood that the buffers/routers within the WAN core will experience a very low utilisation compared to the network access lines, while the access lines will be highly utilised during busy hours, as they represent the bottlenecks. For this reason we concentrate our attention on these bottleneck access lines.

3.1 Buffering VoIP Traffic

In order to obtain some practical estimate of the accuracy we can expect when using probing, we predict measurement accuracy for 2 standard traffic models, against increasing load. The 1st model we adopt is the standard VoIP model (exponential On and Off periods, mean on time = 0.96 seconds, mean off time = 1.69 seconds), which is very well accepted and widely used. This gives the following parameter values:

- $b = 80 \text{ pkts per sec} * 0.96$ (average ON period in seconds, or activity factor) = 76.8
- $N_o = \text{capacity in pkts per sec} / 80 \text{ pkts per sec}$
- packets of length 100 bytes.

Figure 2 shows the value of the error, δ , on the measurements for 128kbps access lines. This shows that at least 512kbps, or 2Mbps should be used. It can be seen that, for probe packet injection rates corresponding to the use of 1%, 2% and 5% of the 128kbit/sec access channel capacity (i.e. non-trivial bandwidth available for measurements) the error gets very large as the load increases.

Naturally, given constant measurement bandwidth, i.e. 10 times fewer probes when 10 times larger probes are being used, this effect is much more marked when the 1000 byte probes are used. Figure 3 repeats this example for 34Mbit/sec access lines, and here significant error is apparent at high loads and with the larger packet size.

3.2 Traffic More ‘Bursty’ than VoIP

In this section we repeat the previous studies, but this time replace the VoIP traffic model with the burstier ‘data’ model, i.e. burstier than VoIP. To obtain such a model we use the following parameters:

- average source active (On) period = 0.1 seconds
- packet generation rate when in the On periods = 1250 packets /second
- $b = 1250 \text{ packets / second} * 0.1 = 125$
- $N_o = \text{Capacity (in packets per second)} \div b$
- Data packets = 1000 bytes in length

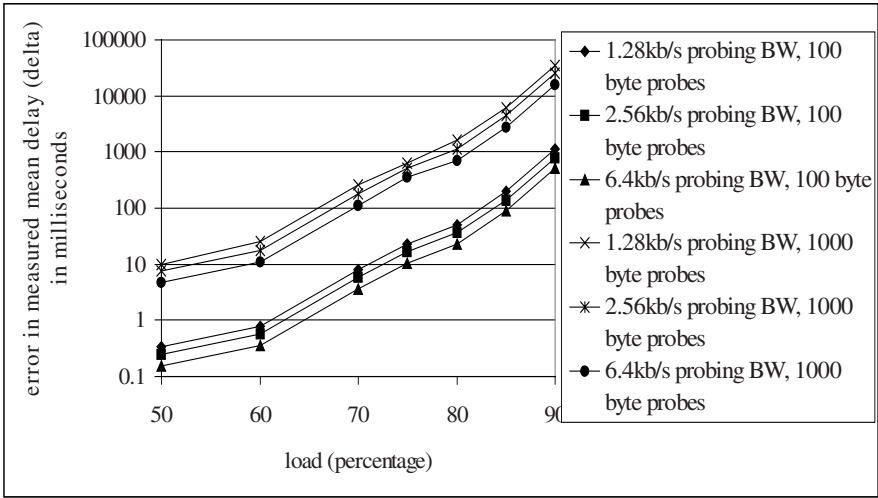


Fig. 2. Error (δ) in the probing measurements for VoIP access scenarios of 128kbit/sec with 100 byte and 1000 byte probing packets

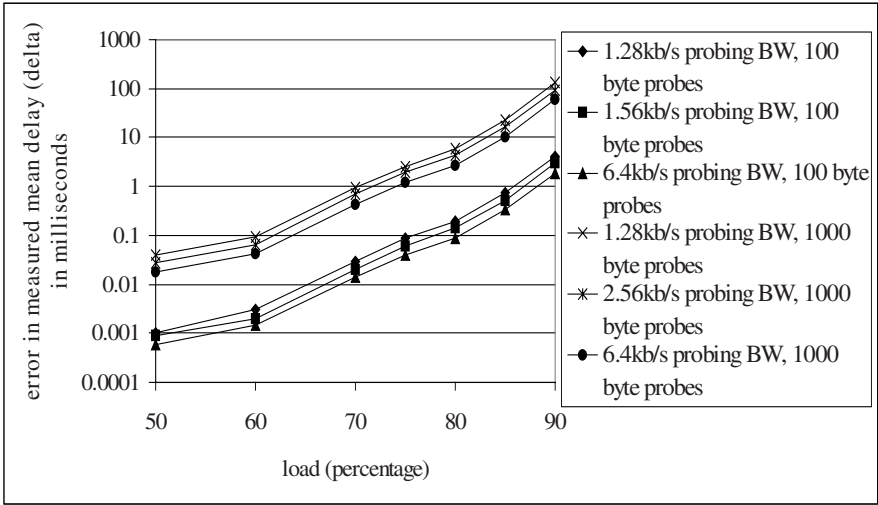


Fig. 3. Error (δ) in the probing measurements for VoIP access scenarios of 34Mbit/sec with 100 byte and 1000 byte probing packets

It can be seen that, for both cases we used (i.e. 128kbps and 34Mb/s access line rate), see Figures 4&5, the absolute error (δ) in the measurements is larger for the data packets than for the VoIP at all utilisations. This is intuitive: the burstier the traffic the more variable the queueing/delay distributions and hence the more probes needed to accurately estimate the mean delay.

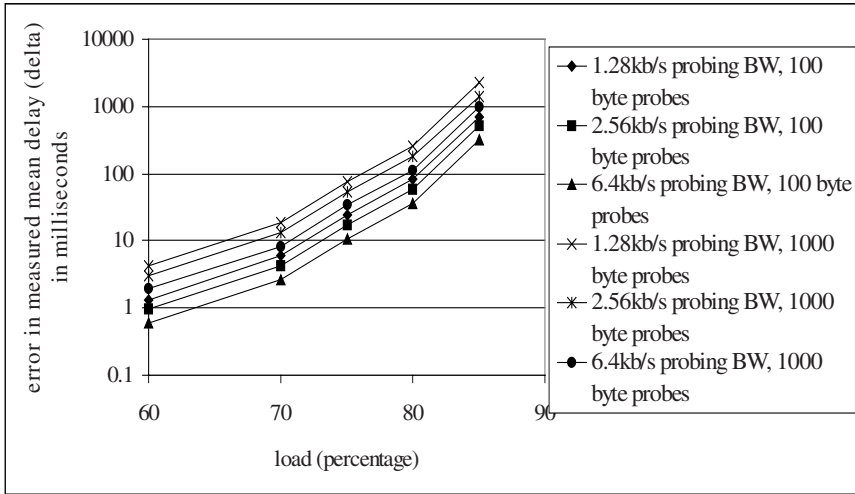


Fig. 4. Error (δ) in the probing measurements for generic ‘data’ access scenarios of 128kbit/sec with 100 byte and 1000 byte probing packets

It is important here to note that the probing bandwidths we have used are actually quite large: 1%, 2% and 5% of 128kbit/sec in all cases. Many organisations are using probing rates equivalent to 1 packet per second, and it is clear that at such low rates measurement accuracy will be significantly worse than predicted here.

3.3 Example Using a 3rd Party Video Trace

In this example we do not use a traffic model but instead use the results of a previous experiment [13] with a real traffic trace of a video sequence (based on the MPEG Star Wars video trace). The parameter values were: channel access link capacity = 17Mbps, video packet size = 424 bytes, and the (measured) utilisation ≈ 0.8 . From these results: $\eta = 0.99924$, and $P_{BS} = 0.01$, and therefore the unconditional variance of the queue length = 34359. (It should be noted that this figure for unconditional variance in queue length is broadly in line with the values generated by the traffic models we used in sections 3.1 and 3.2.) The error when using probing packets is given in Table 1:

Table 1. Error [milliseconds] in the measurements (δ) of mean delay (video trace example)

Probe size = Access link rate =	100 byte		1000 byte	
	128kbps	34Mbps	128kbps	34Mbps
Probe rate = 1.28 kb/s	$\delta=30$ [ms]	$\delta=0.1$ [ms]	$\delta=305$ [ms]	$\delta=1.1$ [ms]
Probe rate = 2.56 kb/s	$\delta=22$ [ms]	$\delta=0.08$ [ms]	$\delta=216$ [ms]	$\delta=0.8$ [ms]
Probe rate = 6.4kb/s	$\delta=14$ [ms]	$\delta=0.05$ [ms]	$\delta=136$ [ms]	$\delta=0.5$ [ms]

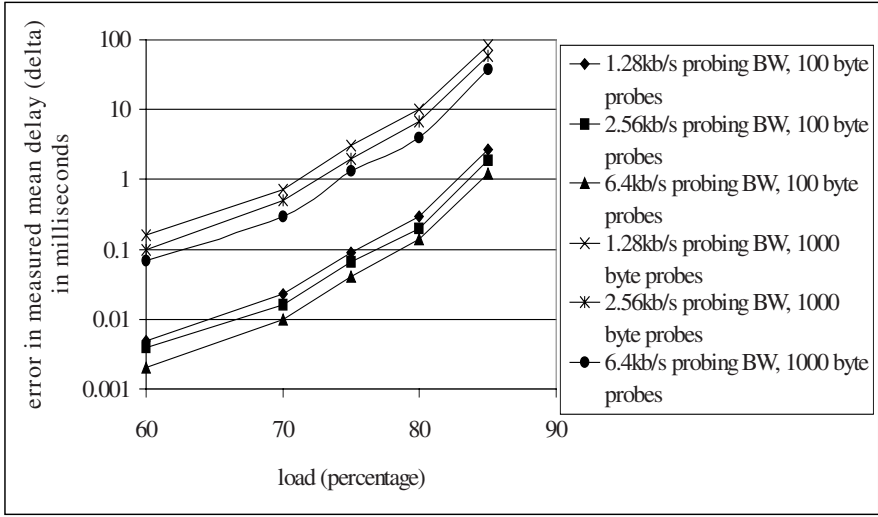


Fig. 5. Error (δ) in the probing measurements for generic 'data' access scenarios of 34Mbit/sec with 100 byte and 1000 byte probing packets

4 Accuracy When Probing Loss

It is important for the effectiveness of network measurement that the number of probes, and hence bandwidth overhead, is kept to a minimum. To do this the stream of probes that are being used to measure delay should also be used to measure loss probability. However this will require the use of probe packets that are of equivalent length to the data packets for which the packet loss probability is to be quantified. To see that this is the case consider loss resulting from overflow in a buffer which is multiplexing large packets, of L bytes, and small packets of S bytes. Perhaps the large packets represent TCP data while the small are TCP ACK packets (they could also be VoIP, however for QoS reasons it is unlikely that such dissimilar services would pass through the same buffers). The total capacity (in bytes) of the buffer is X , where $X = N_L \cdot L = N_S \cdot S$, i.e. N_S , N_L are the number of (respectively) small and large packets that the buffer can hold in the absence of the other type. Clearly once the buffer contains at minimum $((N_L - 1) \cdot L + S)$ bytes it is 'full' as far as the large packets are concerned, and further arrivals of large packets will be lost. However the small packets can still be accepted: these will not be lost until the buffer contains the equivalent in bytes of N_S small packets, i.e. $N_S \cdot S$ bytes. Therefore if the probes are small packets they will not necessarily be lost when large packets are.

So it is easy to see that smaller probe packets cannot accurately be used to measure the loss probability for larger packets. While it is easy to see this intuitively, it is important to discover how significant this inaccuracy will be. To do this, we use the method given in [14], as this allows us to evaluate the ratio:

$$R = \frac{[\text{measured packet loss probability found by using small probes}]}{[\text{Actual packet loss probability}]}$$

We use the results of an analysis of multiplexing large and small packets through a classical M/G/1 queue model with 2 packet sizes, small and large. In order to perform these calculations we have had to make an assumption about the relationship between X, L and S, so for simplicity, and in the absence of further information, we have assumed that the extra space, on average, available to the small packets is $= \lfloor L/2 \rfloor$, which is $\frac{1}{2}$ of the length of a large packet.

Table 2. The value of the error ratio ‘R’ for varying proportions of small packets, with small packets also used as probes

Total load (%)	10% small packets	20% small packets	50% small packets
50	R = 6.00E-10	R = 9.22E-09	R = 3.20E-07
60	R = 1.30E-09	R = 1.87E-08	R = 7.10E-07
70	R = 2.20E-09	R = 3.52E-08	R = 1.23E-06
80	R = 3.90E-09	R = 5.90E-08	R = 2.09E-06
90	R = 6.30E-09	R = 9.60E-08	R = 3.35E-06

These results for the ratio ‘R’ are given in Table 2. Percentage “small packets” refers to the TOTAL traffic carried by small packets. These results show that the measured packet loss probability, in this standard buffer model, will be between 10^{-6} and 10^{-10} SMALLER than the actual packet loss probability for large packets. Clearly measuring packet loss probabilities as being smaller by such large amounts will be a very significant problem. This is the justification for presenting results from Section 3 that use 1000 byte probing packets as well as 100 byte packets.

5 Conclusions

In this paper we have shown that the standard error in the measurements of mean packet delay, for traffic with VoIP type characteristics, or similar bursty data, will be significant: at non-trivial but representative access link load the measurement error may reach many hundreds or even thousands of milliseconds. The effect of this may be to invalidate any SLA based on guaranteeing mean delays within reasonably tight limits.

It is important to note that the probing bandwidths we have used are actually quite large: 1%, 2% and 5% of 128kbps in all cases. Furthermore, when using packet sizes that accurately represent the data packet size, and therefore the packet loss probability of the data traffic, the measurement inaccuracy may be considerable worse (for constant measurement bandwidth).

These results pertain only to measuring the mean delay. Also of clear significance is the achievable accuracy when measuring the delay jitter and packet loss probabilities. It is clear from the work in this paper, and others [10, 11], that either a) much more bandwidth needs to be made available these for these measurements, or b) inaccuracy so severe as to invalidate them must be accepted when attempting their measurement.

In QoS oriented networks there are two more reasons for suspecting that measurement accuracy will be worse than presented here: a) queueing in the WAN will add to the effect of the bottleneck queueing and b) traffic will be divided into (probably about 4) Classes of Service, the effect of which will be to subdivide the available probing bandwidth, meaning a smaller number of measurements is actually available for each calculated average in each CoS class. If, in addition to this, SLAs are to be guaranteed on a per packet size, or per application type basis, then the problems of inaccuracy raised in this paper will become yet more severe.

References

1. Verma, D: "Supporting Service Level Agreements on IP Networks". MacMillan Technical Publishing, ISBN 1-57870-146-5.
2. Cisco Systems – Service Assurance Agent.
www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/%120t5/saaoper.htm, 2000.
3. Service-Level Management: Defining and Monitoring Service Levels in the Enterprise. Cisco White Paper, 2003, available from:
http://www.cisco.com/en/US/products/sw/cscowork/ps2428/products_white_paper09186a0080091ba5.shtml
4. Nevil Brownlee and Chris Loosley: "Fundamentals of Internet Measurement: a tutorial", CMG Journal of Computer Resource Management, Issue 102, 2001. Also available at [www.keynote.com/solutions/assets/applets/Fundamentals of Internet Measurement A Tutorial.pdf](http://www.keynote.com/solutions/assets/applets/Fundamentals_of_Internet_Measurement_A_Tutorial.pdf)
5. Claffy, K.C., Polyzos, G.C., and Braun, H-W: "Application of sampling methodologies to network traffic characterisation". SIGCOMM 1993, 194–203, and also from: citeseer.nj.nec.com/claffy93application.html
6. Hill, J.: "Assessing the Accuracy of Active Probes for Determining Network Delay, Jitter and Loss". MSc Thesis in High Performance Computing, The University of Edinburgh, 2002.
7. Roberts, J. (Editor): "Performance Evaluation and Design of Multiservice Networks". COST224 Final Report, 1991.
8. Schormans, J.A. and Pitts, J.M.: "From Erlangs to Excess Rate", Journal of the IBTE Oct-Dec 2001.
9. Pitts, J.M. and Schormans, J.A.: "Introduction to IP and ATM Design and Performance", Wiley, 2000.
10. Timotijevic, T. and Schormans J.A.: "Bandwidth overhead of probe technology guaranteeing QoS in packet networks". To appear in Electronics Letters, 2003.
11. Schormans, J.A.: "Bandwidth overhead of packet probing". Queen Mary internal report, available from john.schormans@elec.qmul.ac.uk
12. Leung, C.M., Schormans, J.A. and Ma, A.H.I.: "Measurement-based queue length distribution estimation for power law traffic". Electronics Letters, 21st November, 2002, Vol. 38, No. 24, pp. 1608–1610.

13. Choe, J. and Shroff, N.B.: "Queueing Analysis of High-Speed Multiplexers including Long-Range Dependent Arrival Processes". IEEE INFOCOM'99, New York, NY, March 1999, pp. 617–624.
14. Pitts, J.M., Schormans, J.A., Phillips, C.I. and Griffiths, J.M.: "Accurate Delay Bounds For Real-Time IP Services in Presence of Variable Packet Size". Electronics Letters, 1st February 2001, Vol.37, No. 3, pp. 166–167.