In Network of Queues, M/M/1 Can Outperform M/D/1

Mor Harchol-Balter and David Wolfe

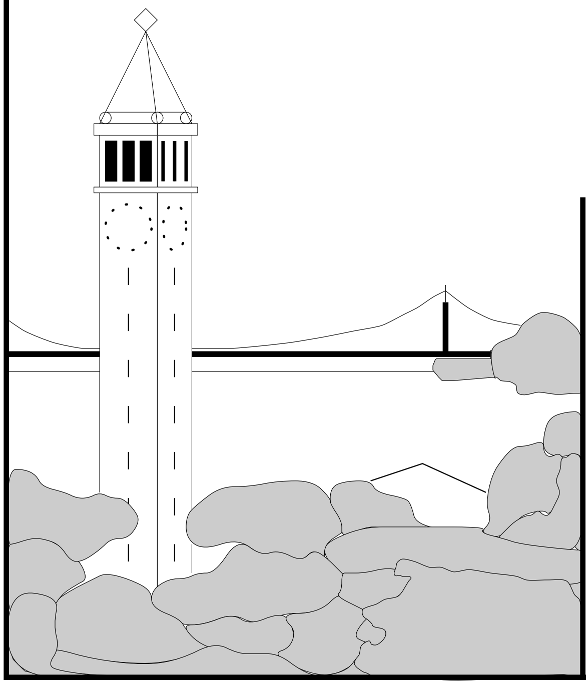

Report No. UCB/CSD-94-841

November 1994

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

In Network of Queues, M/M/1 Can Outperform M/D/1

Mor Harchol-Balter[*] David Wolfe [†]

November 1994

Abstract

Let $\mathcal{N}$ be an open queueing network where the servers have generally distributed service times (with possibly different means) and the outside arrivals to the servers are Poisson. Define $\mathcal{N}_{C,FCFS}$ (respectively, $\mathcal{N}_{E,FCFS}$) to be queueing network $\mathcal{N}$ where each server has a constant (respectively, exponentially distributed) service time with the same mean as the corresponding server in $\mathcal{N}$, and the packets are served in a First-Come-First-Served order.

It has long been conjectured that for all networks $\mathcal{N}$, the average packet delay in $\mathcal{N}_{C,FCFS}$ is upper bounded by the average packet delay in $\mathcal{N}_{E,FCFS}$. In this paper, we present a counterexample to this conjecture.

1 Introduction

1.1 Definitions

Let $\mathcal{N}$ be an open queueing network where the servers have generally distributed service times (with possibly different means) and the outside arrivals to the servers are Poisson. Define $\mathcal{N}_{C,FCFS}$ (respectively, $\mathcal{N}_{E,FCFS}$) to be queueing network $\mathcal{N}$ where each server has a constant (respectively, exponentially distributed) service time with the same mean as the corresponding server in $\mathcal{N}$, and the packets are served in a First-Come-First-Served order.

In general we assume each outside arrival to $\mathcal{N}$ is associated with some class. A packet of class ℓ moves from server i to server j next with probability p_{ij}^{ℓ}. The special case of a *Markovian* network $\mathcal{N}$ is defined as a network with only one class of packets.

Note that $\mathcal{N}_{E,FCFS}$ is a product-form network (more specifically it's a Jackson queueing network, see [BS93]) and therefore the average packet delay is easy to determine for networks of this type

[*]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

(see, for example, [Wal89] [BS93]). On the other hand, it is not known how to compute average packet delay for all but the simplest $\mathcal{N}_{C,FCFS}$ networks.

1.2 Previous Work

All previous work leads one to conjecture that $\mathcal{N}_{E,FCFS}$ has greater average packet delay than $\mathcal{N}_{C,FCFS}$ for all networks $\mathcal{N}$. For example, the average packet delay is an increasing function of of the variance in the service time distribution for each of the following single queue 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. 353–356].

With respect to networks of queues, [HBW94] generalized an earlier result of [ST91] to show that for all Markovian networks $\mathcal{N}$, $\mathcal{N}_{E,FCFS}$ has greater average packet delay than $\mathcal{N}_{C,FCFS}$. There are also empirical studies of several non-Markovian networks $\mathcal{N}$ (i.e. general classed networks) which show that the average packet delay measured is greater for $\mathcal{N}_{E,FCFS}$ than for $\mathcal{N}_{C,FCFS}$ (see empirical results of [HBB94] [MC86] [HC86]).

The above results have led to a general belief that greater variance in service times leads to greater average packet delay [Whi84] [Wal94] [Fer94] [Kle94]. Counterexamples to this theory have only been found in the case where arrivals are not (strictly) Poisson [Wol77] [Ros78]. For example Figure 1 indicates why counterexamples can be found which use *batch* Poisson arrivals such as those in [Wol77]. *In this paper, we demonstrate a counterexample for the case of Poisson arrivals.*

1.3 This Paper

In this paper, we demonstrate an $\mathcal{N}$ for which

$$\text{averagedelay}(\mathcal{N}_{C,FCFS}) > \text{averagedelay}(\mathcal{N}_{E,FCFS})$$

More specifically, define $\mathcal{N}_{C,PS}$ to be the queueing network $\mathcal{N}$ where each server has a constant service time with the same mean as the corresponding server in $\mathcal{N}$ and the service order is Processor Sharing. By [BCMPG75] and [Kel75], we know that the average packet delay in $\mathcal{N}_{C,PS}$ is equal to the average packet delay in $\mathcal{N}_{E,FCFS}$ for all $\mathcal{N}$.[1] In this paper we demonstrate a queueing network $\mathcal{N}$ for which:

$$\text{averagedelay}(\mathcal{N}_{C,FCFS}) > \text{averagedelay}(\mathcal{N}_{C,PS}) = \text{averagedelay}(\mathcal{N}_{E,FCFS})$$

In the next section we describe the network $\mathcal{N}$ and then we prove the inequality.

[1] This powerful theorem is also described more recently in [Wal89] and [Kle76].

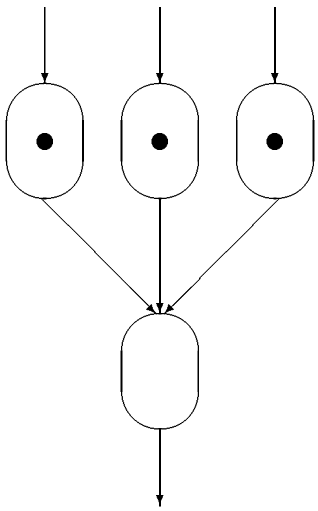

Figure 1: *Non-Poisson (in this case batch-Poisson) arrivals can favor more variance in service distributions. For example, if three packets arrive in a batch (as shown here serving in the top three servers), they'll collide at the next server unless their service-completions are staggered.*

2 Network Description and Analysis

2.1 Network Description

Let $\mathcal{N}$ be the queueing network shown in Figure 2. The servers in $\mathcal{N}$ either have service time 1 or ϵ, as shown. The only outside arrivals are into the top server. Packets arrive from outside $\mathcal{N}$ according to a Poisson Process with rate $\lambda = \frac{1}{n^3}$, where n is the number of servers in $\mathcal{N}$. Half the arriving packets are of type *solid* and half are of type *dashed* (by "type" we mean class). Packets of type *solid* are routed straight down, only passing through the time 1 servers. Packets of type *dashed* are routed through the dashed edges, i.e. through all the ϵ servers and through every other 1-server.

2.2 Intuition

We will compare the average delay in $\mathcal{N}_{C,FCFS}$ with the average delay in $\mathcal{N}_{C,PS}$, as shown in Figure 3. By PASTA (Poisson Arrivals See Time Averages) it will be sufficient to compare the average delay experienced by an arriving packet p at $\mathcal{N}_{C,FCFS}$ and $\mathcal{N}_{C,PS}$.

The intuition behind the analysis is as follows: Since λ is so low, usually for either network, p will see no other packets during its traversal of the network. In this case $\mathcal{N}_{C,FCFS}$ behaves identically to $\mathcal{N}_{C,PS}$. With some probability, however, one other packet will be present in the network during p's traversal of the network. The expected delay on p in this case is greater for the $\mathcal{N}_{C,FCFS}$ network than for the $\mathcal{N}_{C,PS}$ network. Figure 4 shows us why: Consider first $\mathcal{N}_{C,FCFS}$. Suppose q is of type *solid* and some packet p of type *dashed* enters $\mathcal{N}_{C,FCFS}$ within $\frac{n}{2}$ seconds after q. Then p will eventually catch up to q, and from this point on, q will delay p by one second at every other server throughout the rest of the $\mathcal{N}_{C,FCFS}$. That is, p will be delayed by $\Theta(n)$ seconds. Now observe that the same scenario would only cause a delay of at most 2 seconds in $\mathcal{N}_{C,PS}$, because when p catches up to q, it will only interfere with q for two servers and then p will pass q forever. A worse situation for $\mathcal{N}_{C,PS}$ is the case where p meets up with another packet of the same type as p during its traversal (since in that case p is clearly delayed by $\Theta(n)$). Observe, however, that this scenario can only happen if the two packets both arrived at $\mathcal{N}_{C,PS}$ within a second of each other. This occurs with such low probability for our choice of small λ that the scenario's affect on average delay is negligible.

Lastly, we have to consider the case that two or more packets are present in the network during p's traversal of the network. The expected delay on p in this case is greater for the $\mathcal{N}_{C,PS}$ network than for the $\mathcal{N}_{C,FCFS}$ network, but this case occurs with such low probability that its effect on p's delay is also negligible.

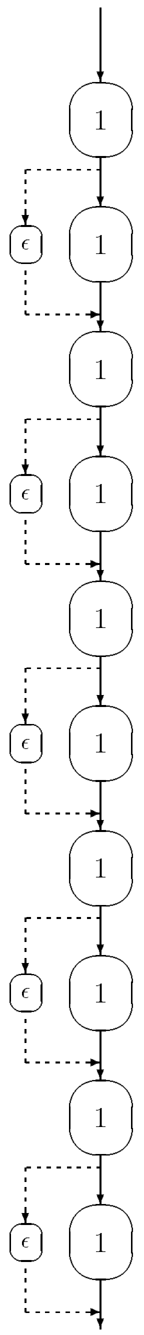

Figure 2: *Network $\mathcal{N}$. Packets arrive at the top; half follow the dashed route, while half follow the solid route. The service times at the servers are either 1 or ϵ as shown.*

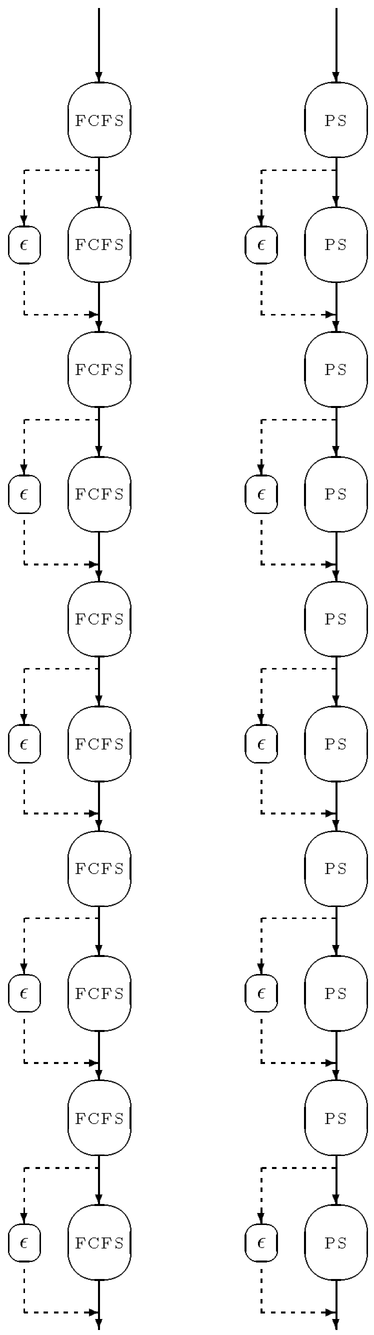

Figure 3: $\mathcal{N}_{C,FCFS}$ *and* $\mathcal{N}_{C,PS}$ *networks*

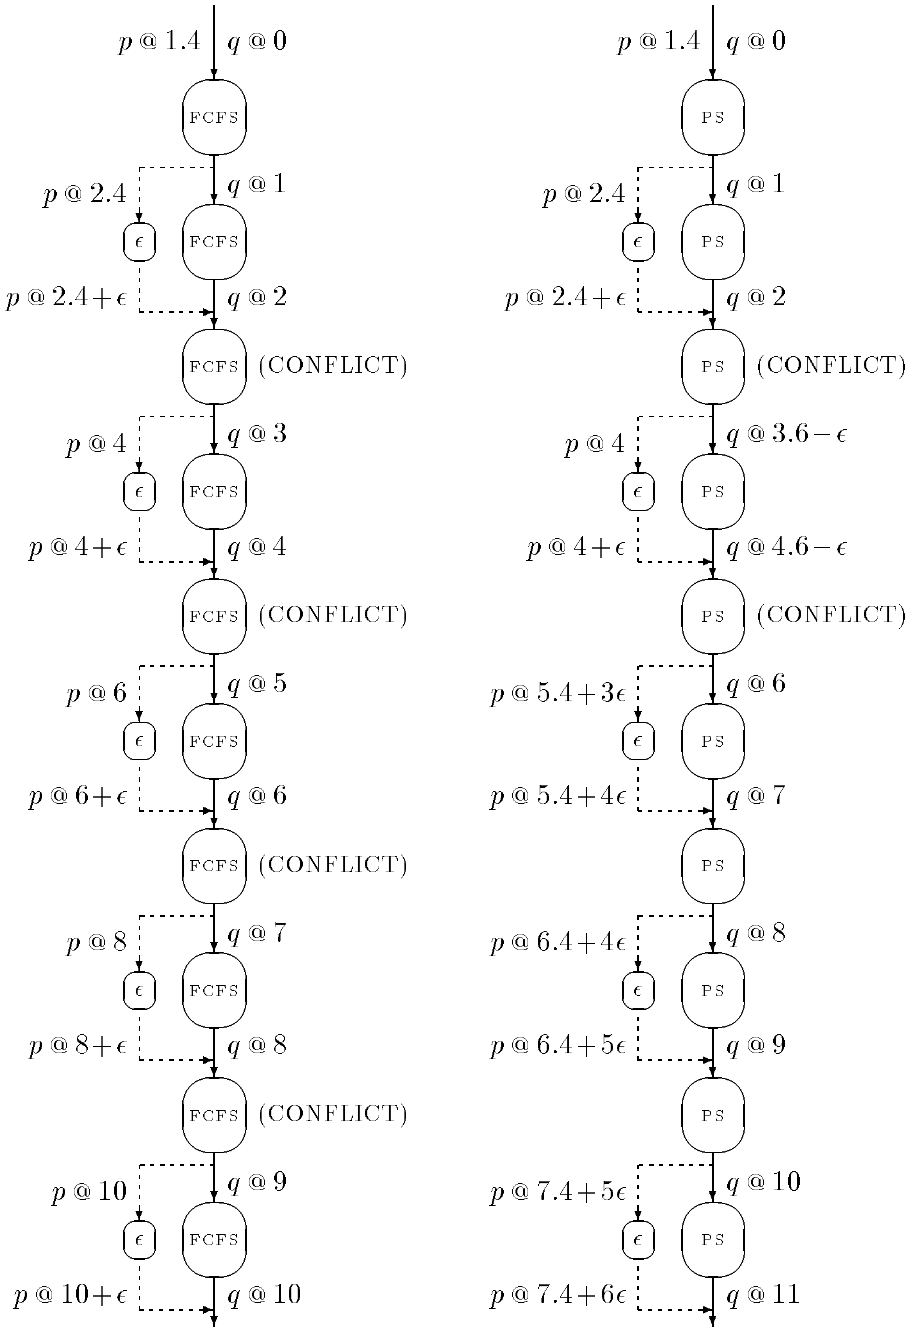

Figure 4: *Example illustrating how a packet, p, of type dashed and a packet, q, of type solid clash repeatedly in $\mathcal{N}_{C,FCFS}$, but only twice in $\mathcal{N}_{C,PS}$.*

2.3 The details

Define

$$P_i^{kn} = \mathbf{Pr}\{\text{there are } i \text{ arrivals during the last } kn \text{ seconds}\}.$$

Recall that arrivals are Poisson with rate $\lambda = 1/n^3$, where n is the number of servers. So,

$$\begin{aligned} P_i^{kn} &\approx \frac{k^i}{n^{2i}} \\ P_i^{kn} &= \Theta\left(\frac{1}{n^{2i}}\right), \text{ for fixed } i, k \end{aligned}$$

By PASTA, the expected delay a newly arriving packet experiences is equal to the average packet delay for the network. We will compute an upper bound on the delay an arrival experiences in $\mathcal{N}_{C,PS}$ and a lower bound on the delay an arrival experiences in $\mathcal{N}_{C,FCFS}$. We will show

$$\text{lowerbound}(\mathbf{E}\{\text{Delay on arrival in } \mathcal{N}_{C,FCFS}\}) > \text{upperbound}(\mathbf{E}\{\text{Delay on arrival in } \mathcal{N}_{C,PS}\}).$$

2.3.1 upperbound on $\mathbf{E}\{$Delay on arrival in $\mathcal{N}_{C,PS}\}$

Let p represent an arriving packet in $\mathcal{N}_{C,PS}$.

Clearly, p may only be delayed by packets which are in $\mathcal{N}_{C,PS}$ during the time p is in $\mathcal{N}_{C,PS}$. Note that if k packets are in $\mathcal{N}_{C,PS}$, they may take up to time kn to clear the system. So, if we call p's arrival time 0, packet p *may* be delayed if one of the following occur:

- 1 other packet arrives during $(-n, n)$.
- 2 other packets arrives during $(-2n, 2n)$.
- 3 other packets arrives during $(-3n, 3n)$.
- etc.

We will compute the expected delay on p due to each of the above events, and then we'll sum these. This will be an overcount, but that's o.k. because we're just upperbounding.

$$
\begin{aligned}
& \mathbf{E}\{\text{Delay on } p \text{ caused by 1 other packet arriving in } (-n,n)\} \\
= \; & \mathbf{E}\{\text{Delay on } p \text{ caused by 1 other packet of same type arriving in } (-n,n)\} \\
& +\mathbf{E}\{\text{Delay on } p \text{ caused by 1 other packet of opposite type arriving in } (-n,n)\} \\
= \; & \mathbf{Pr}\{\text{1 other same type arrival in } (-n,n)\} \cdot \mathbf{E}\{\text{Delay on p} \mid \text{one other same type arrival in } (-n,n)\} \\
& +\mathbf{Pr}\{\text{1 other opp. type arrival in } (-n,n)\} \cdot \mathbf{E}\{\text{Delay on p} \mid \text{one other opp. type arrival in } (-n,n)\} \\
= \; & \Theta(\frac{1}{n^2}) \cdot \Theta(\frac{1}{n} \cdot n) \qquad \text{(get delay of } n \text{ only if same type packet arrived in } (-1,1)) \\
& +\Theta(\frac{1}{n^2}) \cdot \Theta(1) \qquad \text{(opposite type packet causes at most delay of } \Theta(1)) \\
= \; & \Theta(\frac{1}{n^2}) \cdot O(1) \\
= \; & \Theta(\frac{1}{n^2})
\end{aligned}
$$

$$
\begin{aligned}
& \mathbf{E}\{\text{Delay on } p \text{ caused by 2 other packet arriving in } (-2n,2n)\} \\
\leq \; & \mathbf{Pr}\{\text{2 other packets arrive during } (-2n,2n)\} \cdot \mathbf{Max}\{\text{Delay on p} \mid \text{2 other arrivals in } (-2n,2n)\} \\
= \; & \Theta(\frac{1}{n^4}) \cdot O(2n) \\
= \; & \Theta(\frac{1}{n^3})
\end{aligned}
$$

$$
\begin{aligned}
& \mathbf{E}\{\text{Delay on } p \text{ caused by 3 other packet arriving in } (-3n,3n)\} \\
\leq \; & \mathbf{Pr}\{\text{3 other packets arrive during } (-3n,3n)\} \cdot \mathbf{Max}\{\text{Delay on p} \mid \text{3 other arrivals in } (-3n,3n)\} \\
= \; & \Theta(\frac{1}{n^6}) \cdot O(3n) \\
= \; & \Theta(\frac{1}{n^5})
\end{aligned}
$$

To compute an upper bound on $\mathbf{E}\{\text{Delay on } p\}$, we sum the above terms. From the above computations its clear that the delay on p from the remaining summands not shown above is negligible. We find that

$$\mathbf{E}\{\text{Delay on } p\} = O(\frac{1}{n^2})$$

2.3.2 lowerbound on $\mathbf{E}\{\text{Delay on arrival in } \mathcal{N}_{C,FCFS}\}$

Let p represent an arriving packet in $\mathcal{N}_{C,FCFS}$. Assume p arrives at $\mathcal{N}_{C,FCFS}$ at time 0. To lowerbound the $\mathbf{E}\{\text{Delay on } p \text{ in } \mathcal{N}_{C,FCFS}\}$, we consider only the delay on p caused by 1 packet arriving during $(-n, n)$. Observe that if 1 packet (other than p) arrived during $(-n, n)$, and if the packet was of a different type than p, then p and the packet would meet, and the delay caused to p (if p is dashed) is $\Theta(n)$.

$$
\begin{aligned}
& \mathbf{E}\{\text{Delay on } p \text{ caused by 1 other packet arriving in } (-n,n)\} \\
= \; & \mathbf{Pr}\{\text{1 other packet arrives during } (-n,n)\} \cdot \mathbf{E}\{\text{Delay on p} \mid \text{one other arrival in } (-n,n)\} \\
= \; & \Theta(\frac{1}{n^2}) \cdot \Theta(n) \qquad \text{(see Intuition Section)} \\
= \; & \Theta(\frac{1}{n})
\end{aligned}
$$

Thus,

$$\mathbf{E}\{\text{Delay on } p\} = \Omega(\frac{1}{n})$$

References

[BCMPG75] F. Baskett, K.M. Chandy, R.R. Muntz, and F. Palacios-Gomez. Open, closed, and mixed networks of queues with different classes of customers. *Journal of the Association for Computing Machinery*, 22:248–260, 1975.

[BS93] John A. Buzacott and J. George Shanthikumar. *Stochastic Models of Manufacturing Systems*. Prentice Hall, 1993.

[Fer94] Dominico Ferrari, 1994. Personal Communication.

[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.*

[HC86] Bruce Hajek and Rene L. Cruz. Delay and routing in interconnection networks. In *Flow Control of Congested Networks*, pages 235–242, 1986.

[Kel75] F. P. Kelly. Networks of queues with customers of different types. *Journal of Applied Probability*, 12:542–554, 1975.

[Kle76] Leonard Kleinrock. *Queueing Systems Volume II: Computer Applications*. John Wiley and Sons, New York, 1976.

[Kle94] Len Kleinrock, 1994. Personal Communication.

[MC86] Debasis Mitra and R. Cieslak. Randomized parallel communications. In *International Conference on Parallel Processing*, pages 224–230, 1986.

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

[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.

[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.