Correctness Within a Constant of an Optimal Buffer Allocation Rule of Thumb

Venkat Anantharam and A. J. Ganesh

Abstract—**The problem is to allocate a fixed number of buffers among the nodes of an open network of exponential servers with Bernoulli routing and Poisson 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. In earlier work, we used pathwise probabilistic arguments to derive a simple rule of thumb for this problem: allocate the buffers 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 in the network with infinite buffers. We showed that this rule of thumb is accurate to within a known constant times the logarithm of the number of buffers as the number of buffers to be allocated becomes large.**

In this paper, we use time reversal and Poisson clumping arguments to show that our rule of thumb is, in fact, much better than previously demonstrated. We show that the optimal buffer allocation is within a constant of the rule of thumb as the number of buffers to be allocated becomes large, although now we cannot estimate the constant. In numerical terms, our earlier result reduced the search space for the optimal buffer allocation from $O(N^{J-1})$ to $O((\log N)^{J-1})$, where J denotes the number of nodes and N the number of buffers to be allocated. Our improvement reduces the search space to $O(1)$.

Index Terms—**Buffer allocation, communication networks, Jackson networks.**

I. Introduction

NETWORKS of queues are commonly used as models for the queueing processes taking place in computer networks, communication networks, and manufacturing systems. 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 paper, we address the problem using an asymptotic approach. That is, we assume that the number of buffers to be allocated is relatively large. This seems to be a fairly reasonable assumption, at least in the communication network applications. Our network model will be the skeleton of a Jackson network. Here we have a network of J exponential servers, with Bernoulli routing and Poisson exogenous arrivals. The problem studied is how to allocate a fixed number N of buffer spaces among the nodes of the 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.

It is generally accepted that this problem is analytically intractable. In view of this, the problem of estimating the time to buffer overflow by simulation has been studied by several investigators. Simulation-based approaches to this problem include the large deviation-theory-based ideas of Cottrell *et al.* [5] and Parekh and Walrand [9], [10] and the perturbation analysis technique of Ho *et al.* [6].

In [2], the following rule of thumb for this problem was considered: allocate the buffers 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 there that this rule of thumb is accurate to within a constant times $\log N$ as the number of buffers to be allocated becomes large. That is, if one considers the optimal buffer allocation scheme consisting of the optimal choice of buffer allocations for each N, then there is an explicitly known constant K independent of N such that, for each N, the rule of thumb buffer allocation is within $K \log N$ of the optimal allocation.

Our contribution in this work is to show that the rule of thumb is, in fact, much better than previously demonstrated. We show that there is a fixed constant independent of N such that the rule of thumb allocation is within this constant of the optimal buffer allocation for each N. We also give an explicit estimate for the constant when the criterion of interest is the mean time to buffer overflow. In numerical terms, the earlier result reduced the search space for the optimal buffer allocation from $O(N^{J-1})$ to $O((\log N)^{J-1})$. Our improvement reduces the search space to $O(1)$.

A closely related problem has been considered in [7] and in [8], and solved analytically. The problem considered in [8] is that of estimating the mean time between visits to a "rare" set in the equivalent infinite buffer network in stationarity. The rare set is taken to be the set of states corresponding to an overflow having occurred in the finite buffer network. A solution to this problem using ergodic arguments is presented therein. It is shown that the buffer allocation that maximizes the mean time to

Manuscript received ; revised July 16, 1993. This work was supported by the NSF under Grant NCR 8710840, PYI Award NCR 8857731, and IRI 9005849, by IBM under a Faculty Development Award, and by Bellcore Inc. This paper was presented at the ORSA/TIMS Conference on Applied Probability, Monterey, CA, January 1991.

The authors are with the Department of Electrical Engineering, Cornell University, Ithaca, NY 14853.

IEEE Log Number 9402028.

buffer overflow is within a constant of the rule of thumb. In a recent preprint [7],[1] a general class of Markov processes is considered, and it is shown that the distribution of the hitting time of a rare set, for the process in stationarity, is approximately exponential; the parameter of the exponential distribution is expressed as the eigenvalue of a transition rate kernel, and bounds are presented for the eigenvalue. Network problems are not considered in this paper. However, we can specialize the results of this paper to Jackson networks and derive a result similar to ours. In particular, it can be inferred from the bounds on the hitting time distribution therein that the optimal buffer allocation is within a constant of the rule of thumb.

A difference in our paper from the two mentioned above is that we consider that the time to overflow in the network started empty rather than in stationarity. Also, we deal directly with the finite buffer network rather than approximating the time to buffer overflow in the finite buffer case with the hitting time of a corresponding set in the infinite buffer case. The connection between the two can be seen using Poisson clumping arguments [1]. Consider the infinite buffer network in stationarity; let B denote the set of states corresponding to an overflow in its finite buffer counterpart. We wish to estimate the hitting time of the set B. A typical picture of the path of the process will consist of it starting with a small number of customers initially in the system. This number will eventually build up until the state enters the set B. There may subsequently be a clump of returns to B before the state again comes down to a small number of customers in the system. Our paper and [7] deal with the time to first hit B, and ignore the subsequent returns to B which we expect will take comparatively short times. The work in [8] finds the average time to hit B where the averaging involves both the time to first hit B and the subsequent return times. It will follow from the proof in Section V that the mean number of returns to B, after hitting B and before the network again becomes empty, is bounded above by a constant, uniformly in B. This explains why including the return times in the averaging does not significantly affect the estimate of the hitting time.

Our results are more general than [8] in that they apply to criteria associated with the time to buffer overflow other than the mean. They also give much more detailed information about the pathwise behavior of the network than [7], which does not at all consider network issues directly. Our technique is based on the use of time reversal and pathwise probabilistic arguments to bound the probability of buffer overflow in any cycle, starting with the network empty and ending when the network first empties after having had customers. This probability is then related to the time to buffer overflow by the use of renewal arguments to yield the desired result about optimality.

[1] We thank a referee of the paper for bringing this work to our attention.

In Section II, we introduce some notation and state the problem more formally. We state our main theorem in Section III, and use it to obtain our result about the approximate optimality of the rule of thumb. Section IV and V deal with the proof of the theorem stated in Section III. In Section IV, we introduce the reverse-time network, and reduce the problem of proving the theorem to that of bounding the probability of a certain event in the reverse-time network from below. We obtain the desired lower bound in Section V, completing the proof of the theorem. Section VI contains some concluding remarks and discusses what seem to be important directions for further investigation.

II. Problem Formulation

In this section, we describe our model and state the problem more formally. A Jackson network with J nodes is described by a J vector $(\mu_1,\cdots,\mu_J)$ of service rates at the individual nodes, a $(J+1)\times(J+1)$ routing matrix R that describes the interconnections of the nodes among themselves and with the outside world, and a number γ corresponding to the rate of exogenous arrivals. Customers arrive from the external world according to a Poisson process of rate γ, are routed to node j with probability r_{0j}, and join the queue there to await service. The service discipline at the nodes is assumed to be first-come, first-served (FCFS). Service times are i.i.d. exponential, and independent of the arrival process. When a customer finishes service at node j, he is routed to the queue at node k with probability r_{jk}, and leaves the system with probability r_{j0}. Routing is Bernoulli, independent of the arrival process and the departure times.

We assume that our Jackson network is irreducible, i.e., it is possible for an exogenous arrival to visit any queue before leaving the system. We also assume that it is stable, namely, that the solutions of the flow balance equations

$$\lambda_i = \gamma r_{0i} + \sum_{j\in[J]} \lambda_j r_{ji}, \qquad 1 \le i \le J \tag{2.1}$$

satisfy

$$\lambda_i < \mu_i, \qquad 1 \le i \le J$$

where, by $[J]$, we denote the set $\{1,\cdots,J\}$. This is a natural requirement if the network is to operate for reasonably long periods of time without buffer overflow. We call $\eta_i \triangleq \mu_i/\lambda_i$ the effective service rate at node i.

Let $X_j(t)$ denote the queue length process at node j, and let $X(t) = (X_1(t),\cdots,X_J(t))$ denote the vector queue length process. Then, it is known for Jackson networks that $X(t)$ is a Markov process, with unique stationary distribution π given by

$$\pi(x_1,\cdots,x_J) = \prod_{i\in[J]} \rho_i^{x_i}(1-\rho_i) \tag{2.2}$$

where $\rho_i \triangleq \eta_i^{-1}$ is called the load factor at node i (see, e.g., [11]).

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

$$B = \{(x_1, \cdots, x_J): x_i = N_i + 1, \quad \text{for some } 1 \leq i \leq J$$
$$x_j \leq N_j, \quad \text{for all } j \neq i, 1 \leq j \leq J\}. \tag{2.3}$$

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.

Our main result relates the probability that $X(t)$, started at the origin, hits B before returning to the origin to the stationary probability of the set B. We then relate the probability of hitting B to the distribution of the time to hit B, and use this result to justify the rule of thumb proposed in [2].

III. Main Result

The following rule of thumb was considered in [2]. Given N buffers, allocate roughly $p_i N$ buffers to node i, where the fraction p_i is such that $c \triangleq p_i \log \eta_i$ is a constant independent of i. It was shown there that if $(N_1, \cdots, N_J)$ is an optimal allocation, it must be close to the rule of thumb in the sense that

$$|N_i - p_i N| < K \log N, \qquad 1 \leq i \leq J \tag{3.1}$$

for a constant K independent of N and large enough N. What was shown was that, if any allocation differed from the rule of thumb by more than the above amount along a subsequence of $N \to \infty$, it performed worse than the rule of thumb along that subsequence.

Consider any buffer allocation that satisfies (3.1), and the associated boundary B as defined in Section II. Let α denote the probability that $X(t)$, the Markov chain denoting the vector queue length process starting at the origin, hits B before returning to zero.

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

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

$$\pi(B) = \sum_{x \in B} \pi(x) \tag{3.2}$$

where $\pi(x)$ is given in (2.2).

Our main result can be stated as follows.

Theorem: Consider a Jackson network and a scheme for allocating buffers among its nodes such that the allocations satisfy (3.1). Then, for some constants c_1, c_2 that do not depend on the number of buffers N but may depend on the network parameters, we have

$$c_1 \pi(B) \leq \alpha \leq c_2 \pi(B). \tag{3.3a}$$

Further, for N sufficiently large, we have the explicit estimates

$$c_1 = \frac{\gamma + \sum_{i \in [J]} \mu_i}{\gamma \pi(0)} \cdot \left(\prod_{i \in [J]} (1 - \rho_i) \right) \cdot \min_{j \in [J]} \frac{1}{2} \tilde{\beta}_j \left(1 - \frac{\tilde{\nu}_j^{\{j\}}}{\mu_j} \right). \tag{3.3b}$$

Here, $\tilde{\beta}_j$ is defined just prior to (4.15), which gives a lower bound for it. $\tilde{\nu}_j^{\{j\}}$ corresponds in the time-reversed network to $\nu_j^{\{j\}}$, which is defined in (5.3).

$$c_2 = \frac{\gamma + \sum_{i \in [J]} \mu_i}{\gamma \pi(0)}. \tag{3.3c}$$

As defined above, α is the probability that the network, started empty, suffers a buffer overflow before returning to the empty state, and B is the boundary associated with the buffer allocation, as defined in Section II.

Note: It is enough to assume that the buffer allocation $(N_1, \cdots, N_J)$ satisfies the following: if $\eta_i \geq \eta_j$, then $N_i \leq N_j$ to within a term that grows slower than N. This is a much weaker requirement than (3.1), but since the result of Anantharam, see [2], guarantees that the optimal allocation satisfies (3.1), restricting ourselves to such allocations entails no loss of generality.

We now introduce some notation. σ denotes an infinite subset of the positive integers. We write $\lim_{N \overset{\sigma}{\to} \infty}$ 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 $\lim_{N \overset{\sigma}{\to} \infty} (f(N)/g(N)) = 0$. 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.

Next, we relate α to the time to buffer overflow using pathwise probabilistic arguments. The approach closely follows that in [2].

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. Notice that the quantities T and ν depend on N and the buffer allocation, as do some of the other quantities to be defined in the sequel. However, for convenience, we suppress this dependence in the notation.

Since the process started at 0 hits B before returning to 0 with probability α independent of its past, we see that the number of cycles ν of returns to 0 before hitting B is a geometric random variable with parameter $(1 - \alpha)$, i.e.,

$$P(\nu = k) = \alpha(1 - \alpha)^k \tag{3.4}$$

and

$$E\nu = \frac{1 - \alpha}{\alpha}. \tag{3.5}$$

Let δ have the distribution of the time taken to return to 0, starting from 0 and conditioned on not visiting B. Let $\delta_1, \delta_2, \cdots$ be i.d.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, since the evolution of the process starts afresh each time it hits 0 (by the strong Markov property), it is easy to see that

$$T \stackrel{d}{=} \sum_{k=1}^{\nu} \delta_k + \Delta \tag{3.6}$$

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

$$\begin{aligned} ET &= \frac{1-\alpha}{\alpha} E\delta + E\Delta \\ &= \frac{1}{\alpha}[(1-\alpha)E\delta + \alpha E\Delta]. \end{aligned} \tag{3.7}$$

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 time is stochastically dominated by the mean time to return to 0 starting from 0 in the network with infinite buffers. Call this latter time T_0. Then,

$$ET \leq \frac{1}{\alpha} ET_0. \tag{3.8}$$

Also, 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. Hence, ET_0 is independent of N and $(N_1, \cdots, N_J)$. Thus, by Markov's inequality,

$$P(T \geq \tau(N)) \leq \frac{ET}{\tau(N)} \leq \frac{1}{\alpha \tau(N)} ET_0. \tag{3.9}$$

Also observe that δ stochastically dominates an exponential random variable of mean γ^{-1}. Indeed, if the network is empty, we have to wait at least that long for an arrival. From (3.4) and (3.6), 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 $1 - \alpha/\alpha\gamma$. Hence,

$$P(T \leq \tau(N)) \leq 1 - \exp\left(-\frac{\alpha\gamma\tau(N)}{1-\alpha}\right). \tag{3.10}$$

Notice that α is a function of N and the buffer allocation, although the dependence has been suppressed in the notation. Keeping this in mind, we consider a subsequence σ and a corresponding sequence of buffer allocations. Then, from (3.9), we see that

$$\tau(N) = \omega_\sigma(\alpha^{-1}) \Rightarrow P(T \geq \tau(N)) \underset{\sigma}{\to} 0. \tag{3.11}$$

Likewise, from (3.10), we observe that

$$\tau(N) = o_\sigma(\alpha^{-1}) \Rightarrow P(T \leq \tau(N)) \underset{\sigma}{\to} 0. \tag{3.12}$$

We shall now use these estimates to show that, for any reasonable criterion associated with the time to buffer overflow T, there is a fixed constant independent of N such that the rule of thumb allocation proposed in [2] is within this constant of the optimal buffer allocation for every N.

Rule of Thumb: In allocating N buffers to "maximize the time to buffer overflow," one should allocate roughly a fraction $p_i N$ of the buffers to node i, where p_i is inversely proportional to $\log \eta_i$.

Justification: Let $c \triangleq p_i \log \eta_i$ (independent of i). Now, for any buffer allocation scheme $(N_1, \cdots, N_J)$ such that $\sum_{i \in [J]} N_i = N$, let $g(N) \triangleq \min_{i \in [J]} \eta_i^{N_i}$. Then, from the expression for the stationary probability in (2.2) and the definition of the set B in (2.3), we see that

$$\pi(B) \geq g(N)^{-1}\left(\prod_{i \in [J]} (1-\rho_i)\right) \min_{i \in [J]} \rho_i$$

and also,

$$\pi(B) \leq \sum_{i \in [J]} \rho_i^{N_i+1} \leq J \max_{i \in [J]} \rho_i \cdot g(N)^{-1}.$$

Consequently, we obtain

$$k_1 g(N)^{-1} \leq \pi(B) \leq k_2 g(N)^{-1} \tag{3.13}$$

for constants $k_1 = (\prod_{i \in [J]}(1-\rho_i)) \cdot \min_{i \in [J]} \rho_i$ and $k_2 = J \max_{i \in [J]} \rho_i$, independent of N and the buffer allocation scheme.

Denote $2^{-cN} g(N)$ by $h^2(N)$. Suppose

$$\liminf_{N \to \infty} 2^{-cN} g(N) = 0.$$

Then, there is a subsequence σ such that $h^2(N) = o_\sigma(1)$. Then, also, $h(N) = o_\sigma(1)$. Let T and T^* have the distribution of the time to buffer overflow for this buffer allocation scheme and for the rule of thumb, respectively. Note that, for the buffer allocation according to the rule of thumb, the quantity $g(N)$ defined above takes the value 2^{cN}. We now see that

$$\lim_{N \overset{\sigma}{\to} \infty} P(T \leq 2^{cN} h(N)) = \lim_{N \overset{\sigma}{\to} \infty} P\left(T \leq \frac{g(N)}{h(N)}\right) = 1 \tag{3.14}$$

by (3.3a), (3.11), and (3.13), whereas

$$\lim_{N \overset{\sigma}{\to} \infty} P(T^* \leq 2^{cN} h(N)) = 0 \tag{3.15}$$

by (3.3a), (3.12), and (3.13), and the fact that $\pi(B)^{-1}$, for the allocation according to the rule of thumb, is within a multiplicative constant of 2^{cN}.

From (3.14) and (3.15), it is clear that the buffer allocation scheme is worse than the rule of thumb along σ. This suggests that the optimal buffer allocation rule must satisfy

$$\liminf_{N \to \infty} 2^{-cN} g(N) > \epsilon \tag{3.16}$$

for some $\epsilon > 0$. Hence,

$$\liminf_{N \to \infty}\left(\min_{i \in [J]} N_i \log \eta_i - cN\right) > \log \epsilon$$

or, equivalently,

$$\min_{i \in [J]} N_i \log \eta_i - cN > \log \epsilon, \qquad \text{for all } N$$

with possibly a different ϵ. Substituting $c = p_i \log \eta_i$, it is straightforward to see that this implies

$$|N_i - p_i N| < K \tag{3.17}$$

for all N, where $K = -\log \epsilon \cdot \Sigma_{i \in [J]} (\log \eta_i)^{-1}$. Thus, we see that the optimal buffer allocation is within a constant of the rule of thumb. □

In the above argument, note that it is not possible to evaluate the constant explicitly because of its dependence on the unknown parameter ϵ. However, if the criterion of interest is the mean time to buffer overflow, a simpler argument gives an explicit estimate for how far the rule of thumb is from optimality.

Observe from (3.3a) and (3.8) that

$$ET \le \frac{ET_0}{c_1} \pi(B)^{-1}$$

where ET_0 is the mean time taken to return to 0 starting from 0 in the network with infinite buffers. But ET_0 is a constant independent of N and the buffer allocation, and can be estimated using ergodic arguments. Indeed,

$$\pi(0) = \frac{\gamma^{-1}}{\gamma^{-1} + ET_0}$$

since the mean time spent in 0 in each visit to it is γ^{-1} and the mean time between visits is ET_0. Substituting for ET_0 above now gives

$$ET \le \hat{c}_1 \pi(B)^{-1} \tag{3.18}$$

where

$$\hat{c}_1 = \frac{1}{\gamma c_1} \left(\frac{1}{\pi(0)} - 1 \right)$$

can be explicitly computed from the estimate for c_1 in the statement of the main theorem. Also, it is easy to see from (3.10) that

$$ET \ge \frac{1 - \alpha}{\alpha \gamma}.$$

Hence, it follows from the estimate in (3.3a) that, for large enough N,

$$ET \ge \hat{c}_2 \pi(B)^{-1} \tag{3.19}$$

where $\hat{c}_2 = 1/2\gamma c_2$, say.

Fix N, the total number of buffers to be allocated, sufficiently large. Let B denote the boundary corresponding to a given buffer allocation and B^* the boundary corresponding to the rule of thumb allocation (see (2.3) for the definition of the boundary). Now, from (3.13),

$$\pi(B^*) \le k_2 2^{-cN}$$

and

$$\pi(B) \ge k_1 g(N)^{-1}$$

where the quantities c and $g(N)$ are as defined earlier, and k_1 and k_2 are constants defined in (3.13). Let T and T^* have the distribution of the time to buffer overflow for the above buffer allocation scheme and for the rule of thumb, respectively. Then, it is clear from the above and from (3.18) and (3.19) that

$$ET \le \tilde{c}_1 g(N) \tag{3.20}$$

and

$$ET^* \ge \tilde{c}_2 2^{cN} \tag{3.21}$$

where $\tilde{c}_1$ and $\tilde{c}_2$ are constants that can be computed from above. We get

$$\tilde{c}_1 = \frac{1}{\gamma c_1 k_1} \left(\frac{1}{\pi(0)} - 1 \right)$$

and

$$\tilde{c}_2 = \frac{1}{2\gamma c_2 k_2}$$

where c_1, c_2 are as given in our theorem and k_1, k_2 are given in (3.13).

Let

$$K^* = J \frac{\log(\tilde{c}_1/\tilde{c}_2)}{\min_{i \in [J]} \log \eta_i}.$$

Suppose the buffer allocation above differs from the rule of thumb by more than K^*, in the sense that there is some node where the allocations differ by more than K^*, i.e., $\max_{i \in [J]} |N_i - p_i N| > K^*$. Then, there is some node $j \in [J]$ such that $N_j < p_j N - (K^*/J)$. Now, from the definition of $g(N)$, we get

$$g(N) \le \eta_j^{N_j} < 2^{cN} \eta_j^{-K^*/J}.$$

Hence, from the definition of K^*, it is easy to see that

$$\tilde{c}_1 g(N) < \tilde{c}_2 2^{cN}.$$

Therefore, from (3.20) and (3.21), it follows that

$$ET^* > ET.$$

Hence, the allocation according to the rule of thumb performs better than any allocation which differs from it by more than the quantity K^* defined above when the criterion of interest is the mean time to buffer overflow in the network started empty. Also observe that the quantity K^* can be computed explicitly in terms of the constants c_1 and c_2 appearing in the statement of the main theorem as

$$K^* = \frac{J\left(\log c_2 - \log c_1 + \log(2J) - \Sigma_{i \in [J]} \log(1 - \rho_i) - \min_{i \in [J]} \log \rho_i + \max_{i \in [J]} \log \rho_i\right)}{\min_{i \in [J]} \log \eta_i}$$

where c_1 and c_2 are given in the statement of the main theorem.[2]

In Sections IV and V, we will prove the main theorem stated at the beginning of this section.

IV. Watching the Time-Reversed Network

In this section, we begin the proof of our main theorem, which was stated at the beginning of Section III, and used subsequently in that section to justify the rule of thumb for buffer allocation proposed in [2].

We shall use time-reversal arguments to reduce the problem to one of estimating the probability that the Markov process $X(t)$, starting from a point k on the boundary B, hits 0 before returning to B. We show that, if for some specified points in B this probability is bounded away from zero, uniformly in N, then the theorem is true. This section will deal with the proof of the above statement. Then, in Section V, we will show that this probability is indeed bounded away from zero, as required.

Consider a Jackson network as defined earlier, with vector queue length process $X(t)$. We first consider the discrete-time Markov chain $\hat{X}(n)$ obtained by embedding at the virtual jump times. The virtual jump process is the sum of the exogenous arrival process and the virtual departure processes as the individual nodes. These are independent Poisson processes, with future independent of the current state; hence, the virtual jump process is Poisson (with rate $\gamma + \sum_{i \in [J]} \mu_i$) with future independent of the current state. It is an easy consequence of this that the embedded process $\hat{X}(n)$ has the same stationary distribution π as the original process $X(t)$; see, e.g., [11].

Next, let $\hat{Y}(n)$ be obtained from $\hat{X}(n)$ by watching it in the set $0 \cup B$. Then, $\hat{Y}(n)$ is also a discrete-time Markov chain, with stationary probability $\hat{\pi}$ given by

$$\hat{\pi}(j) = \Big(\sum_{i \in 0 \cup B} \pi(i) \Big)^{-1} \pi(j). \tag{4.1}$$

We are interested in its transition matrix $\hat{P}$. Specifically, we have

$$\alpha = P_0(\text{hit } B \text{ before } 0 | \text{ leave } 0).$$

But

$$\begin{aligned}\hat{P}(0,0) &= P_0(\hat{X}(n) = 0 \text{ before } \hat{X}(n) \in B | \hat{X}(1) \neq 0) \\ &\quad P_0(\hat{X}(1) \neq 0) + P_0(\hat{X}(1) = 0) \\ &= \frac{\gamma}{\gamma + \sum_{i \in [J]} \mu_i} P_0(\text{return to 0 before hitting } B) \\ &\quad + \frac{\sum_{i \in [J]} \mu_i}{\gamma + \sum_{i \in [J]} \mu_i} \\ &= \frac{\gamma}{\gamma + \sum_{i \in [J]} \mu_i}(1-\alpha) + \frac{\sum_{i \in [J]} \mu_i}{\gamma + \sum_{i \in [J]} \mu_i} \\ &= 1 - \frac{\alpha\gamma}{\gamma + \sum_{i \in [J]} \mu_i}.\end{aligned}$$

[2] A sharper pass through our development can give a somewhat improved K^*. The essential point, though, is that there is such a K^* and it can be explicitly estimated.

Therefore,

$$\alpha = \sum_{k \in B} \hat{P}(0,k) \cdot \frac{\gamma + \sum_{i=1}^{J} \mu_i}{\gamma} \tag{4.2}$$

using the stochasticity of $\hat{P}$, which implies that $1 - \hat{P}(0,0) = \sum_{k \in B} \hat{P}(0,k)$. Now, let $\tilde{Y}(n)$ be the time reversal of $\hat{Y}(n)$, so $\tilde{Y}(n)$ is a Markov chain with the same stationary probability $\hat{\pi}$, and its transition matrix $\tilde{P}$ is given by

$$\tilde{P}(i,j) = \frac{\hat{\pi}(j)\hat{P}(j,i)}{\hat{\pi}(i)}, \qquad i,j \in 0 \cup B. \tag{4.3}$$

Thus,

$$\hat{P}(0,k) = \frac{\hat{\pi}(k)\tilde{P}(k,0)}{\hat{\pi}(0)} = \frac{\pi(k)\tilde{P}(k,0)}{\pi(0)}, \qquad k \in B, \tag{4.4}$$

and so

$$\alpha = \frac{1}{\pi(0)} \sum_{k \in B} \pi(k)\tilde{P}(k,0) \cdot \frac{\gamma + \sum_{i=1}^{J} \mu_i}{\gamma} \tag{4.5}$$

by (4.1). Since $\tilde{P}(k,0) \leq 1$, we clearly have

$$\alpha \leq \frac{1}{\pi(0)} \sum_{k \in B} \pi(k) \frac{\gamma + \sum_{i=1}^{J} \mu_i}{\gamma} = \frac{\pi(B)}{\pi(0)} \cdot \frac{\gamma + \sum_{i=1}^{J} \mu_i}{\gamma}. \tag{4.6}$$

Now, consider $k \in B$ of the form $(N_j + 1)e_j$ where e_j is the jth unit vector. this corresponds to a buffer overflow in the jth queue, with all other queues being empty. α is clearly no smaller than the sum in (4.5) restricted to just such terms. Thus,

$$\alpha \geq \frac{1}{\pi(0)} \sum_{j \in [J]} \pi\big((N_j + 1)e_j\big) \cdot \tilde{P}\big((N_j + 1)e_j, 0\big) \cdot \frac{\gamma + \sum_{i=1}^{J} \mu_i}{\gamma}. \tag{4.7}$$

We also have

$$\pi(B) \leq \sum_{j \in [J]} \pi\big((N_j + 1)e_j\big) \cdot \prod_{\substack{i \in [J] \\ i \neq j}} (1 - \rho_i)^{-1}. \tag{4.8}$$

If we can show, for each j, that $\tilde{P}((N_j + 1)e_j, 0)$ is bounded below by a positive constant, which is independent of N and $(N_1, \cdots, N_J)$ which satisfy the assumption (3.1), then it is easy to see, from (4.7) and (4.8), that

$$\alpha \geq c_1 \pi(B). \tag{4.9}$$

Combining this with (4.6), we get

$$c_1 \pi(B) \leq \alpha \leq c_2 \pi(B)$$

which is the claim of our main theorem.

Notice that, for c_2, we have the explicit estimate from (4.6):

$$c_2 = \frac{\gamma + \sum_{i \in [J]} \mu_i}{\gamma \pi(0)}. \tag{4.10a}$$

This completes the proof of the second inequality in the theorem. It remains to prove the first inequality and estimate the constant c_1. Now, from (4.7) and (4.8), it is clear that we can take

$$c_1 = \frac{\gamma + \sum_{i\in[J]}\mu_i}{\gamma\pi(0)} \cdot \left(\prod_{i\in[J]}(1-\rho_i)\right) \cdot \min_{j\in[J]} \tilde{P}\big((N_j+1)e_j, 0\big). \tag{4.10b}$$

The term $\tilde{P}((N_j+1)e_j, 0)$ will be estimated in the course of the proof of the theorem.

Thus, we have reduced the proof of our theorem to showing that, for each $1 \le i \le J$, $\tilde{P}([N_i+1]e_i, 0)$ is bounded away from zero, uniformly in N, and for all buffer allocations allowed by (3.1).

Here, $\tilde{P}$ is the transition matrix of $\tilde{Y}(n)$, which was obtained by embedding $X(t)$ at the virtual jump times, watching the resulting Markov chain in the set $0 \cup B$, and taking the time reversal. For ease of computation, we would like to have the result in terms of $\tilde{X}(t)$, the time reversal of $X(t)$. Define $\hat{Z}(n)$ to be the embedding of $\tilde{X}(t)$ at its virtual jump times, and let $\tilde{Z}(n)$ be $\hat{Z}(n)$ watched in the set $0 \cup B$. Then, since the time reversal of the watching of the embedding is the same as the watching (in the same set) of the embedding of the time reversal, we have $\tilde{Z}(n) = \tilde{Y}(n)$. Thus our result can be stated in terms of $\tilde{Z}(n)$ instead of $\tilde{Y}(n)$.

It is known that the time reversal $\tilde{X}(t)$ of the queue length process $X(t)$ of a Jackson network is the queue length process of a different Jackson network, with the same number of nodes but different parameters as follows; see [11]. Also, if the original network is irreducible and open, then so is its time reversal:

$$\tilde{\gamma} = \gamma \tag{4.11}$$

$$\tilde{\mu}_i = \mu_i, \qquad 1 \le i \le J \tag{4.12}$$

$$\tilde{r}_{ij} = \frac{\lambda_j}{\lambda_i} r_{ji}, \qquad 0 \le i, j \le J. \tag{4.13}$$

Here, γ is the rate of exogenous arrivals, μ_i is the service rate at the ith node, and r_{ij} is the (i,j)th entry in the routing matrix R for the original network, and the same quantities with tildes refer to the corresponding quantities in the reverse-time network. Also, in (4.13), λ_0 is understood to refer to γ. The λ_i's solve the flow balance equations (2.1) in the original network. The solutions to the flow balance equations in the reverse-time network are also the same, i.e.,

$$\tilde{\lambda}_i = \lambda_i, \qquad 1 \le i \le J. \tag{4.14}$$

Thus, our problem is equivalent to showing that the process $\tilde{Z}(n)$, started at $(N_j+1)e_j$, hits 0 before returning to B with probability bounded away from zero, uniformly in N and in the permitted buffer allocations, and for all $1 \le j \le J$. Observe that, by our assumption that the network is irreducible and open, there is a path along which it is possible for a customer in any queue j to leave the system in the time-reversed network. In other words, there is a sequence of distinct nodes $j = j_0, j_1, \cdots, j_{n-1}, j_n = 0$ such that $\tilde{r}_{j_k, j_{k+1}} > 0$ for each $k \in \{0, \cdots, n-1\}$. Define $\tilde{\beta}_j$ to be the probability that $\tilde{Z}(n)$, started at $(N_j+1)e_j$, hits $N_j e_j$ before returning to B. It is then clear that

$$\tilde{\beta}_j \ge \prod_{k=0}^{n-1} \frac{\mu_{j_k}\tilde{r}_{j_k, j_{k+1}}}{\gamma + \sum_{i\in[J]}\mu_i} \tag{4.15}$$

since the right-hand side above is the probability that the leading customer in queue j leaves the network before any new customers enter the network or any other customers move out of queue J. In the event that this happens, the network state reaches the point $N_j e_j$ before hitting B. Notice that the above probability is positive and independent of N and the buffer allocation. Furthermore, $\tilde{Z}(n)$ is in B before 0 after this if, and only if, $\tilde{X}(t)$ is since $\tilde{Z}(n)$ is obtained from $\tilde{X}(t)$ by embedding at its virtual jump times and watching the resulting process in $0 \cup B$. So, our problem is reduced to showing that $\tilde{X}(t)$, started at $N_j e_j$, hits 0 before returning to B with uniformly positive probability. We also have the estimate

$$\tilde{P}\big((N_j+1)e_j, 0\big) \ge \tilde{\beta}_j \tilde{P}(N_j e_j, 0) \tag{4.16}$$

where the quantity $\tilde{\beta}_j$ was defined above and shown to be strictly positive.

Proposition 4.1: Let $X(t)$ be the queue length process of a stable, irreducible, open Jackson network of the form defined earlier. Consider a sequence σ with, for each N in the sequence, a buffer allocation $(N_1, \cdots, N_J), \sum_{i\in[J]} N_i = N$ satisfying the condition in (3.1). Then, for each j such that $1 \le j \le J$, we have

$$\tilde{P}(k_j, 0) > \epsilon$$

for some $\epsilon > 0$ independent of N and $(N_1, \cdots, N_J)$. Here, k_j is the state $N_j e_j$, and $\tilde{P}(k_j, 0)$ denotes the probability that $\tilde{X}(t)$, the time reversal of $X(t)$, started at k_j, hits 0 before B, where B is the boundary associated with the buffer allocation.

If this proposition is true, then, in conjunction with (4.5)–(4.9), it implies that our main theorem is true. The proof of the above proposition is addressed in the next section.

V. Proof of the Main Result

In this section, we discuss the proof of Proposition 4.1. That is to say, we show that the Markov chain $\tilde{X}(t)$, started at $k_j \triangleq N_j e_j$, hits 0 before returning to B with probability larger than some positive real number ϵ which is independent of N and the buffer allocation. The proof consists of two parts. The first part involves showing that the system started with the jth queue full does not undergo a buffer overflow in some other queue before the system becomes empty (with sufficiently high probability). The proof employs a fluid limit approach and is dealt with in Sections V-A–V-C. The second part involves showing that, with probability bounded away from 0, the system empties before suffering another overflow in the jth queue. The proof is contained in Section V-C, and uses a simple martingale argument.

As stated earlier, $\tilde{X}(t)$ evolves like the queue length process of a stable, irreducible, open Jackson network whose parameters are related to those of the original network through (4.11)–(4.13). Observe that the effective service rates η_i at the individual nodes remain the same. Since our rule of thumb for allocating buffers depends only on the effective service rates, it would give rise to the same buffer allocation in the reverse-time network as in the original one. Thus, if a boundary B corresponds to a buffer allocation in the original network satisfying the condition in (3.1), then the same is true in the reverse-time network. We use the above facts to simplify our notation by dropping tildes. It should be noted in the following that all network parameters refer to those in the time reversal of the original network, although this will not be made explicit in the notation. Also, we shall henceforth deal only with the time-reversed network; $X(t)$ will be used to denote its queue-length process.

Consider N going to infinity along a subsequence σ of the positive integers, and a corresponding sequence of buffer allocations $(N_1,\cdots, N_J)$. Define $p_k^N = N_k/N$ for each N in the subsequence and each $k \in [J]$. Observe that, if the buffer allocations satisfy (3.1), then $p_k^N \to p_k$ as $N \to \infty$ for each $k \in [J]$ (where p_i was defined in Section III to be the fraction of buffers allocated to node i by the rule of thumb). If we let $X^N(t)$ denote the queue length process when the Jackson network is started with N_j customers in the jth queue and all the other queues empty, then we are interested in $P(X^N(t) \in B$ before $X^N(t) = 0)$ for large N. We shall bound this probability from above using the fluid limit of the processes $X^N(t)$ as $N \to \infty$.

A. The Fluid Limit

The fluid limit of a sequence of discrete flow processes is described in [3] and [4]. The main results that we shall use in this paper are the following.

The sequence of stochastic processes $X^N(t)$ described above converges to a limit $X(t)$ in the sense that, for any $\epsilon_0 > 0$ and all $\epsilon > \epsilon_0$,

$$\lim_{N\to\infty} P\left(\sup_{0\le t\le T} \left\| \frac{1}{N} X^N(Nt) - X(t) \right\| \geq \epsilon \,\middle|\, \left\| \frac{1}{N} X^N(0) - X(0) \right\| < \epsilon_0 \right) = 0 \quad (5.1)$$

where $\|X\| = \max_{i\in[J]}|X_i|$. This is a consequence of [3, Corollary 1]. The proof involves showing that $X^N(Nt)/N - X(t)$ is a martingale, showing that its quadratic variation goes to zero as $N \to \infty$, and applying a martingale maximal inequality (Lenglart's inequality); see [3] for details. Furthermore, it is shown in [4] that the evolution of $X(t)$ is deterministic, and is described for $t \geq 0$ by the solution of

$$X_i(t) = X_i(0) + \gamma r_{0i} t + \sum_{j\in[J]} (\mu_j t - Y_j(t)) r_{ji} - (\mu_i t - Y_i(t)) \quad (5.2)$$

subject to

$$X(t) \geq 0 \quad (5.2a)$$

$$Y \text{ is nondecreasing with } Y(0) = 0 \quad (5.2b)$$

$$Y_i \text{ is increasing only at times } t \geq 0 \text{ when } X_i(t) = 0, \quad i \in [J]. \quad (5.2c)$$

The intuition behind the above set of equations is the following. Initially, the buffer at node i contains a quantity of fluid $X_i(0)$. Flow enters node i from outside at rate γr_{0i}, and, in addition, a fraction r_{ji} of the flow leaving node j up to time t is routed to node i, all these quantities being deterministic. Each node j is capable of pushing flow through at a maximum rate of μ_j, and does so while it is not empty. If it is empty and the inflow rate into it is smaller than μ_j, then the outflow rate is equal to the inflow rate and the buffer at that node continues to be empty. The quantity $Y_j(t)$ measures the difference between the maximum flow that node j could potentially have pushed through up to time t, which is $\mu_j t$, and the amount of flow it actually pushed through, i.e., it measures the "wasted" throughput at node j. This picture makes it easy to write equations describing the quantities $Y_i(t)$ and $X_i(t)$, which we proceed to do below.

We now consider a fairly general initial condition, where the only condition we impose is that

$$X_i(0) = \lim_{N\to\infty} (1/N) X_i^N(0)$$

exist for each $i \in [J]$. Let I denote the set of nodes i such that $X_i(0)$ is nonzero. Let

$$\tau = \inf\{t > 0: X_i(t) = 0 \text{ for some } i \in I\}$$

be the earliest time (possibly never) that the buffer at some node in I becomes empty. Now, by (5.2b) and (5.2c),

$$Y_i(t) = 0, \qquad \forall\, 0 \leq t \leq \tau, \forall\, i \in I.$$

Let ν^I solve the generalized flow balance equations

$$\nu_k^I = \gamma r_{0k} + \sum_{i\in I} \mu_i r_{ik} + \sum_{i\in[J]-I} (\nu_i^I \wedge \mu_i) r_{ik}, \qquad k \in [J]. \quad (5.3)$$

The ν_k^I have the interpretation of being the rate at which fluid enters node k until some node in I empties out. This consists of γr_{0k} entering from outside, plus a fraction r_{jk} of the fluid leaving node j. If node j is not empty, fluid leaves it at rate μ_j; else fluid leaves it at the minimum of the rate at which it enters and the rate at which node j is capable of pushing flow through, which is μ_j.

The following lemma is a consequence of the idea that, if the flow incident at a node is less than its capacity, the difference is wasted throughput at that node. This enables us to write expressions for the $Y_i(t)$; a simple computation then gives us the $X_i(t)$, i.e., the fluid limit. The computation is shown in the proof of the lemma below.

Lemma 5.1:

$$Y_i(t) = (\mu_i - \nu_i^I)^+ t, \qquad \forall\, 0 \le t \le \tau, i \in [J] - I$$

where $(\mu_i - \nu_i^I)^+ = \max(\mu_i - \nu_i^I, 0)$.

Proof: Substitute

$$Y_i(t) = \begin{cases} 0, & i \in I \\ (\mu_i - \nu_i^I)^+ t, & i \in [J] - I \end{cases}$$

in (5.2) to obtain

$$\begin{aligned} X_j(t) &= X_j(0) + \gamma r_{0_j} t + \sum_{i \in I} \mu_i r_{ij} t \\ &\quad + \sum_{i \in [J]-I} (\mu_i \wedge \nu_i^I) r_{ij} t - (\mu_j \wedge \nu_j^I) t \\ &= X_j(0) + \nu_j^I t - (\mu_j \wedge \nu_j^I) t \quad \text{by (5.3)} \\ &= X_j(0) + (\nu_j^I - \mu_j)^+ t, \qquad 0 \le t \le \tau, j \in [J] - I \end{aligned} \tag{5.4a}$$

and also

$$\begin{aligned} X_j(t) &= X_j(0) + \gamma r_{0_j} t + \sum_{i \in I} \mu_i r_{ij} t \\ &\quad + \sum_{i \in [J]-I} (\mu_i \wedge \nu_i^I) r_{ij} t - \mu_j t \\ &= X_j(0) + (\nu_j^I - \mu_j) t, \qquad 0 \le t \le \tau, j \in I. \end{aligned} \tag{5.4b}$$

It is now easy to check that $(X(t), Y(t))$ satisfy (5.2)–(5.2c) for all $0 \le t \le \tau$, and the lemma follows by uniqueness of the solution, which is also established in [4].

B. Negative Drift of the Fluid Limit

In the last section, we computed the fluid limit of the queue length process, and showed that the actual queue length process stays close to the fluid limit with high probability. The fluid limit is a deterministic process. In this section, we use the assumptions of stability and irreducibility of the network to show that the fluid limit has negative drift. More precisely, $\sum_{j \in [J]} X_j(t)$, the total quantity of fluid in the network, is strictly decreasing at a positive rate as long as the network is not empty. In the next section, this fact will be used to show that a buffer overflow, in the queues other than queue j, cannot occur for the fluid limit, and hence, with high probability, for the actual queue length process, before the system empties. This depends on the *a priori* assumption, made in the statement of Proposition 4.1, that the buffer allocation satisfies (3.1).

We now take up the drift computation. For I as above, define $U_I = \{i \in J \colon \nu_i^I > \mu_i\}$ where $\nu_i^I, i \in [J]$ solve (5.3). U_I denotes the set of nodes where fluid builds up when the nodes in I are putting out flow at their maximum rates. That is, the inflow rates at these nodes exceed their capacity to push flow through. Equivalently, in the original queueing network, the queues in U_I have arrival rates greater than their service rates when the departure processes out of the queues in I are replaced by the corresponding virtual departure processes.

Let $U_I' = I \cup U_I$. This is the set of nodes which have a nonzero fluid level $X_j(t)$ during $(0, \tau)$, and hence have outflow rates equal to their maximum outflow rates during this period. For $i, j \in U_I' \cup 0$, let ρ_{ij} denote the probability that a customer leaving node i visits node j before visiting any other node in $U_I' \cup 0$, where 0 refers to the outside world. Then, for $j \in U_I'$, we can rewrite (5.3) as

$$\nu_j^I = \gamma \rho_{0j} + \sum_{i \in U_I'} \mu_i \rho_{ij}. \tag{5.5}$$

This just says that the flow rate into node j in U_I' is the fraction of flow entering the network from outside which first enters the set U_I' at node j, plus the rate of flow leaving nodes in U_I' which first returns to the set U_I' in node j. The justification for this is that no flow is either created or held back at nodes outside U_I', the inflow rates into these nodes being exactly equal to the outflow rates therefrom. Hence, we get

$$\sum_{j \in U_I'} (\nu_j^I - \mu_j) = \sum_{j \in U_I'} (\gamma \rho_{0j} - \mu_j \rho_{j0}) + \sum_{i \in U_I'} \sum_{j \in U_I'} (\mu_i \rho_{ij} - \mu_j \rho_{ji}) \tag{5.6}$$

where we have written

$$\mu_j = \mu_j \left(\rho_{j0} + \sum_{i \in U_I'} \rho_{ji} \right)$$

for each $j \in U_I'$. Now, observe that

$$\sum_{j \in U_I'} \gamma \rho_{0j} = \sum_{j \in U_I'} \lambda_j \rho_{j0}$$

where the λ's solve the original flow balance equations (2.1). This just says that the net flow rate into U_I' equals the net flow rate out of U_I' in equilibrium. Hence, (5.6) can be rewritten as

$$\sum_{j \in U_I'} (\nu_j^I - \mu_j) = \sum_{j \in U_I'} (\lambda_j - \mu_j) \rho_{j0} \tag{5.7}$$

where $\lambda_j < \mu_j$ for each j by our assumption that the network is stable, and at least one ρ_{j0} is greater than zero by our assumption that it is irreducible. Hence,

$$\sum_{j \in U_I'} (\nu_j^I - \mu_j) < 0. \tag{5.8}$$

We infer from (5.4a) and (5.4b) and the definition of U_I' that

$$\sum_{j \in U_I'} X_j(t) = \sum_{j \in U_I'} \left[p_j + (\nu_j^I - \mu_j)t \right], \qquad 0 \le t \le \tau \tag{5.9a}$$

$$X_j(t) = 0, \quad \forall j \in [J] - U_I', \qquad 0 \le t \le \tau. \tag{5.9b}$$

It follows from (5.8), (5.9a), and (5.9b) that $\sum_{j \in [J]} X_j(t)$ is decreasing at a constant positive rate on $[0, \tau]$. Since $\sum_{j \in [J]} X_j(0) = \sum_{j \in I} p_j$ is finite, it follows that τ is finite.

So far, we have shown negative drift until some node in I, the set of initially nonempty nodes, becomes empty. At this time, the flow rates change, but a similar argument can be used to show that the fluid limit has negative drift with the new rates. Also, the nodes that were stable earlier continue to be stable under the new rates. Hence, nodes not in U_I' continue to remain empty. This argument can be repeated inductively until the time that the total amount of fluid in the system becomes zero. The induction is made precise in the remainder of this subsection.

Let $\tilde{I} = \{i \in [J]: X_i(\tau) > 0\}$. It is clear from (5.9b) that $\tilde{I} \subseteq U_I'$. In fact, it is a proper subset since $X_i(\tau) = 0$ for some $i \in I$ by definition of τ. It is then clear from (5.3) that $\nu^{\tilde{I}} \le \nu^{U_I'}$. Also observe from (5.3) that $\nu^I = \nu^{U_I'}$. It is now an easy inference that $U_{\tilde{I}} \subseteq U_I$ and $U_{\tilde{I}}' \subseteq U_I'$, where the latter inclusion is strict. To see this, observe that if $i \in I - \tilde{I}$ (there is such an i), then $X_i(t)$ was decreasing in $[0, \tau]$, so $i \notin U_I$, consequently $i \notin U_{\tilde{I}}$, and therefore $i \notin U_{\tilde{I}}'$.

We now consider the flow process when started in the initial condition $X(0) = p_j e_j$, where e_j is the jth unit vector. Let $(X(t), t \ge 0)$ denote the resulting flow process specified by (5.2)–(5.2c). Define $I_0 = \{j\}$, $\tau_0 = \inf\{t > 0: X_j(t) = 0\}$. Assume inductively that $I_0, \cdots, I_n$ and $\tau_0, \cdots, \tau_n$ have been defined. Define

$$I_{n+1} = \{i \in [J]: X_i(\tau_n) > 0\}$$

and, if $I_{n+1} \ne \varnothing$, define

$$\tau_{n+1} = \inf\{t > \tau_n: X_i(t) = 0 \text{ for some } i \in I_{n+1}\};$$

else, if $I_{n+1} = \varnothing$, define $\tau_{n+1} = \tau_n$. Finally, define $\tau = \sup_n \tau_n$.

Our arguments above imply that $\tau_{n+1} - \tau_n$ is finite for each n, and that U_{I_n}' is a strictly decreasing sequence of sets until it becomes empty. Since $I_n \subseteq U_{I_n}'$, it follows that only finitely many I_n are nonempty (at most J, in fact), and therefore only finitely many $\tau_{n+1} - \tau_n$ are nonzero. Hence, τ is finite. Also observe from (5.8), (5.9a), and (5.9b), applied to each phase $[\tau_n, \tau_{n+1}]$, that

$$\sum_{i \in U_I'} X_i(t) \le p_j, \qquad \forall\, 0 \le t \le \tau \tag{5.10a}$$

since the total amount of fluid initially in the network is a quantity p_j at node j and zero at the remaining nodes. Also,

$$X_i(t) = 0, \qquad \forall\, 0 \le t \le \tau, \forall\, i \notin U_{\{j\}}'. \tag{5.10b}$$

In fact,

$$X_i(t) = 0, \qquad \forall\, \tau_0 \le t \le \tau, \forall\, i \notin U_{\{j\}}. \tag{5.10c}$$

C. Proof of the Main Result

In this subsection, we complete the proof of the main result. We have obtained above an expression for the fluid limit and have computed its drift. We combine this with (5.1) to obtain bounds on the actual queue length process. Observe that, by (5.1),

$$\lim_{N \to \infty} P\left\{ \sup_{0 \le t \le \tau} \left\| \frac{1}{N} X^N(Nt) - X(t) \right\| \ge \epsilon \right\} = 0$$

for all $\epsilon > 0$, where $\|X\| = \max_{i \in [J]} |X_i|$. Then, (5.10a)–(5.10c) imply that the following statements are true with probability going to one as N goes to infinity:

$$\sum_{i \in [J]} X_i^N(t) < (p_j + \epsilon)N, \qquad \forall\, 0 \le t \le N\tau \tag{5.11a}$$

so the inequality holds, in particular, for each $i \in U_j$:

$$X_i^N(t) < \epsilon N, \qquad \forall\, 0 \le t \le N\tau, \forall\, i \notin U_{\{j\}}' \tag{5.11b}$$

and, in addition,

$$X_j^N(t) < \epsilon N, \qquad \forall\, N\tau_0 \le t \le N\tau \tag{5.11c}$$

for all $\epsilon > 0$. Finally, since $X(\tau) = 0$, by definition of τ, we also have

$$X_i^N(N\tau) < \epsilon N, \qquad \forall\, i \in [J] \tag{5.12}$$

with asymptotic probability one.

We thus have bounds on the queue length until the time $N\tau$ that the fluid limit hits zero. We would like to extend this to the time that the actual queue length process hits zero. In order to do so, we need the following lemma, which is proved in [2].

Lemma 5.2: Consider a stable, irreducible, open Jackson network, started with a total of N customers in the system. Then, the time for the network to empty is stochastically dominated by the sum of N independent, identically distributed random variables of finite mean and variance.

Let $T_N = \inf\{t > 0: X_i^N(t) = 0\ \forall\, i \in [J]\}$. Then, by the above lemma and (5.12), $T_N - N\tau$ is stochastically dominated by the sum of ϵJN i.i.d. random variables of finite mean and variance. Since the exogenous arrival process is Poisson of rate γ, the total number of exogenous arrivals in the period $[\tau N, T_N]$ (taken to be empty if $\tau N > T_N$) is less than a constant times ϵN, with probability going to one as $N \to \infty$. In conjunction with (5.12), this implies that

$$\sum_{i \in [J]} X_i^N(t) < k\epsilon N, \qquad \forall\, N\tau \le t \le T_N \tag{5.13}$$

with asymptotic probability one, where k is some constant independent of N and $\epsilon > 0$ is arbitrary. This allows us to extend the validity of (5.11a)–(5.11c) to the entire period $[0, T_N]$, where T_N is the first time that the system is empty.

In (5.11a)–(5.11c), we have obtained bounds on the queue length process until the first time that it hits zero. We wish to use these bounds to show that there is no buffer overflow during this time. To that end, we make use of the following lemma and its corollary.

Lemma 5.3:

$$\frac{\mu_k}{\lambda_k} < \frac{\mu_j}{\lambda_j}, \qquad \text{for all } k \in U_{\{j\}}.$$

Proof: For a J vector η, define

$$\phi_k(\eta) = \gamma r_{0k} + \sum_{i \in [J]-j} (\eta_i \wedge \mu_i) r_{ik} + \mu_j r_{jk}.$$

Then, the solution $\nu^{\{j\}}$ to the generalized flow balance equations (5.3) with $I = \{j\}$ is the unique fixed point of ϕ. Observe that ϕ is monotone, i.e., if $\eta \geq \xi$, then $\phi(\eta) \geq \phi(\xi)$. Let $\lambda_i, i \in [J]$ solve the original flow balance equations (2.1). Consider

$$\begin{aligned}\phi_k\left(\frac{\mu_j}{\lambda_j}\lambda\right) &= \gamma r_{0k} + \sum_{i \in [J]-j} \left(\frac{\mu_j}{\lambda_j}\lambda_i \wedge \mu_i\right) r_{ik} + \mu_j r_{jk} \\ &\leq \frac{\mu_j}{\lambda_j}\gamma r_{0k} + \sum_{i \in [J]-j} \frac{\mu_j}{\lambda_j}\lambda_i r_{ik} + \frac{\mu_j}{\lambda_j}\lambda_j r_{jk} \\ &= \frac{\mu_j}{\lambda_j}\lambda_k\end{aligned}$$

where we have used (2.1) and the fact that $\mu_j > \lambda_j$. Hence, $\phi((\mu_j/\lambda_j)\lambda) \leq (\mu_j/\lambda_j)\lambda$, and it follows from our arguments above that $(\mu_j/\lambda_j)\lambda \geq \nu^{\{j\}}$. But $\mu_k < \nu_k^{\{j\}}$ for $k \in U_{\{j\}}$ by definition of $U_{\{j\}}$; hence, $\mu_k < (\mu_j/\lambda_j)\lambda_k$ $\forall$ $k \in U_{\{j\}}$, which proves the lemma.

Corollary: $p_k > p_j$ for all $k \in U_{\{j\}}$. This is a simple consequence of the rule of thumb that $p_k \log(\mu_k/\lambda_k)$ is constant and the above lemma.

If $(N_1, \cdots, N_J)$ is a buffer allocation consistent with (3.1), then $|N_k - p_k N| < K \log N$. In conjunction with the corollary above and (5.11a), which was shown earlier to hold for all $0 \leq t \leq T_N$, this implies that $X_i^N(t) < N_i$ $\forall$ $0 \leq t \leq T_N$, $\forall$ $i \in U_{\{j\}}$, with probability going to one as N goes to infinity. Likewise, (5.11b), coupled with the observation that $p_i > 0$ for all $i \in [J]$, implies the same result for queues $i \notin U'_{\{j\}}$. We have thus shown that, for all buffer allocations consistent with (3.1), the network started with N_j customers in the jth queue and all other queues empty has asymptotic probability one of the emptying before suffering a buffer overflow in any queue other than queue j. That leaves us with the task of bounding the probability of a buffer overflow in the jth queue before the system empties.

Let $\tilde{X}^N(t)$ denote the process started in the same initial condition as $X^N(t)$, but with the output of the queues in $U'_{\{j\}}$ replaced by their virtual departure processes. Clearly, the queue length process $\tilde{X}^N(t)$ dominates $X^N(t)$ (to see this, color red the virtual departures from nodes in $U'_{\{j\}}$ that are not actual departures, color blue all other departures from all nodes and exogenous arrivals, give blue customers preemptive service priority, noting that this does not affect the distribution of total number in queue, and observe that $X^N(t)$ is the process of blue customers, while $\tilde{X}^N(t)$ is the process of all customers). It is easy to check that $\tilde{X}^N(t)$ evolves like an unstable Jackson network, the inflow rates into whose queues are given by the solution $\nu^{\{j\}}$ of the generalized flow balance equation (5.3). If we now consider an initial condition where queues outside $U'_{\{j\}}$ are in their stationary distributions, then the inflow process into queue j will be Poisson of rate $\nu_j^{\{j\}}$. This dominates the inflow process into queue j in the process $\tilde{X}^N(t)$, wherein the queues outside $U'_{\{j\}}$ were initially empty (this can also be shown using the coloring technique employed above).

We thus see that $X_j^N(t)$ is dominated by a birth and death process $Y(t)$, of birth rate $\nu_j^{\{j\}}$, death rate μ_j, and started at $Y(0) = N_j$. Now, from (5.8) applied to $I = \{j\}$, we see that $\sum_{i \in U'_{\{j\}}} (\nu_i^{\{j\}} - \mu_i) < 0$, where $U'_{\{j\}} = j \cup U_{\{j\}}$. But $\nu_i^{\{j\}} > \mu_i$ for $i \in U_{\{j\}}$ by definition of $U_{\{j\}}$. Hence, $\nu_j^{\{j\}} < \mu_j$. Now, observing that $(\mu_j/\nu_j^{\{j\}})^{Y(t)}$ is a martingale, it can easily be seen that

$$P\big(Y(t) = N_j + 1 \text{ before } Y(t) = 0\big) < \nu_j^{\{j\}}/\mu_j$$

and since $Y(t)$ dominates $X_j^N(t)$, it follows that

$$P\big(X_j^N(t) = N_j + 1 \text{ before } X_j^N(t) = 0\big) < \nu_j^{\{j\}}/\mu_j. \quad (5.14)$$

That is, queue j empties before overflowing with probability bounded away from zero, uniformly in N.

Let $\tilde{T} = \inf\{t > 0\colon\ Y(t) = 0\}$. Then, since $Y(t)$ is a stable birth and death process,

$$\lim_{N \to \infty} P\big(Y(t) < N_j, \quad \forall\, t \in [\tilde{T}, N\tau_0]\big) = 1$$

since τ_0 is a constant that can be computed from the rates, and a stable birth and death process does not grow by N in time linear in N, with asymptotic probability one. Hence, $X_j^N(t)$ does not overflow before the time $N\tau_0$, with probability bounded uniformly away from zero. Applying (5.11c) to the interval $[N\tau_0, T_N]$, we see that $X_j^N(t) < N_j$ for all t in this interval with probability asymptotically equal to one. Putting these results together, we conclude that the probability of the jth queue suffering a buffer overflow before the system empties is bounded uniformly away from zero as N goes to infinity.

We have thus shown that, for buffer allocations $(N_1, \cdots, N_J)$ consistent with (3.1), the queue length process $X(t)$ with initial state $X(0) = N_j e_j$ satisfies

$$\liminf_{N \to \infty} P(X(t) = 0 \text{ before } X_i(t) = N_i + 1, \quad \text{for any } i \in [J]) > 0.$$

This completes the proof of Proposition 4.1.

It is, in fact, possible to obtain an explicit lower bound in the last equation. Observe from (5.14) and the arguments following it that the probability of a buffer overflow in the jth queue before the system empties is bounded above by $\nu_j^{\{j\}}/\mu_j$ for large enough N. Also, by the arguments involving the fluid limit, it was shown that the probability of a buffer overflow before the system empties, in a queue other than queue j, goes asymptotically to zero. Therefore,

$$\liminf_{N\to\infty} P(X(t) = 0 \text{ before } X_i(t) = N_i + 1, \text{ for any } i \in [J]) \geq 1 - \frac{\nu_j^{\{j\}}}{\mu_j}$$

under the initial condition $X(0) = N_j e_j$. But the above probability is precisely $\tilde{P}(N_j e_j, 0)$. Hence, using (4.10) and (4.16), we obtain the estimate for large enough N,

$$c_1 = \frac{\gamma + \sum_{i\in[J]}\mu_i}{\gamma\pi(0)} \cdot \left(\prod_{i\in[J]}(1-\rho_i)\right) \cdot \min_{j\in[J]} \frac{1}{2}\tilde{\beta}_j\left(1 - \frac{\tilde{\nu}_j^{\{j\}}}{\mu_j}\right)$$

where (4.15) gives a lower bound on $\tilde{\beta}_j$, and we write $\tilde{\nu}_j^{\{j\}}$ instead of $\nu_j^{\{j\}}$ since this quantity refers to the time-reversed network. This completes the proof of the theorem.

VI. Conclusion

The problem of allocating storage in a Jackson network so as to avoid frequent buffer overflows was considered, and it was shown that a rule of thumb proposed for this problem in [2] is within a constant of the optimal allocation.

It would be of considerable interest to extend this result to less restrictive models of queueing networks. This is particularly so in the context of manufacturing systems, where the Jackson network model is not very realistic. Another extension that would be of interest is to multi-class networks, where customers could belong to one of a finite set of classes. It may be remarked here that while the method of allocating buffers that we have considered has given rise to cuboidal boundaries in the corresponding Markov chain, essentially the same techniques developed here could be applied to different boundary shapes arising from a different model.

The results derived in this paper are of an asymptotic nature. It would be useful to have error bounds that are applicable for all finite values of N. This may be easier to accomplish for specific cases of the cost function, such as the mean time to buffer overflow, for instance.

References

[1] D. Aldous, *Probability Approximations via the Poisson Clumping Heuristic*. Berlin: Springer-Verlag, 1989.

[2] V. Anantharam, "The optimal buffer allocation problem," *IEEE Trans. Inform. Theory*, vol. 35, July 1989.

[3] V. Anantharam, P. Heidelberger, and P. Tsoucas, "Analysis of rare events in continuous time Markov chains via time reversal and fluid approximation," IBM Res. Rep. RC 16280(#71858), Oct. 1990.

[4] H. Chen and A. Mandelbaum, "Discrete flow networks: Bottleneck analysis and fluid approximation," *Math. Oper. Res.*, vol. 16, May 1991.

[5] M. Cottrell, J. Fort, and G. Malgouyres, "Large deviations and rare events in the study of stochastic algorithms," *IEEE Trans. Automat. Contr.*, vol. AC-28, Sept. 1983.

[6] 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.*, vol. 17, no. 6, 1979.

[7] I. Iscoe and D. McDonald, "Asymptotic of exit times for Markov jump processes," preprint, Feb. 1991.

[8] S. P. Meyn and M. R. Frater, "Recurrence times of buffer overflows in Jackson networks," *IEEE Trans. Inform. Theory*, to be published.

[9] S. Parekh, "Rare events in networks of queues," Ph.D. dissertation, Univ. California, Berkeley, 1986.

[10] J. Walrand, "Quick simulations of networks of GI/G/1 queues," plenary lecture presented at the 16th Conf. Stochastic Processes and Their App., Stanford Univ., 1987.

[11] ——, *An Introduction to Queueing Networks*. Englewood Cliffs, NJ: Prentice-Hall, 1987.