

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/224500675>

Measurement of packet loss probability by optimal design of packet probing experiments

Article in IET Communications · July 2009

Impact Factor: 0.74 · DOI: 10.1049/iet-com.2008.0075 · Source: IEEE Xplore

CITATIONS

9

READS

51

3 authors:



[Ben M. Parker](#)

University of Southampton

10 PUBLICATIONS 15 CITATIONS

[SEE PROFILE](#)



[Steven G Gilmour](#)

University of Southampton

82 PUBLICATIONS 1,026 CITATIONS

[SEE PROFILE](#)



[John A. Schormans](#)

Queen Mary, University of London

55 PUBLICATIONS 232 CITATIONS

[SEE PROFILE](#)

Measurement of packet loss probability by optimal design of packet probing experiments

Ben M. Parker

Steven G. Gilmour

John Schormans

Queen Mary, University of London

July 31, 2008

Abstract

Packet level measurement is now critical to many aspects of broadband networking, e.g. for guaranteeing Service Level Agreements, facilitating measurement-based admission control algorithms, and performing network tomography. Because it is often impossible to measure all the data passing across a network, the most widely used method of measurement works by injecting probe packets. The probes provide samples of the packet loss and delay, and from these samples the loss and delay performance of the traffic as a whole can be deduced. However measuring performance like this is prone to errors. Recent work has shown that some of these errors are minimized by using a gamma renewal process as the optimal pattern for the time instants at which to inject probes. This leaves the best rate at which to inject probes as the key unsolved problem, and we address this here by using the statistical principles of the Design of Experiments. The experimental design approach allows us

to treat packet level measurements as numerical experiments that can be designed optimally. Modelling the overflow of buffers as a 2-state Markov chain, we deduce the system’s likelihood function, and from this we develop a technique (using the Fisher information matrix) to determine the upper-bound on the optimal rate of probing. A generalization of this method accounts for the effect of the probed observations interfering with the experiment. Our numerical results focus on VoIP traffic, allowing us to show how this methodology would be used in practice. One application of this is in measurement-based admission control algorithms, where our technique can be used to provide an upper-bound on the rate at which probes should be injected to monitor the loss performance of the target network, prior to making an admit / don’t admit decision.

1 Introduction

Simulation and analysis are routinely used to aid the design of broadband packet networks. However, when in operation packet level measurement is critical to ensure that the desired performance metrics are being achieved. Examples of the uses of packet level measurement include guaranteeing Service Level Agreements (SLAs), and facilitating measurement based admission control[\[13\]](#) and network tomography (including end-to-end measurements to infer specific resource congestion [\[19\]](#)).

Many UK measurement research projects have focused on the technological aspects of measurement. UKLight/MASTs, 46PaQ and ESLEA looked at monitoring in IP WAN packet networks and concentrated on the challenges arising from the scale of the technology itself: bandwidth in Gbps, large networks and storage space, and complex inter-working of protocols and algorithms. Considerable recent work also concentrated on developing measurement techniques, e.g. [\[15\]](#), [\[23\]](#), and [\[14\]](#). However, much of the basic technology is now ma-

ture and is described in textbooks, e.g. [5], [4]. The research focus has now shifted to the evaluation and the performance of measurement techniques, and the optimisation of their operation.

The most common technique for monitoring packet level performance uses the injection of probe packets. Probing packet-level measurements, often called active measurement, infers the overall traffic performance from the loss and delay of the probe packets. This involves the sampling of stochastic processes, where the sampling is constrained by the resources available (primarily bandwidth). Whatever the precise nature of the algorithm:

- sending packets into an already congested network increases the congestion;
- any measurement may be altered by performing it, [17], [21], and it is hard to know in advance at what rate and at what times the probe packets should be injected into the network [18];
- the accuracy of the overall performance inferred is subject to significant errors [9].

There is a growing awareness of the significance of these errors, e.g. see the presentation by Filsfils, [7], (a Cisco Systems engineer) at a Cisco hosted symposium on Measuring Internet Quality. Filsfils reports that the magnitude of the measurement error is very poorly understood, and that simple approaches to improving it are still focused on straightforward adaptations of existing ideas. Despite a wide variety of measurement tools having been reported in the literature, "...there has been very little analysis of the accuracy of these tools or their impact on the network ..." (quoted from [24] SIGCOMM05).

Related to the problem addressed in this paper, and widely examined in the literature, are the challenges associated with accurate flow-sampling. In

[10] the authors develop a flow-sampling method (concentrating on stationary queue models) for the distribution of the number of packets per flow, and the spectral density of the packet arrival process. Earlier related work can be found in [6].

More closely related to the work we present here is research on measuring packet delays via probing, and some papers report the level of accuracy achieved using probing, e.g. in [22] and [25]. At the ACM SIGMETRICS'05, Roughan [17] investigated the inaccuracies caused by correlations between successive probes and derived fundamental limitations. A series of recent papers by Schormans et al., has also explored the potential inaccuracy inherent in packet level measurement, [22], [26], [25], [9], [12], [20], [9]. It has been discovered that even for simple buffering scenarios there are practical load limits beyond which measurement accuracy degrades very rapidly. These load limits are likely to be within the normal operating specification of packet networks, e.g. 70% load on a VoIP access link. Results show that significant levels of expected absolute error are found in sampled results, even for a very simple network. For example, mean delays measured in a VoIP access link have been found to exhibit an absolute error that can rise to hundreds of msec when a link at around 80% utilisation is probed over a busy hour. It has been found that absolute error in measured mean delay increases with the number of end-to-end links, load and traffic burstiness, and is generally (approximately) inversely proportional to bandwidth.

Compared with measuring mean delays, it is much harder to measure packet loss probability or delay jitter, as these are associated with “rare events”, e.g. the queue level reaching the limit (in the case of loss). Prior work has evaluated the packet loss behaviour over the Internet, e.g. [24] and [16], and in [27] probed measurements were used to establish packet loss characteristics in wide area

networks via three different aspects of loss rate. In [24] the authors reported that simply increasing the packet probing rate is an unsatisfactory way to attempt to increase the accuracy of measured results, as the probe stream will interfere with the data traffic, skewing the results (this was also found earlier in [21]). In [22] queueing analysis was used to show that, for a tail-drop queue, probes would need to be of very similar length to the user packets or else the measured loss probabilities could easily be in error by many orders of magnitude. One of the significant results of the work reported in [9] is that the burstiness of the loss process also has a critical effect, causing any measurement process to require far more samples than would be expected from a naive calculation based only on the target loss probability. In this paper we focus on loss probing as there has already been significantly more research on probing delays (mean delays in particular).

In [9] a formula was developed that gives the number of probes required to achieve a given (pre-chosen) level of accuracy when measuring the packet loss probability. While [9] evaluated the number of probe packets required, the best pattern by which to inject these probes was investigated in [2] and [3]. The key discovery was that traditional Poisson sampling is not always optimal, and a more general class of distributions (gamma renewal processes) minimizes the mean square error in the sampled estimator. So, together [9], [2] and [3] provide the number of probes required and the best pattern for probe injection. This leaves the optimal probing rate as a salient unsolved question, and in this paper we address this.

Uniquely, we address the problems of probing the packet loss probability by using the statistical principles of the Design of Experiments (DOE). DOE allows us to approach measurement as a numerical experiment that can be designed optimally. Following [9], we model the network essentially by concentrating on

a low bandwidth access link; see Figure 1. This ignores the core network where loss is generally much lower.

[Figure 1 about here.]

Following [9], we also model the buffer as a 2 state Markov chain. In section 2 we demonstrate how to analyze the overflow state of a buffer using this Markov chain approach, deduce the likelihood function, and hence the Fisher information matrix. In section 3 we present results which show the best design for the experiment for different values of input parameters, and demonstrate the usefulness of this approach in section 4 with a VoIP example. The main value of our analysis is to evaluate the best probing rate. When this ideal rate is combined with the required number of samples, it gives the total time required to probe results with the required level of accuracy.

In section 5 we show what happens when we take into account the effect of active probing (observations interfering with the experiment). Finally, we conclude in section 6 and present some ideas for extensions to the work.

2 Methodology and Derivation

During periods of significant load, packet queues in buffers will grow, and decline, in size. Eventually a long enough period of growth will result in packet loss due to buffer overflow. A generic period of such behaviour is represented in Figure 2. A fully detailed representation of such a queue evolution would require a finite, ergodic Markov chain with $(N + 1)$ states. In such a model, state s represents s packets in the buffer, and N would be the maximum number of packets the buffer can hold ($N + 1$ such states are required in the Markov chain as there is a state 0 representing the buffer being empty). This is shown in Figure 3. In practice this would amount to a very large number of states

around which to design probing experiments optimally.

[Figure 2 about here.]

[Figure 3 about here.]

However, earlier work, [9], has shown that this full Markov chain representation can be reduced to a much simpler 2-state model. In the 2-state model we either have spare capacity in the buffer and are therefore in the non-overflow state (state 0), or we have no spare buffer capacity and are in the overflow state (state 1). This is also shown in Figure 3. Figure 4 maps the 2-state Markov chain model against the evolution of the queue state against time.

[Figure 4 about here.]

The relationship of the work in this paper to the earlier work in [9] is as follows. In [9] the authors measured the effectiveness of packet probing for loss probability through the analysis of the variance associated with the packet loss process. This allowed the prediction of the number of probes by using sampling theory (which relates the number of samples required to the variance). So the analysis in [9] was not able to address the question of the rate at which the probes should be injected; only the required number of probes was determined. In contrast the work reported in this paper is aimed at determining an upper bound on the optimal probing rate. This is particularly important in cases where the loss probability needs to be determined in near real-time, i.e. ideally by injecting probes as quickly as possible, for example in certain measurement based admission control systems.

Figure 5 illustrates the final step in the modelling process and shows the sequence of probes against the simplified representation of the buffer (as a 2-state Markov chain). For simplicity we use a discrete time base; however, since

the size of a time slot can be arbitrarily small, this approach can be as precise as required.

[Figure 5 about here.]

The 2-state Markov chain model has transition matrix

$$P = \begin{pmatrix} 1-p & p \\ r & 1-r \end{pmatrix}.$$

As shown in figure 4, p is the rate at which the buffer moves into the overflow state (state 1 in the Markov chain model), and r is the rate at which the buffer moves out of the overflow state, back into state 0 in the Markov chain model.

The stationary distribution such that $\pi P = \pi$ is $\pi = \begin{pmatrix} \frac{r}{p+r} \\ \frac{p}{p+r} \end{pmatrix}$.

We assume that we can find an estimator of the unknown $\theta = \begin{pmatrix} p \\ r \end{pmatrix}$, which is at least asymptotically unbiased and of minimum variance. We observe the chain every k -th period, for integer k ; so for $k = 1$ we observe every period, $k = 2$ every second period, and so on.

We observe the chain at time 0, and we allow a further fixed number of observations, T , of the chain (at times $k, 2k, 3k, \dots, Tk$), so we have $T + 1$ observations in all; we thus obtain an observation \mathbf{y} for the chain, where \mathbf{y} is a vector of length $T + 1$. We take the likelihood function $L(\theta|\mathbf{y})$, differentiate it twice and take the expectation to get the Fisher information matrix¹. The derivation is shown in Appendix A. We find that the Fisher information matrix is

¹We here use the Fisher information matrix, rather than any other measure of information, for example Shannon information. The two forms can be related, see e.g. [11]

$$I(p, r) = \frac{1}{p+r} \begin{pmatrix} \frac{1}{p} - \frac{1}{(p+r)} + \frac{Tr}{p} + \frac{Tr}{1-p} & \frac{-1}{(p+r)} \\ \frac{-1}{(p+r)} & \frac{1}{r} - \frac{1}{(p+r)} + \frac{Tp}{r} + \frac{Tp}{1-r} \end{pmatrix}.$$

We can then find the Cramér-Rao Lower Bound (CRLB)² from the inverse of the Fisher information Matrix, $I^{-1}(p, r)$. Now, we wish to compare the information in the experiment in which we look at every realisation of the Markov chain ($k = 1$) with that in the experiment in which we look at every second realisation of the chain ($k = 2$).

We observe that the 2-step transition matrix for the chain is:

$$P(2) = P^2 = \begin{pmatrix} (1-p)^2 + pr & (1-p)p + p(1-r) \\ r(1-p) + (1-r)r & rp + (1-r)^2 \end{pmatrix}.$$

We can reparameterize this as

$$P(2) = \begin{pmatrix} (1-p^*) & p^* \\ r^* & (1-r^*) \end{pmatrix}$$

where $p^* = (1-p)p + p(1-r)$ and $r^* = r(1-p) + (1-r)r$. We can then form the Fisher information matrix by following the method in Appendix A and compare $I_{k=2}(p, r)$ with $I_{k=1}(p, r)$.

This can be extended to $k > 2$ to give the k -step chain transition probabilities: the matrix of probabilities that the network will change between overflow and non-overflow states in k time periods.

The engineering value of this is as follows: the value of k we determine in this way is the best value of the interval between packet probes (measured in

²The CRLB of any scalar function of θ is a lower bound for the variance for any unbiased estimator of that function of θ .

integer steps of time); i.e. it is the reciprocal of the optimal probing rate. When we apply this approach to a realistic scenario (VoIP in the chosen example in this paper) it becomes clear that in many practical situations optimality can equally be achieved by setting k larger than a value determined by this approach. This means that other issues, e.g. the desire to minimize the probing rate, can determine the actual decision on the probing rate chosen in practice, while still maintaining a rate that optimizes the use of a fixed number of probes to accurately sample the packet loss probability.

3 Results

To best estimate the parameters p and r , and hence the overflow probability, we need to correctly choose our input parameter k , the time between our observations. We need to decide what criteria are important to us in deciding whether one k is better than another. Here we are estimating both of p and r , so we choose the D-criterion, the determinant of the information matrix. Maximizing the D-criterion allows us to maximise the generalised variance when estimating p and r together [1]. The k which produces the highest value of this D-criterion is our optimal k (an optimal design point). The same methodology could be applied to any other optimality criterion.

We present as Table 1, for $T = 10$ and different values of p and r , the ratio of determinants for the $k = 1$ case to the $k = 2$ case. A value greater than 1 implies that the $k = 1$ experiment is better, for that value of p and r and this criterion. At values of $p + r = 1$, denoted ‘-’ in the table, it is interesting to note that the D-criterion for the $k = 2$ case is zero; we get no information in performing the $k = 2$ experiment.

[Table 1 about here.]

Table 1 shows that for some values of p and r , the $k = 2$ case is better; so for a fixed number of samples, it may be better to sample the chain every other time period rather than every time.

We can continue this procedure, and, for any p, r, T , find the best value of k at which to perform the experiment³. We tabulate this as Table 2 for $T=5$.

[Table 2 about here.]

In Appendix B, we show that we can find effective estimators that get close to the CRLB for moderate sample sizes.

4 Example: VoIP

The optimal design depends on unknown values of p and r , and so prior knowledge must be used to be able to find an optimal design. These values are not naturally specified in the form we have described in section 2 for networks. We must derive or approximate them from information we do know. We present an example of how to obtain prior estimates of p and r , and apply the statistical theory of design of experiments to detecting overflow in a VoIP buffer.

Here we use previous work [9] which showed that the overflow periods can be well approximated as exponentially distributed periods, with mean number of packets $\mu_p = \frac{h^2 T_{on} A_p}{(C - A_p)^2}$, where A_p is the mean packet arrival rate and C is the capacity of the network, both in packets per second, and T_{on} is the mean duration of periods of activity of individual VoIP sources. We let N_{Cycle} be the number of packets which are generated in one overflow period and one non-overflow period together. We note that $N_{Cycle} \approx \frac{\mu_p}{PLP}$ where PLP is the proportion of lost packets in the buffer, and can be known from standards, or calculated (per [9]).

³Earlier work (e.g. [9]) has shown that the time between transitions may be very large, particularly with a large buffer. Hence, for example, probing through a whole busy hour may require T to be of the order of 10.

From the stationary distribution $\pi = \left(\frac{r}{p+r}, \frac{p}{p+r} \right)$ we can approximate

$$\frac{p}{p+r} N_{cycle} = \mu_p,$$

$$\frac{r}{p+r} N_{cycle} = N_{cycle} - \mu_p,$$

and we can find p and r by solving these simultaneous equations.

We assume throughout that we have $N = 100$ packet voice sources, producing packets of size 47 bytes at a rate of $h = 170$ packets per second when active, into a buffer of size $X = 100$ packets. A standard On-Off VoIP traffic source model has mean ON time of 0.96 seconds, and mean OFF time 1.69 seconds, giving an activity factor of 0.362. Other recently published models have slight variations to these. As the actual values are not significant to our analysis, we assume a simple starting point of $T_{on} = T_{off} = 6$ sec (so load is fixed at $\rho = 0.5$), and then vary the capacity, C .

We assume that our number of probes, T , is 300, and that we can probe at intervals of k seconds, $k = 1, 2, \dots^4$. We vary the bandwidth in order to find the optimal probing rate k and present the results as Table 3.

[Table 3 about here.]

As the bandwidth is increased so the load is reduced and the analysis of [9] indicates that the number of probes required goes down dramatically. However the analysis presented here indicates that the probing rate should also go down. This means that the naive conclusion of [9], that you can simply measure much faster at lower loads, is not correct: the optimal probing rate and the number of probes should be taken together to determine the total time required to measure

⁴We have assumed a minimum probing interval of 1 second as Cisco IPQoS probing rate is 1 probe per second

the packet loss probability accurately. The technique developed in this paper, taken together with the analysis provided in [9], yields the lower bound on the total time required to measure the loss probability

5 Observations interfering with the experiment

In a real network, we gather information on the state of the overflow by introducing probe packets into the network and noting whether or not they are received by a monitoring device within a pre-determined time. We assume that this probe procedure occurs without error, and the state indicated by each probe packet we introduce to the network is one experimental unit.

By performing the experiment, and introducing probe packets into the network, we introduce more data into the network, and the network is more likely to become congested and reach capacity (state 1); in other words we interfere with the experiment. In this section, we model this interference in a general way by assuming that the transition matrix of the Markov chain becomes

$$P_{\delta,\epsilon} = \begin{pmatrix} 1 - p - \delta & p + \delta \\ r - \epsilon & 1 - r + \epsilon \end{pmatrix}$$

in the time period after we perform the experiment, and as before $P_{0,0} = P$ in a time period after no experiment is performed. δ and ϵ are two non-negative constants particular to the network and the probing procedure used, which model the additional time spent in overflow (state 1) as a result of the introduction of packet probes.

As in the non-interfering case, we compare the D-criterion (the determinant of the information matrix) for varying k . We first form the two step transition

matrix of the Markov chain as:

$$P(2) = P_{\delta,\epsilon}P = \begin{pmatrix} (1-p-\delta)(1-p)+(p+\delta)r & (1-p-\delta)p+(p+\delta)(1-r) \\ (r-\epsilon)(1-p)+(1-r+\epsilon)r & (r-\epsilon)p+(1-r+\epsilon)(1-r) \end{pmatrix}$$

and more generally

$$P(k) = P_{\delta,\epsilon}P^{k-1}.$$

The determinant for the $k = 1$ case is now $I(p+\delta, r-\epsilon)$ and the determinant for the $k = 2$ case is now $I(p_\delta^*, r_\epsilon^*)$ where p_δ^* is the (1,2) element of $P(2)$ and r_ϵ^* is the (2,1) element of $P(2)$. We can take the determinants to calculate the D-criterion and use this to compare designs.

As an example, Table 4 shows the best value of k to take for $T = 5, \delta = 0.03, \epsilon = 0.01$. This table can be compared with the table of best k obtained without penalty for doing the experiment (Table 2). A difference in values in the table is observable, here even for small δ and ϵ ; for some values of p and r , e.g. $p = 0.55, r = 0.05$, the optimal probing rate changes when we account for interference.

[Table 4 about here.]

6 Conclusions

Accurate packet level measurements are critical to the proper functioning of broadband packet networks. Recent work has determined that the optimal pattern for probes is not as a Poisson process, but as a gamma renewal process. However the best rate at which to inject probes remained an unsolved problem, which we have addressed in this paper by using the principles of the statistical DOE. DOE has here allowed us to find a method to determine an upper bound on the optimal probing rate, and a generalization of this has accounted for the

effect of the probed observations interfering with the experiment.

Numerical results have been generated for a realistic networking scenario: a multiplex of standard VoIP traffic sources feeding a finite buffer in which the packet loss probability has to be determined by packet probing. A very clear pattern has emerged. As the bandwidth is increased so the load is reduced and the analysis of [9] indicates that the number of probes required goes down dramatically. However the analysis presented here indicates that the probing rate should also go down. This means that the naive conclusion of [9], that you can simply measure much faster at lower loads, is not correct: the optimal probing rate and the number of probes should be taken together to determine the total time required to measure the packet loss probability accurately. The technique developed in this paper, taken together with the analysis provided in [9], yields the lower bound on the total time required to measure the loss probability.

One potential application is in measurement based admission control, where our work can be used to provide an upper-bound on the rate at which probes should be injected to monitor the loss performance of the target network, prior to making an admit / don't admit decision.

References

- [1] A. Atkinson, A. Donev, and R. Tobias. *Optimum Experimental Designs, with SAS*. Oxford University Press, 2007.
- [2] F. Baccelli, S. Machiraju, D. Veitch, and J. Bolot. The role of PASTA in network measurement. *ACM SIGCOMM Computer Communication Review*, 36(4):231–242, 2006.

- [3] F. Baccelli, S. Machiraju, D. Veitch, and J. Bolot. On optimal probing for delay and loss measurement. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 291–302. ACM New York, NY, USA, 2007.
- [4] B. Claise and R. Wolter. *Network Management: Accounting and Performance Strategies*. Cisco Press, 2006.
- [5] M. Crovella and B. Krishnamurthy. *Internet Measurement: Infrastructure, Traffic and Applications*. John Wiley & Sons, Inc. New York, NY, USA, 2006.
- [6] N. Duffield, C. Lund, and M. Thorup. Estimating flow distributions from sampled flow statistics. *IEEE/ACM Transactions on Networking*, 13(5):933–946, 2005.
- [7] C. Filsfil and D. Ward. Measuring internet quality, <http://www.cisco.com/web/about/ac50/ac207/proceedings/MIQ.ppt>, 2006.
- [8] J. Gani. Some Theorems and Sufficiency Conditions for the Maximum-Likelihood Estimator of an Unknown Parameter in a Simple Markov Chain. *Biometrika*, 42(3/4):342–359, 1955.
- [9] M. Hasib, J. Schormans, and T. Timotijevic. Accuracy of packet loss monitoring over networked CPE. *Communications, IET*, 1(3):507–513, 2007.
- [10] N. Hohn and D. Veitch. Inverting Sampled Traffic. *IEEE/ACM Transactions on Networking*, 14(1):68–80, 2006.
- [11] S. Kullback. *Information Theory and Statistics*. Courier Dover Publications, 1997.

- [12] C. Leung and J. Schormans. Measurement-based end to end latency performance prediction for SLA verification. *The Journal of Systems & Software*, 74(3):243–254, 2005.
- [13] A. Moore. An implementation-based comparison of Measurement-Based Admission Control algorithms. *Journal of High Speed Networks*, 13(2):87–102, 2004.
- [14] A. Moore, J. Hall, C. Kreibich, E. Harris, and I. Pratt. Architecture of a Network Monitor. *Passive & Active Measurement Workshop 2003 (PAM2003)*, 2003.
- [15] A. Pasztor and D. Veitch. A precision infrastructure for active probing. In *Proceedings of Workshop on Passive and Active Networking*, 2001.
- [16] V. Paxson. End-to-end internet packet dynamics. *IEEE/ACM Transactions on Networking (TON)*, 7(3):277–292, 1999.
- [17] M. Roughan. Fundamental bounds on the accuracy of network performance measurements. *Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 253–264, 2005.
- [18] M. Roughan. A Comparison of Poisson and Uniform Sampling for Active Measurements. *Selected Areas in Communications, IEEE Journal on*, 24(12):2299–2312, 2006.
- [19] D. Rubenstein, J. Kurose, and D. Towsley. Detecting shared congestion of flows via end-to-end measurement. *IEEE/ACM Transactions on Networking*, 10(3):381–395, 2002.

- [20] J. Schormans and C. Leung. *Measurement for guaranteeing QoS in broadband multiservice networks*, volume ‘Mobile and wireless systems beyond 3G: managing new business opportunities’. 2005.
- [21] J. Schormans and J. Pitts. So you think you can measure IP QoS? *Telecommunications Quality of Services: The Business of Success, 2004. QoS 2004. IEE*, pages 151–155, 2004.
- [22] J. Schormans and T. Timotijevic. Evaluating the accuracy of active measurement of delay and loss in packet networks. In *6th IFIP/IEEE Conference on Management of Multimedia Networks and Services (MMNS)*, pages 409–421. Springer, 2003.
- [23] S. Shalunov, B. Teitelbaum, A. Karp, J. Boote, and M. Zekauskas. A One-way Active Measurement Protocol (OWAMP). *Internet Draft, May, 2003*.
- [24] J. Sommers, P. Barford, N. Duffield, and A. Ron. Improving accuracy in end-to-end packet loss measurement. *Proceedings of the 2005 conference on applications, technologies, architectures, and protocols for computer communications*, pages 157–168, 2005.
- [25] T. Timotijevic, C. Leung, and J. Schormans. Accuracy of measurement techniques supporting QoS in packet-based intranet and extranet VPNs: Superhighways technology and broadband VPN. *IEE proceedings. Communications*, 151(1):89–94, 2004.
- [26] T. Timotijevic and J. Schormans. Bandwidth overhead of probe technology guaranteeing QoS in packet networks. *IEE Electronics Letters*, 39(10):816–818, 2003.

- [27] Y. Zhang and N. Duffield. On the constancy of internet path properties. *Proceedings of the First ACM SIGCOMM Workshop on Internet Measurement*, pages 197–211, 2001.

A Obtaining the Information Matrix

We have a Markov chain with transition matrix,

$$P = \begin{pmatrix} 1-p & p \\ r & 1-r \end{pmatrix}$$

depending on unknown parameter $\theta = \begin{pmatrix} p \\ r \end{pmatrix}$. We assume that the chain has been running for a suitably long period such that the start state is random. Thus the probability distribution of the response at time zero, Y_0 , is the stationary distribution $\pi = \begin{pmatrix} \frac{r}{p+r} \\ \frac{p}{p+r} \end{pmatrix}$. For a realisation of the chain \mathbf{y} , we calculate the likelihood as

$$\begin{aligned} L(\theta|\mathbf{y}) &= P(Y_0 = y_0, Y_1 = y_1, \dots, Y_T = y_T) \\ &= P(Y_0 = y_0)P(Y_1 = y_1|Y_0 = Y_0) \dots P(Y_T = y_T|Y_{T-1} = y_{T-1}, \dots, Y_0 = Y_0), \end{aligned}$$

which by the Markov Principle simplifies to

$$\begin{aligned}
L(\theta|\mathbf{y}) &= P(Y_0 = y_0)P(Y_1 = y_1|Y_0 = y_0)\dots P(Y_T = y_T|Y_{T-1} = y_{T-1}) \\
&= P(Y_0 = y_0) \prod_{j=0}^{T-1} P(Y_{j+1} = y_{j+1}|Y_j = y_j) \\
&= \frac{y_0 p + (1 - y_0)r}{p + r} \prod_{j=0}^{T-1} h_{j,j+1},
\end{aligned}$$

where $h_{j,j+1} = (1 - y_j)[y_{j+1}p + (1 - y_{j+1})(1 - p)] + y_j[y_{j+1}(1 - r) + (1 - y_{j+1})r]$.

Taking logarithms:

$$\log(L(\theta|\mathbf{y})) = \log[y_0 p + (1 - y_0)r] - \log(p + r) + \sum_{j=0}^{T-1} \log h_{j,j+1},$$

where we must assume that $0 < p < 1, 0 < r < 1$. Differentiating gives:

$$\begin{aligned}
\frac{\partial \log(L(\theta|\mathbf{y}))}{\partial p} &= \frac{y_0}{y_0 p + (1 - y_0)r} - \frac{1}{p + r} \\
&\quad + \sum_{j=0}^{T-1} \frac{(1 - y_j)y_{j+1} - (1 - y_j)(1 - y_{j+1})}{h_{j,j+1}} \quad (1)
\end{aligned}$$

and differentiating again gives

$$\begin{aligned}
&\frac{\partial^2 \log(L(\theta|\mathbf{y}))}{\partial p^2} \\
&= \frac{-(y_0)^2}{(y_0 p + (1 - y_0)r)^2} + \frac{1}{(p + r)^2} - \sum_{j=0}^{T-1} \frac{((1 - y_j)y_{j+1} - (1 - y_j)(1 - y_{j+1}))^2}{h_{j,j+1}^2} \\
&= \frac{-(y_0)^2}{(y_0 p + (1 - y_0)r)^2} + \frac{1}{(p + r)^2} - \sum_{j=0}^{T-1} \frac{(1 - y_j)^2 (2y_{j+1} - 1)^2}{h_{j,j+1}^2}.
\end{aligned}$$

Now $y_0 = 0$ or 1 so $(y_0)^2 = y_0$. Similarly $y_j = 0$ or 1 and so $(1 - y_j)^2 = (1 - y_j)$ and $y_{j+1} = 0$ or 1 , so $(2y_{j+1} - 1) = \pm 1$. Therefore

$$\frac{\partial_2 \log(L(\theta|\mathbf{y}))}{\partial p^2} = \frac{-y_0}{(y_0 p + (1 - y_0)r)^2} + \frac{1}{(p + r)^2} - \sum_{j=0}^{T-1} \frac{(1 - y_j)}{h_{j,j+1}^2}.$$

When $y_j = 1, 1 - y_j = 0$ and the final term on the RHS is zero, so we can simplify $h_{j,j+1}^2$ so that the expression becomes

$$\frac{\partial_2 \log(L(\theta|\mathbf{y}))}{\partial p^2} = \frac{-y_0}{p^2} + \frac{1}{(p + r)^2} - \sum_{j=0}^{T-1} \frac{(1 - y_j)}{\{(1 - y_j)[y_{j+1}p + (1 - y_{j+1})(1 - p)]\}^2}.$$

For a random \mathbf{Y} , taking expectations:

$$\begin{aligned} & E\left[\frac{\partial_2 \log(L(\theta|\mathbf{Y}))}{\partial p^2}\right] \\ &= -\frac{E(Y_0)}{p^2} + \frac{1}{(p + r)^2} - \sum_{j=0}^{T-1} E\left[\frac{(1 - Y_j)}{\{(1 - Y_j)[Y_{j+1}p + (1 - Y_{j+1})(1 - p)]\}^2}\right] \\ &= \frac{\left(\frac{-p}{p+r}\right)}{p^2} + \frac{1}{(p + r)^2} - TE\left[\frac{(1 - Y_0)}{\{(1 - Y_0)[Y_1p + (1 - Y_1)(1 - p)]\}^2}\right] \\ &= \frac{\left(\frac{-p}{p+r}\right)}{p^2} + \frac{1}{(p + r)^2} - T[P(Y_0 = 0, Y_1 = 1)\frac{1}{(1 - p)^2} + P(Y_0 = 0, Y_1 = 0)\frac{1}{p^2} \\ &\quad + P(Y_0 = 1, Y_1 = 0)0 + P(Y_0 = 1, Y_1 = 1)0] \\ &= \frac{\left(\frac{-p}{p+r}\right)}{p^2} + \frac{1}{(p + r)^2} - T\frac{r}{p + r}\left[\frac{1}{(1 - p)^2}(1 - p) + \frac{1}{p^2}p\right] \\ &= \frac{-1}{p(p + r)} + \frac{1}{(p + r)^2} - \frac{Tr}{p + r}\left(\frac{1}{p} + \frac{1}{1 - p}\right). \end{aligned}$$

Using (1) we can find:

$$\begin{aligned} \frac{\partial_2 \log(L(\theta|\mathbf{y}))}{\partial p \partial r} &= -\frac{(y_0)(1 - y_0)}{(y_0 p + (1 - y_0)r)^2} + \frac{1}{(p + r)^2} \\ &\quad - \sum_{j=0}^{T-1} \frac{(1 - y_j)(2y_{j+1} - 1) - y_j(1 - 2y_{j+1})}{h_{j,j+1}^2} \end{aligned}$$

and it is clear that $y_j(1 - y_j) = 0$, so this simplifies to

$$\frac{\partial_2 \log(L(\theta|\mathbf{y}))}{\partial p \partial r} = \frac{1}{(p+r)^2}.$$

Using symmetry in p and r, we can find the Fisher information matrix⁵

$$\begin{aligned} I(p, r) &= \begin{pmatrix} E\left(-\frac{\partial_2 \log(L(\theta|\mathbf{Y}))}{\partial p^2}\right) & E\left(-\frac{\partial_2 \log(L(\theta|\mathbf{Y}))}{\partial p \partial r}\right) \\ E\left(-\frac{\partial_2 \log(L(\theta|\mathbf{Y}))}{\partial p \partial r}\right) & E\left(-\frac{\partial_2 \log(L(\theta|\mathbf{Y}))}{\partial r^2}\right) \end{pmatrix} \\ &= \frac{1}{p+r} \begin{pmatrix} \frac{1}{p} - \frac{1}{(p+r)} + \frac{Tr}{p} + \frac{Tr}{1-p} & \frac{-1}{(p+r)} \\ \frac{-1}{(p+r)} & \frac{1}{r} - \frac{1}{(p+r)} + \frac{Tp}{r} + \frac{Tp}{1-r} \end{pmatrix}. \end{aligned}$$

B Attaining the Cramer-Rao Lower Bound

We have shown how to find the optimal value for k, assuming that we have a minimum variance unbiased estimator, and here we demonstrate that doing this is useful, and that we can in practice find estimators for p and r with a variances that are close to the CRLB, the minimum variance obtainable for unbiased estimators.

We know that the variance of any unbiased estimator is bounded below by the CRLB. We generally want to find information about a network as quickly as we can, so we are usually interested in sampling over a small time period, i.e. for small values of T, and we need to show that we can obtain an estimator of θ which comes close to this lower bound.

We can simplify the differentiated likelihood equations (see (1) in Appendix A) to give

$$\frac{\partial \log(L(\theta|\mathbf{y}))}{\partial p} = \frac{y_0}{p} - \frac{1}{p+r} + \sum_{j=0}^{T-1} \frac{(1-y_j)(2y_{j+1}-1)}{y_{j+1}p + (1-y_{j+1})(1-p)}. \quad (2)$$

⁵Strictly, we must show first that $\int \frac{\partial_2 L(\theta|\mathbf{y})}{\partial p^2} dy = \int \frac{\partial_2 L(\theta|\mathbf{y})}{\partial r^2} dy = 0$ to use the Fisher information in this form. We can use a symmetry argument, or prove from first principles.

If, for data \mathbf{y} of length $T + 1$, we let a be the number of transitions from 0 to 0 (the number of times 0 follows 0), b the number from 0 to 1, c the number from 1 to 0, and d the number from 1 to 1, then we let $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ be a count of the transitions of the Markov chain. Clearly then, we can express (2) as

$$\frac{\partial \log(L(\theta|\mathbf{y}))}{\partial p} = \frac{y_0}{p} - \frac{1}{p+r} - \frac{a}{1-p} + \frac{b}{p}. \quad (3)$$

Then $\tilde{p} = \frac{b}{a+b}$, is an unbiased estimator of p : conditioning on the chain being in state 0, which occurs $a + b$ times, it leaves each time with probability p , and this happens b times. Similarly $\tilde{r} = \frac{c}{c+d}$ is an unbiased estimator of r .

Another approach would be to use the maximum likelihood estimators. Setting (3) and the equivalent differentiated log likelihood equation for r equal to zero we get the following estimators:

$$\begin{aligned} \frac{y_0 + b}{\hat{p}} - \frac{1}{\hat{p} + \hat{r}} - \frac{a}{1 - \hat{p}} &= 0, \\ \frac{1 - y_0 + c}{\hat{r}} - \frac{1}{\hat{p} + \hat{r}} - \frac{d}{1 - \hat{r}} &= 0. \end{aligned}$$

We can rearrange this and simplify to get:

$$\begin{aligned} \hat{p}^2[-b - y_0 + 1 - a] + \hat{p}[(b + y_0)(1 - \hat{r}) - 1 - a\hat{r}] + (b + y_0)\hat{r} &= 0 \\ \hat{r}^2[-c + y_0 - d] + \hat{r}[(c + 1 - y_0)(1 - \hat{p}) - 1 - d\hat{p}] + (c + 1 - y_0)\hat{p} &= 0, \end{aligned}$$

a pair of simultaneous quadratic equations in \hat{p} and \hat{r} . We can find a solution for \hat{p} simply in terms of \hat{r} , and vice versa, and we find that these solutions converge quickly given an appropriate start-point (e.g $\hat{p} = 0.5$). It is clear that two quadratic curves with negative second derivative can only meet at two points, and we always in practice find one point where $0 < \hat{p} < 1$, and $0 < \hat{r} < 1$, so we

do find a unique solution. Maximum likelihood theory tells us that the variance of $\hat{\theta}$ is asymptotically equal to the CRLB.

We simulated 1000 realisations of the Markov chain for fixed $T=10$ to compare the estimators, and in order to demonstrate their effectiveness for small T . The results are shown in tables 5 and 6. Experimentally by simulation, we find that the MLEs \hat{p} and \hat{r} are biased, although as they are MLEs they are asymptotically unbiased and must also converge to the CRLB, and indeed we find this as we increase T . It is interesting to note that the biased estimator \hat{p} tends to have a smaller bias and variance than \tilde{p} , presumably as it also uses the value of y_0 to provide a better estimate. However, the overall difference is minimal.⁶

[Table 5 about here.]

[Table 6 about here.]

One problem occurs when calculating both estimators \hat{p} and \tilde{p} where each transition in the chain starts entirely in one state (either 0 or 1) for the whole of the period in which we are sampling. So if $a + b = 0$ or $c + d = 0$, we gather no inference for p or r respectively. We can deal with this when simulating our experiment by ignoring this particular simulation and re-sampling. Another method would be to assign a random value between zero and one to the inestimable parameter, or to base its value on some prior belief. The probability of realising such a chain decreases rapidly as T increases, so this is not a major problem in practice.⁷ We use the first method here, which adds to the bias of the estimators.

⁶If less bias is desired, we could use a bootstrapping technique to remove bias. However this may increase the variance of the estimators.

⁷This problem has been considered by Gani [8], who notes some difficulties in obtaining maximum likelihood estimators for Markov chains, and presents sufficient and necessary conditions for the existence of an MLE. However, this only applies to Markov chains that have one parameter, and is not directly applicable here.

In summary, we have shown that we can find two effective estimators that get close to the CRLB for moderate sample sizes.

List of Figures

1	Overview of model	27
2	Generic behaviour of Markov chain model	28
3	Full $N+1$ state model	29
4	Mapping the 2-state Markov chain model against queue state evolution over time	30
5	Sequence of probes in simplified buffer representation	31




Figure 1: Overview of model

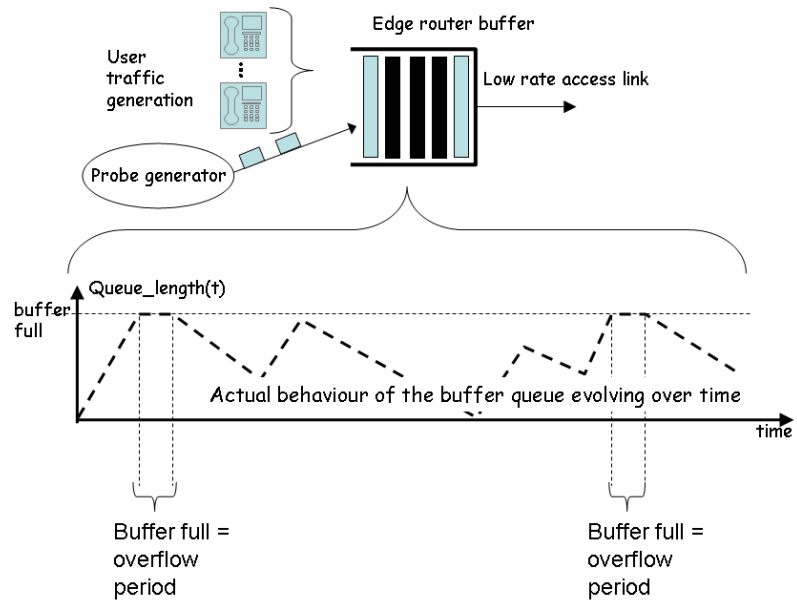


Figure 2: Generic behaviour of Markov chain model

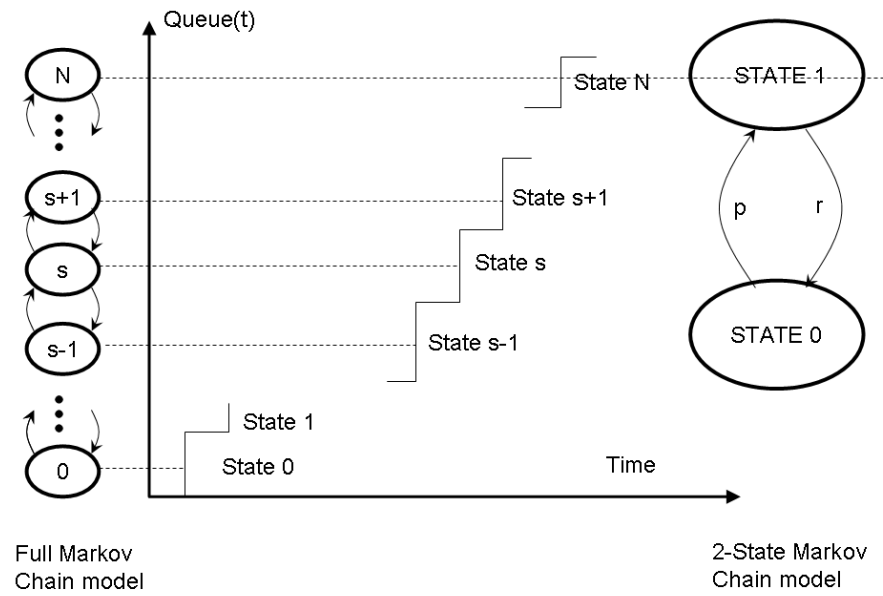


Figure 3: Full $N+1$ state model




Figure 4: Mapping the 2-state Markov chain model against queue state evolution over time




Figure 5: Sequence of probes in simplified buffer representation

List of Tables

1	Comparing $k=1$ with $k=2$ for fixed $T=10$	33
2	Best value of k for $T=5$	34
3	Optimal k for increasing capacity in VoIP	35
4	$T = 5$: Best value of k for varying p and r with penalties for performing experiment	36
5	Accuracy and precision of the simple estimators, \tilde{p}, \tilde{r}	37
6	Accuracy and precision of the maximum likelihood estimators, \hat{p}, \hat{r}	38

Table 1: Comparing $k=1$ with $k=2$ for fixed $T=10$

p/r	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
0.1	0.42	0.47	0.56	0.70	0.97	1.56	3.38	15.11	-
0.2	0.47	0.58	0.77	1.12	1.90	4.31	19.03	-	47.79
0.3	0.56	0.77	1.16	2.05	4.74	20.71	-	35.83	18.2
0.4	0.70	1.12	2.05	4.87	21.38	-	32.29	12.27	11.61
0.5	0.97	1.90	4.74	21.38	-	31.04	10.59	7.47	9.22
0.6	1.56	4.31	20.71	-	31.05	10.17	6.44	5.96	8.52
0.7	3.38	19.03	-	32.29	10.59	6.44	5.35	5.76	9.01
0.8	15.11	-	35.83	12.27	7.48	5.96	5.76	6.78	11.14
0.9	-	47.79	18.28	11.61	9.22	8.53	9.02	11.15	18.25

Table 2: Best value of k for T=5

p/r	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95
0.05	10	7	6	4	4	3	3	2	2	2	2	1	1	1	1	1	1	1	1
0.1	7	5	4	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1
0.15	6	4	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1
0.2	4	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1
0.25	4	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1
0.3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.35	3	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.4	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.45	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.5	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.55	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.6	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.65	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.7	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.75	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.85	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.9	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.95	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

Capacity (Mbps)	Values of D Criterion for $k = 1, 2, \dots, 10$	Optimal k
3.4	378.72, 667.03, 867.01, 985.62, 1033.50, 1023.64, 969.99, 886.32, 785.11, 676.87	k=5
3.6	49.77, 71.38, 72.68, 61.69, 45.57, 29.78, 17.43, 9.28, 4.57, 2.12	k=3
3.8	27.64, 33.94, 28.84, 19.63, 10.90, 4.92, 1.84, 0.59, 0.17, 0.05	k=2
4.0	21.84, 23.54, 17.31, 10.05, 4.60, 1.59, 0.41, 0.09, 0.02, 0.00	k=2
4.2	19.77, 18.86, 12.23, 6.29, 2.56, 0.75, 0.15, 0.02, 0.00, 0.00	k=1

Table 3: Optimal k for increasing capacity in VoIP

Table 4: $T = 5$: Best value of k for varying p and r with penalties for performing experiment

p/r	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95
0.05	10	8	6	5	4	3	3	3	2	2	2	2	1	1	1	1	1	1	1
0.1	7	5	4	3	3	3	2	2	2	1	1	1	1	1	1	1	1	1	1
0.15	5	4	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1
0.2	4	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1
0.25	4	3	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.3	3	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.35	3	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.4	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.45	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.5	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.55	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.6	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.65	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.7	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.75	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.85	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.9	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0.95	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

Table 5: Accuracy and precision of the simple estimators, \tilde{p}, \tilde{r}

p	r	$(bias(\tilde{p}))^2$	$var(\tilde{p})$	$CRLB(p)$	$(bias(\tilde{r}))^2$	$var(\tilde{r})$	$CRLB(r)$
0.2	0.2	0.0050	0.076	0.027	0.0023	0.062	0.027
0.2	0.4	0.0021	0.045	0.018	0.0011	0.107	0.076
0.2	0.6	0.0005	0.026	0.016	0.0001	0.121	0.105
0.2	0.8	0.0003	0.026	0.014	0.0052	0.132	0.085
0.4	0.2	0.0034	0.110	0.076	0.0010	0.040	0.018
0.4	0.4	0.0031	0.059	0.045	0.0045	0.060	0.045
0.4	0.6	0.0011	0.046	0.036	0.0026	0.065	0.060
0.4	0.8	0.0006	0.040	0.032	0.0002	0.052	0.048
0.6	0.2	0.0006	0.115	0.105	0.0012	0.032	0.016
0.6	0.4	0.0030	0.067	0.060	0.0016	0.042	0.036
0.6	0.6	0.0004	0.047	0.046	0.0018	0.047	0.046
0.6	0.8	0.0009	0.042	0.040	0.0016	0.031	0.037
0.8	0.2	0.0011	0.115	0.085	0.0004	0.025	0.014
0.8	0.4	0.0012	0.043	0.048	0.0008	0.041	0.032
0.8	0.6	0.0014	0.031	0.037	0.0010	0.043	0.040
0.8	0.8	0.0004	0.029	0.032	0.0006	0.027	0.032

Table 6: Accuracy and precision of the maximum likelihood estimators, \hat{p}, \hat{r}

p	r	$(bias(\hat{p}))^2$	$var(\hat{p})$	$CRLB(p)$	$(bias(\hat{r}))^2$	$var(\hat{r})$	$CRLB(r)$
0.2	0.2	0.0050	0.071	0.027	0.0026	0.057	0.027
0.2	0.4	0.0016	0.042	0.018	0.0018	0.102	0.076
0.2	0.6	0.0004	0.025	0.016	0.0002	0.120	0.105
0.2	0.8	0.0002	0.026	0.014	0.0049	0.131	0.085
0.4	0.2	0.0036	0.106	0.076	0.0011	0.036	0.018
0.4	0.4	0.0028	0.057	0.045	0.0046	0.058	0.045
0.4	0.6	0.0012	0.046	0.036	0.0029	0.064	0.060
0.4	0.8	0.0006	0.039	0.032	0.0003	0.049	0.048
0.6	0.2	0.0007	0.112	0.105	0.0012	0.030	0.016
0.6	0.4	0.0035	0.065	0.060	0.0015	0.041	0.036
0.6	0.6	0.0005	0.047	0.046	0.0018	0.046	0.046
0.6	0.8	0.0010	0.041	0.040	0.0016	0.031	0.037
0.8	0.2	0.0010	0.114	0.085	0.0005	0.024	0.014
0.8	0.4	0.0012	0.043	0.048	0.0008	0.039	0.032
0.8	0.6	0.0013	0.031	0.037	0.0012	0.042	0.040
0.8	0.8	0.0004	0.028	0.032	0.0006	0.027	0.032