MATHEMATICS OF OPERATIONS RESEARCH
Vol. 29, No. 1, February 2004, pp. 182–190
ISSN 0364-765X | EISSN 1526-5471 | 04 | 2901 | 0182
DOI 10.1287/moor.1030.0067

The Mean Number-in-System Vector Range for Multiclass Queueing Networks

B. Curtis Eaves, Peter W. Glynn
Department of Management Science and Engineering, Stanford University,
Stanford, California 94305 {bceaves@stanford.edu, glynn@stanford.edu}

Uriel G. Rothblum
Faculty of Industrial Engineering and Management, Technion–Israel Institute of Technology,
Haifa 32000, Israel, rothblum@ie.technion.ac.il

In a multiclass network of queues, natural sufficient conditions are given for the mean number-in-system vector to cover the nonnegative orthant as the arrival rate vector varies over the stable domain.

Key words: networks of servers; surjections; onto multifunctions; queueing
MSC2000 subject classification: Primary: 90B22
OR/MS subject classification: Primary: Queues/networks
History: Received August 10, 1998; revised February 15, 2003, and June 11, 2003.

1. Introduction. Our system is a network of servers with customers belonging to different classes; we view this network as a "blackbox." Customers enter the system, are processed by servers (stations), and leave the system. In the system, for example, customers might follow a path (possibly with cycles) of servers where the path depends on the customer class. The service time at each server is random and depends on the customer class. For the system, components of the arrival-rate vector λ correspond to the arrival rate of customers in each class. The stable arrival-rate vectors are those with a steady state in which the expected number of customers in each class in the system is finite. For a stable arrival-rate vector λ, let the components of the mean number-in-system vector function $f(\lambda)$ correspond to the steady-state expected number of customers in each class in the system. Natural and intuitive sufficient conditions are established for the range of the mean number-in-system vector function to cover all nonnegative vector values as arrival-rate vectors vary over the stable set. Viewing the arrival-rate vector as a "control" and the mean number-in-system vector as the "target," the result establishes conditions under which any desired mean number-in-system vector can be attained by choosing the appropriate value of the control.

Our perspective is motivated by results of Eaves and Rothblum (1993). We apply path-following analysis to provide conditions for a real function to cover the nonnegative orthant; see Eaves (1976) and Eaves and Rothblum (1993). We apply queueing analysis to transform these conditions to the context of queueing networks; see Glynn and L'Ecuyer (1995).

We start with some preliminaries in §2. In §3, we state the covering theorem and interpret it in the queueing context. The covering theorem is proved in §4. In §5, tools for verifying assumptions of the covering theorem for a network of queues are developed.

2. Preliminaries. Let $\nu \equiv \{1, \ldots, n\}$ for a positive integer n. Let R^n denote *n-dimensional Euclidean space* and let $R^n_\oplus$ be the set of nonnegative vectors in R^n. The *maximum norm* in R^n will be denoted $\|\ \|$; that is, $\|\cdot\|: R^n \to R$ with the norm of $x = (x_1, \ldots, x_n)$ as $\|x\| = \max_{i\in\nu} |x_i|$. For a sequence $x_1, x_2, \ldots$ of vectors in R^n, we abbreviate $x = \lim_{k\to+\infty} x_k$ by $x_k \to x$. By $x_k \to \infty$ we mean $\|x_k\| \to +\infty$; that is, for every $m > 0$ there is an integer k such that $\|x_h\| > m$ for all $h > k$; in this case we sometimes say that the sequence is *blowing up*. In our system, a network of servers, we index our customer classes with i in $\nu = \{1, \ldots, n\}$. The arrival process is a parameterized family of stochastic processes that is characterized by a nonnegative n-vector $\lambda = (\lambda_1, \ldots, \lambda_n)$, referred to as the *arrival-rate vector*. For example, this is the case when the interarrival sequences of customers of different types follow exponential distributions where the mean interarrival time for customers of class i is λ_i^{-1}.

We say that a vector $\lambda \in R^n_\oplus$ is *weakly stable* if, under the arrival process corresponding to λ, a steady-state distribution exists for the number of customers of each class in the system. If, in addition, the steady-state number-in-system vector has finite mean, then we say that λ is *stable*. In this case, the vector of mean number of customers of each class is denoted $f(\lambda)$ and is referred to as the *mean number-in-system vector associated with* λ. Let S_0 be the set of weakly stable arrival-rate vectors, and let S be the set of stable arrival-rate vectors. For the pure infinite capacity case, S is $R^n_\oplus$, and for the pure finite capacity case S is bounded (by "pure," we mean that either all stations in the network have an infinite number of servers, or all have a finite number of servers). The mean number-in-system function f maps S into $R^n_\oplus$. The main results of this paper are natural sufficient conditions guaranteeing that the function f covers $R^n_\oplus$; that is, every nonnegative n-vector is realizable as the mean number-in-system vector $f(\lambda)$ under the arrival process corresponding to a stable arrival-rate vector λ.

3. Covering theorem and queueing interpretation. In the current section, we state a covering result that is tailor-made for our queueing system, viewed as a blackbox of servers; we interpret the conditions for the result in queueing terminology. The covering result is formally developed in the next section.

Let S be a subset of $R^n_\oplus$. We say that S is *open* in $R^n_\oplus$ if S is open in the relative topology of $R^n_\oplus$; of course, openness in $R^n_\oplus$ does not imply openness in R^n. We say that S is *star-shaped at the origin* if for any x in S we have that γx is in S for every $0 \le \gamma \le 1$. A sequence of vectors in S is defined to *fade from* S if all its cluster points, if any, are outside S; such a sequence is called a *fading sequence* in S. We note that a sequence of vectors in S fades from S if and only if for every compact subset C of S there exists an integer k such that $x_h \in S\backslash C$ for all $h \ge k$. A function $f: S \to R^n_\oplus$ is called *coercive on* S if $f(x_k) \to \infty$ for every fading sequence $x_1, x_2, \ldots$ in S. If S is star-shaped at the origin, a function $f: S \to R^n_\oplus$ is called *ray-coercive on* S if for every $x \in S\backslash\{0\}$, $f(\theta_k x) \to \infty$ for every sequence $\theta_1, \theta_2, \ldots$ of nonnegative numbers with $\theta_k \uparrow \sup\{\theta \in R_\oplus : \theta x \in S\}$.

The proof of Lemma 1 is left to the reader.

Lemma 1. *Let S be a subset of $R^n_\oplus$ and let $f: S \to R^n_\oplus$ be a continuous function. Then, f is coercive on S if and only if for every $a \in R_\oplus$ the set $\{x \in S : \|f(x)\| \le a\}$ is compact.*

Lemma 2. *Let S be a star-shaped subset of $R^n_\oplus$ that is open in $R^n_\oplus$ and let $f: S \to R^n_\oplus$. If f is coercive on S, then f is ray-coercive on S.*

Proof. Let $x \in S\backslash\{0\}$, $\theta^* \equiv \sup\{\gamma \in R_\oplus : \gamma x \in S\}$, and let $\theta_1, \theta_2, \ldots$ be a sequence of nonnegative numbers with $\theta_k \uparrow \theta^*$. The sequence $\theta_1 x, \theta_2 x, \ldots$ has a cluster point only if θ^* is finite, and in this case the only cluster point is $\theta^* x$. In this case, $\theta^* x$ is in the closed set $R^n_\oplus\backslash S$ as it is the limit of a sequence of elements in $R^n_\oplus\backslash S$ (say the sequence $y_k = (\theta^* + 1/k)x$). Therefore, no cluster point of the sequence $\theta_1 x, \theta_2 x, \ldots$ is in S; that is, the sequence fades from S. It now follows from the coerciveness of f on S that $f(\theta_k x) \to \infty$. □

The inverse of Lemma 2 is false—ray-coerciveness does not imply coerciveness.

The following covering theorem is central to our queueing results; the proof is provided in §4.

Theorem 1 (Covering Theorem). *Let S be a nonempty subset of $R^n_\oplus$ and let $f: S \to R^n_\oplus$. If*

(a) *S is open in $R^n_\oplus$,*

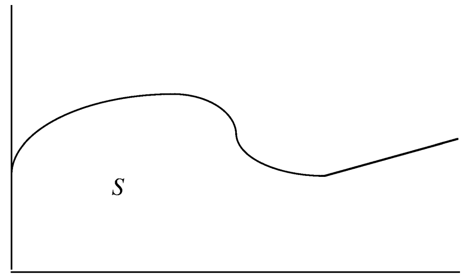

FIGURE 1. A set S satisfying the assumptions of the covering theorem.

(b) *S is star-shaped at the origin,*
(c) *f is continuous,*
(d) *f is coercive on S, and*
(e) *$x \in S$ and $x_i = 0$ implies $f(x)_i = 0$ for i in ν,*
then the range of f covers $R^n_\oplus$.

A possible set $S \subseteq R^2_\oplus$ satisfying the assumptions of the covering theorem is illustrated in Figure 1.

We interpret the assumptions of the covering theorem in a queueing context. The set S represents the stable arrival-rate vectors, and $f(\lambda)$ represents the mean number-in-system vector corresponding to the stable arrival-rate vector λ.

S is open in $R^n_\oplus$. This assumption means that if the queueing system is stable for an arrival-rate vector λ, then it remains stable for all sufficiently small changes in the vector.

S is star-shaped at the origin. This assumption means that if the queueing system is stable for arrival-rate vector λ, then the system remains stable for arrival-rate vectors which are reduced proportionately.

f is continuous. This assumption means that small changes in the arrival-rate vector λ yield small changes in the mean number-in-system vector.

f is coercive on S. This assumption means that the mean number-in-system vector blows up for a fading sequence of stable arrival-rate vectors.

$\lambda_i = 0$ implies $f(\lambda)_i = 0$. This assumption means that if no customers of a class i arrive, then there will be no customers of that class in the system.

For a "normal" queueing system, f could be expected to meet the conditions of the covering theorem. Nonnormal systems would have properties like: no steady state (in which case f is not well defined), discontinuous behavior, bounded system loads with ever-increasing arrival rates, or customer types that are created within the system (for example, after a certain sequence of station transitions). We also note that the covering theorem does not apply to situations in which the components of the number-in-system vector are the mean number at each of the stations of the network. If λ is the vector of external arrival rates to the network, then (c) would generally be violated, for example. Of course, one can easily construct "nonnormal" systems and, further, even embrace their value.

As special cases of the covering theorem, we have the following three corollaries. In the first corollary, $S = R^n_\oplus$. This corollary is designed to deal with queueing networks in which each station has infinite capacity, so that all arrival-rate vectors are typically stable.

COROLLARY 1. *Let $f: R^n_\oplus \to R^n_\oplus$. If*
(a) *f is continuous,*
(b) *$x_k \to \infty$ implies $f(x_k) \to \infty$, and*
(c) *$x \in R^n_\oplus$ and $x_i = 0$ for some i in ν implies $f(x)_i = 0$,*
then the range of f is $R^n_\oplus$.

A second corollary addresses networks in which each station has finite capacity.

Corollary 2. *Let S be a nonempty bounded set in $R^n_\oplus$ and let $f: S \to R^n_\oplus$. If*
(a) *S is open in $R^n_\oplus$,*
(b) *S is star-shaped at the origin,*
(c) *f is continuous,*
(d) *$x_1, x_2, \ldots$ is a sequence of vectors in S converging to a point in $R^n_\oplus \backslash S$, then $f(x_k) \to \infty$, and*
(e) *$x \in S$ and $x_i = 0$ for some i in ν implies $f(x)_i = 0$,*
then the range of f is $R^n_\oplus$.

The third corollary was originally proved in Eaves and Rothblum (1993) and was motivated by the work of Mehrotra (1992).

Corollary 3. *Let U be a nonnegative $n \times p$ matrix with no zero row or column, let $S \equiv \{x \in R^n_\oplus: (x^T U)_j < 1$ for $j = 1, \ldots, p\}$, and let $f: S \to R^n_\oplus$. If*
(a) *f is continuous,*
(b) *$x_1, x_2, \ldots$ is a sequence of vectors in S with $\|x_k^T U\| \to 1$, then $f(x_k) \to \infty$, and*
(c) *$x \in S$ and $x_i = 0$ for some i in ν implies $f(x)_i = 0$,*
then the range of f is $R^n_\oplus$.

Here we can interpret x_i as the rate at which customers of class i enter the network. Customers of class i require, on average, U_{ij} units of processing time effort at station j. Each station is assumed to serve customers at unit rate. Then, we can expect, in great generality, that S (as described above) is the stability region for the network, and that the mean number-in-system function f satisfies the corollary's requirements.

4. Proof of the covering theorem. Our proof of the covering theorem is based on piecewise linear path following, which requires some definitions about triangulations and piecewise linear maps.

For a collection of sets T in a Euclidean space and a subset S of that space, we denote by $T|_S$ the subcollection of sets in T which meet S. The collection T is called *locally finite* if for every point in the underlying Euclidean space, there is a neighborhood N of x with $T|_N$ finite. A *triangulation* of $R^n_\oplus \times R_\oplus$ is a collection T of simplices in $R^n_\oplus \times R_\oplus$ such that

(i) the collection covers $R^n_\oplus \times R_\oplus$,
(ii) the collection is locally finite,
(iii) any two simplices in T are either disjoint or meet in a common face, and
(iv) faces of elements of T are in T.

A function F on $R^n_\oplus \times R_\oplus$ into any Euclidean space is called *piecewise linear*, if for some triangulation T of $R^n_\oplus \times R_\oplus$, F is linear on each $(n+1)$-dimensional cell of T; in this case we also say that F is *T-piecewise linear*. A *piecewise linear path* in $R^n_\oplus$ is a piecewise linear map from $R_\oplus$ to $R^n_\oplus$.

The *diameter* of a set σ in a Euclidean space is the supremum of $\|x - y\|$ as x and y range over σ. For a collection of sets in a Euclidean space, the *grid size* is the supremum of the diameters of its sets. The *projected diameter* of a subset σ of $R^n_\oplus \times R_\oplus$ is the diameter of the projected set $\pi(\sigma) \equiv \{x: (x, t) \in \sigma\} \subseteq R^n_\oplus$. For a collection of subsets of $R^n_\oplus \times R_\oplus$, the *projected grid* is the supremum over the projected diameters of its elements. A triangulation T of $R^n_\oplus \times R_\oplus$ is called *refining*, if it has a finite grid size and if the projected grid size of $T|_{R^n_\oplus \times [k, +\infty)}$ tends to zero as k tends to $+\infty$. For the existence of such a triangulation, see Eaves (1984).

The next result is needed for our development. It determines sufficient conditions for the inverse image of a piecewise linear function from $R^n_\oplus \times R_\oplus$ into $R^n_\oplus$ to contain a fading piecewise linear path. While the arguments we use rely on now-standard methodology, the result itself is new and of independent interest.

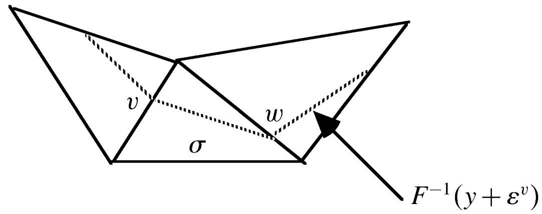

FIGURE 2. The intersection of $F^{-1}(y+\varepsilon^{\nu})$ with an $(n+1)$-cell σ of the subdivision.

THEOREM 2. *Let*
(a) *T be a refining triangulation of $R^n_{\oplus} \times R_{\oplus}$,*
(b) *F_0 be a function mapping the vertices of T into $R^n_{\oplus}$ such that for vertex (v, r) of T, $F_0(v, r) = v$ if $r = 0$ and $[F_0(v, r)]_i = 0$ if $v_i = 0$ and $r > 0$, and*
(c) *F be a T-piecewise linear extension of F_0 to $R^n_{\oplus} \times R_{\oplus}$.*
Then, for every $y \in R^n_{\oplus}$, there exists a piecewise linear path $p = (p_1, p_2)\colon R_{\oplus} \to F^{-1}(y) \subseteq R^n_{\oplus} \times R_{\oplus}$ with $p(0) = (y, 0)$ and $p(t) \to \infty$.

PROOF. Let $\varepsilon^{\nu} \equiv (\varepsilon^1, \ldots, \varepsilon^n)$ be the n-vector of powers of ε where ε is regarded as a positive infinitesimal relative to the real field. The target point y is replaced by the perturbed positive point $y + \varepsilon^{\nu}$. It can be shown that $F^{-1}(y + \varepsilon^{\nu})$ does not meet any simplex of dimension less than n in T (more generally, no simplex of dimension less than n with vertices in R^{n+1}). The intersection of $F^{-1}(y+\varepsilon^{\nu})$ with any $(n+1)$-simplex of σ of T is a line segment $[v, w]$ where the components of v and w are polynomials in the infinitesimal ε. The endpoints v and w lie in the relative interiors of two distinct n-faces of σ, and v and w lie in at most two $(n+1)$-simplices of T (they lie in one only if they are in the boundary of $R^n_{\oplus} \times R_{\oplus}$); see Figure 2.

With $\sigma(1)$ as the unique $(n+1)$-dimensional simplex of T containing $(y+\varepsilon^{\nu}, 0)$, we can follow line segments to build a path p^{ε}. The elements of $p^{\varepsilon}(t)$ for each member $t \in R_{\oplus}$ will be polynomials of degree, at most, n in ε. There are no forks in the path. The path cannot return to the facet $R^n_{\oplus} \times \{0\}$, as the only point of $F^{-1}(y+\varepsilon^{\nu})$ in it is $(y+\varepsilon^{\nu}, 0)$. The path cannot meet any facet $\{(x, t) \in R^n_{\oplus} \times R_{\oplus}\colon x_i = 0\}$ for $i = 1, \ldots, n$, as all points (x, t) on such a facet $F(x, t)_i = 0 \neq (y+\varepsilon^{\nu})_i$. It follows that for the constructed path, $p^{\varepsilon}(t) \to \infty$. Now, as the vertices of this path are mapped into polynomials in the infinitesimal, one can replace ε with 0 to obtain the desired path p^0 mapping $R_{\oplus}$ into $F^{-1}(y)$ with $p^0(t) \to \infty$. The path p^0 may have loops and its range may intersect the vertical facets of $R^n_{\oplus} \times R_{\oplus}$; nevertheless, $p(t) = p^0(t) \to \infty$. □

The treatment of the infinitesimal ε in the above proof is informal. A formal argument requires consideration of an extension of the reals to the ordered field of rational functions in the infinitesimal ε; see Eaves and Rothblum (1994).

We are now ready to prove the covering theorem.

PROOF OF THEOREM 1. Fix y in $R^n_{\oplus}$, and we will construct a vector x in S with $f(x) = y$. As $f(0) = 0$, we can assume that $y \neq 0$. Define

$$\delta \equiv \inf\{\|x\|\colon x \in S,\ \|f(x)\| \geq 3n\|y\|\}.$$

As f is coercive on S and S is open in $R^n_{\oplus}$, the set $\{x \in S\colon \|f(x)\| \geq 3n\|y\|\}$ is clearly nonempty and δ is finite. Furthermore, as $f(0) = 0$ and f is continuous, δ is positive. In summary, $0 < \delta < +\infty$.

Select $\theta > 3n\delta^{-1}\|y\|$. Let $\tilde{S} \equiv \theta S = \{\theta x\colon x \in S\}$ and $\tilde{f}\colon \tilde{S} \to R^n_{\oplus}$ with $\tilde{f}(\theta x) = f(x)$. The conditions of the theorem continue to hold when (S, f) is replaced by $(\tilde{S}, \tilde{f})$. Furthermore,

$$\begin{aligned}\tilde{\delta} &\equiv \inf\{\|z\|\colon z \in \tilde{S},\ \|\tilde{f}(z)\| \geq 3n\|y\|\} = \inf\{\|\theta x\|\colon \theta x \in \tilde{S},\ \|\tilde{f}(\theta x)\| \geq 3n\|y\|\}\\ &= \theta[\inf\{\|x\|\colon x \in S,\ \|f(x)\| \geq 3n\|y\|\}] = \theta\delta > 3n\|y\|.\end{aligned}$$

As $\tilde{f}(z) = y$ if and only if $f(\theta^{-1}z) = y$, we can obtain a solution x for $f(x) = y$ from any solution z of $\tilde{f}(z) = y$. Thus, without loss of generality, we can and do assume that $\delta > 3n\|y\|$; that is, $\|f(x)\| \geq 3n\|y\|$ implies $\|x\| \geq \delta > 3n\|y\|$.

Let B be the (bounded) ball $\{x \in R^n_\oplus \colon \|x\| \leq 3n\|y\|\}$ and let C be the preimage of B under f; that is, $C = \{x \in S \colon \|f(x)\| \leq 3n\|y\|\}$. As f is coercive on S, the set C is compact (Lemma 1). We first argue that $S \cap B \subseteq C$. Indeed, if $x \in S \cap B$, then $x \in S$ and $\|x\| \leq 3n\|y\| < \delta$; the definition of δ then implies that $\|f(x)\| < 3n\|y\|$ and so $x \in C$. Our next claim is that $B \subseteq S$. Indeed, let x be a vector in $R^n_\oplus \backslash S$, and we will show that $x \notin B$. As S is nonempty, star-shaped at 0, and open in $R^n_\oplus$, and as f is ray-coercive on S (by Lemma 2), we have that $x \neq 0$ and for some $0 < \gamma < 1$, $\gamma x \in S$ and $\|f(\gamma x)\| > 3n\|y\|$. Thus, $\gamma x \notin C$ and as $S \cap B \subseteq C$, $\gamma x \notin B$. As B is star-shaped at the origin, we conclude that $x \notin B$. As $S \cap B \subseteq C$, $B \subseteq S$, and $C \subseteq S$, we have $B \subseteq C \subseteq S$.

Let T be a refining triangulation of $R^n_\oplus \times R_\oplus$. Now on vertices (v, s) of T, let

$$F_0(v, r) \equiv \begin{cases} v & \text{if } r = 0, \\ v & \text{if } r > 0 \text{ and } v \notin S, \text{ and} \\ f(v) & \text{if } r > 0 \text{ and } v \in S, \end{cases}$$

and extend F_0 linearly on each of the $(n+1)$-dimensional simplices of T to obtain a piecewise linear function $F \colon R^n_\oplus \times R_\oplus \to R^n_\oplus$; in particular, for $i = 1, \ldots, n$, $[F(v, r)]_i = 0$ whenever $r_i = 0$. It now follows from Theorem 2 that there is a piecewise linear path $p = (p_1, p_2) \colon R_\oplus \to F^{-1}(y) \subseteq R^n_\oplus \times R_\oplus$ with $p(0) = (y, 0)$ and $p(t) \to \infty$ as $t \to \infty$.

Let (v, r) be a vertex of T with $v \in R^n_\oplus \backslash C = (S \backslash C) \cup (R^n_\oplus \backslash S)$. If $v \in S \backslash C$, then $v \in S$ with $\|f(v)\| > 3n\|y\|$ and the assertion $\delta > 3n\|y\|$ assures that $\|v\| > 3n\|y\|$; alternatively, if $v \notin R^n_\oplus \backslash S$, then the inclusion $B \subseteq S$ assures $v \notin B$, namely, $\|v\| > 3n\|y\|$. In either case, we have from the definition of F_0 that $\|F_0(v, r)\| > 3n\|y\|$. With this observation we show that if σ is an $(n+1)$-dimensional simplex in T containing $p(t)$ for some $t \in R_\oplus$, then the projection of some vertex of σ lies in C. Indeed, suppose this assertion is false. Let (v_i, r_i) for $i = 0, \ldots, n+1$ with $v_i \in R^n_\oplus \backslash C$ be the vertices of σ. For some nonnegative coefficients $\theta_0, \ldots, \theta_{n+1}$ that sum to 1, $p(t) = \sum_i \theta_i (v_i, r_i)$ and $y = F[p(t)] = \sum_i \theta_i F_0(v_i, r_i)$. Obviously, one of the θ_is, say θ_j, is larger than or equal to $(n+2)^{-1}$. As $v_j \in R^n_\oplus \backslash C$, we have $\|F_0(v_j, r_j)\| > 3n\|y\|$. From the nonnegativity of the $F_0(v_i, r_i)$s we conclude that $\|y\| \geq \theta_j \|F_0(v_j, r_j)\| > (n+2)^{-1} 3n\|y\| > \|y\|$. We have a contradiction; hence, some vertex (v_i, r_i) has v_i in C.

As T has a finite grid and each $p(t)$ lies in a simplex of T with a projected vertex in the compact set $C \subseteq S$, we have that the collection $\{p_1(t) \colon t \in R_\oplus\}$ is bounded. As $p(t) \to \infty$ and $U \equiv \{p_1(t) \colon t \in R_\oplus\}$ is bounded, we conclude that $p_2(t) \to +\infty$ as $t \to \infty$. Let x be any cluster point of the bounded sequence $p_1(1), p_1(2), \ldots$, in particular, for some positive distinct integers $t_1, t_2, \ldots$, $\lim_{q \to +\infty} p_1(t_q) = x$; from $p_2(t) \to +\infty$ as $t \to \infty$, we have that $p_2(t_q) \to +\infty$ as $q \to +\infty$. We will show that $f(x) = y$.

As $C \subseteq S$, C is compact and $R^n_\oplus \backslash S$ is closed, the distance $\min\{\|c - d\| \colon c \in C,\ d \in R^n_\oplus \backslash S\}$ is positive, hence, for some positive u, the set $C_u = \{d \in R^n_\oplus \colon \|d - c\| \leq u\}$ is contained in S. An earlier argument assures that for each q, any $(n+1)$-dimensional simplex in T containing $p(t_q)$ has a vertex in C. As T is refining and $p_2(t_q) \to \infty$ as $q \to \infty$, for sufficiently large q the projected diameter of such a simplex is less than u; as one of the projected vertices of the simplex must be in C, it follows that all the projected vertices are in C_u. Therefore, for sufficiently large q, $p(t_q)$ is contained in an $(n+1)$-dimensional simplex of T, all of whose projected vertices are in $C_u \subseteq S$. As f is uniformly continuous on the compact set C_u and $\lim_{q \to +\infty} p_1(t_q) = x$, it follows from the definition of F that $F[p(t_q)] - f(x) \to 0$ as $q \to +\infty$. As $F[p(t_q)] = y$ for each q, we conclude that $y = f(x)$. $\square$

5. Verifying the covering conditions. Our plan is to apply the covering theorem to show that the mean number-in-system vector function is onto in a network of queues. For that purpose, verification of the conditions of the covering theorem is required when the blackbox system is in fact a network of servers, as described in §§1 and 2. In applying the covering theorem to general queueing networks, the principal difficulty is that the steady-state distribution is typically analytically intractable, and in particular, f is not known in closed form. As a consequence, one must rely on "indirect" methods to verify the hypotheses of the covering theorem. The standard way to proceed is to first make the queueing network Markov by adding appropriate supplementary variables to the state description. One can then apply the rich set of tools associated with the theory of Markov processes to indirectly verify the hypotheses of the covering theorem.

In particular, let $X = \{X_i : i \geq 0\}$ be the Markov chain corresponding to the queueing network. (The random variables X_i may, for example, correspond to the state descriptor of the system at the ith transition epoch of the number-in-system process.) By applying the Lyapunov criterion associated with Theorem 11.3.4 of Meyn and Tweedie (1993), one can potentially establish that a given set W is a subset of the set S_0 of weakly stable arrival-rate vectors by proving that for each $\lambda \in W$, the associated Markov chain X is positive Harris recurrent (see Meyn and Tweedie 1993 for the definition of positive Harris recurrence—this form of recurrence implies the existence of a steady state and much more). Further, Lyapunov criteria exist that permit one to establish that for $\lambda \notin W$, the corresponding Markov chain is either null recurrent or transient (see Theorems 8.4.2, 9.5.6, and 11.5.1 of Meyn and Tweedie 1993 for such results). Thus, the current state-of-the-art in Lyapunov methodology provides the tools necessary to identify S_0 with some known subset W. Certain refined Lyapunov methods (see, for example, p. 337 of Meyn and Tweedie 1993) permit one to verify that whenever $\lambda \in W = S_0$, then λ is strongly stable. This, in turn, permits one to conclude that W equals the set S of stable arrival-rate vectors. Hence, one can reasonably expect to be able to compute the region S for a variety of complex queueing systems, including those with nonexponential interarrival times and processing times. A great deal of progress has been made recently in following the above program; see, for example, Sigman (1990), Dai (1995), and Kumar and Meyn (1996). In any case, with an explicit description of S, the assertion that it is open in $R^n_\oplus$ and star-shaped at the origin (conditions (a) and (b) of Theorem 1) will typically be readily verifiable.

To prove continuity of f (condition (c) of Theorem 1), we note that general machinery for establishing the stronger assertion of differentiability of f in the Markov process setting has been developed by Glynn and L'Ecuyer (1995); the key condition again involves constructing an appropriate Lyapunov function. Next, condition (e) of Theorem 1 will usually be trivial to verify in the generalized Kelly-type networks (Kelly 1979) that we have in mind.

We are left with establishing the coerciveness of f (condition (d) of Theorem 1). To show that the mean number-in-system blows up as the sequence fades in the set of stable arrival-rate vectors S, one approach involves a proof by contradiction. The idea is to prove that if the mean number-in-system does not blow up, then S_0 contains points on the boundary of S. For this purpose, we can use Theorem 3. However, we argued earlier that we can typically expect to be able to conclude that $S_0 = S$. The fact that S contains points on its boundary and is assumed to be open then yields the required contradiction.

To precisely state Theorem 3, let P_n for $1 \leq n \leq +\infty$ be a transition kernel on a complete separable metric space Γ. The function $V: \Gamma \to R_\oplus$ is said to be *normlike* if the sublevel sets $\{x \in \Gamma : V(x) \leq r\}$ are precompact for each $r > 0$.

Theorem 3. *Suppose that*

(a) *for each bounded continuous function* $g: \Gamma \to R$,

$$\int_\Gamma g(y) P_n(x_n, dy) \to \int_\Gamma g(y) P_{+\infty}(x, dy) \quad \textit{whenever } x_n \to x \in \Gamma,$$

(b) *for* $1 \le n < +\infty$, *there exists a probability measure* π_n *that is a stationary measure of* P_n, *and*
(c) *there exists a normlike function* V *such that* $\sup_{n\ge 1} \int_\Gamma \pi_n(dx)V(x) < +\infty$.
Then,
(i) $P_{+\infty}$ *has a stationary probability measure* $\pi_{+\infty}$, *and*
(ii) $\int_\Gamma \pi_{+\infty}(dx)V(x) < +\infty$.

Proof. We start by noting that (c) forces $\{\pi_n\colon 1 \le n < +\infty\}$ to be a tight sequence of probability measures; see Meyn and Tweedie (1993, p. 522). Hence, by Prohorov's theorem, $\{\pi_n\colon 1 \le n < +\infty\}$ is relatively compact in the topology of weak convergence of probability measures defined on Γ. Consequently, for every subsequence $\{\pi_{n_k}\colon k \ge 1\}$, we can extract a further subsequence $\{\pi_{n'_k}\colon k \ge 1\}$ for which there exists a probability measure $\pi_{+\infty}$ such that

$$\pi_{n'_k} \Rightarrow \pi_{+\infty} \quad \text{as } k \to +\infty, \tag{1}$$

where $\Rightarrow$ denotes weak convergence. Next we prove that the limit measure $\pi_{+\infty}$ is, in fact, a stationary distribution of $P_{+\infty}$, yielding part (i) of the theorem.

For $1 \le n \le +\infty$, let $\tilde{E}_n(\cdot)$ be the expectation operator defined on the path space of X under which the chain has transition kernel P_n and is initiated with distribution π_n. Then, for each bounded continuous function g, stationarity implies that for $1 \le n \le +\infty$, we have

$$\tilde{E}_n[g(X_0)] = \tilde{E}_n[g(X_1)] = \tilde{E}_n[\tilde{E}_n[g(X_1) \mid X_0]] = \tilde{E}_n[h_n(X_0)], \tag{2}$$

where $h_n(x) = \int_\Gamma P_n(x, dy)g(y)$. Now, the Skorohod representation theorem guarantees the existence of a probability space (Ω^*, P^*, F) supporting a sequence of random variables $\{X_n^*\colon 1 \le n \le +\infty\}$ such that $\tilde{P}_n(X_0 \in \cdot) = P^*(X_n^* \in \cdot)$ with $X_{n'_k}^* \to X_{+\infty}^*$ almost surely as $k \to +\infty$. Note that (a) then ensures that $h_{n'_k}(X_{n'_k}^*) \to h_{+\infty}(X_{+\infty}^*)$ almost surely as $k \to +\infty$. Hence, the bounded convergence theorem then yields the convergence relation

$$\tilde{E}_{n'_k} h_{n'_k}(X_0) = E^* h_{n'_k}(X_{n'_k}^*) \to E^* h_{+\infty}(X_{+\infty}^*) = \tilde{E}_{+\infty} h_{+\infty}(X_0) \quad \text{as } k \to +\infty. \tag{3}$$

On the other hand, by the definition of weak convergence, (1) implies that

$$\tilde{E}_{n'_k} g(X_0) \to \tilde{E}_{+\infty} g(X_0) \quad \text{as } k \to +\infty. \tag{4}$$

Combining (2), (3), and (4), we then get the equalities

$$\int_\Gamma \pi_{+\infty}(dx)g(x) = \tilde{E}_{+\infty} g(X_0) = \tilde{E}_{+\infty} h_{+\infty}(X_0) = \int_\Gamma \pi_{+\infty}(dx) \int_\Gamma P_{+\infty}(x, dy)g(y). \tag{5}$$

Because (5) holds for all bounded continuous g, this implies that $\pi_{+\infty}$ is stationary for $P_{+\infty}$.
To prove (ii), note that Fatou's lemma implies that (c) yields

$$E^* V(X_{+\infty}^*) \le \liminf_{k\to+\infty} E^* V(X_{n'_k}^*) = \liminf_{k\to+\infty} \tilde{E}_{n'_k} V(X_0) \le \sup_{k\ge 1} \int_\Gamma \pi_{n'_k}(dx)V(x) < +\infty. \quad \square$$

In Theorem 3, if $P_{+\infty}$ is known also to be a Ψ-irreducible transition kernel (as would typically occur in most queueing settings), then existence of a stationary probability $\pi_{+\infty}$ guarantees that the Markov chain X associated with $P_{+\infty}$ will be positive Harris recurrent on some subset of its state space Γ (thereby establishing that $P_{+\infty}$ corresponds to a point in S_0); see Theorems 10.1.1 and 9.0.1 of Meyn and Tweedie (1993) for details.

One potential difficulty in applying Theorem 3 to queueing networks is the verification of assumption (a) of Theorem 3. To obtain a Markov representation of such a system, one typically adds to the number-in-system process state variables that record the residual time, for each event possible, to the next such event. Unfortunately, the continuity-type hypothesis (a) is frequently violated at points $x \in \Gamma$ in which at least two events (say, an arrival and departure) have the same residual "lifetime" readings. However, one can then use the following relaxed version of (a). Suppose that $P_{+\infty}$ is a Ψ-irreducible transition kernel with maximal irreducibility measure μ (see Meyn and Tweedie 1993, pp. 86–88 for the definitions). Then we can replace (a) with

(a′) *For each bounded continuous function* $g: \Gamma \to R$, $\mu(C(g)^c) = 0$, *where* $C(g) = \{x \in \Gamma: \int_\Gamma g(y) P_n(x_n, dy) \to \int_\Gamma g(y) P_{+\infty}(x, dy)$ *whenever* $x_n \to x\}$.

To verify (a′), one needs to compute μ; such methods are available in Henderson and Glynn (1998). The proof of Theorem 3 with (a) replaced by (a′) follows very similar lines and is therefore omitted.

Acknowledgments. This research was supported in part by the National Science Foundation under Grant DMS 95-02892, by the Army Research Office under Contract DAAL-04-94-G-0021, and by the Office of Naval Research under Grant N00014-92-J1142.

References

Dai, J. 1995. On positive Harris recurrence for multiclass queueing networks: A unified approach via fluid limit models. *Ann. Appl. Probab.* **5** 49–77.

Eaves, B. C. 1976. A short course in solving equations with PL homotopies. *SIAM-AMS Proc.* **IX** 73–143.

Eaves, B. C. 1984. *A Course in Triangulations for Solving Equations with Deformations. Lecture Notes in Economics and Mathematical Systems*, No. 234. Springer Verlag, New York.

Eaves, B. C., U. G. Rothblum. 1993. A class of "onto" multifunctions. *Math. Programming* **61** 327–343.

Eaves, B. C., U. G. Rothblum. 1994. Formulation of linear problems and solution by a universal machine. *Math. Programming* **65** 263–310.

Glynn, P. W., P. L'Ecuyer. 1995. Likelihood ratio gradient estimation for stochastic recursions. *Adv. Appl. Probab.* **27** 1019–1053.

Henderson, S., P. W. Glynn. 2001. Regenerative steady-state simulation of discrete-event systems. *ACM Trans. Model. Comput. Simulation* **11** 313–345.

Kelly, F. P. 1979. *Reversibility and Stochastic Networks.* John Wiley and Sons, New York.

Kumar, P. R., S. P. Meyn. 1996. Duality and linear programs for stability and performance analysis of queueing networks and scheduling policies. *IEEE Trans. Automatic Control* **41** 4–17.

Mehrotra, V. 1992. An approximation procedure for general closed multiclass queueing networks with deterministic routing. Ph.D. dissertation, Department of Operations Research, Stanford University, Stanford, CA.

Meyn, S. P., R. L. Tweedie. 1993. *Markov Chains and Stochastic Stability.* Springer Verlag, New York.

Sigman, K. 1990. The stability of open queueing networks. *Stochastic Processes Their Appl.* **35** 11–25.

Walrand, J. 1988. *An Introduction to Queueing Networks*. Prentice Hall, Englewood Cliffs, NJ.