Queueing Syst (2009) 61: 255–272
DOI 10.1007/s11134-009-9107-z

Determining an adequate probe separation for estimating the arrival rate in an M/D/1 queue using single-packet probing

A. Novak · R. Watson

Received: 27 October 2007 / Revised: 2 January 2009 / Published online: 21 March 2009

Abstract We present a technique to estimate the arrival rate from delay measurements, acquired using single-packet probing. With practical applications in mind, we investigate a lower bound on the value of probe separation, to obtain a satisfactory estimate in a fixed amount of time. This leads to a problem: how long does it take for an $M/D/1$ queue to converge to its steady state as a function of the load? We examine this problem using three independent approaches: the time until the autocovariance of the transient workload process becomes negligible; the time it takes for the first transient moment of the workload process to get close to its stationary limit; and the convergence rate of the transient workload distribution to stationarity. These approaches yield different, yet strikingly similar results. We conclude with a recommendation for the probe separation threshold, and a practical approach to obtaining an arrival rate estimate using single-packet probing.

Keywords Active probing · Single packet · $M/D/1$ · Threshold · Autocovariance · Mean · Convergence rate

Mathematics Subject Classification (2000) 60K25 · 60M20 · 90B22

1 Introduction

In this paper we present a method for estimating the arrival rate on the basis of delay measurements, acquired by using single-packet probing. We consider the case where all packets are assumed to be the same size and the arrival process is Poisson, so that a single hop-network can be modelled by an $M/D/1$ queue. Although the assumption of Poisson arrivals is inaccurate, since a Poisson process cannot include the

A. Novak · R. Watson (✉)
Department of Mathematics and Statistics, University of Melbourne, Melbourne, Australia
e-mail: rayw@ms.unimelb.edu.au

Fig. 1 Probing stream

traffic burstiness of any kind, or describe long-range dependence [1], it nevertheless provides valuable insight into network behaviour. Thus, the problem is equivalent to estimating the mean waiting time (and hence the arrival rate) by observing the system at consecutive time points d time units apart.

We investigate a threshold value for probe separation, above which the assumption of independent measurements holds to a reasonable approximation. To this end we consider three different approaches: the time until the autocovariance of the transient workload process becomes negligible; the time it takes for the first transient moment of workload process to get close to its stationary limit; and the convergence rate of the transient workload distribution to stationarity. Even though these three approaches yield quite different equation forms, the threshold values are very close.

In the case of an M/D/1 queue, delay can be modelled explicitly by transient and/or stationary workload distributions. However, the transmission and propagation delays are fixed per route, and do not convey information about the underlying network dynamics.

1.1 Single-packet experiment

A standard method for investigation of Internet traffic is active probing. We consider a single-packet experiment, which consists of injecting a sequence of probes into the network, such that the time separation between probes is some constant value d (see Fig. 1) and the probe service time is x_p. It is common in active probing to choose small probes, so that the invasiveness to the network is minimised. In this paper we comply with this practise, by selecting the probe service time to be much smaller than the cross-traffic service time. The input of the experiment is a sequence of arrival-time stamps $\{t_1, t_2, t_3, \dots, t_n\}$ or $\{t_1, t_1 + d, t_1 + 2d, \dots, t_1 + (n-1)d\}$, and the outcome of the experiment is a sequence of departure-time stamps, $\{t_1^*, t_2^*, t_3^*, \dots, t_n^*\}$. The objective is to use these raw measurement data to estimate the arrival rate in the single-hop network.

We combine the input and the output raw data to give a sequence of differences

$$\{t_1^* - t_1, t_2^* - t_2, t_3^* - t_3, \dots, t_n^* - t_n\}, \tag{1}$$

yielding a new sequence which is equivalent to probe delays $\{\delta_1, \delta_2, \delta_3, \dots, \delta_n\}$. The delay of the ith packet is defined as

$$\delta_i \equiv t_i^* - t_i = w_i + x_i + D,$$

where w_i is the waiting time, x_i is the service time and D is the transmission and propagation delay combined, of the ith packet. The transmission delay, D, predominantly occurs as a result of asynchronous clocks of the sender and receiver, and the propagation delay is dependent on the physical medium and the distance. However, in the case of the equisized cross-traffic packets, this problem can be removed applying a linear-based algorithm [2] to the delay measurements, to remove the clock's skew. In this paper, we begin the analysis after the 'non-queueing' delay D has been eradicated from the delay measurements. Let us denote the new measurement set by $\{\delta^*\}$. Then, assuming negligible service time, we have for the ith packet,

$$\delta_i^* \equiv t_i^* - t_i = w_i + x_i \approx w_i.$$

1.2 Estimation

Active probing can be used to estimate a number of different quantities, such as packet loss across a network path, the capacities of links [3, 4], and the unused capacity or available bandwidth across a path [5].

Existing active probing techniques are typically based on heuristics. There have been a few recent studies on non-empirical methods, aiming to provide theoretical understanding of single-hop bandwidth estimation. In [6] an analytic methodology was developed, featuring intrusion residual analysis in the context of single-hop probing. Also, [7] has shown a stochastic analysis of the problem of estimating capacity/bandwidth of a single congested link. This was the first model to provide an asymptotically accurate estimation for capacity and available bandwidth in the presence of arbitrary cross-traffic.

Outside of the active probing domain, there have been numerous papers regarding parameter estimation in queues. [8] considered parameter estimation for a FIFO queue with deterministic service times and two independent arrival streams of 'observed' and 'unobserved' packets.

Basawa and Prabhu [9, 10] considered moment and maximum likelihood estimation of the model parameters for a GI/G/1 queue. [11] and [12] further investigate maximum likelihood estimation for a GI/G/1 queue. Reference [13] provides an overview of the literature on statistical analysis of queueing systems.

Our approach is simpler: we are interested only in estimating the arrival rate based on restricted information, for which we use moment estimation. Our main interest is the restriction.

1.3 Overview

This paper is organised into four sections. Following this introductory section, in Sect. 2 we apply the single-packet probing technique to a single-hop network, with an additional assumption that all packets are of the same size. We make use of the stationary and transient workload distributions and their respective moments, for an M/D/1 queue, to model time delays observed by the probes. Consequently, we propose an estimation procedure to obtain an arrival rate in a single-hop network.

In Sect. 3 we examine the problem of probe separation from several angles. In each case the aim is to specify a threshold value for probe separation, above which the assumption of independent measurements holds to a reasonable approximation. We begin with the standard approach of examining autocorrelation: when this is effectively zero, then we can assume that the process is close to independent. We apply a result of Abate and Whitt [15] to give a separation threshold. A second approach is to examine the closeness of the mean workload process to its stationary value. An approximate result for the time at which the mean is 95% of its stationary value yields a separation threshold, which is very similar to the autocorrelation-based threshold. The third approach is to consider the convergence rate of the transient distribution to the stationary distribution, following Lund, Meyn and Tweedie [26]. This result gives only an exponential rate of approach. However, the fact that this result matches the other results up to a constant multiple adds further weight to the separation threshold result. These results specify a separation threshold value, which is such that probes further apart than the threshold can be considered to yield approximately independent measurements. The section concludes with a recommendation for the probe separation threshold and describes a practical approach to experimentation using single-packet probing.

Section 4 gives a summary and conclusion.

2 Estimating the arrival rate

2.1 Estimation procedure

We consider an M/D/1 queue: a process with input A_t, an ordinary Poisson process with constant intensity λ, and constant service time b.

Let $W_{t;c}$ denote the workload process, with initial condition $W_{0;c} = c$. Then

$$W_{t;c} = c + bA_t - \int_0^t \mathbf{I}\{W_{s;c} > 0\}\, ds, \quad c \geq 0,\ t \geq 0. \tag{2}$$

To describe the transient behaviour of the queueing system analytically, is a difficult task. It is for this reason, that often in the literature results are given in an implicit way. In the case of an M/D/1 queue, the probability that the workload is zero at time t, given an initial workload of $c \geq 0$ is given by

$$\mathbf{P}(W_{t;c} = 0) = \begin{cases} \sum_{k=0}^{\lfloor \frac{t-c}{b} \rfloor} (1 - \frac{kb}{t}) e^{-\lambda t} \frac{(\lambda t)^k}{k!} & \text{if } t \geq c; \\ 0 & \text{if } t < c. \end{cases}$$

Also, using (2),

$$\mathbf{E}(W_{t;c}) = c + \lambda b t - \int_0^t \mathbf{P}(W_{s;c} > 0)\, ds.$$

In this section we present a method for estimating the arrival rate in single-hop network, using periodic single-packet probing (as described in Sect. 1.1). The set of data we wish to model, $\mathfrak{S}_w = \{t_1^* - t_1 - x_p, t_2^* - t_2 - x_p, \ldots, t_n^* - t_n - x_p\}$, corresponds to a sequence of waiting times that the probes experience in the queue, say $\{w_1, w_2, \ldots, w_n\}$.

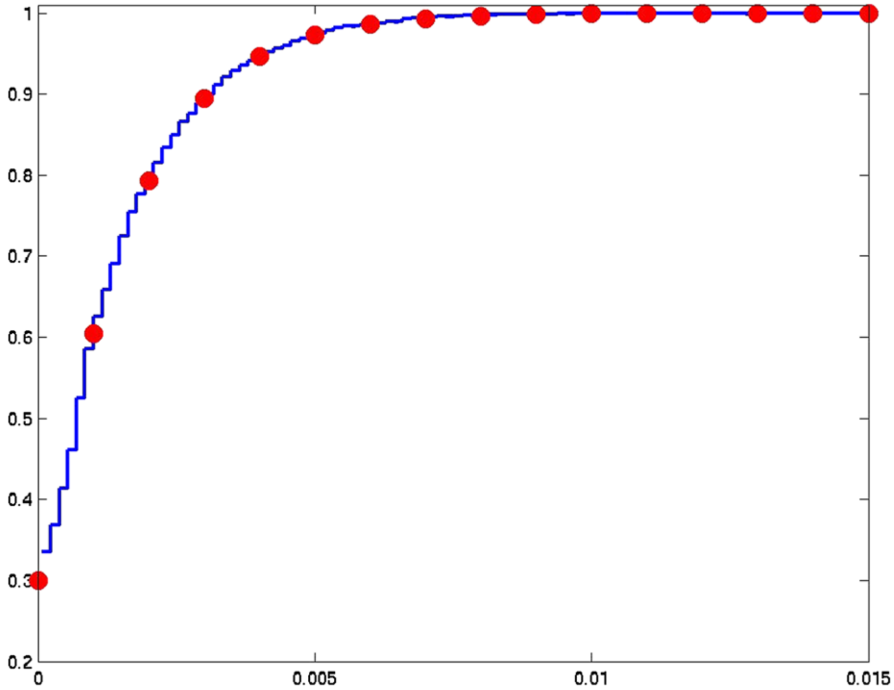

Fig. 2 Stationary workload distribution (theoretical and from simulation). The network was simulated with the following parameters; all cross-traffic packets are the same size and take $b = 0.001$ s to process, probes are very small, $x_p = 0$, the link rate is $\mu = 1000$ pkt/s, cross-traffic packets arrive at rate $\lambda = 700$ pkt/s and $N = 3000$ probes are sent with a period $d = 100b$. The stairs function represents simulation data and the circles are obtained from the stationary workload distribution

If the consecutive probes are spaced sufficiently far apart from each other, then they will each observe a network close to its stationary state (see Fig. 2). In this case, we can model the data in the set $\mathfrak{S}_w$ with the stationary workload distribution for the M/D/1 queue, given by [14, p. 152]

$$\mathbf{P}[W \le x] = (1 - \lambda b) \sum_{j=0}^{\lfloor x/b \rfloor} e^{-\lambda(bj-x)} \frac{[\lambda(bj-x)]^j}{j!}, \quad x \ge 0. \tag{3}$$

The modelling consists of two parts. We construct an estimator given the assumption that cumulative delay measurements follow a stationary workload distribution. The second part is finding a threshold for the probe separation, say α, such that a probe stream constructed with a period $d > \alpha$, will produce close to independent measurements on the stationary distribution.

Due to the simple form of the mean for the stationary workload process, we can use the method of moments to construct an estimator for the arrival rate. That is, we can match the mean of the data in the set $\mathfrak{S}_w$ with the stationary workload mean given by [16, p. 201]

$$m_1 = \mathbf{E}[W] = \frac{\lambda b^2}{2(1-\lambda b)},$$

i.e. we choose $\hat{\lambda}$ so that

$$\bar{w} = \frac{1}{n}\sum_{i=1}^{n} w_i = \frac{\hat{\lambda} b^2}{2(1-\hat{\lambda} b)}.$$

So, a point estimate of λ is given by

$$\hat{\lambda} = \psi(\bar{w}) = \frac{2\bar{w}}{b(2\bar{w}+b)},$$

where ψ is an increasing function, and $\lambda = \psi(m_1)$. Hence

$$\mathbf{P}\left[\psi\left(\bar{w} - 1.96\frac{s_w}{\sqrt{n}}\right) \le \psi(m_1) \le \psi\left(\bar{w} + 1.96\frac{s_w}{\sqrt{n}}\right)\right] \approx 0.95,$$

gives an approximate 95% confidence interval for the arrival rate, λ:

$$\left(\frac{\bar{w} - 1.96\frac{s_w}{\sqrt{n}}}{b(\bar{w} - 1.96\frac{s_w}{\sqrt{n}} + \frac{1}{2}b)}, \frac{\bar{w}+1.96\frac{s_w}{\sqrt{n}}}{b(\bar{w}+1.96\frac{s_w}{\sqrt{n}} + \frac{1}{2}b)}\right). \tag{4}$$

2.2 Iterative scheme

So far we have only stated that the separation between probes, d, has to be 'large enough' to justify the assumption that probes see close to independent events; but 'how large' is the question. The key observation is that seeing independent events depends on the workload process reaching its steady state between two consecutive probes. Thus, the question of how large d has to be, relates to how far the observed process is from its steady state.

Time constraints commonly occur in practice. Thus, when we are talking about the 'best' estimate, it must be put in a context of a limited time-frame.

In order to improve the estimate we need to increase the number of probes (or the sample size). However, for the estimate to be reliable, the measurements have to be sufficiently far apart to be effectively independent, but if we have a limited time-frame, the number of probes is limited. The best we can do then is to choose d/b as close as possible to the threshold value α, that is $d \approx \alpha b$, where $\alpha = \alpha(\rho)$ is specified in Sect. 3.5. But, in order to determine the appropriate α, we need to know the utilisation of the system (or the arrival rate), which is what we want to estimate.

What we propose is an iterative scheme, not unlike a pilot sample commonly used in sampling surveys: see Cochran [17]. However, rather than estimating the variance, here we estimate the required separation. The form of the pilot sample can be readily modified by the user to suit the specific application.

We inject a probe stream with probe separation d_0, where $d_0 = \alpha(\rho_0)b$, so that d_0 is the threshold separation for utilisation ρ_0. One of the user choices is to select d_0.

Table 1 Iterative scheme in the critical region $\rho \in (0.9, 1)$. Network parameters in the experiment are: $d = 50b$, $N = 100$ probes, $x_p \approx 0$ s, $b = 0.001$ s. Only for $\lambda = 900$ pkt/s does the confidence interval contain the true parameter value

λ [pkt/s]	$\hat{\lambda}$ [pkt/s]	CI	Bias [pkt/s]	Bias [%]
900	897.3	(875.7, 912.5)	−2.7	−0.3
950	941.5	(929.7, 949.9)	−8.4	−0.9
970	960.3	(953.1, 965.5)	−9.7	−1.0
990	975.4	(971.4, 978.5)	−14.5	−1.5
999	980.1	(976.9, 982.5)	−18.9	−1.9

We choose $d_0 = 50b$, so that $\rho_0 \approx 0.81$. We use a sample of $N = n_0$ probes. This is another user choice: if n_0 is too small then the result of the pilot sample is unreliable, if it is too large we may have spent too much on a sub-optimal experiment. We choose $n_0 = 100$. The result of this sample then is an estimate of utilisation; call this estimate $\hat{\rho}_1$.

If $\hat{\rho}_1 \le \rho_0$, then we already have a satisfactory (small) experiment. We now perform a second (larger) experiment with $d = d_1$, where $d_1 = \alpha(\hat{\rho}_1)b$; or at least one with $d \ge d_1$.

If $\hat{\rho}_1 > \rho_0$ then we have a problem, since the separation d_0 is not enough to ensure stationarity, and as a result we are likely to have an underestimate of ρ. Table 1 indicates that if a sample with $N = 100$, $d = 50b$ is used for a network with $\rho \ge 0.95$, the result is a significant underestimate.

However, this underestimate can be at least approximately allowed for. To a good approximation, based on simulation, it is found that

$$\mathrm{E}(\hat{\rho}_{100;50} \,|\, \rho) \approx \xi(\rho) = 0.996\rho - 0.4(\rho - 0.8)^2, \quad \rho \ge 0.8. \tag{5}$$

We can then adjust the estimate to $\hat{\rho}_1^* = \xi^{-1}(\hat{\rho}_1)$ and proceed as before. Similar expressions, based on simulation, could be derived for other sample size and separations as required.

Now, we could choose to run another experiment with $d = \alpha(\hat{\rho}_1^*)b$, but as ρ gets close to 1, the value of $\alpha(\rho)$ becomes very large, and the overall benefits of increasing probe separation may be outweighed by the consequent reduction in sample size. This is suggested by Table 1: even with a sample of 100, which is sub-threshold, the bias is not too serious except when ρ gets quite close to 1; and the bias can be approximately corrected for.

In any case, we recommend a choice of separation for the second experiment in the range $d_0 \le d \le \alpha(\hat{\rho}_1^*)b$. We note that if $N \ge n_0$ and $d \ge d_0$ then the bias will be less than that indicated by the approximation (5).

3 Probe separation threshold

3.1 Introduction

We wish to take effectively independent observations on the system, and thus we are concerned with the time between observations so that the observations are effectively independent. In the limit as $T \to \infty$, W_t and W_{t+T} are independent. We seek T so that W_t and W_{t+T} are effectively independent. This is equivalent to determining T so that $W_T \overset{\mathrm{d}}{\approx} W$, i.e. the time for the workload process to have effectively reached its stationary distribution. This time has been called the relaxation time or the settling time for the process.

The relaxation time for a queueing system has been the subject of study for some time. Early results are given in [14] and [16]; see also [18–21]. A more recent overview is given in [22] and [23]. Further, we note that the results of [24] suggest that, at least asymptotically, there is an independence "threshold" in some queueing systems.

3.2 Workload autocorrelation-based threshold

Autocorrelation plots are commonly used for assessing 'pseudo independence' in a time-series: it is a standard practical tool. This approximate independence is ascertained by computing autocorrelations for data values at varying time lags. For example, if a sequence of data is such that we can treat it as if it were independent, then the corresponding autocorrelations should be near zero for any time lag. However, if this is not true, then one or more of the autocorrelations will be significantly non-zero.

For a sample $\{y_1, y_2, \ldots, y_N\}$, an autocorrelation coefficient at lag k can be found by [25]

$$r_k = \frac{C_k}{C_0},$$

where C_k is the autocovariance function

$$C_k = \frac{1}{N-k} \sum_{t=1}^{N-k} (y_t - \bar{y})(y_{t+k} - \bar{y})$$

and C_0 is the variance function

$$C_0 = \frac{1}{N} \sum_{t=1}^{N} (y_t - \bar{y})^2.$$

When an autocorrelation plot is being used to test for independence, the critical values $\pm 2/\sqrt{N}$, where N is the sample size, are recommended [25].

Figure 3 illustrates how pseudo independence of the delay measurements improves with the increase of the probe separation. For example, when $d = 2b$ (probes are very close) the autocorrelation plot of the workload measurements shows strong correlation. At lag 1 the autocorrelation is just under 1, after which it gradually declines to

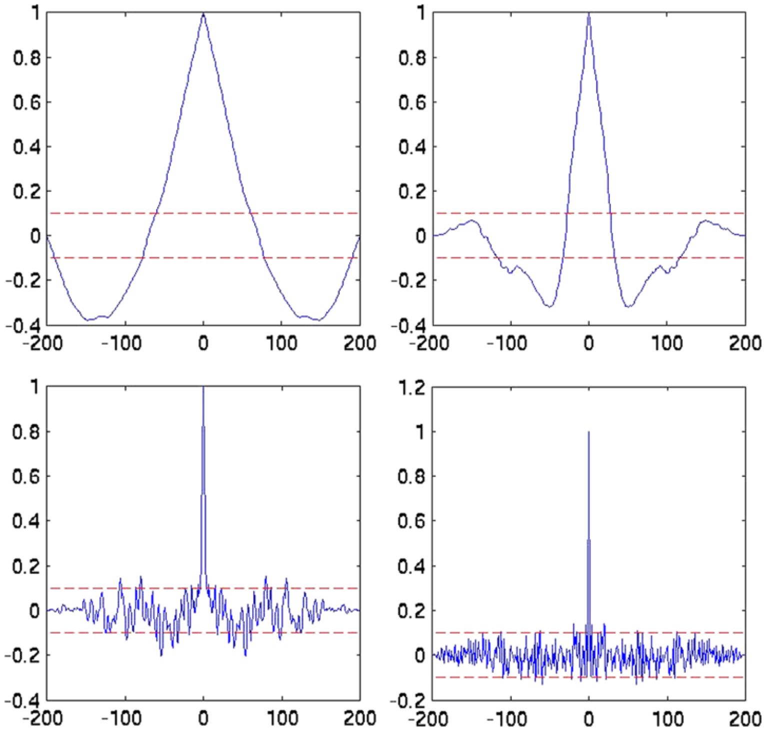

Fig. 3 Autocorrelation of the probe workload measurements. Network parameters: number of probes 200, link rate is 1000 pkt/s, cross traffic arrives at rate 800 pkt/s, probe rate is $1/d$, where d is the probe separation, crosstraffic packet service time is $b = 0.001$ and probe service time is $x_p = 0$

about -0.4, and then it increases back to zero as the lag increases further. A similar sort of pattern is exhibited for $d = 5b$ with the approach to zero less slow.

When $d = 10b$ (probes are further apart) the plot starts with a moderately high autocorrelation at lag 1 (approximately 0.65) and then rapidly decreases. After lag 100 the autocorrelation increases, with a significant noise component, and in this region is effectively zero. The decreasing autocorrelation is approximately linear and as such is a signature of the weak autocorrelation.

On the other hand, when $d = 50b$, with the exception of lag 0 almost all of the autocorrelations are non-significant. The plot exhibits no obvious pattern and so we expect the data to be close to independent. In fact, later on we show that for the utilisation of $\rho = 0.8$, as used in Fig. 3, probe separation needs to be $d > 44b$ to assume an approximate independence of measurements.

In this section we analyse the correlation between periodic probe observations, which are d time units apart, with the assumption that the probe service time $x_p \approx 0$. Under such conditions, we can apply the result on covariance function of the workload process, from Abate and Whitt [15], for a special case of an M/D/1 queue.

Consider the workload process $W_{t;c}$, as defined earlier in (2). Also, let $\{W_t^*: t \geq 0\}$ be a stationary version of the workload process, where $W_0^* \stackrel{\mathrm{d}}{=} W$. In the following, we use results for the first three moments of the stationary workload. These are given by [16, p. 201]

$$m_1 = \mathbf{E}[W] = \frac{\lambda b^2}{2(1-\lambda b)}, \tag{6}$$

$$m_2 = \mathbf{E}[W^2] = 2m_1^2 + \frac{\lambda b^3}{3(1-\lambda b)}, \tag{7}$$

$$m_3 = \mathbf{E}[W^3] = \frac{\lambda b^2}{1-\rho}\left(\frac{3}{2}m_2 + bm_1 + \frac{1}{4}b^2\right). \tag{8}$$

Abate and Whitt [15, p. 752] showed that "a summary description" of the time it takes for the dependence between W_0^* and W_t^* to die out (in the stationary version) is given by

$$u = \frac{\sigma^2}{\mathrm{var}(W_0^*)} = \frac{m_3 - m_2 m_1}{(1-\rho)(m_2 - m_1^2)}, \tag{9}$$

where m_k denotes the kth moment of the stationary workload process. Substituting (6), (7) and (8) for m_1, m_2 and m_3 into (9) and simplifying, we obtain an indication of the time to approximate independence, for an M/D/1 queue

$$u(\rho) = b\frac{3+4\rho-\rho^2}{(4-\rho)(1-\rho)^2}, \quad \text{where } \rho = \lambda b. \tag{10}$$

This then is the autocorrelation-based threshold for separation: the time to approximate independence. In Fig. 4 we illustrate the dependence of the separation threshold, $u(\rho)/b$ given by (10), on the system utilisation, ρ. From the plot we see that, if the utilisation of the system is $\rho = 0.6$, then the probes have to be spaced at least 10 service times apart to assume independence, and if the utilisation is $\rho = 0.8$, then the probe separation needs to be greater than 43 service times. In general, as the utilisation increases, so does the probe separation necessary to assume independence of measurements. This is as expected because higher utilisation means on average longer busy periods, and thus it takes a longer amount of time to erase the memory of the past information. Notice, in the limiting case as $\rho \to 1$ then $u/b \to \infty$.

In order to relate this result back to our problem, we plotted the graph of autocorrelation at lag 1 versus the probe separation, for a simulation performed with utilisation of $\rho = 0.75$ and $N = 1000$ probes (see Fig. 5). Here we observe that the autocorrelation decreases as we increase the separation between probes. According to (10), for the measurements in the sample to be considered independent, the probe separation needs be chosen such that $d > 27b$ (see vertical line in Fig. 5).

3.3 Threshold based on transient mean workload

As $t \to \infty$, the first moment of the transient workload process approaches its steady state value, that is $\mathbf{E}[W_{t;c}] \to \mathbf{E}[W]$. In order to separate the steady state value m_1,

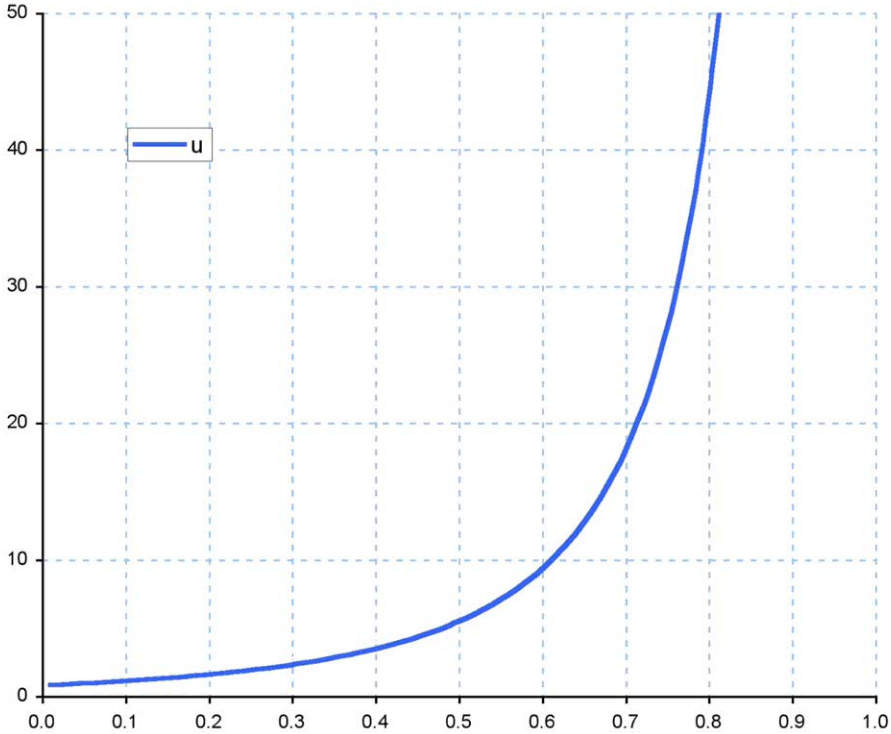

Fig. 4 Separation threshold, $u(\rho)$. Service time: $b = 1$

from the proportion of the steady state value attained at time t, we define a function $H_1(t)$ to be a ratio of the first transient moment (starting empty) and stationary moment [15],

$$H_1(t) = \frac{\mathbf{E}[W_{t;0}]}{\mathbf{E}[W]}. \tag{11}$$

Thus, $H_1(t)$ represents the time dependent moment function starting empty: see Fig. 6.

The function $H_1(t)$ is non-negative, increasing and approaches 1 as $t \to \infty$. Hence, we can interpret it as a distribution function. In the literature [15], it is called the moment cdf. We comply with this convention.

Let h_{11} be the mean of the moment cdf $H_1(t)$, and let $m_k = \mathbf{E}[W^k]$ be the kth moment of the steady state workload cdf. Then, by the corollary to Theorem 6 of Abate and Whitt [15, p. 758],

$$h_{11} = \frac{m_2}{2m_1(1-\rho)} \tag{12}$$

where m_1 is given in (6) and m_2 is given in (7). Substituting these into (12), we get

$$h_{11} = \frac{b(\rho+2)}{6(1-\rho)^2}. \tag{13}$$

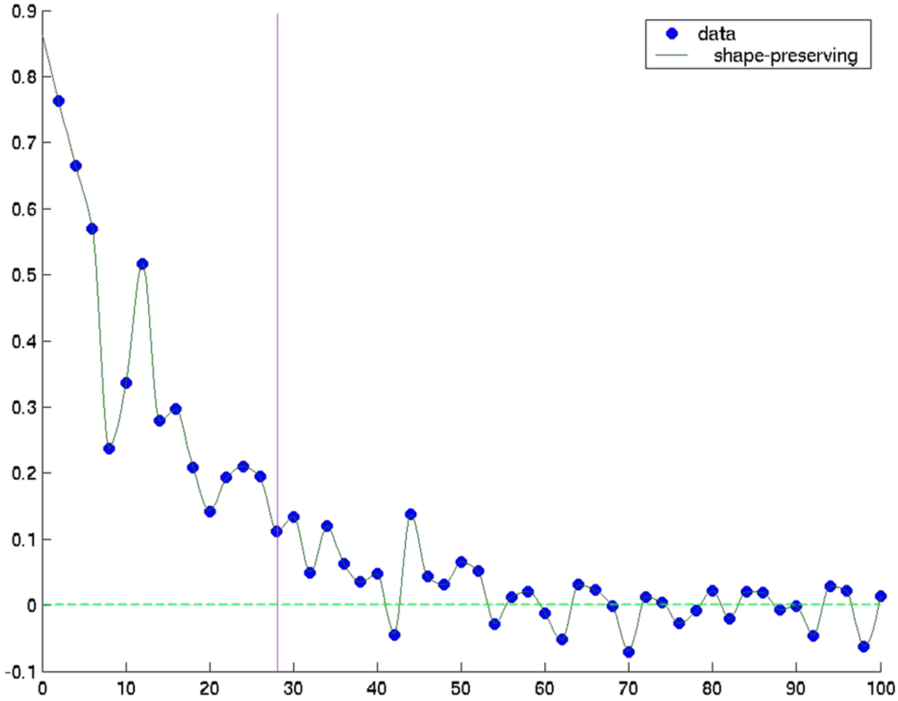

Fig. 5 Autocorrelation at Lag 1 vs. normalised probe separation d/b. Network parameters: number of probes is $N = 1000$, system utilisation $\rho = 0.75$, cross-traffic service time $b = 0.001$, probe service time $x_p = 0$

Similarly, Abate and Whitt [15, p. 758] give the second moment of the cdf H_1 as

$$h_{12} = \frac{1}{(1-\rho)^2}\left(\frac{m_3}{3m_1} + m_2\right) = b^2 \frac{1 + 4\rho + \rho^2}{6(1-\rho)^2},$$

and hence the variance corresponding to H_1 is given by

$$v_{12} = h_{12} - h_{11}^2 = b^2 \frac{2 + 20\rho + 5\rho^2}{36(1-\rho)^4}.$$

In the absence of the effectively intractable 95th percentile of H_1, we use a simple approximation based on the normal quantile:

$$v(\rho) = h_{11}(\rho) + 1.645\sqrt{v_{12}(\rho)}.$$

The graph of $v(\rho)/b$ is shown in Fig. 7. Clearly, there is an element of arbitrariness about the chosen percentile. Despite the approximations, and the difference in functional form, the graphs of $u(\rho)$ and $v(\rho)$ are remarkably similar. This lends support to the threshold value.

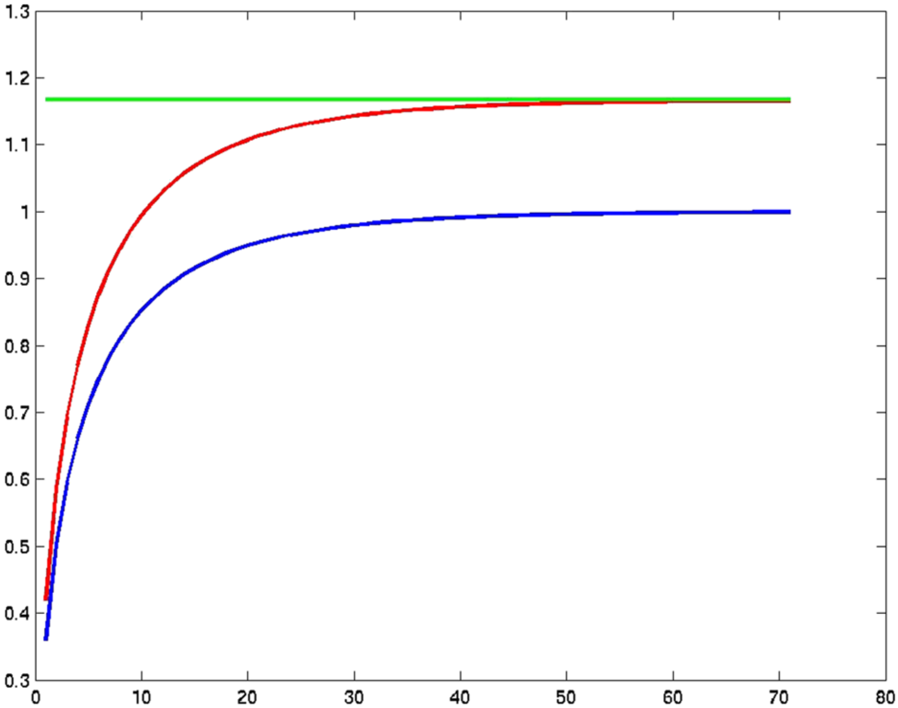

Fig. 6 Moment cdf H_1. Network parameters: $b = 1$, $c = 0$, $\rho = 0.7$

3.4 Convergence rate of the workload distribution

In the previous subsection we examined the approach of the mean transient workload to its stationary limit. In order to get another view of the approach to stationarity, we proceed by evaluating the convergence rate, say β, of the transient workload distribution to the stationary workload distribution. This section applies the results on convergence rate from Lund, Meyn and Tweedie [26] for the special case of an M/D/1 queue. They showed that for many Markov processes (including workloads in queues), the largest possible convergence rate is the radius of convergence of the moment generating function of the first hitting time to zero, and that this radius of convergence is often bounded by 'drift inequalities' based on the generator of the input process.

Consider again the workload process, $W_{t;c}$ given by (2). Let $W_{t;c}$ and $W'_{t;c'}$ be two sample paths of the process, driven by the same sample path of bA_t, but starting from different the initial levels c and c', respectively. If $c < c'$, then by (2) $W'_{u;c'} - W_{u;c} = c' - c$ for all $u \le t$ and therefore $W_{t;c} < W'_{t;c'}$. In other words, $\{W_{t;c}\}$ is a pathwise ordered Markov process.

The process $W_{t;c}$ converges exponentially fast to W, if there exists a $\beta > 0$ such that

$$\lim_{t\to\infty} e^{\beta t} \sup_{\mathcal{A}} \left|\mathbf{P}[W_{t;c} \in \mathcal{A}] - \mathbf{P}[W \in \mathcal{A}]\right| = 0.$$

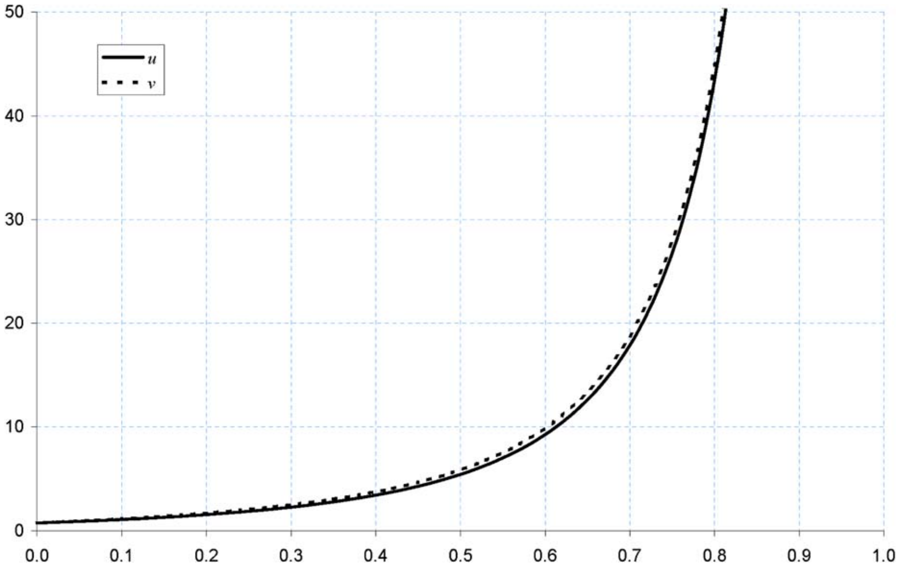

Fig. 7 Plots of the separation thresholds based on autocorrelation and mean workload, $u(\rho)$ and $v(\rho)$. Service time: $b=1$

The best possible exponential rate of convergence, β^* denotes the largest such β.

Define $\tau_{0;c} = \inf\{t \geq 0 : W_{t;c} = 0\}$ as the first hitting time to state zero and $\tau_{0;0} = 0$. Denote by $M_c(s) = \mathrm{E}[e^{s\tau_{0;c}}]$ the moment generating function (MGF) of the first hitting time to zero. From Theorems 2.1 and 2.3 in [26], it follows that the best possible exponential rate of convergence, β^*, is the radius of convergence of M_c for any $c > 0$.

It was shown by Prabhu, Takács and others [14, 22, 26] that

$$M_c(\beta) = e^{c\eta(s)}$$

where $\eta(s)$ is a solution of the functional equation

$$\eta(s) = s + \phi\big(\eta(s)\big),$$

and $\phi(u)$ is the characteristic exponential of the input process

$$\phi(s) = \ln \mathbf{E}\big[e^{sbA_1}\big] = \lambda\big(e^{bs} - 1\big).$$

We are interested in the largest value of s for which $\eta(s) < \infty$. This is, the largest root of the equation

$$f\big(\eta(s)\big) = s,$$

where $f(y) = y - \phi(y)$. For a stable M/D/1 queue $f(y) = y - \lambda(e^{by} - 1)$.

Theorem 4.1 from Lund, Meyn and Tweedie [26] states that the exponential convergence rate β^* is attained for

$$\beta \leq \beta^* = \sup\big\{f(s) : s > 0\big\} = \sup\big\{s - \phi(s) : s > 0\big\}.$$

Thus, equating $f'(s)$ with zero yields the maximum of the function $f(s)$, that is

$$s = \frac{1}{b}\ln\frac{1}{\lambda b} = -\frac{\ln\rho}{b}, \quad \text{and}$$
$$f\left(-\frac{\ln\rho}{b}\right) = \lambda - \frac{1}{b}(1+\ln\lambda b).$$

Therefore, the desired exponential convergence rate is

$$\beta^* = \lambda - \frac{1}{b}(1+\ln\lambda b) = \frac{1}{b}(\rho - 1 - \ln\rho).$$

Now this means that, for any event $\mathcal{A}$

$$\sup_{\mathcal{A}}\left|\mathbf{P}[W_{t;c}\in\mathcal{A}] - \mathbf{P}[W\in\mathcal{A}]\right| < Ke^{-\beta^* t}.$$

But the constant K is unspecified, though bounds can be placed on it. However, in order that

$$\left|\mathbf{P}[W_{t;c}\in\mathcal{A}] - \mathbf{P}[W\in\mathcal{A}]\right| < \epsilon,$$

we require that $Ke^{-\beta^* t} < \epsilon$, and hence

$$t > w(\rho) = \frac{\kappa}{\beta^*(\rho)}.$$

This suggests that the threshold value should take the form

$$w(\rho) = b\frac{\kappa}{\rho - 1 - \ln\rho}, \quad 0 < \rho < 1.$$

This gives a third form for the threshold value, of a different functional form to the other two. But it ought to be a multiple of the others. Choosing $\kappa = 1$ gives a close approximation to the other threshold values: see Fig. 8. Our confidence in the threshold value is further increased.

3.5 Summary and experimental procedure

We choose to adopt as the threshold value

$$\alpha(\rho) = \lceil u(\rho)\rceil,$$

though there is little difference between $u(\rho)$, $v(\rho)$ and $w(\rho)$ (see Fig. 8 and Table 2). The percentage difference between them is greatest when ρ is small, and the threshold is small, in which case it does not matter much.

A recommended experimental procedure is described in Sect. 2.2, though now the detail of the separation threshold $\alpha(\rho)$ can be filled in.

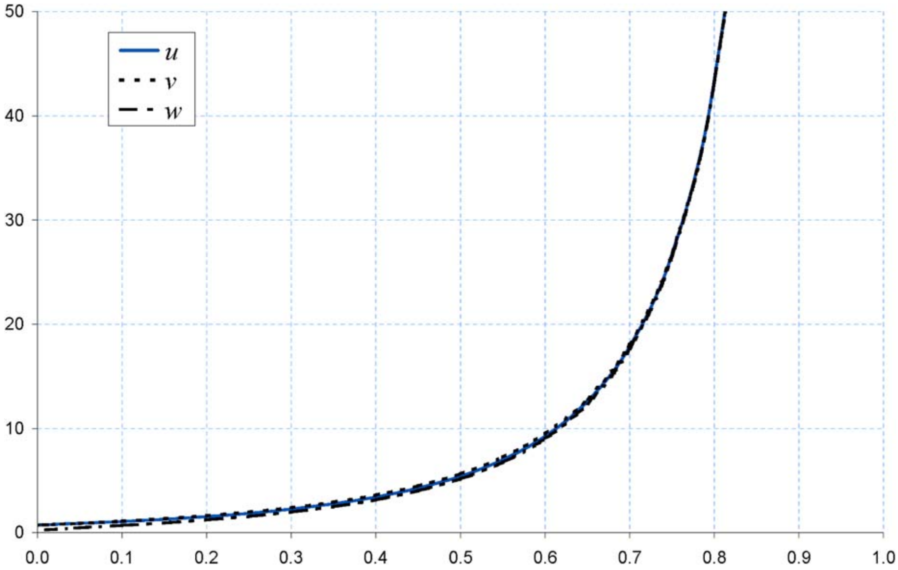

Fig. 8 Plots of the three threshold values: $u(\rho)$, $v(\rho)$ and $w(\rho)$. Service time: $b = 1$

Table 2 Values of u, v, w and the threshold α for a selection of values of ρ

ρ	u	v	w	α
0.10	1.07	1.11	0.71	2
0.20	1.55	1.64	1.24	2
0.30	2.27	2.41	1.98	3
0.40	3.43	3.61	3.16	4
0.50	5.43	5.66	5.18	6
0.55	7.01	7.26	6.76	8
0.60	9.26	9.52	9.02	10
0.65	12.62	12.86	12.38	13
0.70	17.88	18.08	17.64	19
0.75	26.77	26.86	26.54	27
0.80	43.44	43.23	43.20	44
0.85	80.11	79.05	79.87	81
0.90	186.77	182.79	186.53	187
0.95	773.44	750.59	773.14	774

4 Conclusion

In this paper we looked at the problem of estimating the arrival rate for electronic traffic on the basis of delay measurements, acquired using single-packet probing.

Table 3 Summary: separation thresholds

Autocorrelation-based threshold	$u(\rho)/b = \dfrac{3+4\rho-\rho^2}{(4-\rho)(1-\rho)^2}$
Mean-based threshold	$v(\rho)/b = \dfrac{\rho+2+1.645\sqrt{2+20\rho+5\rho^2}}{6(1-\rho)^2}$
Distribution-based threshold	$w(\rho)/b = \dfrac{\kappa}{\rho-1-\ln\rho}$

This led to an investigation of a lower bound on the value of probe separation in order that adequate convergence to the steady state is achieved.

We considered the case of an M/D/1 queue, and used three independent approaches:

1. The time until the autocovariance of the transient workload process (starting empty) becomes negligible.
2. The time it takes for the first transient moment of the workload process to approach within 5% of its stationary limit.
3. The convergence rate of the transient distribution to the stationary workload distribution.

In the first two cases, we evaluated a threshold value for the probe separation, with respect to the system utilisation, ρ. The third case determines a rate of convergence which specifies a threshold only up to a multiplicative constant. The results are summarised in Table 3.

Figure 8 is a visual representation of the results given in Table 3, with $\kappa = 1$. The similarity of these graphs endorses the threshold.

The practical application of these results is a procedure, described in Sect. 2.2 based on a separation threshold, specified in Table 2, which enables sound inference concerning the utilisation to be obtained.

Acknowledgements The authors would like to acknowledge the support of the Australian Research Centre of Excellence in the Mathematics and Statistics of Complex Systems (MASCOS).

The authors are grateful to two anonymous referees who made a number of valuable comments, and whose input has significantly improved the paper.

References

1. Paxson, V., Floyd, S.: Wide area traffic: the failure of Poisson models. IEEE/ACM Trans. Netw. Univ. Mass. **3**, 226–244 (1995)
2. Moon, S.: Measurements and analysis of end-to-end delay and loss in the Internet. Ph.D. Thesis, University of Massachusetts (2000)
3. Dovorlis, C., Ramanathan, P., Moore, D.: Packet-dispersion techniques and capacity-estimation methodology. IEEE Trans. Netw. **12**, 963–977 (2004)
4. Pasztor, A., Veitch, D.: Active probing using packet quartets. In: ACM SIGCOMM Internet Measurement Workshop (IMW-2002), pp. 293–305, Marseille (2002)

5. Riberio, V., Navratil, J., Reidi, R., Cottrell, L.: Efficient available bandwidth estimation for network paths. In: PAM 2003, Passive and Active Measurement Workshop, La Jolla, California (2003)
6. Liu, X., Ravindran, K., Loguinov, D.: Single-hop probing asymptotics in available bandwidth estimation: Sample-path analysis. In: ACM IMC, pp. 300–313 (2004)
7. Kang, S., Liu, X., Dai, M., Loguinov, D.: Packet-pair bandwidth estimation: Stochastic analysis of a single congested node. In: IEEE ICNP, pp. 316–325 (2004)
8. Chen, T.M., Walrand, J., Messerschmitt, D.G.: Parameter estimation for partially observed queues. IEEE Trans. Commun. **42** (1994)
9. Basawa, I.V., Prabhu, N.U.: Estimation in single server queues. Nav. Res. Logist. Q. **28**, 475–487 (1981)
10. Basawa, I.V., Prabhu, N.U.: Large sample inference from single server queues. Queueing Syst. **3**, 289–304 (1988)
11. Basawa, I.V., Bhat, U.N., Lund, R.: Maximum likelihood estimation for single server queues from waiting time data. Queueing Syst. **24**, 155–167 (1996)
12. Acharya, S.K.: On normal approximation for maximum likelihood estimation from single server queues. Queueing Syst., Theory Appl. **31**, 207–216 (1999)
13. Narayan Bhat, U., Subba Rao, S.: Statistical analysis of queueing systems. Queueing Syst., Theory Appl. **1**, 217–247 (1987)
14. Takács, L.: Combinatorial Methods in the Theory of Stochastic Processes. Wiley, New York (1967)
15. Abate, J., Whitt, W.: Transient behaviour of the $M/G/1$ workload process. Oper. Res. **42**, 750–764 (1994)
16. Kleinrock, L.: Queueing Systems, vol. 1. Wiley, New York (1975)
17. Cochran, W.G.: Sampling Techniques, 2nd edn. Wiley, New York (1963), p. 78
18. Blanc, J.P.C., van Doorn, E.A.: Relaxation times for queueing systems. In: de Bakker, J.W., Hazewinkel, M., Lenstra, J.K. (eds.) Mathematics and Computer Science. C.W.I. Monograph 1, pp. 139–162. North-Holland, Amsterdam (1984)
19. Stamoulis, G.D., Tsitsikilis, J.N.: On the settling time of the congested $G/G/1$ queue. Laboratory for Information and Decision Systems Technical Report (1989). http://hdl.handle.net/1721.1/3149
20. Masaaki, K.: On the relaxation time for single-server queues. Oper. Res. Jpn. **32**, 103–111 (1989)
21. Merchant, A.: Settling times for $M/G/1$ queues. Queueing Syst. **8**, 105–110 (1991)
22. Prabhu, N.U.: Stochastic Storage Processes, Queues, Insurance Risk, Dams and Data Communication, 2nd edn. Springer, Berlin (1998)
23. Daigle, J.N.: Queueing theory with applications to packet telecommunication. In: Elementary Continuous-Time Markov-Chain Based Queueing Models. Springer, Berlin (2005). Chap. 3
24. Konstantopoulos, P., Baccelli, F.: On the cut-off phenomenon in some queueing systems. INRIA Research Report 1290 (1990). ftp://ftp.inria.fr/INRIA/publication/publi-pdf/RR/RR-1290.pdf
25. NIST/SEMATECH, E-Handbook of Statistical Methods (2006). http://www.itl.nist.gov/div898/handbook/
26. Lund, R.B., Meyn, S.P., Tweedie, R.L.: Computable exponential convergence rates for stochastically ordered Markov processes. Ann. Appl. Probab. **6**, 218–237 (1996)