Greater Variance Does Not Necessarily Imply Greater Average Delay

Mor Harchol-Balter and David Wolfe

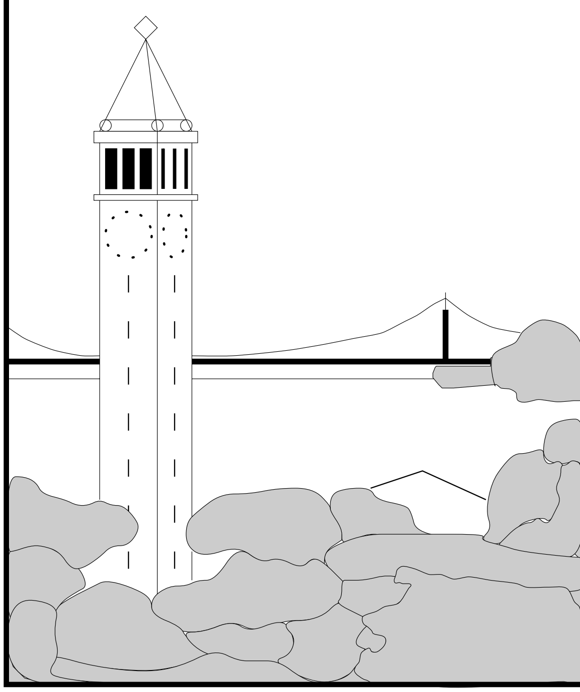

Report No. UCB/CSD-94-821

July 1994

Computer Science Division (EECS)
University of California
Berkeley, California 94720

Greater Variance Does Not Necessarily Imply Greater Average Delay

Mor Harchol-Balter* David Wolfe †

July 1994

Abstract

Real-world packet routing networks differ in many ways from the networks which are analyzable by current-day queueing theory methods. For example the service distributions in real-world networks are constant, whereas the vast majority of queueing theory applies most powerfully to exponential service distributions.

Consequently, it is desirable to at least be able to approximate the behavior of real-world networks by networks which we know how to analyze. Towards this end, much previous research has been done showing that for many networks *greater variance (in service-time distributions and other things) leads to greater congestion*, and therefore we can obtain upper bounds on delays in real-world networks by computing the delay in a related network, having more variance, which we know how to analyze.

The class of networks for which greater variance leads to greater congestion is not known. This paper contributes to determining this classification by demonstrating a network for which increasing the variance in either of two very different ways leads to *smaller* delays.

The arguments we make in this paper are not traditional to the field of queueing theory and are much more in the spirit of discrete mathematics.

*Computer Science Division, UC Berkeley, CA 94720. Supported by National Physical Science Consortium (NPSC) Fellowship. Also supported by NSF grant number CCR-9201092. Email: harchol@cs.berkeley.edu

†Computer Science Division, UC Berkeley, CA 94720. Email: wolfe@cs.berkeley.edu

1 Introduction

Networks of queues for which the service times have a general distribution and/or the arrival process is general are typically very difficult to analyze, with the exception of the case where both the service times and the interarrival times are distributed exponentially. For many well studied networks, previous work shows that the greater the variance in the distributions, the greater is the average time a packet spends in the system (from now on we will denote this by *average packet system time*). In this paper, we show two examples where increasing the variance leads to a **smaller** average packet system time.[1] In this introduction we motivate why we came to consider these two particular examples, and discuss what makes these examples important.

Average packet system time has been shown to be an increasing function of the variance in the service time distribution for the following networks: the M/G/1 queue, the M/G/1 queue with batch arrivals, the M/G/1 queue with priorities, and the M/G/k queue [Whi83] [Whi80] [Ros89, pp. 374–376]. Average packet system time is also an increasing function of the variance in the interarrival times in the G/G/1 queue [Wal89, pp. 353-357]. Even for large classes of complex networks, increased service-time variance leads to higher packet system time. For example, Stamoulis and Tsitsiklis [ST91] [HBW94] show that exponential service times are worse than constant service times for Markovian queuing networks with Poisson arrivals.

In this paper we consider four types of networks of queues. We'll refer to them as Constant Declassed, Constant Classed, Exponential Declassed, and Exponential Classed. (The reason we define these types will become apparent soon). Each of these types is an open network of queues [Ros83, pp. 164–168], where packets arrive according to a Poisson Process. The words Constant/Exponential refer to the distribution of the service time at each server of the network. The words Classed/Declassed refer to the routing policy for the packets in the network. In a Declassed (or Markovian) network, when a packet finishes serving at server i, the probability that it next moves to some server j (or leaves the network) is independent of anything except the fact that the packet just served at i. In particular, all past history is irrelevant. (Thus a Declassed network can simply be described by a directed graph with probabilities on its edges). This is in contrast to

[1] In the case where arrivals are non-Poisson, there are already known counterexamples where increased service-time variance leads to smaller average packet system time. However the case of Poisson arrivals remains open [Wol77] [Ros78] [Whi84].

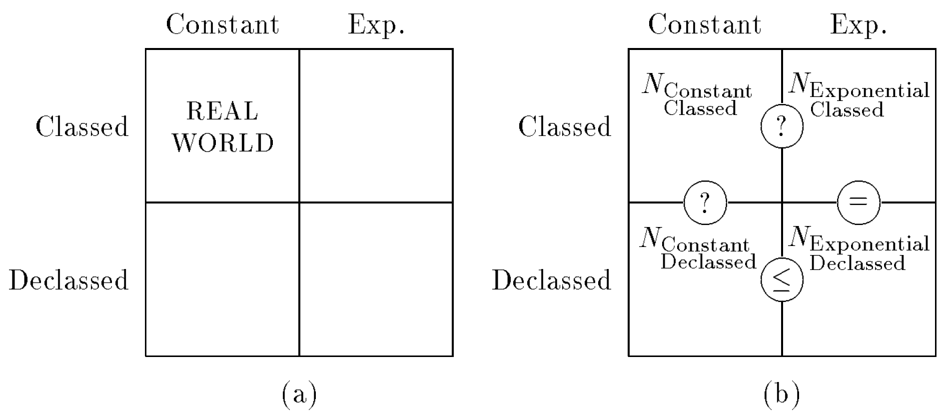

Figure 1: *The four types of networks (a), and the known relationships between them (b). The network type which corresponds to* typical *real-world packet-routing networks is type Constant Classed.*

a Classed (or Non-Markovian) type network, where the probability that a packet next moves to server j after finishing serving at server i is a function of the history of the packet. The name "Classed" comes from the fact that one can imagine a packet being assigned a class at birth, where the class determines the route the packet will follow within the network.

Corresponding to every Classed network, N_{Classed}, there is a related Declassed network, $N_{\text{Declassed}}$, defined as follows: Suppose in N_{Classed}, p proportion of all the packets leaving server i next move to server j. Then in $N_{\text{Declassed}}$, we will define the probability of moving from server i to server j to be p. Thus the Declassed network is a version of the Classed network which no longer depends on history or classes (hence the name). Likewise, corresponding to every Constant type network there is a related Exponential type network, which is identical except that all the servers now have exponentially distributed service times.

Figure 1(a) shows the four types of networks we have defined. The upper left-hand box, type Constant Classed, represents real-world packet-routing networks. This is the type we care most about, but know almost nothing about. Since we know nothing about the upper left-hand box, we'd like to relate it to the other boxes which we know more about.

Figure 1(b) shows what we do know. Define T_N to be the average packet system time for network N. Firstly, we know for any exponential service network N,

$$T_{N_{\text{Exponential-Classed}}} = T_{N_{\text{Exponential-Declassed}}}$$

This follows because the total arrival rate into a server of $N_{\text{Exponential-Classed}}$ is the same as the total arrival rate into the same server in $N_{\text{Exponential-Declassed}}$. *When service times are all exponential*, we know that average packet system time is completely determined by the total arrival rate into each server [HBB94], [Ros83, pp. 164–168]. Since these are the same for both networks, we're done. We also know

$$T_{N_{\text{Constant-Declassed}}} \leq T_{N_{\text{Exponential-Declassed}}}$$

This is due to the result of [ST91], as generalized in [HBW94]. To the best of our knowledge, the other relations indicated in Figure 1(b) are completely open.

What might one hypothesize? Because the Constant distribution has less variance than the Exponential Distribution, it has been conjectured that

$$T_{N_{\text{Constant-Classed}}} = T_{N_{\text{Exponential-Classed}}}$$

One might also conjecture that because the packet routing is "less random" in a Classed Network than in its related Declassed network, that

$$T_{N_{\text{Constant-Classed}}} = T_{N_{\text{Constant-Declassed}}}$$

We will show that both these statements are false, for the general case. (Note, these statements may still be true for special classes of networks.)

The particular network we chose to study was the single-server ring. The ring differs from a single-server feedback network in that in the case of a ring network, packets returning for service rejoin the *end* of the queue, whereas feedback networks are typically defined to return packets to the front of the queue. The single-server feedback network has been well studied by dozens of authors ranging from [Tak63] to [vB91]. Its susceptibility to analysis stems from the fact that it can be modeled as an $M/G/1$ queue. The single-server ring network is, however, more practical because it models real world round-robin applications such as multi-level feedback queue scheduling [CMDD62]. However, the single-server ring network appears far more difficult to analyze, and few formal results are known [Rei88].

Figure 2 shows a Classed Ring and the related Declassed Ring network. In both networks, packets arrive from outside according to a Poisson Process

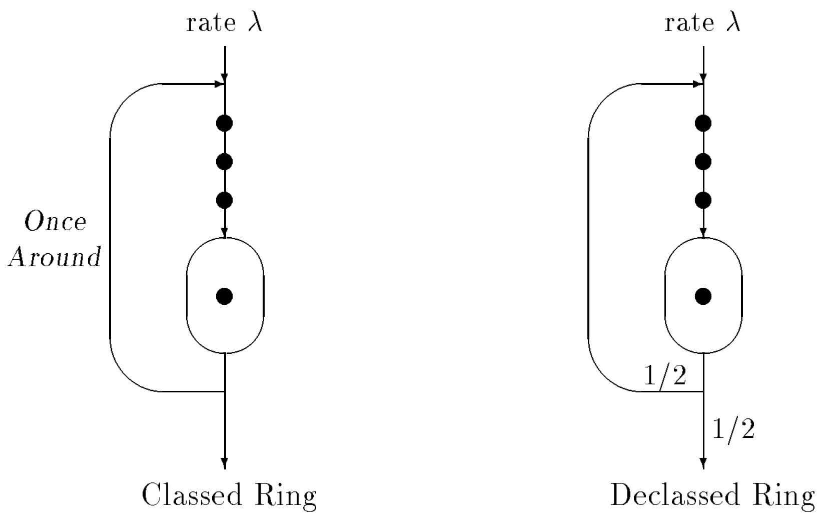

Figure 2: *The related network pair (Classed Ring Network, Declassed Ring Network)*

with rate λ. In the Classed Ring, each packet waits in the queue (if there is one), serves once, then returns *to the end of the queue* to wait again, then serves one more time and leaves the network. This network is Classed, because a packet must carry state with it, namely whether it has served once already or not. Observe that in the Classed Ring, half of all packets which have just served leave the network, and half go back to the end of the queue to serve again. Therefore, in the Declassed Ring a packet has probability $\frac{1}{2}$ of leaving the network after serving once, and probability $\frac{1}{2}$ of returning to serve again. Observe that in both networks the expected number of times a packet serves is 2, *however there is no variance in the Classed Ring, and a variance of* 2 *in the Declassed Ring.*

Our aim in this paper is to show that the Constant Classed Ring has greater average packet system time than the Constant Declassed Ring, and also to show that the Constant Classed Ring has greater average system time than the Exponential Classed Ring. Specifically we prove:

- (Section 2) $T_{\text{Constant-Classed-Ring}} \geq T_{\text{Constant-Declassed-Ring}}$ for the case where all packets arrive in a batch at time zero ($\lambda = \infty$).

  The case of $\lambda = \infty$ is unrealistic, since steady state only occurs for $\lambda < \frac{1}{2}$. However, in Section 3 we will see that $\lambda = \infty$ is indeed an important case, because we'll see that as λ increases, the difference between the Classed Ring and the Declassed Ring becomes more and more pronounced. Therefore studying what happens as $\lambda \to \infty$ is like comparing the two networks under a magnifying glass.

- (Section 3) $T_{\text{Constant-Classed-Ring}} \geq T_{\text{Constant-Declassed-Ring}}$ for the case of 2 packets.
- (Section 4) $T_{\text{Constant-Classed-Ring}} \geq T_{\text{Exponential-Classed-Ring}}$ for the case of 2 packets.

2 $\mathbf{T_{Constant-Classed-Ring} \geq T_{Constant-Declassed-Ring}}$ for $\lambda = \infty$

In this section we consider the case where all packets arrive at time 0 at the Constant Classed Ring and the Constant Declassed Ring. We will prove that the Constant Classed Ring has greater average packet system time than the Constant Declassed Ring, when $\lambda = \infty$, i.e.,

Theorem 1 $T_{\text{Constant-Classed-Ring}} \geq T_{\text{Constant-Declassed-Ring}}$ *for* $\lambda = \infty$

The only difference between the Classed Ring and the Declassed Ring is that in the Classed Ring the service time of each packet is 2 seconds, whereas in the Declassed Ring, the service times of the packets vary (but have mean 2). We describe a step-by-step mapping from the case of varying service times with mean 2 to the case of all service times being exactly 2. We show that each step of this mapping either causes the average packet system time to increase or leaves it unchanged. We therefore conclude that the Classed Ring has greater average packet system time than the Declassed Ring.

We describe the transformation from the Declassed Ring to the Classed Ring via an example. Figure 3 shows service times of packets in the Classed Ring and the Declassed Ring. Each step of the transform consists of:

1. Finding a number in the the Declassed Ring column which is bigger than 2 and decreasing it by one,
2. Finding a '1' in the the Declassed Ring column and increasing it to a '2'.

Classed Ring	Declassed Ring		Classed Ring	Declassed Ring
2	1		2	2
2	5		2	4
2	1	$\Longrightarrow$	2	1
2	2		2	2
2	3		2	3
2	1		2	1
2	1		2	1

Figure 3: *One step of the transformation from the Declassed Ring to the Classed Ring. The columns show service time of packets.*

Since the numbers in the the Declassed Ring column average to 2, after doing enough such transform steps, eventually the the Declassed Ring column will be transformed into the Classed Ring column.[2]

Claim 1 *Each step of the transformation from the Declassed Ring to the Classed Ring causes the average packet system time to either increase or stay the same.*

Proof: Consider the effect on packet system times of the transform step which decreases the service time of packet A from i to $i-1$ ($i > 2$) and increases the service time of packet B from 1 to 2. This transformation step has no effect on the system time of packets whose service times are $\geq i+1$, since the effects of the change in packets A and B are cancelled. The transformation increases the system time of packets whose service times are $< i-1$, since these packets are affected by the change in packet B, but not by the change in packet A. The system time of packets whose service times are i or $i-1$ may increase or may stay the same, depending on the order of the packets. ■

[2]Note that the numbers in the Declassed Ring column may not always average to 2, but their expected average is 2. This is sufficient to prove the Theorem 1, but the averaging argument is omitted.

3 $\mathbf{T}_{\text{Constant}-\text{Classed}-\text{Ring}} \geq \mathbf{T}_{\text{Constant}-\text{Declassed}-\text{Ring}}$ for the case of 2 packets.

In this section we analyze the Constant Classed Ring and the Constant Declassed Ring, where only two packets arrive at each network (according to a Poisson Process with rate λ). Even in the case of just two packets, the calculations are already messy.

3.1 Analysis of $\mathbf{T}_{\text{Constant}-\text{Classed}-\text{Ring}}$ for 2 packets

To determine the average packet system time for the Constant Classed Ring with two packets, we compute the average total (combined) time for the two packets and divide by two. To compute the average total time for two packets, we condition on the arrival time of the second packet, t.

$$\begin{aligned}
&\mathbf{E}\{\text{Total Time for 2 Packets}\} \\
&\quad = \int_{t=0}^{\infty} \mathbf{E}\{\text{Total Time for 2 Packets} \mid \text{interarrival time} = t\} \\
&\qquad \cdot \mathbf{Pr}\left\{2^{\text{nd}} \text{ packet arrives } t \text{ seconds after } 1^{\text{st}}\right\}
\end{aligned}$$

Observe that the expected total time for the two packets depends on whether the second packet arrives after the first packet has completed both its services, one of its services, or neither of its services.

$$\begin{aligned}
&\mathbf{E}\{\text{Total Time for 2 Packets} \mid \text{interarrival time} = t\} \\
&\quad = \begin{cases} 4 & \text{if } t > 2 \\ 6 - t & \text{if } 1 < t < 2 \\ 7 - t & \text{if } 0 < t < 1 \end{cases}
\end{aligned}$$

So,

$$\begin{aligned}
&\mathbf{E}\{\text{Total Time for 2 Packets}\} \\
&\quad = \int_{t=0}^{1} (7-t)\lambda e^{-\lambda t} + \int_{t=1}^{2} (6-t)\lambda e^{-\lambda t} + \int_{t=2}^{\infty} 4\lambda e^{-\lambda t} \\
&\quad = \frac{1}{\lambda} e^{-2\lambda} - e^{-\lambda} - \frac{1}{\lambda} + 7
\end{aligned}$$

Observe that this answer makes sense, since for large λ it evaluates to 7, and for small λ it evaluates to 4.

3.2 Analysis of $\mathbf{T}_{\textbf{Constant-Declassed-Ring}}$ for 2 packets

The average packet system time for the Constant Declassed Ring is easier to compute because the two packets are essentially indistinguishable. In particular, the same total expected delay is incurred if packets are assumed to feedback to the head of the queue rather than the tail for their next service. With this in mind, we'll let m denote the number of times the first packet serves (so, $\mathbf{E}\{\mathrm{m}\} = 2$). Let t denote the time between arrivals of the first and second packets. By conditioning on m and t, we obtain,

$$\begin{aligned}
&\mathbf{E}\{\text{Total time for 2 packets}\} \\
&\quad= 4 + \mathbf{E}\{\text{Overlap time during which both packets are in network}\} \\
&\quad= 4 + \int_t \sum_m \mathbf{E}\{\text{Overlap time} \mid m, t\} \cdot 2^{-m} \lambda e^{-\lambda t} \\
&\quad= 4 + \sum_{m=1}^{\infty} 2^{-m} \int_{t=0}^{m} \lambda e^{-\lambda t}(m-t) \\
&\quad= 6 - \frac{1}{\lambda} + \frac{1}{\lambda(2e^{\lambda}-1)}
\end{aligned}$$

Figure 4 and Figure 5 are plots of the average packet system times for the Classed Ring and the Declassed Ring as a function of λ, in the case of 2 packets, as derived in Section 3.1 and Section 3.2. The average packet system time in both networks is 4 for small λ, and increases to 6 for the Declassed Ring and to 7 for the Classed Ring as $\lambda \to \infty$. For all values of λ, the average packet system time for the Classed Ring is greater than that for the Declassed Ring, and the difference is accentuated as λ is increased.

4 $\mathbf{T}_{\textbf{Constant-Classed-Ring}} \geq \mathbf{T}_{\textbf{Exponential-Classed-Ring}}$ for the case of 2 Packets

In this section we compare the Constant Classed Ring with the Exponential Classed Ring where only two packets arrive at each network, according to a Poisson Process with rate λ. We already computed the average packet system time for the Constant Classed Ring with two packets in Section 3.1. We now look at the Exponential Classed Ring.

We will use t to denote the interarrival time between the two packets. We will use m_1 to denote the time spent by packet 1 on its first service. We will use m_2 to denote the time spent by packet 1 on its second service.

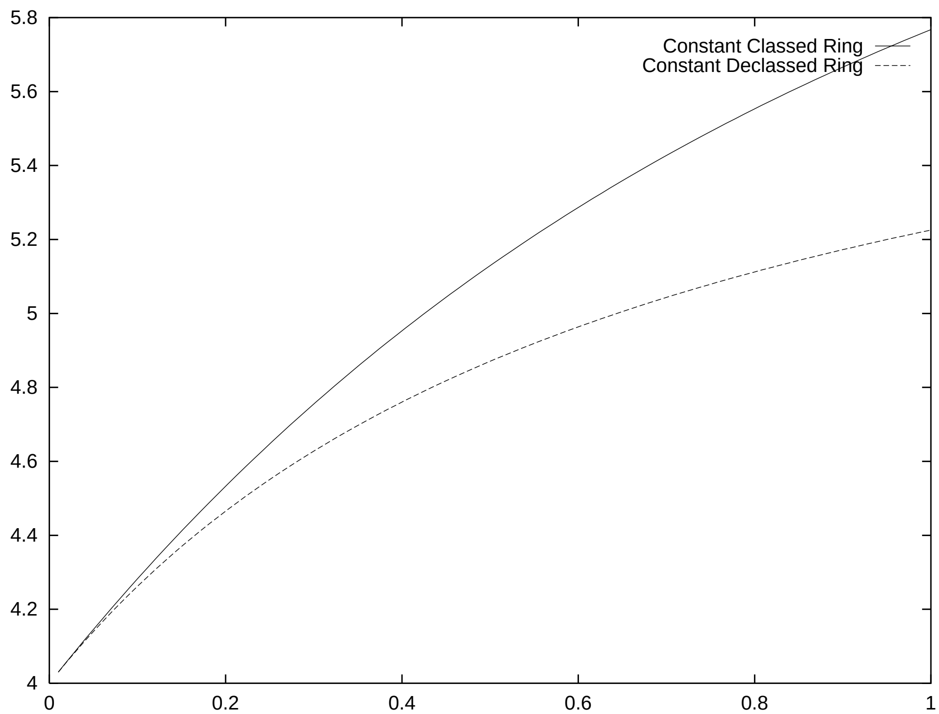

Figure 4: *Expected packet system time for the Constant Classed Ring and the Constant Declassed Ring as a function of* λ*, in the case of* 2 *packets, shown for* $0 \leq \lambda \leq 1$.

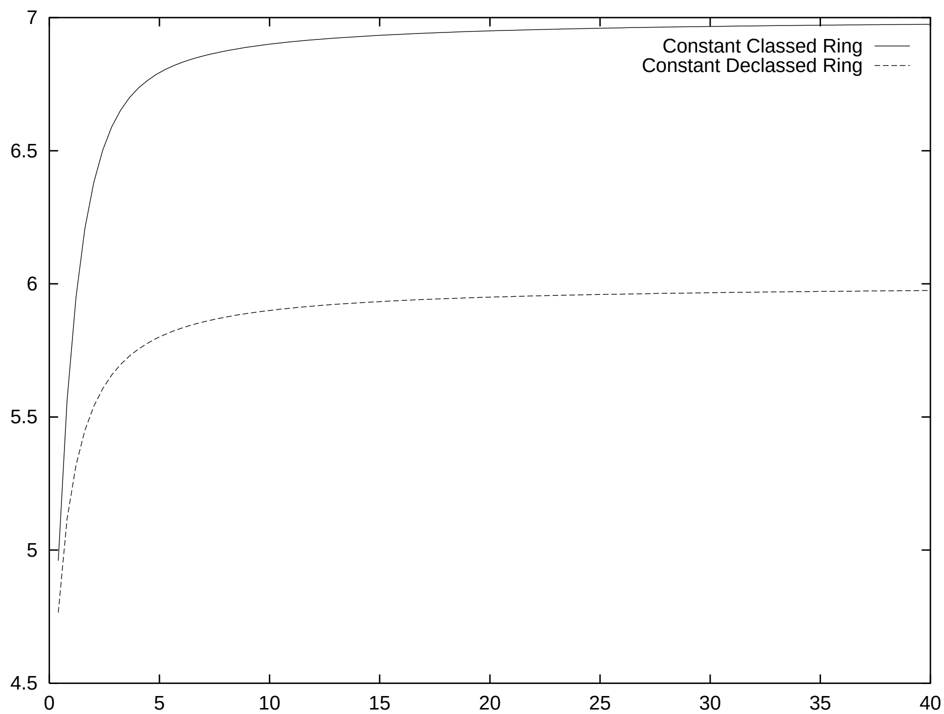

Figure 5: *Expected packet system time for the Constant Classed Ring and the Constant Declassed Ring as a function of* λ*, in the case of* 2 *packets, shown for* $0 \leq \lambda \leq 40$*.*

Conditioning on t we have,

$$\begin{aligned}
&\mathbf{E}\{\text{Total time for 2 packets}\} \\
&\quad = \int_{t=0}^{\infty} \mathbf{E}\{\text{Total time for 2 packets} \mid \text{interarrival} = t\} \\
&\qquad \cdot \mathbf{Pr}\{\text{interarrival time} = t\}
\end{aligned}$$

To determine $\mathbf{E}\{\text{Total time for 2 packets} \mid \text{interarrival time} = t\}$, we condition on m_1 and m_2.

$$\begin{aligned}
&\mathbf{E}\{\text{Total time for 2 packets} \mid t\} \\
&\quad = \int_{m_1=0}^{\infty} \int_{m_2=0}^{\infty} \mathbf{E}\{\text{Total time for 2 packets} \mid t, m_1, m_2\} \\
&\qquad \cdot \mathbf{Pr}\left\{\text{packet 1 takes } m_1 \text{ for } 1^{\text{st}} \text{ service, } m_2 \text{ for } 2^{\text{nd}}\right\}
\end{aligned}$$

To compute the above expression, we consider 3 cases: the second packet may arrive after the first packet is finished with both its services, with one of its services, or with none of its services.

1. $m_1 + m_2 < t \Longrightarrow \mathbf{E}\{\text{Total} \mid t, m_1, m_2\} = m_1 + m_2 + 2$
2. $m_1 < t < m_1 + m_2 \Longrightarrow \mathbf{E}\{\text{Total} \mid t, m_1, m_2\} = 2m_1 + 2m_2 + 2 - t$
3. $t < m_1 \Longrightarrow \mathbf{E}\{\text{Total} \mid t, m_1, m_2\} = 2m_1 + 2m_2 + 3 - t$

We can now rewrite $\mathbf{E}\{\text{Total time for 2 packets} \mid t\}$ as:

$$\begin{aligned}
&\mathbf{E}\{\text{Total time} \mid t\} \\
&\quad = \int_{m_1=0}^{t} \int_{m_2=0}^{t-m_1} (m_1 + m_2 + 2) e^{-m_1} e^{-m_2} \\
&\qquad + \int_{m_1=0}^{t} \int_{m_2=t-m_1}^{\infty} (2m_1 + 2m_2 + 2 - t) e^{-m_1} e^{-m_2} \\
&\qquad + \int_{m_1=t}^{\infty} \int_{m_2=0}^{\infty} (2m_1 + 2m_2 + 3 - t) e^{-m_1} e^{-m_2} \\
&\quad = \frac{-4 + 4\,e^t - 4\,t - t^2}{e^t} \\
&\qquad + \frac{t\,(4 + t)}{e^t}
\end{aligned}$$

$$
\begin{aligned}
&+\frac{7+t}{e^t} \\
=\ & 4 + te^{-t} + 3e^{-t}
\end{aligned}
$$

Now,

$$
\begin{aligned}
&\mathbf{E}\{\text{Total time for 2 packets}\} \\
&= \int_{t=0}^{\infty} (4 + te^{-t} + 3e^{-t}) \cdot \lambda e^{-\lambda t} dt \\
&= 7 - \frac{2}{\lambda+1} - \frac{1}{(\lambda+1)^2}
\end{aligned}
$$

Figure 6 is a plot of the expected packet system time for the Constant Classed Ring versus the Exponential Classed Ring, as a function of λ, in the case of 2 packets. Observe that for small λ, the Constant Classed Ring has smaller average packet system time, however for large λ, the Exponential Classed Ring has smaller average packet system time.

5 Acknowledgement

Mor thanks Manuel Blum for his never-ending encouragement to look for the simplest problem within each problem. It was Manuel who suggested that the simplest ring has only *one* server, and "if you can't handle an infinite number of packets, how about two?" We also thank Jean Walrand for his constant willingness to listen to and critique our ideas.

References

[CMDD62] F. J. Corbato, M. Merwin-Daggett, and R. C. Daley. An experimental time-sharing system. In *Proceedings of the AFIPS Fall Joint Computer Conference*, pages 335–344, May 1962.

[HBB94] Mor Harchol-Balter and Paul E. Black. Queueing analysis of oblivious packet-routing algorithms. In *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 583–592, January 1994.

[HBW94] Mor Harchol-Balter and David Wolfe. Constant beats memoryless for service times in a markovian queueing network. Technical Report UCB//CSD-94-820, University of California, Berkeley, 1994. Also submitted to *Journal of the ACM.*

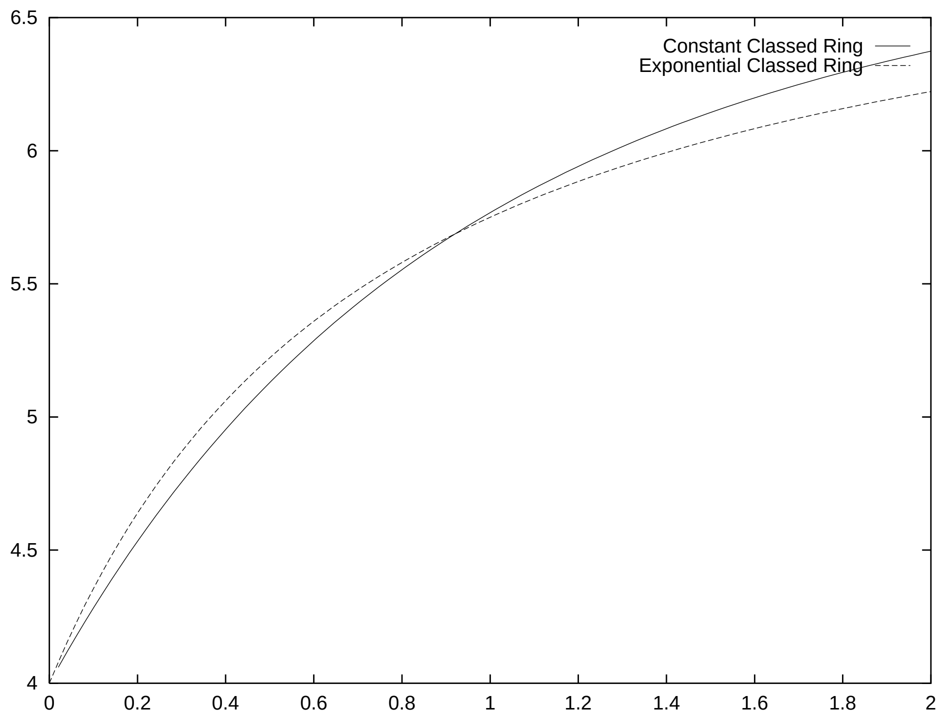

Figure 6: *Expected packet system time for the Constant Classed Ring constant and the Exponential Classed Ring as a function of* λ*, in the case of* 2 *packets, shown for* $0 \leq \lambda \leq 2$*.*

[Rei88] Martin Reiman. A multiclass feedback queue in heavy traffic. *Advances in Applied Probability*, 20(1):179–207, 1988.

[Ros78] Sheldon M. Ross. Average delay in queues with non-stationary Poisson arrivals. *Journal of Applied Probability*, 15:602–609, 1978.

[Ros83] Sheldon M. Ross. *Stochastic Processes*. John Wiley and Sons, New York, 1983.

[Ros89] Sheldon M. Ross. *Introduction to Probability Models*. Academic Press, Inc., Boston, 1989.

[ST91] George D. Stamoulis and John N. Tsitsiklis. The efficiency of greedy routing in hypercubes and butterflies. *Journal of the ACM*, 1991.

[Tak63] L. Takács. A single-server queue with feedback. *The Bell System Technical Journal*, March 1963.

[vB91] J. L. van den Berg and O. J. Boxma. The $M/G/1$ queue with processor sharing and its relation to a feedback queue. *Queueing Sytems*, 9(4):365–401, 1991.

[Wal89] Jean Walrand. *Introduction to Queueing Networks*. Prentice Hall, New Jersey, 1989.

[Wal94] Jean Walrand, 1994. Personal Communication.

[Whi80] Ward Whitt. The effect of variability in the $GI/G/s$ queue. *Journal of Applied Probability*, 17(4):1062–1071, 1980.

[Whi83] Ward Whitt. Comparison conjectures about the $M/G/s$ queue. *Operations Research Letters*, 2(5):203–209, December 1983.

[Whi84] Ward Whitt. Minimizing delays in the $GI/G/1$ queue. *Operations Research*, 32(1):41–51, 1984.

[Wol77] R. W. Wolff. The effect of service time regularity on system performance. In K. M. Chandy and M. Reiser, editors, *International Symposium on Computer Performance Modeling, Measurement, and Evaluation*, pages 297–304, Amsterdam, 1977.