Bounding Delays in Packet-Routing Networks with Light Traffic

Mor Harchol-Balter

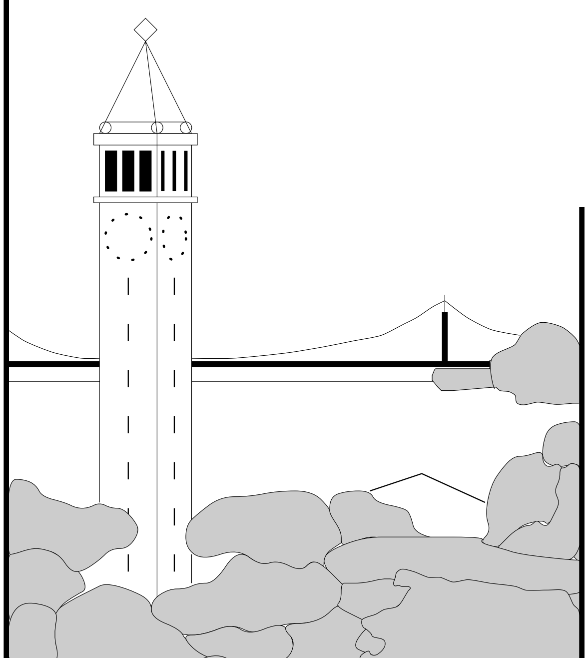

Report No. UCB/CSD-95-885

October 1995

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

Bounding Delays in Packet-Routing Networks with Light Traffic

Mor Harchol-Balter*

Abstract

If $\mathcal{N}$ is a queueing network and c_s is the mean service time at server s of $\mathcal{N}$, define $\mathcal{N}_{C,FCFS}$ (respectively, $\mathcal{N}_{E,FCFS}$) to be the queueing network $\mathcal{N}$ where the service time at server s is a constant c_s (respectively, an independent exponentially distributed random variable with mean c_s) and the packets are served in a first-come-first-served order.

Recently, Harchol-Balter and Wolfe introduced the problem of determining the class $\mathcal{S}$ of queueing networks $\mathcal{N}$ for which $\mathcal{N}_{C,FCFS}$ has smaller average delay than $\mathcal{N}_{E,FCFS}$. This problem has applications to bounding delays in packet-routing networks.

In this paper we consider the same problem, only restricted to the case of light traffic. We define $\mathcal{S}_{Light}$ to be the set of queueing networks $\mathcal{N}$ for which $\mathcal{N}_{C,FCFS}$ has smaller average delay than $\mathcal{N}_{E,FCFS}$ in the case of light traffic. We discover a sufficient criterion to determine whether a network $\mathcal{N}$ belongs to $\mathcal{S}_{Light}$, where this criterion is extremely simple and easy to check. Using this criterion we are able to show that many networks belong to $\mathcal{S}_{Light}$ that were previously not known to belong to $\mathcal{S}$. The significance of this result is that it suggests that many more networks are contained in $\mathcal{S}$ than has already been shown.

1 Introduction

Throughout this paper, whenever we refer to a queueing network, we will have in mind a network of servers where outside arrivals occur according to a Poisson Process and each outside arrival (packet) is born with a path (route) which it follows. Figure 1 illustrates an example of a possible routing scheme: Packets arrive into the network from outside at an average rate of one packet every 5 seconds. With probability 1/2, the packet has the path $a \rightarrow b \rightarrow c \rightarrow$; with probability 1/4 the packet has the path $a \rightarrow b \rightarrow a \rightarrow b \rightarrow$; with probability 1/4 the packet has the path $b \rightarrow c \rightarrow$.

*Supported by National Physical Science Consortium (NPSC) Fellowship. Also supported by NSF grant number CCR-9201092. harchol@cs.berkeley.edu

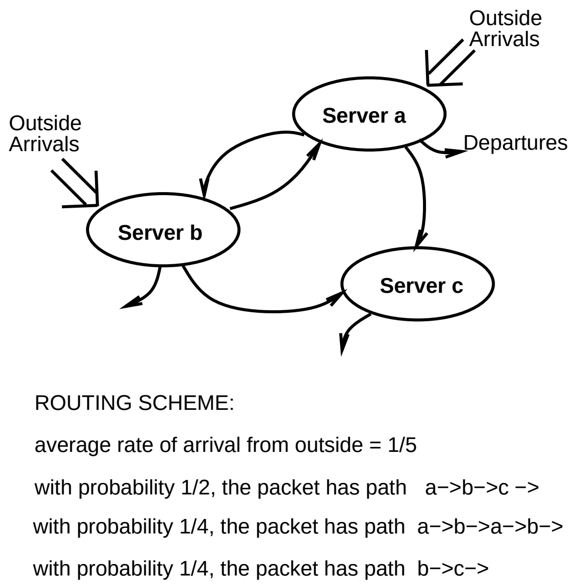

Figure 1: *In this paper, a queueing network denotes a network of servers together with a routing scheme.*

A queueing network is also characterized by the service time distribution associated with each server and the order in which packets are served at a server (the contention resolution protocol). If $\mathcal{N}$ is a queueing network and c_s is the mean service time at server s of $\mathcal{N}$, define $\mathcal{N}_{C,FCFS}$ (respectively, $\mathcal{N}_{E,FCFS}$) to be the queueing network $\mathcal{N}$ where the service time at server s is a constant c_s (respectively, an independent exponentially distributed random variable with mean c_s) and the packets are served in a first-come-first-served order. Likewise, define $\mathcal{N}_{C,PS}$ to be the queueing network $\mathcal{N}$ where the service time at server s is a constant c_s and the contention resolution protocol is processor-sharing.

Harchol-Balter and Wolfe, [3], demonstrate that many real-world packet-routing networks can be modeled by queueing networks of type $\mathcal{N}_{C,FCFS}$. It is therefore desirable to be able to compute the steady-state average packet delay of networks of type $\mathcal{N}_{C,FCFS}$. (The delay of a packet is defined as the total time the packet spends waiting in queues at servers from the time it is born until it reaches its destination.) Unfortunately, it is not known how to compute the average packet delay for all but the simplest $\mathcal{N}_{C,FCFS}$ type networks. However, the corresponding network of type $\mathcal{N}_{E,FCFS}$ is a product-form network (more specifically it can be modeled as a classed Jackson queueing network) and the average packet delay is easy to determine for networks of this type ([6], [2]). Harchol-Balter and Wolfe therefore ask the following question:

Is it possible to bound the average delay of $\mathcal{N}_{C,FCFS}$ (which we care about) by the average delay of $\mathcal{N}_{E,FCFS}$ (which we know how to compute)?

Let $\mathcal{S}$ denote the set of queueing networks $\mathcal{N}$ for which

$$\text{AvgDelay}(\mathcal{N}_{C,FCFS}) \leq \text{AvgDelay}(\mathcal{N}_{E,FCFS}). \tag{1}$$

Harchol-Balter and Wolfe give a simple proof that every network with Markovian routing is contained in $\mathcal{S}$. (In Markovian routing a packet's route is not contained within the packet, but rather there are probabilities on the edges leaving a server which determine all packets' routes. In other words, Markovian routing is classless.) They also demonstrate a network which is not contained in $\mathcal{S}$. They leave as an open problem the question of determining whether more networks are contained in $\mathcal{S}$.

In this paper we approach the problem of determining $\mathcal{S}$ by restricting ourselves to only networks with light traffic. Let $\lambda_{\mathcal{N}}$ denote the outside arrival rate into queueing network $\mathcal{N}$. $\mathcal{S}$ is the set of queueing networks $\mathcal{N}$ that satisfy equation(1) for all (stable) values of $\lambda_{\mathcal{N}}$. Define $\mathcal{S}_{Light}$ to be the set of queueing networks $\mathcal{N}$ that satisfy equation(1) in the case of light traffic, i.e., small $\lambda_{\mathcal{N}}$.

We give a simple sufficient criterion for whether a queueing network is in $\mathcal{S}_{Light}$. This simple criterion enables us to prove many networks belong to $\mathcal{S}_{Light}$ which haven't yet been shown to belong to $\mathcal{S}$.

By definition $\mathcal{S}$ is contained in $\mathcal{S}_{Light}$. However it seems likely that $\mathcal{S}_{Light}$ is also contained in $\mathcal{S}$, since it seems probable that AvgDelay($\mathcal{N}_{E,FCFS}$) should increase at a faster rate than AvgDelay($\mathcal{N}_{C,FCFS}$) as the traffic load is increased. Therefore, the significance of the above result is that it suggests that many more networks are contained in $\mathcal{S}$ than has already been proven.

Section 2 states the sufficient criterion theorem precisely and proves it. In Section 3 we discuss which queueing networks can easily be seen to satisfy the sufficient criterion.

2 Main Theorem

By [1] and [4], 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] Therefore it is equivalent to study for which queueing networks $\mathcal{N}$,

$$\text{AvgDelay}(\mathcal{N}_{C,FCFS}) \leq \text{AvgDelay}(\mathcal{N}_{C,PS}).$$

We will assume this formulation of the problem throughout the rest of the paper, since it simplifies our analysis.

[1] This powerful theorem is also described more recently in [6] and [5].

In this section we see that, speaking loosely, to test whether a queueing network $\mathcal{N}$ belongs to $\mathcal{S}_{Light}$ it is enough to check whether the expected delay created by exactly 2 packets in $\mathcal{N}_{C,FCFS}$ is smaller than the expected delay created by exactly 2 packets in $\mathcal{N}_{C,PS}$.

Theorem 1 *Given a queueing network, $\mathcal{N}$, if $\lambda_{\mathcal{N}} < min(\frac{1}{8e^2km^2}, s)$ then*

$$P_1 E_1^{FCFS} < \textit{AvgDelay}(\mathcal{N}_{C,FCFS}) < P_1(E_1^{FCFS} + \frac{1}{k})$$
$$P_1 E_1^{PS} < \textit{AvgDelay}(\mathcal{N}_{C,PS}) < P_1(E_1^{PS} + \frac{1}{k})$$

where

$$\begin{aligned} E_1^{FCFS} &= \mathbf{E}\{\textit{Delay on packet in } \mathcal{N}_{C,FCFS} \mid \textit{one other packet in } \mathcal{N}_{C,FCFS}\} \\ E_1^{PS} &= \mathbf{E}\{\textit{Delay on packet in } \mathcal{N}_{C,PS} \mid \textit{one other packet in } \mathcal{N}_{C,PS}\} \\ \lambda_{\mathcal{N}} &= \textit{the total arrival rate into } \mathcal{N} \textit{ from outside} \\ P_1 &= \mathbf{Pr}\{1 \textit{ arrival during } (-m, m)\} = e^{-\lambda_{\mathcal{N}} \cdot m}(\lambda_{\mathcal{N}} \cdot m) \\ k &= \textit{a free parameter} \geq 1. \\ s &= \textit{the max. value of } \lambda_{\mathcal{N}} \textit{ which } \mathcal{N} \textit{ can accept and still be stable} \\ m &= \textit{the length of the longest route in } \mathcal{N}\textit{'s routing scheme,} \\ & \quad\ \textit{where length is measured by total mean time in service.} \end{aligned}$$

Corollary 1 *Given a queueing network, $\mathcal{N}$, if $\lambda_{\mathcal{N}} < min(\frac{1}{8e^2km^2}, s)$ then*

$$P_1(E_1^{FCFS} - E_1^{PS}) - P_1\frac{1}{k} < \text{AvgDelay}(\mathcal{N}_{C,FCFS}) - \text{AvgDelay}(\mathcal{N}_{C,PS}) < P_1(E_1^{FCFS} - E_1^{PS}) + P_1\frac{1}{k},$$

where E_1^{FCFS}, E_1^{PS}, $\lambda_{\mathcal{N}}$, P_1, k, s, and m are as defined in the above theorem.

Corollary 2 *Given a queueing network, $\mathcal{N}$, if $\lambda_{\mathcal{N}} < min(\frac{1}{8e^2km^2}, s)$,*

$$\textit{if } E_1^{FCFS} < E_1^{PS} - \frac{1}{k}, \textit{ then}$$

$$\textit{AvgDelay}(\mathcal{N}_{C,FCFS}) < \textit{AvgDelay}(\mathcal{N}_{C,PS})$$

where E_1^{FCFS}, E_1^{PS}, $\lambda_{\mathcal{N}}$, P_1, k, s, and m are as defined in the above theorem.

A few comments on the above corollary before we begin the proof. First, observe that k is a free parameter of $\lambda_{\mathcal{N}}$. Therefore $\frac{1}{k}$ above can be made as small as we wish by decreasing $\lambda_{\mathcal{N}}$. Second, note that E_1^{FCFS} is an average. Therefore, it includes the case where the two packets happen to have the same route and both packets start within one service time of each other. In this case PS clearly does worse than FCFS, and we can make k large enough so that this difference exceeds $\frac{1}{k}$. Thus for light traffic, AvgDelay$(\mathcal{N}_{C,FCFS}) <$

AvgDelay($\mathcal{N}_{C,PS}$) whenever $E_1^{FCFS} < E_1^{PS}$, given that the two packets are on different paths.

Proof of Theorem:

By PASTA (Poisson Arrivals See Time Averages), the expected delay a newly arriving packet experiences is equal to the average packet delay for the network. Let $\mathcal{N}$ be any queueing network. For the case of light traffic (i.e., $\lambda_{\mathcal{N}} < min(\frac{1}{8e^2km^2}, s)$), we will compute upper and lower bounds on the delay an arrival experiences in $\mathcal{N}_{C,FCFS}$. The proof for $\mathcal{N}_{C,PS}$ is identical.

To compute an upper bound on the delay in $\mathcal{N}_{C,FCFS}$, let p represent an arriving packet in $\mathcal{N}_{C,FCFS}$. Clearly, p may only be delayed by packets which are in $\mathcal{N}_{C,FCFS}$ during the time p is in $\mathcal{N}_{C,FCFS}$. Note that if k packets are in $\mathcal{N}_{C,FCFS}$, they may take up to time kn to clear the system. So, if we call p's arrival time 0, packet p can only possibly be delayed if at least one of the following occur:

- exactly 1 other packet arrives during $(-m, m)$.
- exactly 2 other packets arrive during $(-2m, 2m)$.
- exactly 3 other packets arrive during $(-3m, 3m)$.
- 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.

Let

$$P_i = \mathbf{Pr}\{i \text{ arrivals during time } (-im, im)\}$$

Let

$$E_i^{FCFS} = \mathbf{E}\{\text{delay on } p \text{ due to } i \text{ arrivals during } (-im, im) \text{ in } \mathcal{N}_{C,FCFS}\}$$

So

$$\begin{aligned}\mathbf{E}\{\text{delay on } p \text{ in } \mathcal{N}_{C,FCFS}\} &\leq P_1 E_1^{FCFS} + P_2 E_2^{FCFS} + P_3 E_3^{FCFS} + \ldots \\ &\leq P_1 E_1^{FCFS} + P_2(2m) + P_3(3m) + \ldots\end{aligned}$$

where the last inequality is an over-estimate, since we are assuming the worst case where all the packets continually run into each other over and over again during their entire time in the network. By definition of the Poisson Process,

$$P_i = \frac{e^{-\lambda_{\mathcal{N}} \cdot 2im} (\lambda_{\mathcal{N}} \cdot 2im)^i}{i!}$$

For $i \geq 2$, we can express P_i in terms of P_1 as follows:

$$\begin{aligned} P_i(i \geq 2) &= \frac{e^{-\lambda_{\mathcal{N}} \cdot 2im} (\lambda_{\mathcal{N}} \cdot 2im)^i}{i!} \\ &= \frac{i^i}{i!} \cdot e^{-\lambda_{\mathcal{N}} \cdot 2im} (\lambda_{\mathcal{N}} \cdot 2m)^i \\ &< e^i \cdot e^{-\lambda_{\mathcal{N}} \cdot 2m} (\lambda_{\mathcal{N}} \cdot 2m)^i \\ &= P_1 \cdot (\lambda_{\mathcal{N}} \cdot 2m)^{i-1} \cdot e^i \end{aligned}$$

Substituting $\lambda_{\mathcal{N}} = \frac{1}{8e^2km^2}$, we have:

$$\begin{aligned} P_i(i \geq 2) &< P_1 \cdot (\lambda_{\mathcal{N}} \cdot 2m)^{i-1} \cdot e^i \\ &= P_1 \cdot (\frac{1}{4e^2km})^{i-1} \cdot e^i \\ &< P_1 \cdot \frac{1}{k4^{i-1}m} \end{aligned}$$

Now, substituting P_i, $i \geq 2$ into the formula for the expected delay on p, we have:

$$\begin{aligned} \mathbf{E}\{\text{delay on } p \text{ in } \mathcal{N}_{C,FCFS}\} &\leq P_1 E_1^{FCFS} + P_2(2m) + P_3(3m) + \ldots \\ &< P_1 E_1^{FCFS} + \frac{P_1}{2k} + \frac{P_1}{2^2k} + \frac{P_1}{2^3k} + \ldots \\ &= P_1 E_1^{FCFS} + P_1 \frac{1}{k} \\ &= P_1 \left(E_1^{FCFS} + \frac{1}{k} \right) \end{aligned}$$

To derive a simple lower bound for the expected delay in $\mathcal{N}_{C,FCFS}$, again 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 $(-m, m)$.

$$\mathbf{E}\{\text{delay on } p \text{ in } \mathcal{N}_{C,FCFS}\} \geq P_1 E_1^{FCFS}$$

□

3 Characterizing $\mathcal{S}_{Light}$

In Section 2, we found that to check whether $\mathcal{N}$ is in $\mathcal{S}_{Light}$ it is enough to check whether the expected delay created by exactly 2 packets in $\mathcal{N}_{C,FCFS}$ is smaller than the expected delay created by exactly 2 packets in $\mathcal{N}_{C,PS}$, when the packets take different paths. In this section we discuss which networks which satisfy this easy test.

Throughout this section, we will only consider the case where all servers have the same mean service time, 1. At a first glance it may appear that all such networks $\mathcal{N}$ satisfy the test. Here is the *wrong proof* behind that thought: We only need to look at the case of two packets in the network. The two packets don't affect each other at all until they bump into each other. Thus, until the packets bump into each other, $\mathcal{N}_{C,FCFS}$ and $\mathcal{N}_{C,PS}$ behave identically. In $\mathcal{N}_{C,FCFS}$, if the two packets do bump into each other, even if their routes now intersect for a long time, the packets will only interfere with each other for the duration of one server and then one packet will pass the other. Whereas, in $\mathcal{N}_{C,PS}$, when the packets bump into each other, they will slow each other down by the total length of the intersection of their routes.

This proof is mostly correct. The flaw is the claim that in $\mathcal{N}_{C,FCFS}$ the two packets will never see each other once they pass each other. In fact, the two routes may cross repeatedly as shown in Figure 2. Even if the two routes do cross repeatedly, though, it's tough to find examples where $\mathcal{N}_{C,FCFS}$ behaves badly in the case of two packets. The reason is as follows: If two packets interfere with each other at all in $\mathcal{N}_{C,FCFS}$, the next time that they meet (if ever) will be such that they both arrive at a server at the exact same time. When this happens, either packet could end up serving first with equal likelihood. Since we don't know which packet will go first it is difficult to construct routes which force the packets to meet again. A bad instance for $\mathcal{N}_{C,FCFS}$ must force the packets to meet again regardless of which packet served first in the previous collision. Figure 2 shows the worst-case instance for $\mathcal{N}_{C,FCFS}$ that we were able to come up with, yet even here, the expected number of collisions is only $O(\lg n)$, where , is the number of route crossings. Also, it's not clear that the corresponding $\mathcal{N}_{C,PS}$ wouldn't behave just as badly.

Thus for the case where all servers have the same service time, the expected delay given only two packets in the network is usually greater for $\mathcal{N}_{C,PS}$ than $\mathcal{N}_{C,FCFS}$, and when it is greater for $\mathcal{N}_{C,FCFS}$it's not greater by much.

In the case where the servers may have different service times it is possible to construct a network $\mathcal{N}$, where the expected delay in the case of two packets is $O(n)$ for $\mathcal{N}_{C,FCFS}$ and only $O(1)$ for $\mathcal{N}_{C,PS}$, where n is the number of servers in $\mathcal{N}$(see [3]).

4 Future Work

It would be useful to characterize more precisely exactly which networks satisfy the criterion from Section 2.

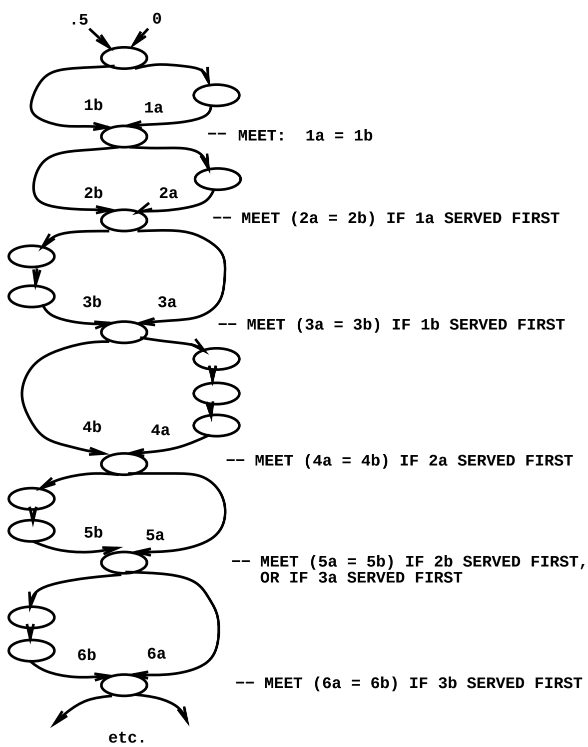

Figure 2: *Bad example for* $\mathcal{N}_{C,FCFS}$ *with 2 packets. The number of collisions is* $O(\lg n)$ *where* n *is the number of route crossings. At each same-time collision, regardless of which packet served first, the packets are guaranteed to meet again some time down the chain.*

References

[1] 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.

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

[3] Mor Harchol-Balter and David Wolfe. Bounding delays in packet-routing networks. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing*, pages 248–257, May 1995.

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

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

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