Adv. Appl. Prob. **22**, 263–267 (1990)
Printed in N. Ireland

ON THE ERGODICITY OF NETWORKS OF $\cdot/GI/1/N$ QUEUES

PANAGIOTIS KONSTANTOPOULOS* AND
JEAN WALRAND,** *University of California, Berkeley*

Abstract

In this paper we consider a network of $\cdot/GI/1/N$ queues with finite-length buffers. A suitable Markov process for the time-evolution of this system is defined. This process is subsequently shown to be ergodic under the conditions of Borovkov (1987).

QUEUEING NETWORKS, REGENERATIVE PROCESSES

1. Introduction

The objective of this paper is the study of the asymptotic stationarity of a network consisting of finitely many queues with finite-length buffers. This is a typical model of a computer communication network where the customers are interpreted as messages or packets. First we shall assume that the network is closed. Each queue is modeled in this paper as a $\cdot/GI/1/N$ queue with first-come first-served queueing discipline. Thus, to each queue i, we assign a sequence $\{\sigma_i^{(k)}, k=1, 2, \cdots\}$ of i.i.d. random variables representing the successive service times in the queue. These service sequences are also mutually independent. When a customer finishes service in queue i it decides to jump queue j with probability p_{ij}, independently of everything else.

It is then expected that under minimal natural assumptions (non-lattice service times with finite means and irreducible routing matrix) the network is ergodic in the sense that a suitable Markovian state process (which at each time summarizes the past evolution) possesses a uniquely defined stationary probability distribution and, furthermore, this process starting from any initial condition converges weakly to its stationary distribution. If the service times are exponential or of phase type (see Walrand (1988), Asmussen (1987)) then the question has a trivial affirmative answer since the network can be described by a finite-state irreducible Markov chain. To answer this question under a more general setup we restrict the class of service time distributions to that of Borovkov (1987) or Gnedenko and Kovalenko (1968) (see definition in the next section). Borovkov (1987) has answered the same question posed above for a network of $\cdot/GI/1/N$ queues with infinite-capacity buffers. The difference in our problem is that it is harder to visualize the movement of the customers in the network when there is blocking and thus the problem seems more complicated. It turns out though that there is an argument (see Konstantopoulos and Walrand (1988)) for a sketch of the proof) that overcomes this difficulty and yields an easy answer. In Section 2 we give a more detailed description of the model and the assumptions and state the theorem. The proof of the

Received 7 March 1989; revision received 25 August 1989.

* Present address: INRIA, Unité de Recherche Sophia-Antipolis, 2004 Route des Lucioles, 06565 Valbonne Cedex, France.

** Postal address: Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, USA.

theorem is given in detail in Section 3. Finally, in Section 4, we discuss some related problems.

2. Preliminaries and the model

We consider a closed network. Let N be the number of queues and n the total number of customers in the network. Let b_i be the length of the buffer of queue i (the population in queue i is restricted to be always smaller than or equal to the number b_i). Let $P=[p_{ij}]$ be the routing matrix which is assumed to be irreducible. Finally, let σ_i denote a typical service time in queue i having distribution F_i and mean $\mu_i^{-1}<\infty$.

Next we discuss the blocking and the conventions associated with it. We say that queue i is blocked by queue j at time t if queue i is non-empty at time t, queue j is full at the same time, and the customer who is at the server in queue i has decided to jump to queue j afterwards. The type of blocking that we consider is called manufacturing blocking and works as follows: if queue i is blocked by queue j when the customer in i finishes service then this customer remains in queue i (while, at the same time, its server remains idle) until the first time an empty spot appears in j. This type of blocking may result in the following deadlocks:

(i) A number of queues are simultaneously blocked by the same queue. Then we assign the first empty spot of this queue at random among the blocked customers of the other queues.

(ii) There is collection of queues $i_1, \cdots, i_k$ such that i_m is blocked by i_{m-1} $(2\leqq m\leqq k)$, i_1 is blocked by i_k and all but i_1 are idle. When the customer in i_1 finishes service we pick at random a queue $i\in\{i_2,\cdots,i_k\}$ for which there is a $j\notin\{i_1,i_2,\cdots,i_k\}$ with $p_{ij}>0$ and redirect the customer of queue i to one of these queues j chosen again at random.

(iii) A queue is blocked by itself. Then we redirect the customer when it finishes service to another queue at random.

Of course, one can think of many other ways of resolving the deadlocks. However, the conventions introduced above are particularly appealing because they simplify the state of the network.

As a state of the network at a given time we may choose the triple

$$x=(\nu,\xi,s)$$

where ν, ξ, s are all N-tuples. The ith component ν_i of ν is the number of customers in queue i, ξ_i is the remaining service time of the customer receiving service in queue i (if queue i is empty then $\xi_i=0$), and s_i is the index of the queue that this customer has decided to visit next (if queue i is empty then we give s_i some dummy value). It is clear that x summarizes the past and the resulting process $\{x(t),t\geqq 0\}$ is Markov.

The following assumptions are made concerning the service time distributions:

(A1) They are non-lattice.

(A2) There is a queue, say queue 1, with unbounded service time.

(A3) The service times are *strongly Cramér* with common parameter α: this means that there is an $\alpha>0$ and an $\varepsilon>0$ such that

$$P(T<\sigma_i\leqq t+\alpha)\geqq\varepsilon P(\sigma_i>t),\quad \text{for all}\quad t\geqq 0\quad\text{and all}\quad 1\leqq i\leqq N.$$

The terminology 'strongly Cramér' is non-standard. It suggests only that if a random variable σ_i has this property then the moment generating function of σ_i exists in a neighborhood of 0 (which is the Cramér condition), its tail distribution $P(\sigma_i>t)$ goes to 0 exponentially fast as $t\to\infty$ and, consequently, has moments of all orders. Intuitively, this condition can be thought of as a very weak memoryless property in the sense that if it is given that σ_i has exceeded t then there is a positive probability, which is bigger than ε, that σ_i will occur in the interval $(t,t+\alpha]$, where this ε is independent of t. See Borovkov (1987) and Gnedenko and Kovalenko (1968) for other uses of this condition.

One of the most common examples of random variables satisfying condition (A3) is that of a bounded random variable. It is indeed trivial to observe that (A3) is satisfied if α is equal to an upper bound of the random variable. We would also like to mention that condition (A2), that is, existence of a service station with unbounded service time, is a reasonable model for a machine with breakdowns. Assume that the service time, when there are no breakdowns, is denoted by τ. Assume that τ is a bounded random variable. The breakdowns occur in a Poisson fashion with rate $\lambda > 0$. Assume also that the repair times are bounded random variables distributed like r. Then the actual service time of this station is $\sigma = \tau + \sum_{k=1}^{N(0,\tau)} r_k$. The r_k's are independent copies of r and $N(0, \tau)$ denotes the number of points of a Poisson process N of rate λ in the interval $(0, \tau)$. (Note that N, τ, r_k, $(k \geqq 1)$ are independent.) It can then be seen that, no matter how small λ is, there is always a positive probability that σ exceeds any finite number.

Let P_x denote the distribution of the Markov process $\{x(t)\}$ when it starts from the state x. The main theorem can then be stated as follows.

Theorem 1. The Markov process $\{x(t)\}$ has a unique stationary distribution π which is also a limiting distribution:

$$P_x(x(t) \in \cdot) \to \pi(\cdot) \quad \text{as} \quad t \to \infty,$$

where the convergence is meant to be weak convergence of probability distributions.

3. Proof of ergodicity

Consider a maximal spanning tree $\boldsymbol{T}$ for the directed graph of the network corresponding to the matrix P. Let the tree $\boldsymbol{T}$ have its root at node 1, which, by assumption (A2), is the node with the unbounded service time. Let also $\boldsymbol{T}$ be directed towards the node 1. Thus, for any node $i \neq 1$ there is a unique path of the tree $\boldsymbol{T}$ connecting i with node 1. Let i^* denote the successor of a node $i \neq 1$ with respect to the direction of the tree. Consider now the set of 'blocking states' associated with the chosen tree, that is, all states x for which node 1 is full with its server having a positive amount of workload, and each node $i \neq 1$ is either empty or blocked by node i^* with its server idle. Formally this is the set

$$B = \{x : v_1 = b_1,\ \xi_1 > 0 \text{ and } s_i = i^*,\ \xi_i = 0, \text{ for all } i \neq 1 \text{ with } v_i > 0\}.$$

If we let T_B be the first positive time at which the state process $x(t)$ is in B, we propose to show, as in Borovkov (1987), that there is a $t^* > 0$ and a $p^* > 0$ such that

$$P_x(T_B \leqq t^*) \geqq p^*, \quad \text{for all } x. \tag{3.1}$$

To do this, we let $\{\sigma_i^{(k)}, k = 1, 2, \cdots\}$ be the i.i.d. sequence of successive service times assigned by the server of queue i after time 0, and $\{s_i^{(k)}, k = 1, 2, \cdots\}$ be the i.i.d. sequence of routing decisions taken in queue i after time 0. Let also $t^* = \alpha + \beta$ where α is as in (A2) and β will be defined below. Then the event $T_B < \alpha + \beta$ is implied by the following event: all customers served at all queues at time 0 finish service before time α, the first new service of queue 1 is sufficiently long (longer than $\alpha + \beta$), the first n routing decisions for each queue $i \neq 1$ are to send the finished customer to queue i^*, and for each queue $i \neq 1$ the sum of its first n service times does not exceed β / N. Hence

$$\begin{aligned} P_x(T_B \leqq \alpha + \beta) \geqq P_x(\xi_i(0) \leqq \alpha, \text{ for all } i) P(\sigma_1^{(1)} > \alpha + \beta) \\ P(s_i^{(1)} = \cdots = s_i^{(n)} = i^*, \text{ for all } i \neq 1) \\ P(\sigma_i^{(1)} + \cdots + \sigma_i^{(n)} \leqq \beta / N, \text{ for all } i \neq 1). \end{aligned} \tag{3.2}$$

We have omitted the subscript x from the last three terms of the right-hand side to show their independence of the initial condition. It follows from (A3) that the first term is bounded below by ε^N and from (A2) that the second term is positive. It is also clear that the third term is positive and that the last one is positive for a suitable choice of β. This proves (3.1) with p^* equal to the right-hand side of (3.2).

The rest of the argument proceeds as follows. Observe that the set B can be decomposed into components B_m, where m ranges over some index set M, and where each component is specified like B with the additional requirement that the number of customers in each queue is fixed. Since there are finitely many customers and queues, the number M of components is finite. (Note that if $n \leqq b_1$ then there is only one component.) Let $T_B^{(k)}$ (or $T_{B_m}^{(k)}$) be the kth time that the state process makes the transition $B^c \to B$ (or $B_m^c \to B_m$). Since $B = \bigcup_m B_m$ we have

$$\frac{1}{t}\sum_k 1(T_B^{(k)} \leqq t) = \sum_m \frac{1}{t}\sum_k 1(T_{B_m}^{(k)} \leqq t). \tag{3.3}$$

Note now that (3.1) implies that each $T_B^{(k+1)} - T_B^{(k)}$ is bounded stochastically from above by an exponential random variable S with positive rate $\lambda = |\log(1-p^*)|/t^*$. This implies that

$$\limsup_{t\to\infty} \frac{1}{t}\sum_k 1(T_B^{(k)} \leqq t) > \lambda > 0. \tag{3.4}$$

Since M is finite, (3.3) and (3.4) imply that there is an m such that

$$\limsup_{t\to\infty} \frac{1}{t}\sum_k 1(T_{B_m}^{(k)} \leqq t) := \lambda_m > 0. \tag{3.5}$$

Fix this m and let $S_m^{(k)}$ be the first time after $T_{B_m}^{(k)}$ that the state process leaves the set B_m. Observe that the process regenerates at these times. The points $\{S_m^{(k)}\}$ form a renewal process and hence the $\lim_{t\to\infty} 1/t \sum_k 1(S_m^{(k)} \leqq t)$ exists and is clearly equal to λ_m as in (3.5). Standard theory of regenerative processes (Asmussen (1987)) shows then that there is a unique stationary distribution for the process $\{x(t)\}$ given by the formula

$$\pi(A) = \lambda_m E_x \int_{S_m^{(1)}}^{S_m^{(2)}} 1(x(t) \in A)\, dt.$$

Finally, to show weak convergence, we have to show that the distribution of a typical increment $S_m^{(k+1)} - S_m^{(k)}$ is not supported on a lattice. This is done by contradiction. Suppose that its distribution is supported on a lattice $\{0, l, 2l, \cdots\}$. Observe that the number $N_m[rl, rl+\delta]$ of renewals $S_m^{(k)}$ on the any interval $[rl, rl+\delta]$ (where r is a positive integer and δ a positive real) is at most equal to the number $N_1[rl, rl+\delta]$ of service completions by queue 1 over the same interval. If we take $\delta < l$ then $\lim_{r\to\infty} EN_m[rl, rl+\delta] = \lambda_m$, while $\limsup_{r\to\infty} EN_1[rl, rl+\delta] \leqq \mu_1\delta$, since the service processes are non-lattice, by (A1). We then have $\lambda_m \leqq \mu_1\delta$ for all $\delta < l$, which means that $\lambda_m = 0$. This contradicts (3.5).

4. Additional comments

4.1. *Convergence in total variation—rates of convergence.* Suppose that the service time distributions are spread-out, i.e., they have a component with density. Then the distribution of a typical renewal increment $S_m^{(k+1)} - S_m^{(k)}$ is also spread-out. Indeed, let $\sigma_1^{(k)}$ denote the last service time of server 1 on the interval $(S_m^{(k)}, S_m^{(k+1)})$. Observe then that $S_m^{(k+1)} - S_m^{(k)}$ is written as the sum of $\sigma_1^{(k)}$ plus an independent time $Y^{(k)}$. The sequence $\{Y^{(k)}\}$ is i.i.d. because the system regenerates at the time $S_m^{(k)}$. Hence the distribution of $S_m^{(k+1)} - S_m^{(k)}$ is the convolution of the distribution of $\sigma_1^{(k)}$ and $Y^{(k)}$. As the distribution of $\sigma_1^{(k)}$ has a component with density, so does the distribution of $S_m^{(k+1)} - S_m^{(k)}$ (see Feller (1971), Chapter V).

The theory of regenerative processes guarantees that, in this case, the convergence of $x(t)$ towards its stationary distribution is in the sense of total variation, i.e.,

$$\sup_A |P_x(x(t) \in A) - \pi(A)| \xrightarrow[t\to\infty]{} 0, \tag{4.1}$$

where the supremum is taken over all measurable sets of states A.

Although condition (A3) is a strong condition on the distributions of the service times, it has the advantage that when it holds it gives exponential rates of convergence in (4.1). For a proof of this see Borovkov (1987).

4.2. *Existence of stationary distribution without the conditions* (A1), (A2) *and* (A3). If we assume only that the service time distributions have finite means then the analysis of Section 3 is not valid. However, if we are interested only in the existence of a stationary distribution of the x-process, then we can obtain it by reasoning as in Whitt (1980). First, we replace the service time distributions F_i, $1 \leqq i \leqq N$ by phase-type distributions $F_i^{(n)}$, $1 \leqq i \leqq N$ and we thus obtain the nth approximation to the original network. The phase-type distributions are chosen so that $\lim_{n\to\infty} F_i^{(n)} = F_i$ (weak convergence) and $\lim_{n\to\infty} \int tF_i^{(n)}(dt) = \int tF_i(dt)$ (i.e., the means converge to the original ones). As was mentioned in the introduction, the stationary distribution $\pi^{(n)}$ of the nth approximation exists. It then remains to show that the sequence $\{\pi^{(n)}\}$ is tight. Any limit point of the sequence will then be a stationary distribution for the original network. For a proof of this, in a more general context, see Whitt (1980).

4.3. *Ergodicity without the conditions* (A2) *and* (A3). While it seems plausible that Theorem 1 holds under general finite-mean service time distributions, provided that they are non-lattice and that the routing matrix is irreducible, the problem remains open. For the case of a closed network with infinite-buffer queues a similar restricted ergodicity result has been obtained in Sigman (1989). The restriction comes from the fact that the author assumes that $p_{ij} > 0$ for all i and j.

4.4. *Ergodicity of open networks of* $\cdot/GI/1/N$ *queues.* Consider finally an open network with one arrival stream and finite buffers in all queues. Assume that conditions (A1) and (A3) hold and make the additional assumption that the interarrival time in an unbounded and strongly Cramér random variable. Since the total number of customers in the network can never exceed the sum of all buffer sizes, we can easily see that the proof of ergodicity for such a network is the same as that of Section 3 where σ_1 now represents the interarrival time.

References

ASMUSSEN, S. (1987) *Applied Probability and Queues.* Wiley, New York.

BOROVKOV, A. A. (1987) Limit theorems for queueing networks I. *Theory Prob. Appl.* **31,** 413–427.

FELLER, W. (1971) *An Introduction to Probability Theory,* Vol. II. Wiley, New York.

GNEDENKO, B. AND KOVALENKO, I. N. (1968) *An Introduction to Queueing Theory.* Israel Program for Scientific Translations, Jerusalem.

KONSTANTOPOULOS, P. AND WALRAND, J. (1988) Ergodicity of networks with blocking. *Proc. 27th CDC.*

SIGMAN, K. (1989) Notes on the stability of closed queueing networks. *J. Appl. Prob.* **26,** 678–682.

WALRAND, J. (1988) *An Introduction to Queueing Networks.* Prentice-Hall, Englewood Cliffs, NJ.

WHITT, W. (1980) Continuity of generalized semi-Markov processes. *Math. Operat. Res.* **5,** 494–501.