Bounding Delays in Packet-Routing Networks with Light Traffic

Mor Harchol-Balter

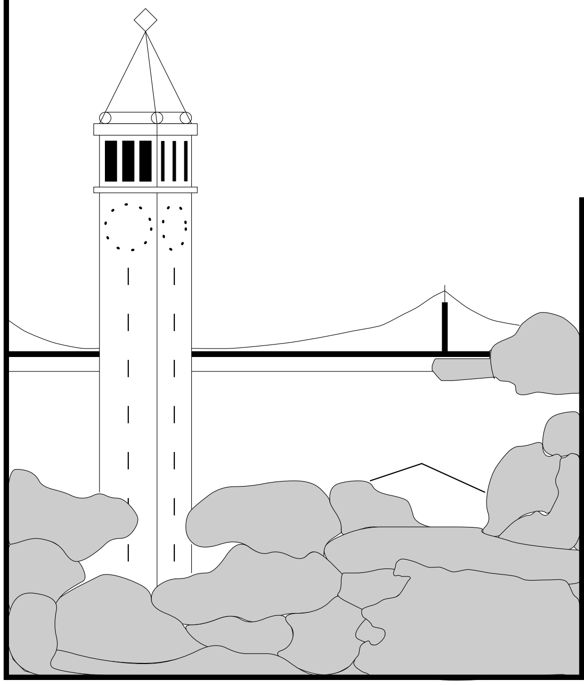

Report No. UCB//CSD-95-885

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 $a \to b \to c \to$; with probability 1/4 the packet has the path $a \to b \to a \to b \to$; with probability 1/4 the packet has the path $b \to c \to$.

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

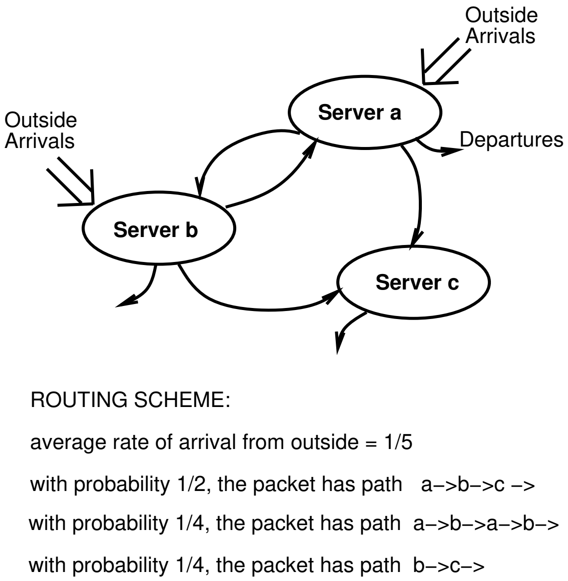

Figure 1: *In this paper, a queueing network denotes a network of servers together with a routing scheme.*

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 ([8], [2]). Harchol-Balter and Wolfe therefore pose 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.

Reiman and Simon ([6], [7]) prove a light traffic result which is somewhat similar to ours, as well as many other light traffic results. Their technique involves computing derivatives of the delay at the point of zero arrival rate. The light traffic proof techniques we use are much simpler, and will hopefully lead to other simple light traffic analysis.

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 $\mathcal{N}$,

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

[1] This powerful theorem is also described more recently in [8] and [5].

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

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 $\lambda_{\mathcal{N}} < \frac{1}{8e^2km^2}$ then*

$$\begin{array}{rcl} P_1 D_1^{FCFS} < & \textit{AvgDelay}(\mathcal{N}_{C,FCFS}) & < P_1(D_1^{FCFS} + \frac{1}{k}) \\ P_1 D_1^{PS} < & \textit{AvgDelay}(\mathcal{N}_{C,PS}) & < P_1(D_1^{PS} + \frac{1}{k}) \end{array}$$

where

$$\begin{array}{rcl} D_1^{FCFS} & = & \mathbf{E}\{\textit{Delay on packet in } \mathcal{N}_{C,FCFS} \mid \textit{one other packet in } \mathcal{N}_{C,FCFS}\} \\ D_1^{PS} & = & \mathbf{E}\{\textit{Delay on a packet in } \mathcal{N}_{C,PS} \mid \textit{one other packet in } \mathcal{N}_{C,PS}\} \\ \lambda_{\mathcal{N}} & = & \textit{the total arrival rate into } \mathcal{N} \textit{ from outside} \\ P_1 & = & \mathbf{Pr}\{1 \textit{ outside arrival during } (-m,m)\} = e^{-\lambda_{\mathcal{N}} \cdot m}(\lambda_{\mathcal{N}} \cdot m) \\ k & = & \textit{a free parameter} \geq 1. \\ m & = & \textit{the length of the longest route in } \mathcal{N}\textit{'s routing scheme,} \\ & & \textit{where length is measured by total mean time in service.} \end{array}$$

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

$$P_1(D_1^{FCFS} - D_1^{PS}) - P_1\frac{1}{k} < \mathrm{AvgDelay}(\mathcal{N}_{C,FCFS}) - \mathrm{AvgDelay}(\mathcal{N}_{C,PS}) < P_1(D_1^{FCFS} - D_1^{PS}) + P_1\frac{1}{k},$$

where D_1^{FCFS}, D_1^{PS}, $\lambda_{\mathcal{N}}$, P_1, k, and m are as defined in the above theorem.

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

$$\textit{if } D_1^{FCFS} < D_1^{PS} - \frac{1}{k}, \textit{ then}$$

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

where D_1^{FCFS}, D_1^{PS}, $\lambda_{\mathcal{N}}$, P_1, k, 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 D_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 unit of each other. For this particular case PS clearly does worse than FCFS; call the difference in delay for this case δ. Since we can always choose k in $\lambda_{\mathcal{N}}$ such that $\delta > \frac{1}{k}$, we can state a simple consequence of Corollary 2, namely:

For light enough traffic, AvgDelay($\mathcal{N}_{C,FCFS}$) $<$ AvgDelay($\mathcal{N}_{C,PS}$) whenever $D_1^{FCFS} < D_1^{PS}$, 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., $\lambda_{\mathcal{N}} < \frac{1}{8e^2km^2}$), 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 i packets are in $\mathcal{N}_{C,FCFS}$, they may take up to time im to clear the system. So, denoting p's arrival time by 0, if packet p is delayed, at least one of the following must occur:

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

Define

E'_i : the event that *at least* i packets arrive during $(-im, im)$

E_i : the event that *exactly* i packets arrive during $(-im, im)$

Now p can only be delayed is at least one of the E'_i occur, i.e., if $\bigcup E'_i$ is true. However $\bigcup E'_i$ can only occur if $\bigcup E_i$ occurs (See footnote[2] for a proof of this subtle point), so p can only be delayed if at least one of the following events occurs:

- 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.

[2] Let $x(i)$ denote the number of arrivals during $(-im, im)$. Observe that $x(i)$ is a non-decreasing integer-valued function of i. Let $\mathcal{L}$ be the line $x(i) = i$. Since $\mathbf{E}\{x(i)\}$ is less than 1, if x(i) is ever above $\mathcal{L}$, with probability 1 it must eventually cross $\mathcal{L}$ and come below $\mathcal{L}$ (by the Law of Large Numbers). Thus if $\bigcup E'_i$ is true, then so is $\bigcup E_i$.

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

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

Let

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

So

$$\begin{aligned}\mathbf{E}\{\text{delay on } p \text{ in } \mathcal{N}_{C,FCFS}\} &\leq P_1 D_1^{FCFS} + P_2 D_2^{FCFS} + P_3 D_3^{FCFS} + \ldots \\ &\leq P_1 D_1^{FCFS} + P_2(2m) + P_3(3m) + \ldots\end{aligned}$$

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

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

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

$$\begin{aligned}\mathbf{E}\{\text{delay on } p \text{ in } \mathcal{N}_{C,FCFS}\} &\leq P_1 D_1^{FCFS} + P_2(2m) + P_3(3m) + \ldots \\ &< P_1 D_1^{FCFS} + \frac{P_1}{2k} + \frac{P_1}{2^2 k} + \frac{P_1}{2^3 k} + \ldots\end{aligned}$$

$$= \quad P_1 D_1^{FCFS} + P_1 \frac{1}{k}$$
$$= \quad P_1 \left(D_1^{FCFS} + \frac{1}{k} \right)$$

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 $(-m, m)$.

$$\mathbf{E}\{\text{delay on } p \text{ in } \mathcal{N}_{C,FCFS}\} \geq P_1 D_1^{FCFS}$$

□

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. We make two points:

1. We demonstrate a class of networks with *non-Markovian* routing which satisfies the test for being in $\mathcal{S}_{Light}$. (Because the routing is not Markovian, the result in [3] on networks in $\mathcal{S}$ does not apply to these networks.)
2. In [3], a network $\mathcal{N}$ is constructed for which the expected delay in the case of two packets is significantly greater in $\mathcal{N}_{C,FCFS}$ than in $\mathcal{N}_{C,PS}$. This property is highly dependent on the fact $\mathcal{N}$'s servers have different service times. We therefore pose the question of whether *all* networks where the servers have the *same* service time are in $\mathcal{S}_{Light}$ (or almost in $\mathcal{S}_{Light}$). We give some arguments for why it is likely that the answer to this question is yes.

Consider the class of networks (with non-Markovian routing) with the following property: any two routes which intersect and then split up can never subsequently rejoin. It's easy to see that these networks are all in $\mathcal{S}_{Light}$: 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. The first time the two packets do bump into each other will therefore occur at exactly the same time and will be at exactly the same location in $\mathcal{N}_{C,FCFS}$ and $\mathcal{N}_{C,PS}$. The delay incurred on the packets during the period of intersection will be smaller in $\mathcal{N}_{C,FCFS}$ than in $\mathcal{N}_{C,PS}$. After the routes split up, the two packets will never again see each other in either network. Thus the delay in every case of two

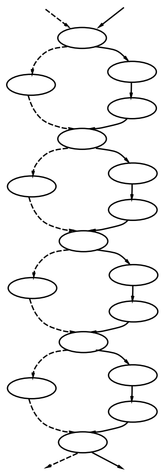

Figure 2: *Example where $\mathcal{N}_{C,FCFS}$ behaves slightly worse on average than $\mathcal{N}_{C,PS}$, given exactly two packets are in the network, and all service times are equal. Packets follow the dashed route or the solid route with equal probability. In both networks two packets on opposite routes are likely to meet because the dashed-route packet will likely catch up to the solid-route packet regardless of the original arrival times of the packets. When the two packets do catch up, in $\mathcal{N}_{C,PS}$ the packets will interfere with each other at most for the duration of one server. In $\mathcal{N}_{C,FCFS}$, with probability $\frac{1}{2}$, the the packets meet once more and with probability $\frac{1}{4}$ they meet twice more, etc., thus the expected delay for $\mathcal{N}_{C,FCFS}$ is slightly higher than for $\mathcal{N}_{C,PS}$.*

packets is smaller in $\mathcal{N}_{C,FCFS}$ than $\mathcal{N}_{C,PS}$, so certainly the average delay for the case of two packets is smaller in $\mathcal{N}_{C,FCFS}$ than $\mathcal{N}_{C,PS}$.

If in the above proof the routes had crossed repeatedly, it is possible that the two packets in $\mathcal{N}_{C,FCFS}$ might meet repeatedly with high probability and the two packets in $\mathcal{N}_{C,PS}$ might never meet again, resulting in greater average delay for $\mathcal{N}_{C,FCFS}$. [3] shows such an example involving servers with different service times, where the expected delay in the case of two packets is $\Theta(n)$ for $\mathcal{N}_{C,FCFS}$ and only $\Theta(1)$ for $\mathcal{N}_{C,PS}$, where n is the number of servers in $\mathcal{N}$. We now question whether this could still happen in the case where all servers have the same mean service time, 1.

For networks where all servers have service time 1, even if the routes do cross repeatedly, it's tough to find examples where $\mathcal{N}_{C,FCFS}$ behaves badly in the case of two packets. The reason is twofold:

1. To guarantee that the *average* delay is high in the case of two packets for $\mathcal{N}_{C,FCFS}$, we need to be able to guarantee that with high probability the two packets will meet at least once.

2. 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 probability. 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 repeatedly regardless of which packet served first in the previous collision.

We were able to come up with a network which addressed issue 1 above but didn't address issue 2 well, see Figure 2. We were also able to come up with a network which addressed issue 2 above, but not issue 1 (i.e., we constructed an $\mathcal{N}$ such that, given a collision, the expected delay in $\mathcal{N}_{C,FCFS}$ was $O(\lg n)$, where n was the number of servers in $\mathcal{N}$, but the probability of two packets colliding in the first place was very low). We hypothesize that in the case where all packets have the same service time $\mathcal{N}_{C,FCFS}$ never behaves significantly worse than $\mathcal{N}_{C,PS}$.

4 Future Work

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

5 Acknowledgements

We thank David Wolfe for his help in clarifying some subtle points in the analysis. Thanks also to Micah Adler for many useful discussions.

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] Martin I. Reiman and Burton Simon. Light traffic limits of sojourn time distributions in markovian queueing networks. *Communications in Statistics - Stochastic Models*, 4(2):191–233, 1988.

[7] Martin I. Reiman and Burton Simon. Open queueing systems in light traffic. *Mathematics of Operations Research*, 14(1):26–59, 1989.

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