

Optimal buffer allocation in tandems of last come first served queues

V. Anantharam^{a,1} and A. Ganesh^{b,2}

^a*Department of Electrical Engineering and Computer Science,
University of California, Berkeley, CA 94720, USA*

^b*Department of Computer Science, University of Edinburgh,
Edinburgh EH9 3JZ, UK*

Received 23 September 1994; revised 21 March 1995

A large fixed number of buffer spaces is given. We consider the problem of allocating these spaces among the nodes of a tandem of last-come-first-served queues with general service time distributions and Poisson external arrivals so as to optimize some performance criterion associated with the time to buffer overflow, such as maximizing its mean or maximizing the probability that it exceeds some value. Consider the following rule of thumb: allocate the buffer spaces in inverse proportion to the logarithms of the effective service rates at the nodes. Here effective service rate denotes the ratio of the service rate to the stationary arrival rate. We prove that this rule of thumb achieves a nearly optimal buffer allocation under the assumption that the service time distributions satisfy an exponential tail condition. This problem has been studied earlier in the context of Jackson networks, where it was shown that the same rule of thumb achieves an allocation that is close to optimal. The technique of proof here is similar, but there are important differences. Both Jackson networks and the LCFS tandems considered here are product form networks (with infinite buffers). Optimism should lead us to expect that the near optimality of this rule of thumb holds much more generally for product-form networks, but this remains a conjecture at present.

Keywords: Buffer allocation, LCFS queues, queueing networks, tandem queues.

1. Introduction

Networks of queues are used as models for the queueing processes taking place in manufacturing systems, computer networks, and communication networks. An important problem in these applications is how best to allocate buffer spaces among the nodes of the network so as to avoid frequent overflows. Indeed, buffer overflow often has catastrophic consequences in the applications above.

In this note we address the problem using an asymptotic approach. That is, we assume that the number of buffer spaces to be allocated is relatively large.

¹ Research supported by NSF under NCR 8857731, by AT&T, and by Bellcore Inc.

² Research supported by IBM under a graduate fellowship.

Our network model will be a tandem of last-come-first-served (LCFS) queues with general independent and identically distributed (i.i.d.) service times and Poisson input. Tandems of LCFS queues with infinite buffers are product form queueing networks (see Kelly [5]), admitting a particularly simple form of stationary probability distribution. The problem studied is how to allocate a fixed amount N of total buffer space among the nodes of such a network so as to optimize some performance criterion associated with the time to buffer overflow, such as maximizing its mean or maximizing the probability that it exceeds some value.

In the context of the Jackson network, which is also a product form model, this problem has been studied by several researchers. A rule of thumb for that case is proposed in Anantharam [1], and analyzed there and in Anantharam and Ganesh [2]. Several simulation based approaches have also been tried, see, for example Cottrell et al. [3], Ho et al. [4], and Parekh and Walrand [6, 7]. Here, we make a start in the direction of extending the results in [1] to non-Jacksonian product form queueing networks.

In [1] the following rule of thumb for this problem was considered: allocate the buffer space in inverse proportion to the logarithms of the effective service rates at the nodes, where effective service rate means the ratio of the service rate to the stationary arrival rate in the network with infinite buffers. It was shown that this rule of thumb is accurate to within a constant times $\log N$ as the amount N of buffer space to be allocated becomes large. That is, for any reasonable cost function associated with the time to buffer overflow, the allocation that is optimal for this cost function is within a constant times $\log N$ of the rule of thumb. One way to interpret this result is that the size of the search space to determine the actual buffer allocation is reduced from N^{d-1} to $(\log N)^d$, where d denotes the number of nodes in the network.

Our contribution in this work is to obtain the identical result in the case of tandems of $M/G/1$ -LCFS queues, that is, queues with general i.i.d. service times, last-come-first-served queueing discipline, and Poisson external arrivals into the first queue in the tandem. We need to assume, however, that the service time distributions satisfy an exponential tail condition.

In section 2 we introduce some notation and state the problem more formally. We state our main theorem in section 3 and use it to obtain our result about the approximate optimality of the rule of thumb. Section 4 deals with the proof of the theorem stated in section 3.

The reason that the problem of extending the validity of the rule of thumb to general product-form networks is significantly different from the case of Jackson networks is that in general such networks lack an important monotonicity property. Increasing the number of customers initially in the system does not necessarily increase the amount of work done in the system upto an arbitrary time. For example, in the case of LCFS networks, the extra customers may cause extra pre-empts and thus prevent work from moving smoothly through the system. Hence, the techniques developed for Jackson networks do not carry over in a simple

way. In particular, the method used to prove lemma 1 in section 4 of this paper is very different from that used in Anantharam [1] to establish a similar lemma for Jackson networks.

2. Problem formulation

The model we consider in this paper is a collection of d nodes or service centers, connected in series with no feedback. Customers arrive from the external world into the first queue according to a Poisson process of rate λ . On completing service at node i , $1 \leq i \leq d - 1$, the customer enters node $i + 1$. On completing service at node d , the customer leaves the system. Service times are assumed to have a general distribution, i.i.d. at each node, independent from node to node, and also independent of the arrival process. The mean service time at node i is denoted $1/\mu_i$. We define the quantity $\rho_i \triangleq \lambda/\mu_i$ and assume that $\rho_i < 1$ for all i . This is a natural requirement if the network is to operate for reasonably long periods of time without buffer overflow. The service discipline at the nodes is assumed to be last-come-first-served (LCFS). In particular, if a customer enters a queue while a service is taking place, then the service in progress is interrupted and the server starts work on the new customer. After completing service on new arrivals, the server resumes service on the interrupted customer, that is, the service already delivered to that customer is not wasted.

Let $X(t) = ((Q_1(t), B_1(t)), \dots, (Q_d(t), B_d(t)))$ denote the state process. Here $Q_i(t)$ is the queue length at node i including the customer undergoing service, and $B_i(t)$ is the vector of service times already delivered to the customers currently at node i . We construct the stochastic process $X(t)$ on a sample space consisting of a Poisson arrival processes and i.i.d. sequences of service requirements at the individual nodes which are handed out to customers in the order of their arrival and assumed to be independent from node to node and also independent of the arrival process. We define the filtration \mathcal{F}_t to be the sigma algebra generated by $(X_s, 0 \leq s \leq t)$. It follows that the residual service time vector $Y_i(t)$ for the customers in queue i at time t satisfies

$$E[Y_i(t)|\mathcal{F}_t] = E[Y_i(t)|B_i(t)]. \tag{2.1}$$

It also follows from the above construction that $((X(t), \mathcal{F}_t), t > 0)$ is a Markov process.

It is known for this model that the network is stable under the assumption $\rho_i < 1$ for all i , $1 \leq i \leq d$. Further, the stationary distribution of the number of customers in the system has a particularly simple form (see Kelly [5]), and is given by

$$\pi(q_1, \dots, q_d) = \prod_{i=1}^d \rho_i^{q_i} (1 - \rho_i). \tag{2.2}$$

We shall need the following assumption about the residual service time distribution:

$$E[Y_i(t)|B_i(t)] \leq \beta, \quad 1 \leq i \leq d, \quad (2.3)$$

where β is a deterministic constant that does not depend on $B(t)$. Here $B_i(t)$ denotes the amount of service already delivered to a customer in queue i , and $Y_i(t)$ the residual service that needs to be delivered to that customer. In other words, the mean residual service time is uniformly bounded irrespective of the amount of service already delivered. It can be shown that this assumption is equivalent to the assumption that the moment generating function of the service time distribution is finite in an open neighborhood of 0.

We are interested in the problem of assigning N buffer spaces to the nodes of this network so as to optimize some cost function associated to the time to buffer overflow of the network started empty. Call (N_1, N_2, \dots, N_d) an allocation if $\sum_{i=1}^d N_i = N$ and each N_i is a positive integer. Given an allocation, define the set B as follows:

$$B = \{(x_1, \dots, x_d) : q_i = N_i + 1 \text{ for some } 1 \leq i \leq d, \\ q_j \leq N_j \text{ for all } j \neq i, 1 \leq i \leq d\}, \quad (2.4)$$

where $x_i = (q_i, b_i)$ for $1 \leq i \leq d$. We call B the boundary for this allocation. Then, the time for a buffer overflow to occur in the system started empty is the same as the time for the Markov chain $X(t)$, started at the origin, to hit the set B . We shall show that, for *any* performance measure associated with this time, the rule of thumb stated below is approximately optimal in the sense that it lies within a constant times $\log N$ of the exact optimal allocation.

3. Main result

The following rule of thumb was considered in Anantharam [1]: Given N buffer spaces, allocate roughly $p_i N$ spaces to node i , where the fraction p_i is such that $c \triangleq p_i \log \rho_i^{-1}$ is a constant independent of i , and $\sum_{i=1}^d p_i = 1$. In other words, $p_i = -c / \log \rho_i$ where $c = -(\sum_{i=1}^d \log \rho_i)^{-1}$. It was shown that if (N_1, \dots, N_d) is an optimal allocation, it must be close to the rule of thumb in the sense that

$$|N_i - p_i N| < K \log N, \quad 1 \leq i \leq d, \quad (3.1)$$

for a constant K independent of N , and large enough N . More precisely, if any allocation differs from the rule of thumb by more than the above amount along a subsequence of $N \rightarrow \infty$, it performs worse than the rule of thumb along that subsequence. We show the identical result to be true for our model as well.

We now introduce some notation. σ denotes an infinite subset of the positive integers. We write $\lim_{N \rightarrow \infty}^\sigma$ for N going to ∞ along the subsequence σ . Given two functions $f(N)$ and $g(N)$ on the positive integers, we write $f(N) = o_\sigma(g(N))$ if $\lim_{N \rightarrow \infty}^\sigma f(N)/g(N) = 0$. We write $f(N) = \omega_\sigma(g(N))$ if $g(N) = o_\sigma(f(N))$. Finally, all logarithms are assumed to be to the base 2.

Let T have the distribution of the time to buffer overflow when the system is started empty. This depends on N and the buffer allocation but this dependence will be suppressed from the notation. In the next section, we prove the following theorem. Here a *buffer allocation scheme* consists of a buffer allocation (N_1, \dots, N_d) for each N .

THEOREM 1

For any buffer allocation scheme $(N_1, \dots, N_d), \sum_{i=1}^d N_i = N$, let $g(N) \triangleq \min_i (1/\rho_i)^{N_i}$. Then, for any subsequence σ and any $\tau(N)$ such that $\tau(N) = o_\sigma(g(N))$, we have

$$\lim_{N \rightarrow \infty}^\sigma P(T \leq \tau(N)) = 0. \tag{3.2}$$

On the other hand, for any $T(N)$ such that $T(N) = \omega_\sigma(Ng(N))$, we have

$$\lim_{N \rightarrow \infty}^\sigma P(T \leq T(N)) = 1. \tag{3.3}$$

Here we use this theorem to justify the rule of thumb stated above. This short argument is identical to that in Anantharam [1] but is repeated for the convenience of the reader.

For any buffer allocation (N_1, \dots, N_d) , denote $2^{-cN}Ng(N)$ by $h^2(N)$. Here c is the constant defined in the statement of the rule of thumb. Suppose $\liminf_{N \rightarrow \infty} 2^{-cN}Ng(N) = 0$. Then there is a subsequence σ such that $2^{-cN}Ng(N) = o_\sigma(1)$. Then also $h(N) = o_\sigma(1)$. If T and T^* have the distribution of the time to buffer overflow for this buffer allocation scheme and for the rule of thumb respectively, then we have

$$\lim_{N \rightarrow \infty}^\sigma P(T \leq 2^{cN}h(N)) = \lim_{N \rightarrow \infty}^\sigma P\left(T \leq \frac{Ng(N)}{h(N)}\right) = 1 \tag{3.4}$$

by (3.3), whereas

$$\lim_{N \rightarrow \infty}^\sigma P(T^* \leq 2^{cN}h(N)) = 0 \tag{3.5}$$

by (3.2). Clearly the buffer allocation scheme is eventually worse than the rule of thumb along σ . Modifying the scheme along σ by replacing it with the rule of thumb

is an improvement for large enough N . This suggests that any reasonable buffer allocation scheme must satisfy

$$\liminf_{N \rightarrow \infty} 2^{-eN} Ng(N) > \epsilon$$

for some $\epsilon > 0$. It is a straightforward calculation to see that this implies that

$$|N_i - p_i N| < K \log N, \tag{3.6}$$

for all sufficiently large N , for all $i = 1, \dots, d$, where $K \triangleq \sum_{i=1}^d (\log(1/\rho_i))^{-1}$. In this sense, the rule of thumb is approximately optimal.

4. Proof of theorem 1

Consider any buffer allocation (N_1, \dots, N_d) , $\sum_{i=1}^d N_i = N$, and the associated boundary B as defined in (2.4). Recall that $X(t) = (Q(t), B(t))$ was defined as the state process, where $Q(t)$ is the vector queue length process. Let α denote the probability that the Markov process $X(t)$, starting at the origin, hits B before returning to zero.

$$\alpha \triangleq P_0(X(t) \text{ hits } B \text{ before returning to } 0).$$

Let $\pi(B)$ denote the stationary probability of the set B , i.e.

$$\pi(B) = \sum_{\mathbf{q}: \mathbf{x}=(\mathbf{q}, \mathbf{b}) \in B} \pi(\mathbf{q}), \tag{4.1}$$

where $\pi(\mathbf{q})$ is given in (2.2) and we have used the fact that the definition of the set B involves only the queue lengths, \mathbf{q} . We now relate α to the time to buffer overflow using pathwise probabilistic arguments. The approach closely follows that in Anantharam [1].

Consider a path of the process starting at 0 and ending when it hits B for the first time. Call this time T . The path consists of a certain number of cycles, ν , where the process returns to 0 without hitting B , and a last segment where it hits B before returning to 0. The ν cycles defined above are i.i.d., independent of ν , and ν is a geometric random variable; these assertions are a simple consequence of the fact that the network evolution regenerates when it hits 0 because of the assumption of Poisson arrivals. Thus, we have

$$P(\nu = k) = \alpha(1 - \alpha)^k \tag{4.2}$$

and

$$E\nu = \frac{1 - \alpha}{\alpha}. \tag{4.3}$$

Let δ have the distribution of the time taken to return to 0, starting from 0, having had customers, and conditioned on not visiting B . Let $\delta_1, \delta_2, \dots$ be i.i.d with the distribution of δ . Let Δ have the distribution of the time to hit B , starting from 0 and conditioned on not returning to 0. Also assume that Δ is independent of $(\delta_n, n \geq 1)$. Then, we also have

$$T \stackrel{d}{=} \sum_{k=1}^{\nu} \delta_k + \Delta, \tag{4.4}$$

where T is the time $X(t)$ first hits B , and $\stackrel{d}{=}$ denotes equality in distribution. In particular,

$$\begin{aligned} ET &= \frac{1 - \alpha}{\alpha} E\delta + E\Delta \\ &= \frac{1}{\alpha} [(1 - \alpha)E\delta + \alpha E\Delta]. \end{aligned} \tag{4.5}$$

It is easy to see that $(1 - \alpha)E\delta + \alpha E\Delta$ is the mean time taken to either return to 0 or visit B , starting from 0. This is less than the mean time to return to 0 starting from 0 and after having had customers in the network with infinite buffers. Call this latter time T_0 . Clearly, the distribution of T_0 is independent of N and the buffer allocation, and it is easy to verify that T_0 has finite mean and variance. Then, from (4.5),

$$ET \leq \frac{1}{\alpha} ET_0 \tag{4.6}$$

and so, by Markov's inequality,

$$P(T \geq \tau(N)) \leq \frac{ET}{\tau(N)} \leq \frac{1}{\alpha\tau(N)} ET_0. \tag{4.7}$$

Also observe that δ stochastically dominates an exponential random variable of mean λ^{-1} . Indeed, if the network is empty, we have to wait at least that long for an arrival. From (4.2) and (4.4), and since an independent geometric sum of independent exponential random variables is exponential, it follows that T stochastically dominates an exponential random variable of mean $(1 - \alpha)/\alpha\lambda$. Hence

$$P(T \leq \tau(N)) \leq 1 - \exp\left(-\frac{\alpha\lambda\tau(N)}{1 - \alpha}\right). \tag{4.8}$$

Notice that α is a function of N and the buffer allocation though the dependence has been suppressed in the notation. Keeping this in mind, we consider a subsequence $\sigma \rightarrow \infty$ of natural numbers. Then, from (4.7) we see that

$$\tau(N) = \omega_\sigma(\alpha^{-1}) \Rightarrow P(T \geq \tau(N)) \xrightarrow{\sigma} 0. \quad (4.9)$$

Likewise, from (4.8) we observe that

$$\tau(N) = o_\sigma(\alpha^{-1}) \Rightarrow P(T \leq \tau(N)) \xrightarrow{\sigma} 0. \quad (4.10)$$

It now remains to relate α to the quantity $g(N)$ defined in the last section.

Consider a cycle started with the network empty and ending when the network empties again after having had customers. Notice that at the end of such a cycle the network regenerates because of the assumption of Poisson arrivals. In other words, the future evolution of the network after it becomes empty is independent of its entire past. It then follows that the paths of the system during such cycles are i.i.d., and hence, by a standard regenerative argument, we have

$$\pi(B) = \frac{\text{mean time spent in } B \text{ in a busy cycle}}{E_0 T_0}, \quad (4.11)$$

where $E_0 T_0$ denotes the mean time for the network to return to 0 after having had a customer. But the mean time spent in B in such a cycle is the product of α and the mean time spent in B until the network empties, starting in the conditional hitting distribution on B . Here α is as defined earlier.

Observe that the set B is entered because of an arrival into some queue, and can be left only when the number in this queue decreases by one. Because of the LCFS service discipline, this takes time at least equal to the total service time of the entering customer. Therefore, the mean time spent in B starting from the conditional hitting distribution until the system empties, is at least equal to the mean service time of the customer entering the queue that undergoes an overflow. Hence,

$$\pi(B) \geq \frac{\alpha \min_{i=1}^d \mu_i^{-1}}{E_0 T_0}. \quad (4.12)$$

Observe also that the mean time spent in B starting from the conditional hitting distribution, until the system empties, is dominated by the mean time for the system to empty starting from the conditional hitting distribution. In the lemma below, we show that, if the system is started with an arbitrary initial number of customers M , then the mean time for it to empty is dominated by a constant times

M , where the constant will depend on the network parameters. When the queue length process hits B , it is clear from the definition of B that the number of customers in the system is at most $N + 1$. Using the lemma below, it then follows that the mean time to empty starting from the conditional hitting distribution on B is at most a constant times $N + 1$. Consequently,

$$\pi(B) \leq \frac{\alpha k N}{E_0 T_0} \tag{4.13}$$

for some constant k .

Since $E_0 T_0$ is some finite constant which doesn't depend on n , it follows from (4.12) and (4.13) that

$$\frac{k_1 \pi(B)}{N} \leq \alpha \leq k_2 \pi(B), \tag{4.14}$$

where k_1, k_2 are constants that don't depend on N . It is also easily seen from (2.2) and the definition of B that

$$k_1 g(N)^{-1} \leq \pi(B) \leq k_2 g(N)^{-1},$$

for some constants k_1, k_2 . Substituting this in (4.14) gives

$$k_1 g(N) \leq \alpha^{-1} \leq k_2 N g(N), \tag{4.15}$$

for some constants k_1 and k_2 . Finally, the theorem is easily seen to follow from (4.9), (4.10) and (4.15).

LEMMA 1

The mean time for a stable tandem of LCFS queues with Poisson arrivals to empty, when started with a total of N customers initially in the system, is bounded above by a constant times N , provided the exponential tail condition on the service time distribution stated in (2.3) holds. The constant may depend on the network parameters.

Proof

The proof of the lemma is by induction on the number of nodes in the tandem. We first establish the basis case, which is $d = 1$.

Color the customers initially in queue blue, and color new arrivals red. It is clear from the LCFS nature of service discipline that blue customers are worked on only when there are no red customers in the queue. Thus, the process of red customers is seen to evolve like the process of the same LCFS queue, started empty. Since the queue was assumed to be stable, the process of red customers has regenerative i.i.d. busy cycles

of finite mean length. Every busy period is followed by a random exponential time of mean $1/\lambda$ during which no red customers arrive, and therefore blue customers are served. Define $S_0 = 0$. For $k \geq 1$, let T_k denote the first time after S_{k-1} that a red customer enters the system, and S_k the first time after T_k that the system is empty of red customers. Let $t_k = T_k - S_{k-1}$ and $s_k = S_k - T_k$. Define

$$\kappa = \inf \left\{ k > 0 : \sum_{i=1}^k t_i \geq \sum_{i=1}^{Q(0)} Y_i \right\},$$

where Y_i is the residual service time needed for the i th customer initially in queue. Let τ denote the time for the queue to empty. Then it is clear from above that

$$\tau = \sum_{i=1}^{Q(0)} Y_i + \sum_{i=1}^{\kappa-1} s_i,$$

where, furthermore, the s_i are i.i.d., distributed like the length of a busy period, and independent of κ . Therefore

$$E[\tau | \mathcal{F}_0] = \sum_{i=1}^{Q(0)} E[Y_i | \mathcal{F}_0] + E[\kappa - 1 | \mathcal{F}_0] E[s_i]. \tag{4.16}$$

Now condition on $\sum_{i=1}^{Q(0)} Y_i = y$, and observe that

$$P(\kappa > k) = P\left(\sum_{i=1}^k t_i < y\right) \leq e^{\theta y} E[e^{-\theta t_i}]^k = e^{\theta y} \left(\frac{\lambda}{\lambda + \theta}\right)^k$$

for all $\theta > 0$, where the inequality is by Chernoff's inequality. Summing over k , and choosing θ to be $1/y$ for simplicity, we get

$$E[\kappa | \mathcal{F}_0] \leq e \left(1 + \lambda E \left[\sum_{i=1}^{Q(0)} Y_i | \mathcal{F}_0 \right] \right).$$

Finally, by our assumption on the residual service times, it follows that

$$E[\kappa | \mathcal{F}_0] \leq e(1 + \lambda \beta Q(0)).$$

Substituting this in (4.16) and using the assumption on the residual service times once more gives

$$E[\tau | \mathcal{F}_0] \leq eE[s_i] + \beta(1 + \lambda eE[s_i])Q(0).$$

Observing that $\tau = 0$ if $N = 0$ and that $E[s_i]$, the mean busy period, is some finite constant independent of N now gives the desired result, viz.,

$$E[\tau|\mathcal{F}_0] \leq cQ(0)$$

for some constant c . The total number of arrivals during the period $[0, \tau]$ is the sum of $(\kappa - 1)$ i.i.d. random variables, independent also of κ , each of which is the number of arrivals during some busy period. Hence, it is easily seen that the number of arrivals is also bounded in mean by a constant times N .

We now proceed with the induction step. Suppose that the result is true when there are d or fewer queues in the tandem. Consider a tandem of $d + 1$ queues. Define

$$\tau_d = \inf \left\{ t > 0 : \sum_{i=1}^d Q_i(t) = 0 \right\}$$

to be the first time that the first d queues are all empty. Since the subsystem consisting of the first d queues is a stable LCFS tandem with Poisson input, the induction hypothesis implies that,

$$E[\tau_d|\mathcal{F}_0] \leq c \sum_{i=1}^d Q_i(0) \tag{4.17}$$

for some constant c . Furthermore,

$$Q_{d+1}(\tau_d) \leq \sum_{i=1}^{d+1} Q_i(0) + A(\tau_d), \tag{4.18}$$

where $A(t)$ denotes the total number of arrivals upto time t . But, by the induction hypothesis again,

$$E[A(\tau_d)|\mathcal{F}_0] \leq c \sum_{i=1}^d Q_i(0) \tag{4.19}$$

for some constant c . At time τ_d , the system consists of some number $Q_{d+1}(\tau_d)$ of customers in queue $d + 1$ and the remaining queues are empty. Color these customers blue and color new arrivals red. Consider the following modified system, where blue customers are worked on only when there are no red customers in any of the queues. If the modified process is constructed on the same sample space as the original process, that is, corresponding to the same external arrivals, service time assignments, and the same residual service times assigned to customers originally in the system at time τ_d , then it is easy to see that the original system becomes empty no later than the modified system (and possibly earlier). If we let τ denote the time to empty in the original system, and τ' in the modified system, then $\tau \leq \tau'$. We shall

now estimate $\tau' - \tau_d$. As in the single queue case, this can be written as a sum

$$\tau' - \tau_d = \sum_{i=1}^{Q_{d+1}(\tau_d)} Y_i + \sum_{i=1}^{\kappa-1} s_i, \tag{4.20}$$

where

$$\kappa = \inf \left\{ k > 0 : \sum_{i=1}^k t_i \geq \sum_{i=1}^{Q_{d+1}(\tau_d)} Y_i \right\}. \tag{4.21}$$

Here s_i is the length of the i th busy period of the entire system of $d + 1$ queues, and t_i the length of the i th idle period, which is exponentially distributed with mean λ^{-1} . Consequently, the s_i are i.i.d. and independent of κ . Furthermore, since the tandem is known to be stable, the s_i have finite mean. Also, the t_i are i.i.d., independent of the s_i . Hence, exactly the same arguments as in the case of a single queue suffice to show that

$$E[\tau' - \tau_d | \mathcal{F}_{\tau_d}] \leq c Q_{d+1}(\tau_d). \tag{4.22}$$

It now follows from (4.18), (4.19), and (4.22) that

$$E[\tau' - \tau_d | \mathcal{F}_0] \leq c \sum_{i=1}^{d+1} Q_i(0),$$

and since $\tau \leq \tau'$, that

$$E[\tau - \tau_d | \mathcal{F}_0] \leq c \sum_{i=1}^{d+1} Q_i(0). \tag{4.23}$$

Finally, putting (4.17) and (4.23) together gives

$$E[\tau | \mathcal{F}_0] \leq c \sum_{i=1}^{d+1} Q_i(0). \tag{4.24}$$

It is also easy to see that the total number of arrivals during $[\tau_d, \tau')$ is the sum of $(\kappa - 1)$ i.i.d. random variables of finite mean, and hence that the mean number of arrivals during $[\tau_d, \tau')$ is bounded by a constant times $\sum_{i=1}^{d+1} Q_i(0)$ (since the last is an upper bound on the mean of κ , analogously to the single queue case). Combining this with (4.19) and using the fact that $\tau \leq \tau'$ gives

$$E[A(\tau) | \mathcal{F}_0] \leq c \sum_{i=1}^{d+1} Q_i(0), \tag{4.25}$$

for some constant c . Finally note that $N = \sum_{i=1}^{d+1} Q_i(0)$. This completes the proof of the induction step, and hence the theorem is proved.

Acknowledgements

The authors would like to thank Professors Pravin Varaiya and Jean Walrand for their warm hospitality at the University of California at Berkeley, where this work was done.

References

- [1] V. Anantharam, The optimal buffer allocation problem, *IEEE Trans. Inf. Theory* 35 (4) (1989).
- [2] V. Anantharam and A. Ganesh, Correctness within a constant of an optimal buffer allocation rule of thumb, *IEEE Trans. Inf. Theory* 40 (1994) 871–882.
- [3] M. Cottrell, J. Fort and G. Malgouyres, Large deviations and rare events in the study of stochastic algorithms, *IEEE Trans. Autom. Contr.* AC-28 (9) (1983).
- [4] Y. C. Ho, M. A. Eyler and T. T. Chien, A gradient technique for general buffer storage design in a production line, *Int. J. Prod. Res.* 17 (6) (1979).
- [5] F. Kelly, *Reversibility and Stochastic Networks* (Wiley, New York, 1979).
- [6] S. Parekh, Rare events in networks of queues, Doctoral Dissertation, University of California, Berkeley (1986).
- [7] J. Walrand, Quick simulations of networks of $GI/G/1$ queues, Plenary lecture delivered at the *16th Conf. on Stochastic Processes and Their Applications*, Stanford University (1987).