THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 10, December 1983
Printed in U.S.A.

Equivalent Queueing Networks and Their Use in Approximate Equilibrium Analysis

By A. KUMAR*

(Manuscript received March 14, 1983)

Most Markovian queueing networks that arise as models of stochastic congestion systems (e.g., communication networks and multiprogrammed computer systems) do not have a product form in their stationary probability distributions, and hence are not amenable to the simplicity of product-form analysis. In this paper we suggest an approach for systematically examining the validity of a class of approximation schemes that is based on the idea of equivalent networks and is used for the approximate equilibrium analysis of nonproduct-form networks. We study equivalent networks, and prove a generalization of the so-called "Norton's" Theorem for closed product-form networks in order to study and generalize the equivalent flow method for the approximate analysis of nonproduct-form queueing networks. We then present the results of a study of the approximation scheme as applied to a type of network model (called a central-server model) that arises frequently in modeling multiprogrammed computer systems. In this model the central server uses a priority discipline, so the resulting network is nonproduct form. This study demonstrates the situations under which the approximation can be expected to do well or poorly and the kinds of errors it introduces.

I. INTRODUCTION

Mathematical modeling of stochastic systems frequently gives rise to models in a class referred to as Markovian queueing networks—specifically, queueing networks whose time evolution can be described by a discrete-state, regular Markov stochastic process. Markovian

*Bell Laboratories.

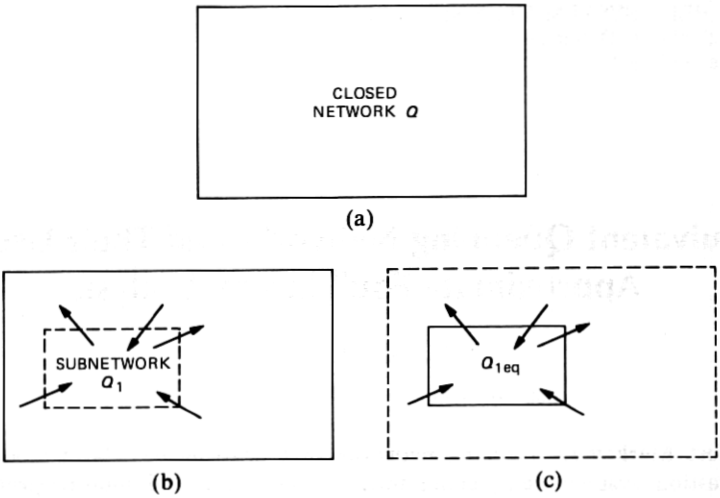

Fig. 1—Notion of an equivalent network; (a) original network, (b) with arrows indicating flows between Q_1 and its complements, and (c) with arrows indicating models of flows between Q_1 and its complement.

queueing network models, known as product-form networks, have been widely studied, owing primarily to their well-understood stochastic behavior, and the simplicity of their analysis in equilibrium. However, the class of product-form queueing network models is far from adequate for modeling many simple real-world congestion systems. The exact equilibrium analysis of nonproduct-form queueing networks is, in most cases, computationally, and often fundamentally, intractable. Much effort has, therefore, been directed towards devising approximation schemes that attempt to reconcile the conflicting requirements of modeling fidelity and the simplicity of product-form analysis. One such class of approximation schemes is based on the idea of equivalent networks. In this paper we systematically study this approximation.

By an equivalent network we mean the following (cf. Fig. 1). Consider a closed queueing network Q constructed from the set of nodes M and the subnetwork Q_1 consisting of nodes $M_1 \subset M$. Let Q_{1eq} be a network constructed from M_1 such that the joint equilibrium (probability) distribution of Q_{1eq} is the same as the marginal joint equilibrium distribution of Q_1 in Q. The network Q_{1eq} is then said to be equivalent to Q_1.* Clearly, to study Q_1 in isolation, one needs to account for the

* This notion of equivalence may appear unduly restrictive. Why not establish a more detailed stochastic equivalence? For the calculation of many performance analysis criteria, the present notion is adequate. However, it is easy to see that the equivalent networks we later identify yield equality in distribution for the entire process in equilibrium.

influence of the nodes in $M - M_1$ on the nodes in Q_1. When Q is product form, the influence of the complementary network on Q_1 takes an especially simple form, and can be determined by analyzing a modified version of the complementary network in isolation! For the case where M_1 consists of a single node, this fact was first recognized by Chandy et al.,[1] who called the equivalent network so obtained a "Norton" equivalent, because of the similarity of this equivalence to Norton equivalence in electrical circuits.

In Section II we study equivalent networks and demonstrate the simplifications that arise for product-form networks. The development yields a generalization of Norton's Theorem to multinode subnetworks of closed product-form networks. Essentially, the same extension to the entire class of closed product-form networks has been obtained independently and concurrently by Kritzinger et al.[2] and Balsamo et al.,[3] through an approach based on verification via detailed computations from the product-form solution. Our approach is substantially different, in that it derives Norton's Theorem directly as a special case of a general result for stochastically equivalent networks. This approach is concise, conceptually and intuitively appealing, gives the result a probabilistic interpretation, and shows up clearly the role played by the product-form solution. It also seems to be the natural approach for the purposes of this study.

This generalization of Norton's Theorem motivates the following approximation scheme. Suppose now that Q is a nonproduct-form network, but for the purposes of studying the subnetwork Q_1, we follow the equivalence procedure for product-form networks. Suppose also that in doing so we find that the version of the complementary network we have to analyze, in order to determine the latter's influence on Q_1, is product form. Let the equivalent network thus obtained be $\hat{Q}_{1eq}$. The approximation scheme, referred to above, approximates the equilibrium distribution of Q_1 with that of $\hat{Q}_{1eq}$ (i.e., approximates Q_{1eq} by $\hat{Q}_{1eq}$). The effort to determine and analyze $\hat{Q}_{1eq}$ will, in general, be considerably less than the effort to exactly analyze Q_1 in Q.

This approximation scheme is an extension of one (often referred to in the literature as an equivalent flow approximation) that has been utilized by several workers, in the field of network performance analysis, with remarkably accurate results. Sauer and Chandy,[4] and Chow and Yu[5] use this idea as the basic step in iterative schemes for approximating central-server models in which the central server is not of product-form type. Schwartz[6] uses the basic scheme directly to approximately analyze a model for a multiple-access communication system. In Section III we draw upon the theoretical development in Section II to study the validity of the approximation scheme when it is applied to a simple test-bed model.

II. EQUIVALENT NETWORKS

Consider a closed Markovian queueing network Q consisting of M congestion nodes. In this section we study the problem of the equilibrium analysis of a subnetwork Q_1 (embedded in Q). To simplify the discussion we shall limit our considerations to networks of First In, First Out (FIFO) nodes. It is easily recognized that the ideas in this section can be extended to apply to more general networks. In Section 2.3 we establish Theorem 1, which explicates the structure of equivalent subnetworks of the networks described in Section 2.1. By combining this result with Theorem 2, we get a generalization of Norton's Theorem.

2.1 Network specifications

Q is a closed queueing network consisting of M FIFO nodes (indexed by $i \in \{1, \cdots, M\}$). There are R classes/types of customers (indexed by $r \in \mathbf{R} = \{1, \cdots, R\}$) with N_r customers in the rth class. The M x M matrix $P^{(r)} = [p_{ij}^{(r)}]$ is the routing probability matrix of type r customers; customers do not change class as they move from node to node. For each r in $\{1, \cdots, R\}$, $P^{(r)}$ is a stochastic matrix which, when considered as a transition probability matrix, leads to a Markov chain, on the state space $\{1, \cdots, M\}$, with a single positive, communicating class.

Throughout the following discussion, the network state process is assumed to be in equilibrium. The state of the ith node (denoted by S^i) is a finite string drawn from the set $\mathbf{R}$. Given a state vector S^i, $r \in \mathbf{R}$ appearing in the kth position in the string S^i denotes that a customer of type r is in the kth position, in FIFO order, at the node i. Thus, by definition, the customer in service is in the first position. $S = (S^1, \cdots, S^M)$ denotes the state of the entire network Q. The ith node is equipped with an exponential server which, when the state of the network is S, serves a customer of class r at the rate $\nu_{ir}(S)$.

Q_1 is a subnetwork of Q, consisting of M_1 ($<N$) nodes (indexed by $i \in \{1, \cdots, M_1\}$). Q_2 is the complementary network consisting of $M_2 = M - M_1$ nodes (indexed by $i \in \{M_1 + 1, \cdots, M\}$).

Some additional notation is inevitable; this we proceed to describe in the next subsection.

2.2 Notation

$\mathbf{N} = (N_1, \cdots, N_R)$ is the population vector of the network Q, where N_r is the number of customers of class r, $r \in \mathbf{R}$.

Let $\mathbf{R}^* = (\cup_{n\geq 1}\{1, \cdots, R\}^n) \cup \Theta$ where Θ denotes the empty string. For any $s \in \mathbf{R}^*$, denote by $N_r(s)$ the population of class r in the string s, and let

$$N(s) = (N_1(s), \cdots, N_r(s), \cdots, N_R(s)).$$

For K, a positive integer, let

$$S_{\mathbf{N}}^{K} = \{(s^1, \cdots, s^K): (\text{for every } i, 1 \leqslant i \leqslant K, s^i \in \mathbf{R}^*)$$

$$\text{and } \sum_{i=1}^{K} N(s^i) = \mathbf{N}\}.$$

As stated in Section 2.1, $S = (S^1, \cdots, S^M)$ denotes the Q-network state. Let $S_1 = (S^1, \cdots, S^{M_1})$ denote the Q_1-network state and $S_2 = (S^{M_1+1}, \cdots, S^M)$ denote the Q_2-network state. Let $\mathbf{F}_{\mathbf{N}}^{Q}$, and $\mathbf{F}_{\mathbf{N}}^{Q_1}$ and $\mathbf{F}_{\mathbf{N}}^{Q_2}$ denote respectively the sets of feasible states, in equilibrium, of the state process of the networks Q, Q_1, and Q_2, respectively.

A network $Q_{1\text{eq}}$ constructed from the nodes $\{1, \cdots, M_1\}$ is said to be *equivalent* to Q_1 if the joint equilibrium (probability) distribution of the state processes of $Q_{1\text{eq}}$ is the same as the marginal joint equilibrium distribution of the state process of Q_1 in Q.

2.3 Construction of $Q_{1\text{eq}}$

Let π: $\mathbf{F}_{\mathbf{N}}^{Q} \rightarrow (0, 1)$ be the equilibrium distribution of the state process of the network Q. Let (for every $(1 \leqslant i \leqslant M)(1 \leqslant r \leqslant R)$) (for every $S \in \mathbf{F}_{\mathbf{N}}^{Q}$) $\nu_{ir}(S) = \nu_{ir}(S_1)$ and (for every $S_1 \in \mathbf{F}_{\mathbf{N}}^{Q_1}$)

$$\rho_{ir}^{S_1} \triangleq \pi\ \{\text{A customer of type } r \text{ is in service at node}$$

$$i/Q_1\ i \text{ is in state } S_1\}.$$

Construct a network $Q_{1\text{eq}}$ from the nodes $\{1, \cdots, M_1\}$ as follows:

1. The routing between the nodes in $Q_{1\text{eq}}$ is the same as in Q_1 (self loops around nodes in Q_1 are included in $Q_{1\text{eq}}$).

2. When the state of $Q_{1\text{eq}}$ is S_1, node $j(1 \leqslant j \leqslant M_1)$ receives an exogenous arrival stream of class r customers with (state dependent) rate $\sum_{i=M_1+1}^{M} \rho_{ir}^{S_1} \nu_{ir}(S_1) p_{ij}^{(r)}$.

3. A customer of class r, after completing service at node i $(1 \leqslant i \leqslant M_1)$, leaves the network $Q_{1\text{eq}}$ with probability $\sum_{j=M_1+1}^{M} p_{ij}^{(r)}$.

Theorem 1: Q_{1eq} as constructed above is equivalent to the subnetwork Q_1 of Q.

Proof: The intuitive appeal of the construction is manifest. In Step 2 of the construction, for every $i(M_1 + 1 \leqslant i \leqslant M)$, $\rho_{ir}^{S_1} \nu_{ir}(S_1)$ is the conditional throughput of type r customers through node i, when Q_1 is in the state S_1. A fraction $p_{ij}^{(r)}$ of this flow through i finds its way into node j of Q_1.

A simple detailed proof can be obtained by summing the Kolmogorov equilibrium equations for Q over the set $\{S \in \mathbf{F}_{\mathbf{N}}^{Q}: S_1 \text{ fixed}\}$, and observing that the resulting equations are exactly the equilibrium equations for $Q_{1\text{eq}}$ described above (cf. Ref. 7). □

Remarks: But for the explosion in notation that occurs in setting up a detailed proof, it is clear that the construction of Q_{1eq} described above extends easily to networks other than those described in Section 2.1. In this work, however, we continue to restrict our attention to networks of the latter type.

We now turn to the subclass of product-form networks of the class of networks described in Section 2.1. Since we are concerned here with FIFO nodes, the service rates cannot be class dependent. We further assume that the service rates are not state dependent in any way, i.e., we now have

$$(\text{for every } (1 \leq i \leq M)(1 \leq r \leq R) \quad \text{and} \quad S \in \mathbf{F}_{\mathbf{N}}^{Q})\nu_{ir}(S) = \nu_i.$$

Let (for every $r \in \mathbf{R}$) $C^{(r)} \subset \{1, \cdots, M\}$ be the subset of nodes of Q that communicate under $P^{(r)}$ (i.e., in queueing-network terminology, the chain corresponding to class r). Let $\mathbf{R}_2 = \{r \in \mathbf{R} : C^{(r)} \cap \{M_1 + 1, \cdots, M\} \neq \emptyset\}$ be the set of customer types that visit Q_2. Let $\|\mathbf{R}_2\| = R_2$, $\|\mathbf{R} - \mathbf{R}_2\| = R_1$ (where $\| \ \|$ denotes set cardinality), and reindex $\mathbf{R}$ so that the elements of $\mathbf{R}_2$ receive the highest indices. Let $\mathbf{N}^2 = (N_{R_1+1}, \cdots, N_R)$ and if s is a string in $\mathbf{R}^*$, let $N^2(s) = (N_{R_1+1}(s), \cdots, N_R(s))$, i.e., $N^2(s)$ is the population vector, of the string s, restricted to the classes in $\mathbf{R}_2$.

For every $\mathbf{N}' = (N'_{R_1+1}, \cdots, N'_R) \leq \mathbf{N}^2$, consider the network $Q_2'(\mathbf{N}')$ obtained from Q by replacing all servers in Q_1 with infinite speed servers (i.e., by short-circuiting the nodes in Q_1), and placing $\mathbf{N}'$ customers in the resulting network. Let $\pi_{\mathbf{N}'}$ be the equilibrium distribution of the state process of the network $Q_2'(\mathbf{N}')$. Define for every $(M_1 + 1 \leq i \leq M)$, $r \in \mathbf{R}_2$,

$$\xi_{ir}^{\mathbf{N}'} \triangleq \pi_{\mathbf{N}'} \{\text{A customer of type } r \text{ is in service at node } i \text{ in } Q_2'(\mathbf{N}')\}.$$

2.4 The product-form case

Theorem 2: If Q is a product-form network then

$$\textit{for every } (M_1 + 1 \leq i \leq M)(r \in \mathbf{R}_2) \textit{ and every } S_1 \in \mathbf{F}_{\mathbf{N}}^{Q_1}$$
$$\rho_{ir}^{S_1} = \xi_{ir}^{\mathbf{N}^2 - N^2(S_1)} \qquad (\rho_{ir}^{S_1} \textit{ as defined earlier}).$$

(Note: it is obvious that (for every $(M_1 + 1 \leq i \leq M)$, $r \notin \mathbf{R}_2$ and $S_1 \in \mathbf{F}_{\mathbf{N}}^{Q_1})\rho_{ir}^{S_1} = 0$.)

Proof: The proof utilizes a simple lemma and is outlined in the appendix. □

Remarks: Theorem 2, when combined with Theorem 1, yields a generalization of Norton's Theorem[1] to multinode subnetworks. Even though the previous development is specific to the class of networks described in Section 2.1, it is clear that the same approach can be used to extend Norton's Theorem to the entire class of closed product-form networks. The product-form solution continues to play the same role as it does in Theorem 2, i.e., it allows the rates of the external arrival streams in Q_{1eq} to be computed from an analysis of Q_2' for all possible customer populations in Q_2'.

III. AN APPROXIMATION SCHEME

In an IBM Research Report, Chow and Yu[5] suggest a somewhat ad hoc, iterative approximation scheme for a class of central-server models, with a priority discipline at the central server. As mentioned earlier, the scheme relies on an inexact application of Norton's Theorem to such networks. In Section I we described a natural generalization of this so-called equivalent flow approximation scheme to more general nonproduct-form networks. In this section, we present the results of a detailed study of the application of this approximation to a simple, test-bed, central-server network.

3.1 The test-bed model

Consider the two-node network Q shown in Fig. 2. There are two customer classes, namely 1 and 2, with N_1 and N_2 customers, respectively (i.e., $\mathbf{N} = (N_1, N_2)$). At node 1, the customers of class 1 (high-priority) have preemptive priority over class 2 (low-priority) customers; after being preempted by a class 1 customer, when a class 2 customer reaches the service station again, it resumes service where it left off; class 1 and 2 customers have exponential service times with rates ν_{11} and ν_{12}, respectively. Such a service discipline is commonly referred to as a *preemptive resume* discipline. At node 2, there is no priority; customers are served in the order in which they arrive (irrespective of class), at the class independent exponential service rate ν_2. Customers alternately seek service at nodes 1 and 2 and stay in the network forever. This model belongs to a class of central-server networks that arise as models of computer systems.

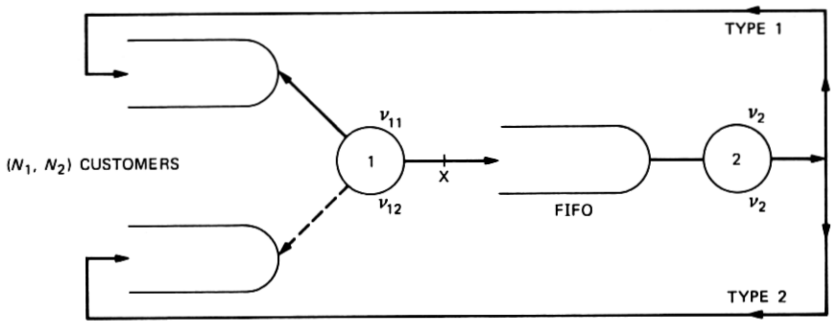

Fig. 2—Q.

3.2 Approximating the test-bed network

The network described in Section 3.1 is nonproduct form because of the preemptive resume discipline at node 1. In order to approximate the equilibrium behavior of node 1, we first increase the service rates at node 1 to infinity, thus effectively short circuiting the node. Denote the resulting network by Q_2' (Fig. 3). Then for each $(k_1,k_2) \leqslant (N_1,N_2)$ analyze Q_2' with k_1 and k_2 customers of types 1 and 2, respectively, in the network. Let (cf. Thm. 2) (for every $(k_1,k_2) \leqslant (N_1,N_2)$) (for every $r \in \{1, 2\}$) $\xi_{2r}^{(k_1,k_2)}$ = Prob {A customer of type r is in service at node 2 when (k_1,k_2) customers are in Q_2'}.

This probability will not depend on the sequence in which the (k_1,k_2) customers are placed in Q_2'. Since the service rate at node 2 is class independent, it is clear that, in equilibrium, all possible states, for any arrangement of the customers, are equally likely. From this we can directly conclude that

$$\text{(for every } (0, 0) < (k_1,k_2) \leqslant (N_1,N_2)\text{) (for every } r \in \{1, 2\}\text{)},$$

$$\xi_{2r}^{(k_1,k_2)} = \frac{k_r}{k_1 + k_2}.$$

Now consider the open network $\hat{Q}_{1eq}$ consisting of the node 1 in isolation. The service rates and discipline remain the same as in Q. When there are n_1 customers of type 1 and n_2 customers of type 2 in $\hat{Q}_{1eq}$ then customers of type r ($\in \{1, 2\}$) enter the network at the rate $\hat{\lambda}_r^{(n_1,n_2)}$ where

$$\text{(for every } (n_1,n_2) \leqslant (N_1,N_2))\hat{\lambda}_r^{(n_1,n_2)} = \xi_{2r}^{\mathbf{N}-(n_1,n_2)}\nu_2.$$

When a customer finishes service in $\hat{Q}_{1eq}$, it leaves the network (Fig. 4).

The evolution of the network $\hat{Q}_{1eq}$ can be described by a regular Markov process on the state space $\{(n_1,n_2) : (n_1,n_2) \leqslant (N_1,N_2)\}$. The idea is to approximate the equilibrium distribution of customers at node 1 in Q with the equilibrium distribution of customers in $\hat{Q}_{1eq}$.

At first glance, the approximation technique described above may seem rather ad hoc. However, we can draw upon the development in

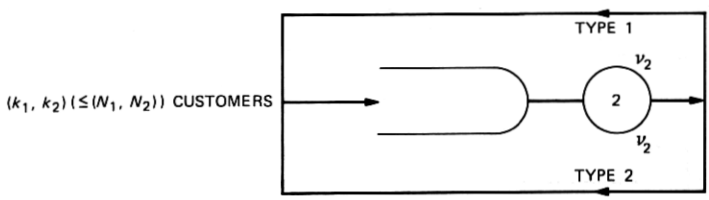

Fig. 3—Q_2.

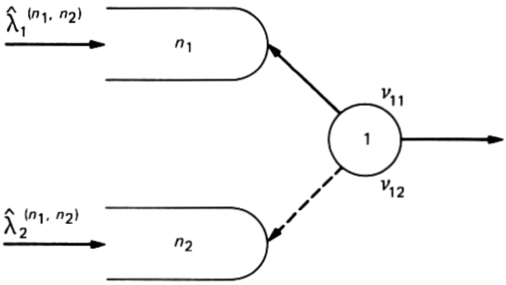

Fig. 4—$\hat{Q}_{1\,\mathrm{eq}}$.

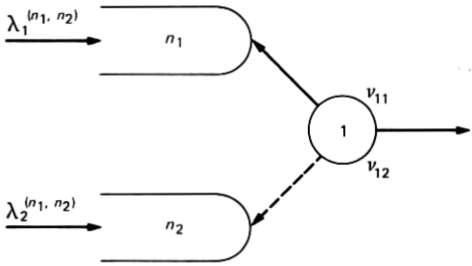

Fig. 5—$Q_{1\,\mathrm{eq}}$.

Section II to understand the inner workings of the test-bed model, and to show that, at least in principle, the approximation scheme is not altogether unreasonable.

It is clear that we can think of node 1, in the test-bed network Q, as comprising two FIFO nodes with service rates that depend just on the joint state of these two nodes. Consider the subnetwork Q_1 of Q consisting only of node 1. Theorem 1 can now be invoked to determine the exact equivalent network $Q_{1\mathrm{eq}}$. Let

(for every $(n_1,n_2) \leqslant (N_1,N_2)$) (for every $r \in \{1, 2\}$),

$$\rho_{2r}^{(n_1,n_2)} = \pi\{\text{A customer of type } r \text{ is in service at node } 2/(n_1,n_2) \text{ customers in } Q_1\}.$$

$Q_{1\mathrm{eq}}$ is then an open network consisting of node 1. When there are n_1 customers of type 1 and n_2 customers of type 2 in $Q_{1\mathrm{eq}}$, then customers of type r ($\in \{1, 2\}$) enter the network at the rate $\lambda_r^{(n_1,n_2)}$ where

$$(\text{for every } (n_1,n_2) \leqslant (N_1,N_2))\lambda_r^{(n_1,n_2)} = \rho_{2r}^{(n_1,n_2)}\nu_2.$$

When a customer finishes service in $Q_{1\mathrm{eq}}$ it leaves the network (Fig. 5).

The equilibrium distribution of customers in $Q_{1\mathrm{eq}}$ is exactly the same as the equilibrium distribution of customers in Q_1. Observe, though, that the form of Q_{1eq} is the same as that of $\hat{Q}_{1eq}$, the difference lying in the state-dependent input rates. It is in this sense that the approximation scheme is reasonable. The idea now is to compare the exact state-dependent input rates, $\rho_{2r}^{(n_1,n_2)}\nu_2$, with the approximate state-dependent rates,

$$\xi_{2r}^{\mathbf{N}-(n_1,n_2)}\nu_2 = \left(\frac{N_r - n_r}{N_1 - n_1 + N_2 - n_2}\,\nu_2\right),$$

i.e., to compare $\rho_{2r}^{(n_1,n_2)}$ with

$$\frac{N_r - n_r}{N_1 - n_1 + N_2 - n_2}$$

for all $(n_1,n_2) \lneqq (N_1,N_2)$ and for $r \in \{1, 2\}$.

3.3 Qualitative evaluation of the approximation

In this section, we present a qualitative evaluation of the approximation scheme as applied to the test-bed model.

Observe that if the service rates for the two FIFO queues comprising node 1, in Q_1, were not state dependent (in the priority scheme they are state dependent), then Q_1 would, in fact, be a product-form network. Theorem 2 would then lead us to conclude that $\hat{Q}_{1eq}$ and Q_{1eq} were the same. Consider what happens if, in Q_1, ν_{11} is allowed to go to infinity. Then, effectively, the high-priority customers do not interfere with the low-priority customers at node 1. With $\nu_{11} = \infty$, the network becomes the one shown in Fig. 6, which is a product-form network. Thus according to our observation above, for values of ν_{11} that are large, compared to ν_{12} and ν_2, the approximation can be expected to yield very good results.

In order to discover the situations in which the approximation can be expected to behave poorly, one needs to understand what aspects

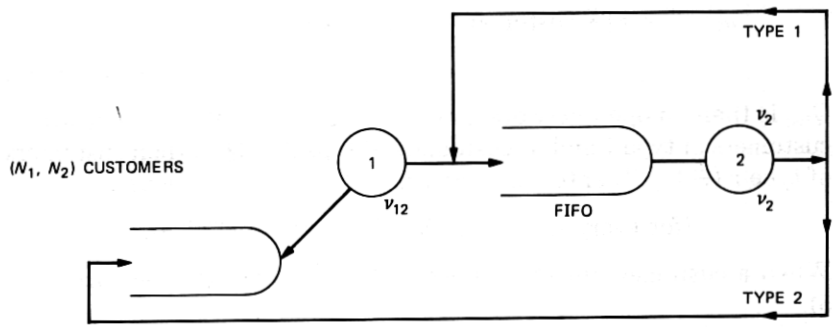

Fig. 6—"Lim" $Q^{\prime}_{11\to\infty}$.

of the exact network the approximation fails to capture. If Q_1 were a product-form network, then, given that $(k_1,k_2)(\leqslant(N_1,N_2))$ customers were in node 2, all arrangements of customers within the node would be equally likely. As it stands, however, at node 1, priority 1 customers can preempt customers of priority 2. This suggests that given $(k_1,k_2)(\leqslant N_1,N_2)$ customers in node 2, some arrangements of customers would be more likely than others. In fact, we conjecture that priority 1 customers are more likely to be ahead of priority 2 customers, leading to the (conjectured) conclusion that

$$(\text{for every } (0,\,0) < (k_1,k_2) \leqslant (N_1,N_2))\,\rho_{21}^{\mathbf{N}-(k_1,k_2)} \geq \frac{k_1}{k_1 + k_2},$$

and

$$\rho_{22}^{\mathbf{N}-(k_1,k_2)} \leqslant \frac{k_2}{k_1 + k_2}.$$

Thus $\hat{Q}_{1eq}$ uses smaller (resp. larger) state-dependent input rates for type 1 (resp. type 2) customers than the exact equivalent Q_{1eq}. This idea is suggestive, but it is difficult to draw any immediate conclusions from this conjecture as to the relationship between exact and approximate performance measures of the network.

Another approach to discovering the direction in which the approximation can be expected to err is to observe that if node 2 in Q is replaced by a processor-sharing node, with class-independent service rate ν_2, then $\hat{Q}_{1eq}$ becomes the exact equivalent of Q_1 (cf. Fig. 7). (This follows because when node 2 is processor sharing, if $(k_1,k_2)(\leqslant (N_1,N_2))$ customers are present at node 2, then the rate of flow of class $r(\in \{1, 2\})$ customers into node 1 is $(k_r/(k_1 + k_2))\nu_2$.) To fix ideas consider the case $N_1 = n(\geqslant 1)$ and $N_2 = 1$. The throughput of the class 2 customer is simply the reciprocal of the mean successive passage times of the (single) class 2 customer through the point X (cf. Figs. 2 and

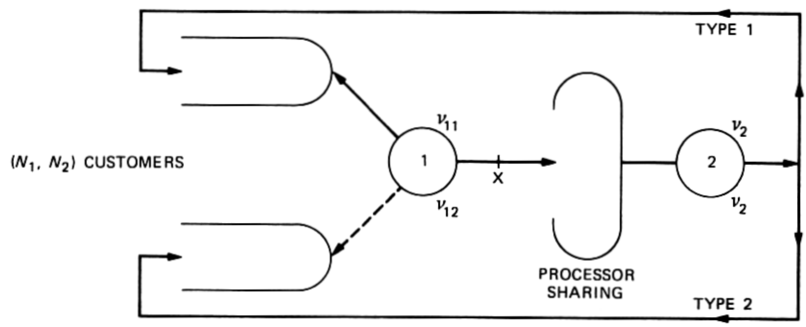

Fig. 7—Q with node 2 processor sharing.

7). In either network, when the class 2 customer crosses the point X to enter node 2, it finds all the class 1 customers receiving service at this node. In the original network, since node 2 is FIFO, the class 2 customer will have to wait for full service completion of the n class 1 customers before it can leave node 2 (and subsequently, at some future time instant, cycle back through X). Thus, the mean sojourn time of the type 2 customer, in node 2 of the original network, is $(n + 1)/\nu_2$. However, if node 2 is processor sharing, then on entering node 2, the customer of type 2 starts receiving service immediately at the rate $\nu_2/(n + 1)$, and continues to receive service at a rate $\nu_2/(k + 1)$ $(0 \leqslant k \leqslant n)$ until it finally leaves. Thus, in this case, the sojourn time of the class 2 customer at node 2 is stochastically dominated by an exponentially distributed random variable with mean $n + 1/\nu_2$, and hence has a mean smaller than $(n + 1)/\nu_2$. We further expect, intuitively, that, after completing service at node 2, when the type 2 customer returns to node 1, it expects to find more type 1 customers at node 1 when node 2 is FIFO than when it is processor sharing. Given that the type 2 customer finds $k(0 \leqslant k \leqslant n)$ type 1 customers on its arrival at node 1, its sojourn time at node 1 does not depend on whether node 2 is FIFO or processor sharing, and increases with increasing k. Thus, we expect that the mean sojourn time of the type 2 customer at node 1 will be larger if node 2 is FIFO than when it is processor sharing. The conclusion is that the mean passage time of the type 2 customer, through the point X, is larger in the original network than in the approximating network.

To see the magnitude of the error this effect could cause, let $N_1 = 1$ and allow $\nu_{12} \rightarrow \infty$, $\nu_2 \rightarrow \infty$, and $\nu_2/\nu_{12} \rightarrow 0$. Under these assumptions, in the original network, the class 2 customer will be blocked once (and only once) at node 1 each time it cycles through the point X. The average time it spends in the blocked condition is $1/\nu_{11}$. The rest of the time in each cycle tends to 0. Hence, the mean response time for the class 2 customer in the original network is $1/\nu_{11}$. If node 2 is replaced by a processor sharing node, then, each time the class 2 customer cycles through X, it is blocked once (and only once) with probability 1/2. Hence, the mean response time for the class 2 customer in the approximating network is $1/2\nu_{11}$, thus yielding an error of 100 percent.

We do not yet have a simple but rigorous argument that would allow us to say conclusively that the approximation yields higher throughputs for low-priority customers. However, the arguments presented above do make the conclusion plausible.

3.4 Numerical examples

To examine how the approximation works with specific examples, we wrote a FORTRAN program to solve the equilibrium equations for $\hat{Q}_{1eq}$ using a simple recursive technique.[8] The program was somewhat more general, in that it could accept arbitrary state-dependent input rates and output rates. Thus, the same program could be used to solve the network exactly, if it were given the exact values of $\rho_{12}^{(n_1,n_2)}$ and $\rho_{22}^{(n_1,n_2)}$ for the various feasible (n_1,n_2).

It is not hard to calculate exactly the probabilities $\rho_{21}^{(n_1,n_2)}$ and $\rho_{22}^{(n_1,n_2)}$ for some simple cases. Of course for $(n_1,n_2) \neq (N_1,N_2)$, $\rho_{22}^{(n_1,n_2)} = 1 - \rho_{21}^{(n_1,n_2)}$. Consider, for the purpose of illustration, the case $N_1 = 1$, $N_2 = 1$. Note that the state of the network Q is completely described by the state of node 2. The epochs of entry into the state $S^2 = (12)$ are renewal epochs. The next state is, inevitably, $S^2 = (2)$. The next state is $S^2 = (21)$ with probability $\nu_{11}/(\nu_{11} + \nu_2)$ and $S^2 = (\Theta)$ with probability $\nu_2/(\nu_{11} + \nu_2)$. Because of the preemptive discipline, the next state to be entered in the set $\{(12), (21)\}$ will be $S^2 = (12)$, thus completing a renewal cycle. Since the expected holding time in each state in $\{(12), (21)\}$ is $1/\nu_2$, therefore

$$\rho_{21}^{(0,0)} = \frac{\dfrac{1}{\nu_2}}{\dfrac{1}{\nu_2} + \dfrac{\nu_{11}}{\nu_{11} + \nu_2} \cdot \dfrac{1}{\nu_2}} = \frac{1}{1 + \dfrac{\nu_{11}}{\nu_{11} + \nu_2}},$$

and, of course,

$$\rho_{21}^{(1,0)} = 0 \quad \text{and} \quad \rho_{21}^{(0,1)} = 1.$$

In Table I we list the exact expressions for $\rho_{21}^{(n_1,n_2)}$, for all $(n_1,n_2) \leqslant (N_1,N_2)$, for some values of (N_1,N_2). These were computed in the same fashion as in the above example.

In Tables II(a), (b), and (c), we give several numerical examples of exact and approximate solutions of the test-bed network. The exact solutions and the approximate solutions were obtained using the FORTRAN program described above. The program yields the equilibrium joint-probability distribution of queue lengths at node 1. In Tables II(a), (b), and (c), we display these joint probabilities and the node 1 utilizations.

The following observations are immediate and summarize our conclusions regarding the performance of the approximation scheme when applied to the test-bed network.

1. The numerical computations support our earlier observations that if ν_{11} is large, then the approximation can be expected to yield excellent results [cf. case (1) in each of Tables II(a), (b), and (c)].

2. The low-priority utilizations are consistently higher, again supporting our earlier observations regarding the direction in which the approximation can be expected to err.

Table I—Exact expressions for $\rho_{21}^{(n_1,n_2)}$ in the test-bed network Q. (cf. Thm. 1 and Fig. 2)

N_1	N_2	(n_1,n_2)	$\rho_{21}^{(n_1,n_2)*}$	Comparison with $\hat{\rho}_{21}^{(\cdot,\cdot)}$
1	1	(1,0)	0	
		(0,1)	1	
		(0,0)	$\dfrac{1}{1+\dfrac{\nu_{11}}{\nu_{11}+\nu_2}}$	$\geq\frac{1}{2}$
2	1	(2,0)	0	
		(0,1), (1,1)	1	
		(1,0)	$\dfrac{1}{1+\dfrac{\nu_{11}}{\nu_{11}+\nu_2}+\dfrac{\nu_2}{\nu_{11}+\nu_2}\cdot\dfrac{\nu_{11}}{\nu_{11}+\nu_2}}$	$\geq\frac{1}{2}$
		(0,0)	$\dfrac{1+\dfrac{\nu_{11}}{\nu_{11}+\nu_2}}{1+\dfrac{\nu_{11}}{\nu_{11}+\nu_2}+\left(\dfrac{\nu_{11}}{\nu_{11}+\nu_2}\right)^2+\dfrac{\nu_2}{\nu_{11}+\nu_2}\cdot\left(\dfrac{\nu_{11}}{\nu_{11}+\nu_2}\right)^2}$	$\geq\frac{2}{3}$
1	2	(1,0), (1,1)	0	
		(0,2)	1	
		(0,1)	$\dfrac{1}{1+\dfrac{\nu_2}{\nu_{12}+\nu_2}\cdot\dfrac{\nu_{11}}{\nu_{11}+\nu_2}+\dfrac{\nu_{12}}{\nu_{12}+\nu_2}\cdot\left[\dfrac{\nu_{11}}{\nu_{11}+\nu_2}+\dfrac{\nu_2}{\nu_{11}+\nu_2}\cdot\dfrac{\nu_{11}}{\nu_{11}+\nu_2}\right]}$	$\geq\frac{1}{2}$
		(0,0)	$\dfrac{1}{1+\dfrac{\nu_{11}}{\nu_{11}+\nu_2}+\dfrac{\nu_{11}}{\nu_{11}+\nu_2}\cdot\dfrac{\nu_{12}}{\nu_{12}+\nu_2}\left[1+\dfrac{\nu_2}{\nu_{12}}+\dfrac{\nu_2}{\nu_2+\nu_{11}}\right]}$	$\geq\frac{1}{3}$

* $(\rho_{22}^{(n_1,n_2)} = 1 - \rho_{21}^{(n_1,n_2)}$ if $(n_1,n_2) \neq (N_1,N_2)$; $\rho_{21}^{(N_1,N_2)} = \rho_{22}^{(N_1,N_2)} = 0)$.

3. When ν_{11}, ν_{12} and ν_2 are comparable, then the approximation yields good results with errors in the utilizations in the neighborhood of 10 percent.

4. Considerable errors in the low-priority utilizations can arise, however. Witness case 3 in each of the Tables II(a), (b), and (c). With a very large low-priority service rate at node 1, the approximate low-priority utilization suffers from an error of 20 to 50 percent.

5. For the range of examples studied, the equilibrium probabilities

Table II—Numerical comparisons of exact and approximate solutions of the test-bed network

Case No.	ν_{11}	ν_{12}	ν_2	Equilibrium State Probabilities at Node 1: State (n_1,n_2)	Exact Probability	Approximate* Probability	Node 1 Utilizations: Class	Exact Utilization	Approximate* Utilization
				(a)	($N_1 = 1, N_2 = 1$)				
1	10	1	1	(0,0)	0.614	0.606			
				(1,0)	0.029	0.028	1	0.064	0.063
				(0,1)	0.322	0.331	2	0.322	0.331
				(1,1)	0.035	0.036			
2	1	.5	1	(0,0)	0.231	0.222			
				(1,0)	0.077	0.056	1	0.462	0.444
				(0,1)	0.308	0.333	2	0.308	0.333
				(1,1)	0.385	0.389			
3	1	100	2	(0,0)	0.393	0.488			
				(1,0)	0.196	0.163	1	0.601	0.504
				(0,1)	0.0059	0.0081	2	0.0059	0.0081
				(1,1)	0.405	0.341			
				(b)	($N_1 = 2, N_2 = 1$)				
1	10	1	1	(0,0)	0.681	0.676			
				(1,0)	0.044	0.043			
				(2,0)	0.002	0.002	1	0.0754	0.0748
				(0,1)	0.244	0.249	2	0.244	0.249
				(1,1)	0.027	0.027			
				(2,1)	0.003	0.003			
2	1	.5	1	(0,0)	0.179	0.171			
				(1,0)	0.083	0.065			
				(2,0)	0.024	0.016	1	0.631	0.618
				(0,1)	0.191	0.211	2	0.191	0.211
				(1,1)	0.250	0.260			
				(2,1)	0.274	0.276			
3	1	100	2	(0,0)	0.170	0.217			
				(1,0)	0.116	0.109			
				(2,0)	0.050	0.036	1	0.827	0.780
				(0,1)	0.0022	0.0033	2	0.0022	0.0033
				(1,1)	0.187	0.188			
				(2,1)	0.474	0.447			
				(c)	($N_1 = 1, N_2 = 2$)				
1	10	1	1	(0,0)	0.464	0.457			
				(1,0)	0.015	0.014			
				(0,1)	0.316	0.319	1	0.0495	0.0487
				(1,1)	0.016	0.016	2	0.487	0.494
				(0,2)	0.170	0.175			
				(1,2)	0.019	0.019			
2	1	.5	1	(0,0)	0.116	0.109			
				(1,0)	0.028	0.108			
				(0,1)	0.177	0.182	1	0.437	0.418
				(1,1)	0.070	0.055	2	0.447	0.473
				(0,2)	0.270	0.291			
				(1,2)	0.340	0.346			
3	1	100	2	(0,0)	0.473	0.584			
				(1,0)	0.167	0.130			
				(0,1)	0.008	0.010	1	0.517	0.404
				(1,1)	0.115	0.090	2	0.0101	0.0123
				(0,2)	0.002	0.002			
				(1,2)	0.235	0.184			

*Based on Norton's Equivalent.

are never drastically wrong, and follow trends similar to the exact values.

The test-bed model is hard to analyze exactly for population sizes larger than the ones considered. We have run detailed simulations of the test-bed model for larger population sizes, and the results of these simulations continue to support the qualitative observations we have made above.

IV. CONCLUSIONS

We have demonstrated an approach for systematically analyzing the equivalent flow approximation. Our investigations have (1) revealed the conceptual basis for the approximation scheme, and (2) led to an understanding of the reasons for, and directions of, the errors that such an approximation scheme could introduce when applied to a class of prioritized central-server models. The approximation as described in the paper is of more general applicability, and much work remains to be done to discover its validity (accuracy and computational tractability) for more complicated, nonproduct-form networks. Our work, we think, provides the theoretical understanding and motivation for pursuing more detailed investigations.

V. ACKNOWLEDGMENTS

We are indebted to J. S. Kaufman for introducing us to the problem. We are grateful to him and to B. T. Doshi, B. Melamed, and B. Sengupta for making themselves available for several helpful discussions.

REFERENCES

1. K. M. Chandy, U. Herzog, L. Woo, "Parametric Analysis of Queueing Networks," IBM J. Res. and Develop. *19,* No. 1 (January 1975), pp. 36–42.
2. P. S. Kritzinger, S. Van Wyck, A. E. Krzesinski, "A Generalization of Norton's Theorem for Multiclass Queueing Networks," Performance Evaluation, *2* (July 1982), pp. 98–107.
3. S. Balsamo and G. Iazeolla, "An Extension of Norton's Theorem for Queueing Networks," IEEE Trans. Software Eng., *SE-8,* No. 4 (July 1982), pp. 298–305.
4. C. H. Sauer and K. M. Chandy, "Approximate Analysis of Central Server Models," IBM J. Res. Develop., *19,* No. 3 (May 1975), pp. 301–13.
5. W. Chow and P. S. Yu, "An Approximation Technique for Central Server Queueing Models with a Priority Dispatching Rule," IBM Res. Rpt. (1980), RC8163 (No. 35479).
6. M. Schwartz, "Performance Analysis of the SNA Virtual Route Pacing Control," IEEE Trans. Commun., *COM-30,* No. 1, Part II (January 1982), pp. 172–84.
7. A. Kumar, unpublished work.
8. J. S. Kaufman, unpublished work.
9. F. Baskett, K. M. Chandy, R. R. Muntz, and F. G. Palacios, "Open, Closed and Mixed Networks of Queues with Different Classes of Customers," J. ACM, *22,* No. 2 (April 1975), pp. 248–60.

APPENDIX

Proof of Theorem 2: Index the nodes of Q_2' in the same order in which they were indexed in Q. We need the following lemma.

Lemma: Let $T^{(r)}$ denote the class r routing probability matrix for the network Q_2'. Partition $P^{(r)}$ as follows:

$$P^{(r)} = \begin{array}{c} \\ M_1 \{ \\ M_2 \{ \end{array} \begin{array}{c|c} \overbrace{\phantom{P_{11}^{(r)}}}^{M_1} & \overbrace{\phantom{P_{12}^{(r)}}}^{M_2} \\ P_{11}^{(r)} & P_{12}^{(r)} \\ \hline P_{21}^{(r)} & P_{22}^{(r)} \end{array}.$$

Then

(1) $r \in \mathbf{R}_2 \Rightarrow T^{(r)} = P_{22}^{(r)} + P_{21}^{(r)}[I - P_{11}^{(r)}]^{-1}P_{12}^{(r)}$

(2) *If $\Lambda^{(r)}$ solves $\Lambda^{(r)}P^{(r)} = \Lambda^{(r)}$ then, partitioning $\Lambda^{(r)}$ as*

$$\Lambda^{(r)} = \left[\underbrace{\Lambda_1^{(r)}}_{M_1} \vdots \underbrace{\Lambda_2^{(r)}}_{M_2} \right]$$

(a) *if $r \notin \mathbf{R}_2$ then $\Lambda_2^{(r)} = 0$*

(b) *if $r \in \mathbf{R}_2$ then $\Lambda_2^{(r)} = \Lambda_2^{(r)}T^{(r)}$.*

Proof of Lemma: Conclusion 1 follows readily from the fact that, for each $r \in \mathbf{R}$, $P^{(r)}$ is the transition probability matrix of a finite Markov chain with a single, positive communication class that has a nonempty intersection with $\{M_1 + 1, \cdots, M\}$. For details see Ref. 7.

Conclusion 2 follows directly from Conclusion 1. □

Returning to the proof of Theorem 2, we let π: $\mathbf{F}_{\mathbf{N}}^{Q} \to (0, 1)$ be the equilibrium distribution of the state process of network Q.* It is now well known (cf. Ref. 9) that $\pi(\cdot)$ is of the form

$$\pi(S) = \frac{1}{G} \prod_{i=1}^{M} f_i(S^i),$$

where G is a normalization constant and, for each $i \in \{1, \cdots, M\}$, f_i depends only $N(S^i)$, ν_i and (for every r, $(1 \leqslant r \leqslant R)$)$\lambda_i^{(r)}$, where $\Lambda^{(r)} = (\lambda_i^{(r)}, \cdots, \lambda_M^{(r)})$ is any solution of $\Lambda^{(r)}P^{(r)} = \Lambda^{(r)}$.

Hence, (for every $S_1 \in \mathbf{F}_{\mathbf{N}}^{Q_1}$) (for every i, r, $(M_1 + 1 \leqslant i \leqslant M)$, $r \in \mathbf{R}_2$)

$$\rho_{ir}^{S_1} = \frac{\displaystyle\sum_{\{S:S\in\mathbf{F}_{\mathbf{N}}^{Q},(S^1,\cdots,S^{M_1})=S_1,S^i(1)=r\}} \prod_{j=1}^{M} f_j(S^j)}{\displaystyle\sum_{\{S:S\in\mathbf{F}_{\mathbf{N}}^{Q},(S^1,\cdots,S^{M_1})=S_1\}} \prod_{j=1}^{M} f_j(S^j)}$$

$$= \frac{\displaystyle\sum_{\{S_2:S_2\in\mathbf{F}_{\mathbf{N}}^{Q_2}\cap S_{\mathbf{N}^2-N^2(S_1)}^{M_2},S_2^i(1)=r\}} \prod_{j=M_1+1}^{M} f_j(S_2^j)}{\displaystyle\sum_{\{S_2:S_2\in\mathbf{F}_{\mathbf{N}}^{Q_2}\cap S_{\mathbf{N}^2-N^2(S_1)}^{M_2}\}} \prod_{j=M_1+1}^{M} f_j(S_2^j)},$$

*For notation see Section 2.2.

(where $\mathbf{F}_{\mathbf{N}}^{Q_2} \cap S_{\mathbf{N}^2 - N^2(S_1)}^{M_2}$ is the set of feasible states of Q_2 when

$$\mathbf{N}^2 - N^2(S_1) \text{ customers are in } Q_1),$$

which, using the above lemma and the fact that the equilibrium distribution of the Q_2' network state process is still product form,

$$= \xi_{ir}^{\mathbf{N}^2 - N^2(S_1)}.$$

Remarks: Some care is needed in asserting the last equality in the case where there are classes r_1 and r_2, such that the submatrices of the communicating classes under $T^{(r_1)}$ and $T^{(r_2)}$ are the same permutation matrices (i.e., members of classes r_1 and r_2 cannot overtake each other). In this case the equality follows because for each $\mathbf{N}'$, $\xi_{ir_1}^{\mathbf{N}'}$ and $\xi_{ir_2}^{\mathbf{N}'}$ are independent of the order in which members of these classes circulate in the network Q_2'. □

AUTHOR

Anurag Kumar, B.Tech (Electrical Engineering), 1977, Indian Institute of Technology, Kanpur, India; Ph.D. (Major: Electrical Engineering, Minor: Operations Research/Math), 1981, Cornell University; Bell Laboratories, 1981—. Mr. Kumar joined the Performance Analysis Department in December 1981, and has since been involved in two projects: a study of a class of approximation schemes for queueing-network analysis; and the identification, analysis, and evaluation of possibilities for giving priority to critical user traffic in the Public Switched Network.