Multiclass Multiserver Threshold-Based Systems: A Study of Noninstantaneous Server Activation

Cheng-Fu Chou, *Member, IEEE Computer Society*, Leana Golubchik, and John Chi-Shing Lui

Abstract—In this paper, we consider performance evaluation of a system which shares K servers (or resources) among N heterogeneous classes of workloads, where server allocation and deallocation for class i is dictated by a class specific threshold-based policy with hysteresis control. In particular, the server activation time for class i is *noninstantaneous*. There are many systems and applications where a multiclass threshold-based queueing system can be of great use. One important utility of using threshold-based approaches is in situations where applications may incur server usage costs. In these cases, one needs to consider not only the performance aspects but also the resulting cost/performance ratio. The motivation for using hysteresis control is to reduce the unnecessary cost of server setup (or activation) and server removal (or deactivation) whenever there are momentary fluctuations in workload. Moreover, servers in such systems and applications are often needed by multiple classes of workloads and, hence, it is desirable to find good approaches to sharing server resources among the different classes of workloads, preferably without statically partitioning the server pool among these classes. An important and distinguishing characteristic of our work is that we consider the modeling and analysis of a multiclass system with *noninstantaneous* server activation. The main contributions of this work are 1) in developing an efficient approximation method for solving such models, 2) in verifying the convergence of our iterative method, and 3) in evaluating the resulting accuracy of the technique for computing performance measures of interest, which can subsequently be used in making system design choices.

Index Terms—Multiclass multiserver queueing system, hysteresis control, performance, evaluation, system design.

1 INTRODUCTION

IN this paper, we consider performance evaluation of a multiclass multiserver system, in which K servers are shared among N heterogeneous classes of workloads and $K \geq N$. In this multiclass multiserver system, servers (resources) are needed by multiple classes of workloads (applications) and we also note that not all classes of requests are operating at high load at the same time. That is, when the traffic loading of class i is low, it is not desirable to operate unnecessarily many servers for that class, due to the incurred usage costs as well as due to the performance consequences of that class under-utilizing the servers while other classes are (possibly) experiencing high traffic workloads. On the other hand, it is also not desirable for a system to exhibit very long delays, which can result from lack of servers under heavy loads. Therefore, it is an important issue for the system to figure out a good approach to use few resources to serve those requests from all N classes and still attain a good cost/performance ratio instead of statically partitioning the server pool among the classes. To deal with the above issue and

efficiently utilize system resources for such multiclass multiserver system, we propose a threshold-based approach to dynamically assign servers to services of different class requests, i.e., how to allocate or deallocate servers to different classes is governed by a set of thresholds.

The motivation for using a threshold-based approach in a system is that applications may incur server usage costs. Thus, one not only needs to consider the performance but also the cost/performance ratio. In this work, we use “multiclass multiserver threshold-based system” to refer to our proposed threshold-based system, which is able to adaptively share the servers among different workloads without statically partitioning the server pool among the classes. Note that in many cases, a “simple” threshold-based system may not suffice since it is prone to workload oscillations. Such workload oscillations coupled with nonnegligible server setup and removal costs can result in a poor cost/performance ratio of a system. Therefore, hysteresis control is used to reduce the unnecessary cost by such oscillation behavior. The goal of this work is to look for efficient analysis techniques of threshold-based queueing systems with hysteresis control.

There are many applications where threshold-based resource management policies can be employed and, thus, performance evaluation of such systems through analysis of *multiclass* threshold-based queueing systems with hysteresis control can be of great use. For example, the Novell file server maintains a memory pool wherein a fraction of it is used for communication buffers and a fraction is used for file buffers, where threshold-based policies are implemented in order to make decisions about when to increase the number of network buffers and when to decrease it; the threshold values are based on perceived packet losses due to increases in network traffic activity. Similarly, OS design has been moving toward maintaining a common buffer space pool that can be

- C.-F. Chou is with the Department of Computer Science and Information Engineering, National Taiwan University, Room 517, CSIE Building, #1 Roosevelt Rd. Sec. 4, Taipei, Taiwan, R.O.C. E-mail: ccf@csie.ntu.edu.tw.
- L. Golubchik is with the Department of Computer Science, University of Southern California, SAL 226, 941 W. 37th Place, Los Angeles, CA 90089-0781. E-mail: leana@cs.usc.edu.
- J.C.-S. Lui is with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Rm 1028 Ho Sing-Hang Engineering Building, Shatin, N.T., Hong Kong. E-mail: cslui@cse.cuhk.edu.hk.

Manuscript received 1 June 2004; revised 1 Apr. 2005; accepted 30 Dec. 2005; published online 28 Nov. 2006.

Recommended for acceptance by K.-J. Lin.

For information on obtaining reprints of this article, please send e-mail to: tpds@computer.org, and reference IEEECS Log Number TPDS-0134-0604.

dynamically managed between the various I/O processes. Another example application is server replication for different classes of Internet services for an overlay network. As the number of requests for a particular class of service increases, the number of servers needed to maintain an acceptable level of quality-of-service guarantees also increases. The use of a threshold-based approach can result in a cost-controlled creation/deletion of servers based on the changes in the workload for a particular class of request. Thus, the model presented in this paper and its efficient solution will be beneficial for many systems and applications.

1.1 Overview of the Problem

Now, we begin to give an overview of the multiclass multiserver threshold-based system, which has a total of K servers. In particular, the number of servers employed for servicing class i customers, $i \in \{1, \dots, N\}$, is governed by a *forward threshold* vector $F_i = [F_i(1), F_i(2), \dots, F_i(K_i - 1)]$ (where $F_i(1) < F_i(2) < \dots < F_i(K_i - 1)$) and a *reverse threshold* vector $R_i = [R_i(1), R_i(2), \dots, R_i(K_i - 1)]$ (where $R_i(1) < R_i(2) < \dots < R_i(K_i - 1)$), where K_i is the maximum number of servers that can be allocated to serve class i customers (i.e., the system includes hysteresis control). The service time of class i customers is represented by an exponential random variable with mean μ_i^{-1} . The server activation time for class i is *noninstantaneous*, and it is represented by an exponential random variable with mean β_i^{-1} . In general, $\mu_i \neq \mu_j$ and $\beta_i \neq \beta_j$ for $i, j \in \{1, \dots, N\}$. Next, we explain how to allocate or deallocate servers to classes in the system as follows. Initially, each class is allocated a minimum of one server.

When a class i customer arrives to an empty system (i.e., when there are no other class i customers), this newly arrived class i request is served by a single server. Arrival of a class i customer when there are already $F_i(j)$ class i customers in the system (with j servers already allocated to serve class i), causes an attempt to allocate one additional server to class i , where $j = 1, \dots, K_i - 1$. Departure of a class i customer which leaves behind $R_i(j)$ class i customers (with $j + 1$ servers already allocated to this class prior to this departure event), causes a deallocation of a server from class i , where $j = 1, \dots, K_i - 1$. In other words, this forces the return of a server, which was earlier allocated to class i , back to the pool of “free” servers which are available for allocation to all classes of customers. Therefore, all N classes of applications share a common pool of K servers, with *dynamic* allocation of servers to classes governed by a set of thresholds with hysteresis behavior.

Note that when $\sum_{i=1}^N K_i \leq K$, then the classes do not “interfere” with each other since the total peak resource demand is less than or equal to the total number of resources in the system. Of course, a more interesting and challenging case is when $\sum_{i=1}^N K_i > K$. In other words, we want to investigate the performance of each class of workload when the number of common servers is less than the total peak resource demand of all classes. By taking advantage of the fact that not all classes are operating at high load at the same time, one may use fewer resources (than with static resource partitioning) to serve requests from all N classes and still achieve a good cost/performance ratio.

1.2 Related Work and Our Contribution

Given the above motivation for the use of threshold-based systems with hysteresis control, we present an efficient

technique for solving the corresponding analytical models and computing various performance measures of interest, in the context of *noninstantaneous* server activation. We begin with a very brief survey of some of the existing literature on the topic. A two-server system is considered in [13], [14], [20]. An approximate solution for solving a degenerate form of this problem (where all thresholds are set to zero) is presented in [7], [9]; an approximate solution for a system that employs (nonzero) thresholds is presented in [21] (but without hysteresis). In [8], the authors solve a multiserver threshold-based queueing system with hysteresis, using the Green’s function method [6], [10], [11]. In [16], we give a solution of several forms of the single class, multiserver threshold-based queueing system with hysteresis using stochastic complementation [17]. Techniques for computation of bounds for performance measures of single class, multiserver threshold-based queueing systems with hysteresis and noninstantaneous server activation are given in [3]. Last, [4] provides a solution technique for multiclass, multiserver threshold-based system with hysteresis and *instantaneous* server activation. In this paper, we extend and generalize that work to *noninstantaneous* server activation; the noninstantaneous server activation can have a significant impact on the system performance (as will be illustrated in Section 5) and, hence, is an important model characteristic to consider. Specifically, in this work we consider and solve a *multiclass*, multiserver threshold-based queueing system with hysteresis control and *noninstantaneous* server activation.

The *contributions* of this work are as follows: To the best of our knowledge, *none* of the works described above give an efficient analytical solution technique for analyzing this model. Since, in many applications, different types of workloads compete for a pool of resources where server activation time is nonnegligible (e.g., it takes a nonzero time to replicate and activate a video server in an overlay network), we consider it an important and distinguishing characteristic of our work. In this paper, we present an iterative solution technique which solves the multiclass model by “breaking” it up into N single class models, “coupled” through a set of model parameters which capture the interaction between classes. We also illustrate the accuracy of our approach, which efficiently computes performance measures of interest, through a set of numerical results. Furthermore, we give a proof and discussion about the convergence of our iterative method to solve the approximate model. This is important since we can get further understanding of our analytic model such that we could construct more precise and efficient analytic model for the multiclass multiserver system. We also believe that the efficiency and accuracy of our iterative approach provides an important step in finding *optimal* threshold values¹ for a multiclass, multiserver threshold based system with hysteresis and *noninstantaneous* server activation. Finally, we note that a variety of iterative approaches have been used in numerous approximation techniques (e.g., refer to [2]). For instance, an iterative technique for a somewhat different control scheme for dynamic resource sharing between multiple classes is employed in [18], [19].

1. In general, the question of optimal values for the forward and reverse threshold vectors is a difficult problem and is outside the scope of this paper. This is because the “good” threshold values are a function of many factors, such as the server setup, usage, and removal costs, characteristics of the arrival process and the service rates, as well as the possible “interaction” between the different classes of workloads.

The remainder of this paper is organized as follows: In Section 2, we give a detailed description of our model. Section 3 describes our iterative solution approach for this model. The convergence of the iterative method to solve the approximate model is presented in Section 4. The quality of this approach, i.e., its accuracy and utility in system design and evaluation is discussed in Section 5 through the use of numerical results. Finally, our conclusions are given in Section 6.

2 SYSTEM MODEL

The Markovian model for our *multiclass, multiserver threshold-based* queueing system with *hysteresis* control has an infinite state space which can be described as follows. There are K servers in the system with $K \geq N$ where N is the total number of classes of customers. The service time of different classes can be different, and the service time requirements of a class i customer are exponentially distributed with parameter μ_i . The customer arrival process is Poisson with rate λ , where with probability α_i an arriving customer is of class i and $\sum_{i=1}^N \alpha_i = 1$ and $1 \leq i \leq N$. The addition and removal of servers for serving customers of class i is governed by the forward and the reverse threshold vectors $F_i = [F_i(1), F_i(2), \dots, F_i(K_i - 1)]$ and $R_i = [R_i(1), R_i(2), \dots, R_i(K_i - 1)]$, where $F_i(j) < F_i(j+1)$ for $1 \leq j \leq K_i - 2$, $R_i(j) < R_i(j+1)$ for $1 \leq j \leq K_i - 2$, and $R_i(j) < F_i(j)$ for $1 \leq j \leq K_i - 1$. Note that, unlike in [4], [5], the activation of a server for class i is *noninstantaneous* where the server activation time is exponentially distributed with mean β_i^{-1} . As mentioned in Section 1, this is motivated by the fact that in many applications addition of a new server takes a nonnegligible amount of time.

Each of these K servers is able to serve a customer of any class. Each class i starts out with one server and may attempt to obtain at most K_i servers. These servers are allocated for service of class i customers and returned to the pool of available servers based on the number of class i customers currently in the system (as stated more formally below). In general, $\sum_{i=1}^N K_i$ may be greater than, equal to, or less than K ; although the more interesting and challenging case is where $\sum_{i=1}^N K_i > K$. We model this system as a Markovian process $\mathcal{M}$, using two different variations. In the first variation, we constrain the number of class i customers when the number of servers allocated to class i is less than K_i (we motivate this variation below). The Markovian model for this variation is referred to as $\mathcal{M}^a$. In the second variation, we do not use such a constraint, and the Markovian model for this variation is referred to as $\mathcal{M}^n$. We now give a more detailed description of each of these models.

2.1 $\mathcal{M}^a$ with Constraint Vector a

The Markovian process $\mathcal{M}^a$, with a constraint vector a , has the following state space $\mathcal{S}^a$:

$$\mathcal{S}^a = \{(n_1, s_1, l_1, \dots, n_N, s_N, l_N) \mid n_i \geq 0, l_i \in \{1, \dots, K_i\}, \sum_{i=1}^N l_i \leq K, l_i \geq s_i, F_i(l_i) \leq n_i \leq F_i(l_i) + a_i^j, s_i \in \{1, 2, \dots, K_i\}, i = 1, \dots, N\},$$

where n_i is the number of class i customers in the system, s_i is the number of “busy” (or active) servers currently serving class i customers, and l_i is the number of servers allocated to class i , not all of which may currently be available for service of class i customers since server *activation* process is

noninstantaneous. Upon an arrival of a class i customer, if $F_i(j) \leq n_i \leq F_i(j) + a_i^j$ where $a_i^j \geq 0$ and $j = l_i$, the system attempts to allocate an additional server for service of class i customers, which is possible *only* when the system has sufficient amount of resources, i.e., if $\sum_{i=1}^N l_i < K$. Note that in a system where $\sum_{i=1}^N K_i > K$, it may not always be possible to allocate another server since it is possible that all K servers may have already been allocated. In this case, the arriving class i customer joins the queue of class i requests as long as $F_i(j) \leq n_i < F_i(j) + a_i^j$ (where $a_i^j \geq 0$ and $j = l_i$). When $n_i = F_i(j) + a_i^j$, the arriving class i customer is rejected by the system if there is no server available for allocation to class i (i.e., if $\sum l_i = K$). For *correctness*, we assume the following constraint on all a_i^j :

$$F_i(j) + a_i^j < F_i(j+1) + a_i^{j+1} \text{ for } i = 1, 2, \dots, N \text{ and } j = 1, 2, \dots, K_i - 1.$$

We also assume that $a_i^{K_i} = \infty$; hence, we have no restrictions on queue length when the maximum number of servers that may be needed by class i have been allocated (i.e., when $l_i = K_i$). The limitation on queue length when $l_i < K_i$ can be motivated by system design considerations. For example, if the system reaches a point where its design dictates that another server be allocated for class i workload but a server is not available, then one may assume that the system is temporarily overloaded and rejection of customers is a reasonable approach to dealing with overload conditions. Of course, a “real” system will not also have an infinite queue length, when the maximum number of servers (K_i) for class i has been allocated. In this case, we may either 1) use a finite queue length model (i.e., $a_i^{K_i}$ is finite) and study the system’s performance under a given queue size limitation, or 2) allow an infinite queue length (i.e., $a_i^{K_i} = \infty$) and use the model to study queue length requirements of the corresponding system. Our solution methodology (refer to Section 3) allows for either type of a model, but for simplicity of exposition, in the remainder of the paper, we will focus our discussion on the infinite queue version (i.e., where a_i^j is finite, for $j = 1, \dots, K_i - 1$, and $a_i^{K_i} = \infty$).

We now give a detailed description and formal structure for the transitions of $\mathcal{M}^a$. The transitions corresponding to an arrival of a class i customer fall into one of the following categories:

C_1 : no need to allocate another server to class i . The conditions for this category could be 1) the number of class i customers does not cross a corresponding forward threshold, or 2) there are already K_i servers allocated to class i , or 3) there are no available servers in the system due to resource contention among the different classes.

C_2 : a need to allocate another server to class i . The condition for this category is that the number of allocated server for class i is less than K_i and there is an available server in the system and the number of class i customers crosses a forward threshold.

The formal structure for these arrival transitions is as follows:

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{\lambda \alpha_i} (n_1, s_1, l_1, \dots, n_i + 1, s_i, l_i, \dots, n_N, s_N, l_N) \text{ if } C_1 \quad (1)$$

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{\lambda \alpha_i} (n_1, s_1, \dots, n_i + 1, s_i, l_i + 1, \dots, n_N, s_N, l_N) \text{ if } C_2, \quad (2)$$

where conditions C_1 and C_2 are

$$C_1 = \left((l_i < K_i) \wedge (n_i < F_i(l_i)) \right) \vee \left(l_i = K_i \right) \\ \vee \left((l_i < K_i) \wedge \left(\sum_{j=1}^N l_j = K \right) \wedge (F_i(l_i) \leq n_i < F_i(l_i) + a_i^{l_i}) \right) \\ C_2 = \left(l_i < K_i \right) \wedge \left(\sum_{j=1}^N l_j < K \right) \wedge \left(F_i(l_i) \leq n_i \leq F_i(l_i) + a_i^{l_i} \right).$$

The transitions corresponding to a departure of a class i customer fall into one of the following categories:

C_3 : no need to deactivate a server. The conditions for this category are either 1) only one server is allocated to class i or 2) the number of class i customers does not drop below a reverse threshold.

C_4 : a need to deactivate a server. The condition for this category is that more than one server is allocated to class i and the number of class i customers drops below a corresponding reverse threshold.

The formal structure for these departure transition is as follows:

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{s_i \mu_i} (n_1, s_1, l_1, \dots, n_i - 1, s_i, l_i, \dots, n_N, s_N, l_N) \text{ if } C_3 \quad (3)$$

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{s_i \mu_i} (n_1, s_1, l_1, \dots, n_i - 1, \min(s_i, l_i - 1), l_i - 1, \dots, n_N, s_N, l_N) \text{ if } C_4, \quad (4)$$

where conditions for C_3 and C_4 are

$$C_3 = ((n_i > 0) \wedge (l_i = 1)) \vee ((n_i > 0) \wedge (n_i - 1 > R_i(l_i - 1)) \wedge (l_i > 1))$$

$$C_4 = ((n_i > 0) \wedge (n_i - 1 = R_i(l_i - 1)) \wedge (l_i > 1)).$$

Last, there are transitions corresponding to class i server activations. The condition for these transitions is that the number of active servers is less than the number of allocated servers. The formal structure for these activation transitions is as follows:

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{(l_i - s_i) \beta_i} (n_1, s_1, l_1, \dots, n_i, s_i + 1, l_i, \dots, n_N, s_N, l_N) \text{ if } C_5, \quad (5)$$

where condition $C_5 = (l_i > s_i)$.

2.2 $\mathcal{M}^n$ without Constraint a

The second model variation is $\mathcal{M}^n$, which represents a Markovian process without constraints on the number of class i customers. It has the following state space, $\mathcal{S}^n$:

$$\mathcal{S}^n = \{(n_1, s_1, l_1, \dots, n_N, s_N, l_N) | n_i \geq 0, l_i \in \{1, \dots, K_i\}, l_i \geq s_i, s_i \in \{1, \dots, K_i\}, \sum s_i \leq K, i = 1, \dots, N\},$$

where n_i is the number of class i customers in the system, s_i is the number of “busy” (or active) servers currently servicing class i customers, and l_i is the number of servers “expected” to be allocated/activated for class i use—more specifically, according to the threshold vectors, l_i servers *should be* in use by class i customers but may not be, because 1) multiple classes of customers are competing for these servers and 2) in our model server activation is *noninstantaneous*. Hence, a major difference between $\mathcal{M}^a$ and $\mathcal{M}^n$ is that we use a constraint vector a to limit the queue length when $l_i < K_i$ in $\mathcal{M}^a$ while we do not limit the queue length in $\mathcal{M}^n$. We give a comparison study between these two models in Section 5.

We now give a detailed description and formal structure for the transitions of $\mathcal{M}^n$. The transitions corresponding to arrivals of class i customers fall into one of the following categories:

C_1 : no need to increase the number of expected servers (according to the forward threshold vector) for class i .

The conditions for this category could be 1) the number of class i customers does not cross a corresponding forward threshold or 2) the number of expected servers is equal to K_i .

C_2 : a need to increase the number of expected servers for class i . The condition for this category is that the number of expected servers for class i is less than K_i and the number of class i customers crosses a forward threshold.

The formal structure for these arrival transitions is as follows:

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{\lambda \alpha_i} (n_1, s_1, l_1, \dots, n_i + 1, s_i, l_i, \dots, n_N, s_N, l_N) \text{ if } C_1 \quad (6)$$

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{\lambda \alpha_i} (n_1, s_1, \dots, n_i + 1, s_i, l_i + 1, \dots, n_N, s_N, l_N) \text{ if } C_2, \quad (7)$$

where conditions C_1 and C_2 are

$$C_1 = ((l_i < K_i) \wedge (n_i < F_i(l_i))) \vee (l_i = K_i)$$

and

$$C_2 = (l_i < K_i) \wedge (F_i(l_i) = n_i).$$

The transitions corresponding departures of class i customers fall into one of the following categories:

C_3 : no need to deactivate a server. The condition for this category is that either 1) the system has only one server for class i or 2) the number of class i customers does not drop below a reverse threshold.

C_4 : a need to deactivate a server. The condition for this category is that there is more than one server for class i and the number of class i customers drops below a corresponding reverse threshold.

The formal structure for these departure transitions is as follows:

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{s_i \mu_i} (n_1, s_1, l_1, \dots, n_i - 1, s_i, l_i, \dots, n_N, s_N, l_N) \text{ if } C_3 \quad (8)$$

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{s_i \mu_i} (n_1, s_1, l_1, \dots, n_i - 1, \min(s_i, l_i - 1), l_i - 1, \dots, n_N, s_N, l_N) \text{ if } C_4, \quad (9)$$

where conditions C_3 and C_4 are

$$C_3 = ((n_i > 0) \wedge (l_i = 1)) \vee ((n_i > 0) \wedge (n_i - 1 > R_i(l_i - 1)) \wedge (l_i > 1))$$

$$C_4 = ((n_i > 0) \wedge (n_i - 1 = R_i(l_i - 1)) \wedge (l_i > 1)).$$

Last, there are the transitions corresponding to class i server activation. The condition for these transitions is that 1) the number of active servers is less than the number of “expected” servers and 2) there is an available server in the system. The formal structure for these server activation transitions is as follows:

$$(n_1, s_1, l_1, \dots, n_i, s_i, l_i, \dots, n_N, s_N, l_N) \xrightarrow{\beta_i} (n_1, s_1, l_1, \dots, n_i, s_i + 1, l_i, \dots, n_N, s_N, l_N) \text{ if } C_5, \quad (10)$$

where the condition $C_5 = ((l_i > s_i) \wedge (\sum_{j=1}^N l_j < K))$.

3 ITERATIVE METHOD

In this section, we describe an *iterative* approach to solving the models presented in Section 2. As described in Section 2, the corresponding Markov process,² $\mathcal{M}$, is infinite in multiple dimensions. One can choose to solve this model by 1) simulating the Markovian process $\mathcal{M}$ or 2) looking for special structure, or 3) looking for efficient approximation techniques. Because $\mathcal{M}$ appears to lack sufficient structure for an efficient exact solution technique (e.g., such as the matrix-geometric technique), we describe an approximate iterative solution technique for solving this model. The use of an approximation is motivated by the desire to construct an efficient solution approach (and simulation can be significantly slower than analytical solutions) as well as an accurate one (and iterative techniques can often produce fairly accurate results).

3.1 Basic Approach

Let us first describe the basic approach to solving the above defined Markovian model. We first break up the original model $\mathcal{M}$ into N single class Markovian submodels, namely, $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_N$ (see Section 3.2 for a more detailed description of the $\mathcal{M}_i$'s). These N Markovian models are “coupled” via a set of blocking probabilities. Specifically, the interaction between classes occurs when class i requires allocation of another server (due to the crossing of a forward threshold), and no servers are available in the system (i.e., all K servers have already been allocated) due to the workload of other classes. Therefore, in general, there is a nonzero probability that class i , which has already 1) allocated s_i servers in the case of $\mathcal{M}^a$ or 2) expected to be allocated/activated s_i servers in the case of $\mathcal{M}^n$, is not able to add a server upon the forward threshold crossing. Let us refer to this as a “blocking” probability P_{i,s_i} , which *approximately* captures this interaction between classes. Note that, $s_i = l$ (number of allocated servers for class i) in $\mathcal{M}^a$ while $s_i = l$ (number of expected to be allocated/activated servers for class i) in $\mathcal{M}^n$.

2. In the remainder of the paper, we use $\mathcal{M}$ to represent either $\mathcal{M}^a$ or $\mathcal{M}^n$, for simplicity of exposition.

We now formally describe our iterative approach. Let $\mathcal{M}_i^{(n)}$ be the Markovian process corresponding to the individual class i model at iteration n with a corresponding steady state probability vector $\tilde{\pi}_i^{(n)}$. The parameters of each $\mathcal{M}_i^{(n)}$ are computed as a function of blocking probabilities, $P_i^{(n)} = \{P_{i,1}^{(n)}, P_{i,2}^{(n)}, \dots, P_{i,K_i}^{(n)}\}$, which are in turn computed as a function of the steady state probability vector, $\tilde{\pi}_i^{(n-1)}$, obtained during the previous iteration. (We give the details of the construction of $\mathcal{M}_i^{(n)}$ and the computation of $\tilde{\pi}_i^{(n)}$ below.³) Then, a high level description of our iterative approach is as follows (a more detailed and formal description is given in Section 3.3):

1. Construct $\mathcal{M}_1^{(0)}, \mathcal{M}_2^{(0)}, \dots, \mathcal{M}_N^{(0)}$; set $n = 0$ (this is iteration 0).
2. Solve $\mathcal{M}_1^{(n)}, \mathcal{M}_2^{(n)}, \dots, \mathcal{M}_N^{(n)}$, i.e., compute the corresponding steady state probabilities to obtain $\tilde{\pi}_1^{(n)}, \tilde{\pi}_2^{(n)}, \dots, \tilde{\pi}_N^{(n)}$; set $n = n + 1$.
3. Use these steady state probabilities to compute $P_1^{(n)}, P_2^{(n)}, \dots, P_N^{(n)}$.
4. Use these blocking probabilities to update the individual class models, i.e., construct $\mathcal{M}_1^{(n)}, \mathcal{M}_2^{(n)}, \dots, \mathcal{M}_N^{(n)}$, where for each $i = 1, \dots, N$, parameters of $\mathcal{M}_i^{(n)}$ are computed as functions of $P_i^{(n)}$ (but not $P_j^{(n)}$ where $j \neq i$).
5. Continue the iterative process (i.e., go back to step 2) until the values of all P_i 's converge.

3.2 Individual Class Model

Since our iterative approach involves solution of individual class models ($\mathcal{M}_i$ s), we now briefly describe the class i model, which can be defined as follows: We have K_i servers each with an exponential service rate μ_i . Customer arrivals are governed by a Poisson process with rate $\lambda_i = \alpha_i \lambda$. The addition and removal of servers is governed by the forward and the reverse threshold vectors, namely, $F_i = [F_i(1), F_i(2), \dots, F_i(K_i - 1)]$ and $R_i = [R_i(1), R_i(2), \dots, R_i(K_i - 1)]$, where $R_i(j) < F_i(j)$ and $1 \leq j \leq K_i - 1$. And, server activation time is exponentially distributed with rate β_i .

3.2.1 Individual Class Model for $\mathcal{M}^a$

Given a K_i -server *single* class threshold-based queueing system with hysteresis control and constraint vector a_i , we model it as a Markov process $\mathcal{M}_i$ with the following state space $\mathcal{S}_i$:

$$\mathcal{S}_i = \{(k, j, l) | k \geq 0; j, l \in \{1, 2, \dots, K_i\}, l \geq j, F_i(l) \leq k < F_i(l) + a_i^l, 1 \leq l < K_i\},$$

where k is the number of customers in the class i queueing system, j is the number of busy (active) servers, and l is the number of allocated, not all of which may currently be activated due to the *noninstantaneous* nature of server activation in our model. Fig. 1 illustrates the state transition diagram for such a system where $K_i = 2$. Formally, the transition structure of $\mathcal{M}_i$ can be specified as shown in

3. Note that there are multiple approaches to constructing $\mathcal{M}_i^{(0)}$ s, i.e., multiple ways to start the iteration; we give details of one such approach below.

Fig. 1. State transition diagram of $\mathcal{M}^a$ for a class i system with $K_i = 2$.TABLE 1
Description of State Transition for $\mathcal{M}^a$

Next State	Rate	Condition for transition
$(k+1, j, l)$	λ_i	$(1 \leq l < K_i) \wedge (k < F_i(l))$
$(k+1, j, l)$	λ_i	$l = K_i$
$(k+1, j, l)$	$\lambda_i P_{i,l}$	$(1 \leq l < K_i) \wedge (F_i(l) \leq k < F_i(l) + a_i^l)$
$(k+1, j, l+1)$	$\lambda_i (1 - P_{i,l})$	$(1 \leq j < K_i) \wedge (F_i(l) \leq k < F_i(l) + a_i^l)$
$(k-1, j, l)$	$j\mu_i$	$(k \geq 1) \wedge (1 < l \leq K_i) \wedge (k-1 > R_i(l-1))$
$(k-1, \min(j, l-1), l-1)$	$j\mu_i$	$(k \geq 1) \wedge (1 < l \leq K_i) \wedge (k-1 = R_i(l-1))$
$(k-1, j, l)$	μ_i	$(l = j = 1) \wedge (k \geq 1)$
$(k, j+1, l)$	$(l-j)\beta_i$	$(l > j)$

(11)

Fig. 2. State transition diagram of $\mathcal{M}^n$ for a class i system with $K_i = 2$.

Table 1,⁴ where all transitions are from state (k, j, l) , with the state description given above.

3.2.2 Individual Class Model for $\mathcal{M}^n$

Given a K_i -server single class threshold-based queueing system with hysteresis control, we model it as a Markov process $\mathcal{M}_i$ with the following state space $\mathcal{S}_i$:

$$\mathcal{S}_i = \{(k, j, l) | k \geq 0; j, l \in \{1, 2, \dots, K_i\}, l \geq j\},$$

where k is the number of customers in the class i queueing system, j is the number of busy (active) servers, and l is the

4. Note that the transition rates described here are a function of the blocking probabilities, $P_{i,l}$, which change from iteration to iteration, as outlined above; however, for simplicity of notation, we do not indicate the iteration step number in the description of the transition structure of a class i model.

number of servers “expected” to be allocated and activated. Fig. 2 illustrates the state transition diagram for such a system where $K_i = 2$. Formally, the transition structure of $\mathcal{M}_i$ can be specified as follows, where all transitions are from state (k, j, l) , with the state description given in Table 1.

Let us now proceed to a more detailed description of the iterative solution technique for the multiclass system. We do this under the assumption that, given P_i , we know how to construct $\mathcal{M}_i$ (using Tables 1 and 2) and compute $\tilde{\pi}_i$, the steady state probability vector corresponding to $\mathcal{M}_i$. The procedure for computing $\tilde{\pi}_i$, is given in Section 3.5.

3.3 Iterative Computation

In this section, we describe the framework for the iterative procedure. This iterative procedure is similar to our work in

TABLE 2
Description of State Transition for $\mathcal{M}^n$

Next State	Rate	Condition
$(k+1, j, l)$	λ_i	$(1 \leq l < K_i) \wedge (k < F_i(l))$
$(k+1, j, l)$	λ_i	$(l = K_i)$
$(k+1, j, l+1)$	λ_i	$(1 \leq j < K_i) \wedge k = F_i(l)$
$(k-1, j, l)$	$j\mu_i$	$(k \geq 1) \wedge (1 < l \leq K_i) \wedge (k-1 > R_i(l-1))$
$(k-1, \min(j, l-1), l-1)$	$j\mu_i$	$(k \geq 1) \wedge (1 < l \leq K_i) \wedge (k-1 = R_i(l-1))$
$(k-1, j, l)$	μ_i	$(l = j = 1) \wedge (k \geq 1)$
$(k, j+1, l)$	$(1 - \mathcal{P}_{i,j})\beta_i$	$(l > j)$

(12)

[4] but we extend it to handle the case wherein the server activation event is noninstantaneous. First, note that in general, there are two cases to consider here:

Case 1: $\sum_{i=1}^N K_i \leq K$; that is, we have a “trivial” case, where the classes do not interfere with each other, and we can solve each individual class model once (i.e., no need for iteration) using the procedure given in Section 3.5 with $\mathcal{P}_{i,s_i} = 0, \forall i, s_i$.

Case 2: $\sum_{i=1}^N K_i > K$, where it is possible that an attempt at server allocation for class i may fail because all K servers in the system are currently allocated. As described above, in this case a form of blocking occurs and we solve the model using the iterative approach outlined in Section 3.1 whose details are now presented below.

Note also that the main difficulty in the iterative technique outlined in Section 3.1 is in determining an appropriate procedure for computing the blocking probabilities which capture the class interaction, i.e., the probabilities that, upon a forward threshold crossing, it is not possible to allocate another server to class i . Recall that, during the n th iteration ($n \geq 0$), $\mathcal{P}_{i,s_i}^{(n)}$ is the blocking probability of class i ($1 \leq i \leq N$) to which s_i servers already have been 1) allocated in the case of $\mathcal{M}^a$ and 2) expected to be allocated/activated in the case of $\mathcal{M}^n$. Before we proceed, let us state the following definitions.

Definition 1. Let $\mathcal{X}$ and $\mathcal{Y}$ be two nonnegative random variables having values in $\{1, 2, \dots\}$ and let $\pi_{\mathcal{X}}$ and $\pi_{\mathcal{Y}}$ be their respective probability mass functions. Let $\mathcal{Z}$ be another nonnegative random variable where $\mathcal{Z} = \mathcal{X} + \mathcal{Y}$; then, $\pi_{\mathcal{Z}} = \pi_{\mathcal{X}} \otimes \pi_{\mathcal{Y}}$ where $\otimes$ is the convolution operator.

Definition 2. Let $\mathcal{X}$ be a nonnegative random variable having values in $\{1, 2, \dots\}$ and let $\pi_{\mathcal{X}}$ be its probability mass function. Let

$$\mathcal{X}' = \begin{cases} \mathcal{X} & \text{if } L_1 \leq \mathcal{X} \leq L_2 \\ 0 & \text{otherwise.} \end{cases}$$

Then, the probability mass function of $\mathcal{X}'$, denoted by $\pi_{\mathcal{X}'}$, is equal to $g(\pi_{\mathcal{X}}, L_1, L_2)$ where function g is defined such that:

$$g(\pi_{\mathcal{X}}, L_1, L_2)[k] = \pi_{\mathcal{X}'}[k] = \begin{cases} \frac{\pi_{\mathcal{X}}[k]}{\sum_{m=L_1}^{L_2} \pi_{\mathcal{X}}[m]} & \text{if } L_1 \leq k \leq L_2 \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Let $\tilde{\pi}_i^{(n)}[k, j, l]$ be the steady state probability of class i having k customers ($k \geq 0$) in the system with j activated servers and l target server allocations (with $1 \leq j \leq l \leq K_i$), computed during the n th iteration. Let $\pi_i^{(n)}$ denote the steady state

probability vector of the number of servers allocated to class i , where $\pi_i^{(n)}[l]$ denotes the steady state probability of l servers having been target allocated to class i , as computed during the n th iteration. Thus, we have:

$$\pi_i^{(n)}[l] = \sum_{j=1}^{K_i} \sum_k \tilde{\pi}_i^{(n)}[k, j, l]. \quad (14)$$

Finally, let $Q_i^{(n)}$ be the transition rate matrix corresponding to the class i model $\mathcal{M}_i^{(n)}$, during the n th iteration, which is computed using the transition structure of $\mathcal{M}_i^{(n)}$ given in Table 1 and $\mathcal{P}_{i,s_i}^{(n-1)}$, where $1 \leq s_i \leq K_i - 1$. Then, the iterative procedure is as follows:

1. *Initialization step:* set $n = 0$ and set $\mathcal{P}_{i,s_i}^{(0)} = 0$ for $1 \leq s_i < K_i$. Given these initial values of blocking probabilities, for each class i , we can construct $Q_i^{(0)}$ using the transition structure given in (11) and then compute $\tilde{\pi}_i^{(0)}$ using the procedure given in Section 3.5. Once we compute the steady state probability vector $\tilde{\pi}_i^{(0)}$ for each class i , we can then compute their respective server allocation probability vectors, $\pi_i^{(0)}$ s, using (14). The $\pi_i^{(0)}$ s are in turn needed in the computation of the blocking probabilities, $\mathcal{P}_{i,s_i}^{(1)}$ s (step 2 below).
2. *Updating of blocking probabilities step:* $n = n + 1$, and

$$\mathcal{P}_{i,s_i}^{(n)} = \begin{cases} 0 & \text{if } K \geq \sum_{j=1}^N K_j \\ 0 & \text{if } K - s_i > \sum_{j=1, j \neq i}^N K_j \\ \Gamma(i, s_i, n) & \text{otherwise.} \end{cases} \quad (15)$$

The first condition in (15) indicates that the system has a sufficient number of servers for all classes (we include this for completeness). The second condition indicates that the system has sufficient resources to allocate at least one more server to class i without affecting the maximum possible server allocation of other classes. In the last condition, the Γ function is used to compute the blocking probability, at iteration n , for class i which has s_i servers already 1) allocated to it in the case of $\mathcal{M}^a$ or 2) expected to be allocated/activated in the case of $\mathcal{M}^n$.

$\Gamma(i, s_i, n)$ can be computed as follows: Let $\mathcal{A}_m(i, s_i, n)$ be the random variable, at iteration n , denoting server allocation of class m , when class i has been 1) allocated s_i servers in the case of $\mathcal{M}^a$ or 2) expected to be allocated/activated s_i servers in the

case of $\mathcal{M}^n$. Let $\Upsilon_m(i, s_i, n)$ be the probability mass function of $\mathcal{A}_m(i, s_i, n)$. Then, we have:

$$\Upsilon_m(i, s_i, n) = g(\pi_m^{(n-1)}, 1, L_m) \quad (16)$$

for $m = \{1, 2, \dots, i-1, i+1, \dots, N\}$, where function g is defined through (13) and L_m is as follows:

$$L_m = \begin{cases} K_m & \text{if } K - s_i - (N - 2) \geq K_m \\ K - s_i - (N - 2) & \text{otherwise} \end{cases} \quad (17)$$

and $\pi_m^{(n-1)}$ in (16) is computed using (14). The normalization in (16) is used to account for the fact that if we know that the system already 1) allocated s_i servers to class i in the case of $\mathcal{M}^n$ or 2) expected to be allocated/activated s_i servers to class i in the case of $\mathcal{M}^n$, then the system only has $(K - s_i)$ servers remaining. Out of these $(K - s_i)$ remaining servers, the system needs to allocate $(N - 2)$ to customers that are neither in class i nor in class m (i.e., the system allocates at least one server to each class). Therefore, if the system potentially has at least K_m available servers, then $\mathcal{A}_m(i, s_i, n)$ can have values in $\{1, \dots, K_m\}$; otherwise, the random variable $\mathcal{A}_m(i, s_i, n)$ can only take on values in $\{1, 2, \dots, K - s_i - (N - 2)\}$. Let $\mathcal{B}(i, s_i, n)$ be a non-negative random variable, at iteration n , denoting the server allocation of all classes *except* class i , where class i already has s_i servers 1) allocated to it in the case of $\mathcal{M}^n$ or 2) expected to be allocated/activated in the case of $\mathcal{M}^n$. Let $\Psi(i, s_i, n)$ be the probability mass function of $\mathcal{B}(i, s_i, n)$. Then, we have:

$$\Psi(i, s_i, n) = g((\Upsilon_1(i, s_i, n) \otimes \dots \otimes \Upsilon_{i+1}(i, s_i, n) \dots \otimes \Upsilon_N(i, s_i, n)), N - 1, K - s_i). \quad (18)$$

The normalization in (18) is used to account for the fact that if the system has already 1) allocated s_i servers to class i in the case of $\mathcal{M}^n$, or 2) expected to be allocated/activated s_i servers to class i in the case of $\mathcal{M}^n$, then the number of servers that have been allocated to other classes can only range in $\{N - 1, N, \dots, K - s_i\}$.

Last, $\Gamma(i, s_i, n)$, the function used to compute blocking probabilities, at iteration n , corresponding to class i with 1) s_i allocated servers in the case of $\mathcal{M}^n$ or 2) s_i expected to be allocated/activated servers in the case of $\mathcal{M}^n$ is:

$$\Gamma(i, s_i, n) = \Psi(i, s_i, n; K - s_i), \quad (19)$$

where $\Psi(i, s_i, n; K - s_i) = \text{Prob}[\mathcal{B}(i, s_i, n) = K - s_i]$ and $\Psi(i, s_i, n)$ is computed using (18).

3. *Updating of individual class models step:* Given the blocking probabilities $\mathcal{P}_{i,s_i}^{(n)}$ of class i in (15), we can compute the new rate matrix $\mathcal{Q}_i^{(n)}$ (based on the transition structure given in Table 1) and then compute the corresponding steady state probabilities $\tilde{\pi}_i^{(n)}$ (using the procedure given in Section 3.5) as well as $\pi_i^{(n)}$, the probability vector of server allocation of class i (using (14)). (The $\pi_i^{(n)}$ s will in turn be needed in the updating of the blocking probabilities, $\mathcal{P}_{i,s_i}^{(n+1)}$'s (step 2 above).)

4. *Test of convergence step:* If $|\mathcal{P}_{i,s_i}^{(n)} - \mathcal{P}_{i,s_i}^{(n-1)}| \leq \epsilon$ for each class i , $1 \leq i \leq N$, and each s_i , $1 \leq s_i \leq K_i - 1$, then stop. Otherwise, go to step 2 and continue iterating.

3.4 Computation of Performance Measures

In this section, we briefly discuss computation of performance measures. Due to lack of space, we only present the derivation for model $\mathcal{M}^n$ and we could use the same approach for the derivation for model $\mathcal{M}^a$. Given the steady state probabilities $\tilde{\pi}_i, i = 1, \dots, N$, computed using the iterative approach described above, we can compute various performance measures of interest. More specifically, for each class i we can compute performance measures which can be expressed in the form of a Markov reward function, $\mathcal{R}_i$, where

$$\mathcal{R}_i = \sum_{k,j,l} \tilde{\pi}_i[k, j, l] R_i(k, j, l)$$

and $R_i(k, j, l)$ is the reward for state (k, j, l) of class i . Some useful performance measures include:

1. expected number of customers of class i ,
2. expected response time for customers of class i ,
3. probability of dropping a customer of class i upon its arrival,
4. throughput of class i customers, and so on.

For instance, let $E[N_i]$ and $E[T_i]$ denote the expected number of customers and the expected response time, respectively, of the class i model, corresponding to the Markov process $\mathcal{M}_i$. Then, $E[N_i]$ can be expressed as $\sum_{k,j,l} k \tilde{\pi}_i[k, j, l]$. Of course, using Little's result [15], we have $E[T_i] = \frac{1}{\lambda_i^*} E[N_i]$, where λ_i^* is the class i throughput. To compute λ_i^* , we need to account for the customers that are dropped from the system (see Section 2). Hence, $\lambda_i^* = \lambda_i (1 - \sum_{l=1}^{K_i} \sum_{j=1}^l \mathcal{P}_{i,j} \tilde{\pi}_i[F_i(j) + a_{i,j}^j, j, l])$.

We believe that the more interesting performance measures are those computed on a per class basis, since a useful part of studying performance of multiclass threshold-based systems is to discover the effect that the various classes have on one another. Therefore, we have concentrated on per class performance measures here. However, we can also use these to compute overall system performance measures, for instance, as a weighted average of the individual class performance measures. For example, we can compute the expected system response time, $E[T]$, as follows:

$$E[T] = \frac{\lambda_1^*}{\lambda^*} E[T_1] + \frac{\lambda_2^*}{\lambda^*} E[T_2] + \dots + \frac{\lambda_N^*}{\lambda^*} E[T_N],$$

where $\lambda^* = \sum_{i=1}^N \lambda_i^*$.

3.5 Analysis of the Individual Class Model

In this section, we briefly summarize the solution technique for the individual class model which was defined in Section 3.2. Specifically, we use the single class solution technique we derived in [16] with some modifications needed to account for the structure of the multiclass model. Since these modification are mostly straightforward, we only summarize the solution technique in this section.

The general approach is as follows: As already stated, we model the class i queueing system as a Markov process, $\mathcal{M}_i$, where: 1) the main goal is to compute the steady state

probabilities of the Markov process and use these to compute various performance metrics of interest and 2) the main difficulty is that the Markov process is infinite (see Section 3.2) and, thus, “difficult” to solve using a “direct” approach.⁵

As is often done in these cases, we need to look for special structure that might exist in the Markov process; specifically, we take advantage of the stochastic complementation technique [17]. The basic approach to computing the steady state probabilities of the Markov process and the corresponding performance measures is as follows: First, we construct an upper bound model, $\mathcal{M}_i^u$, for the original Markov process $\mathcal{M}_i$, while trying to satisfy the criteria that the new model will: 1) provide (hopefully a tight) upper bound on the desired performance measures and 2) be a “simpler” model to solve. Therefore, the upper bound model transitions that replace the original transition can be specified as follows: In Tables 1 and 2, we replace $(k-1, \min(j, l-1), l-1)$ with $(k-1, 1, l-1)$, where the transition rate is $j\mu_i$. Next, we partition the state space of the original Markov process $\mathcal{M}_i$ ⁶ into disjoint sets. Using the concept of stochastic complementation, for each set, we compute the conditional steady state probability vector, given that the original Markov process $\mathcal{M}_i$ is in that set. (A relatively simple construction of the stochastic complement is possible due to the special structure that exists in the individual class models; specifically, we exploit the “single entry” structure as in [3].) By applying the state aggregation technique [1], we aggregate each set into a single state and then compute the steady state probabilities for the aggregated process, i.e., the probabilities of the system being in any given set. Last, we apply the disaggregation technique [1] to compute the individual (unconditional) steady state probabilities of the original Markov process $\mathcal{M}_i$. These can in turn be used to compute various performance measures of interest.

4 CONVERGENCE OF THE ITERATIVE METHOD

As described in Section 3, we break up the original model $\mathcal{M}$ into N single class Markovian submodels, i.e., $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_N$. We use a set of blocking probabilities to describe the interaction among these N Markovian submodels. In particular, the blocking probabilities P_i for class i are functions of all steady state probabilities except the steady state probability of class i , i.e., $\tilde{\pi}_i$. Therefore, we could use N groups of simultaneous equations to represent those N Markovian submodels since the blocking probabilities could be represented by the function of those steady state probabilities. Note that for each class i , the number of unknown variables is $|\tilde{\pi}_i|$ and there are $\sum_{i=0}^N |\tilde{\pi}_i|$ variables for the whole system.

Let $\mathcal{A}_i$ be the transition rate matrix for the class i except all the entries in the first column are equal to 1. In addition, b_i is a row vector, in which all the elements are 0 but the first element is 1, and its size is the same as $|\tilde{\pi}_i|$. Thus, we can use the following nonhomogeneous system of (nonlinear) equations to represent the class i Markovian submodel $\mathcal{M}_i$:

$$\tilde{\pi}_i \mathcal{A}_i = b_i.$$

5. We could consider finite versions of the model or truncation of the infinite version [12]; however, in either case, the Markov process would still be very large and the computational complexity of a direct solution for a reasonable size system still high.

6. For simplicity, we use $\mathcal{M}_i$ instead of $\mathcal{M}_i^u$ in the rest of this paper.

To use the iterative method to solve N groups of simultaneous equations, we split the matrix $\mathcal{A}_i$ into two submatrices $\mathcal{A}_{i,n}$ and $\mathcal{A}_{i,b}$, where $\mathcal{A}_{i,n}$ is the initial nonsingular, nonblocking matrix (with all blocking probabilities are equal to 0) and $\mathcal{A}_{i,b} = \mathcal{A}_i - \mathcal{A}_{i,n}$. Given a splitting with nonsingular $\mathcal{A}_{i,n}$, we have

$$\begin{aligned} \tilde{\pi}_i (\mathcal{A}_{i,n} + \mathcal{A}_{i,b}) &= b_i, \\ \tilde{\pi}_i &= b_i \mathcal{A}_{i,n}^{-1} - \tilde{\pi}_i \mathcal{A}_{i,b} \mathcal{A}_{i,n}^{-1}, \end{aligned} \quad (20)$$

which leads to our iterative procedure.

$$\tilde{\pi}_i^k = b_i \mathcal{A}_{i,n}^{-1} - \tilde{\pi}_i^{k-1} \mathcal{A}_{i,b} \mathcal{A}_{i,n}^{-1}. \quad (21)$$

We note that $\mathcal{A}_{i,b}^k$ is updated based on the steady state probabilities $\tilde{\pi}_j^{k-1}$, where $1 \leq j \leq N$ and $j \neq i$, in the previous $k-1$ round. From the above discussion, one can compute the error vector of class i , e.g., e_i^k after the k th round, we have:

$$\begin{aligned} e_i^0 &= \tilde{\pi}_i^0 - \tilde{\pi}_i = b_i \mathcal{A}_{i,n}^{-1} - (b_i \mathcal{A}_{i,n}^{-1} - \tilde{\pi}_i \mathcal{A}_{i,b} \mathcal{A}_{i,n}^{-1}) \\ &= -\tilde{\pi}_i \mathcal{A}_{i,b} \mathcal{A}_{i,n}^{-1}, \end{aligned} \quad (22)$$

$$e_i^1 = \tilde{\pi}_i^1 - \tilde{\pi}_i = \tilde{\pi}_i^1 \mathcal{A}_{i,b} \mathcal{A}_{i,n}^{-1} - \tilde{\pi}_i \mathcal{A}_{i,b} \mathcal{A}_{i,n}^{-1}, \quad (23)$$

$$e_i^2 = \tilde{\pi}_i^2 - \tilde{\pi}_i = \tilde{\pi}_i^2 \mathcal{A}_{i,b}^2 \mathcal{A}_{i,n}^{-1} - \tilde{\pi}_i \mathcal{A}_{i,b} \mathcal{A}_{i,n}^{-1}, \quad (24)$$

⋮

$$e_i^k = \tilde{\pi}_i^k - \tilde{\pi}_i = \tilde{\pi}_i^k \mathcal{A}_{i,b}^k \mathcal{A}_{i,n}^{-1} - \tilde{\pi}_i \mathcal{A}_{i,b} \mathcal{A}_{i,n}^{-1}. \quad (25)$$

If our iterative method is able to converge, the sufficient condition is that e_k becomes very close to 0 after finite k iterations. Since both $\tilde{\pi}_i^k$ and $\tilde{\pi}_i$ are stationary probability vectors, we can see that as $\mathcal{A}_{i,b}^k$ approaches to the actual $\mathcal{A}_{i,b}$, e_i^k will approach 0. To illustrate the above argument, we give the proof of a 2-class system and show the convergence of the proposed iterative method.

Example. Consider a 2-class, 3-server threshold-based system as follows: $K = 3$, $K_1 = K_2 = 2$, F_1 , R_1 , F_2 , a , and R_2 . We assume that there exists a stationary distribution for this system.

Let $\tilde{\pi}_{i,j}$ be the steady state probability for class i with j servers and then $\tilde{\pi}_1 = (\tilde{\pi}_{1,1}, \tilde{\pi}_{1,2})$, and $\tilde{\pi}_2 = (\tilde{\pi}_{2,1}, \tilde{\pi}_{2,2})$. According to our method to compute the blocking probabilities, we have $\mathcal{P}_{1,1} = \tilde{\pi}_{2,2}$ and $\mathcal{P}_{2,1} = \tilde{\pi}_{1,2}$. Therefore, the goal is to solve two groups of nonhomogeneous of linear equations, i.e., $\tilde{\pi}_1 \mathcal{A}_1 = b_1$ and $\tilde{\pi}_2 \mathcal{A}_2 = b_2$, where $\mathcal{P}_{1,1} = \tilde{\pi}_{2,2}$ and $\mathcal{P}_{2,1} = \tilde{\pi}_{1,2}$.

- **Step 0:** We set $\mathcal{P}_{1,1}^0 = 0$ and $\mathcal{P}_{2,1}^0 = 0$.
 $\mathcal{P}_{1,1}^0 = 0 < \mathcal{P}_{1,1} \rightarrow \tilde{\pi}_{1,2}^0 > \tilde{\pi}_{1,2} = \mathcal{P}_{2,1}$ and $\tilde{\pi}_{1,1}^0 < \tilde{\pi}_{1,1}$.
 $\mathcal{P}_{2,1}^0 = 0 < \mathcal{P}_{2,1} \rightarrow \tilde{\pi}_{2,2}^0 > \tilde{\pi}_{2,2} = \mathcal{P}_{1,1}$ and $\tilde{\pi}_{2,1}^0 < \tilde{\pi}_{2,1}$.
- **Step 1:** We set $\mathcal{P}_{1,1}^1 = \tilde{\pi}_{2,2}^0$ and $\mathcal{P}_{2,1}^1 = \tilde{\pi}_{1,2}^0$.
 $\mathcal{P}_{1,1}^1 > \mathcal{P}_{1,1} \rightarrow \tilde{\pi}_{1,2}^1 < \tilde{\pi}_{1,2} = \mathcal{P}_{2,1}$ and $\tilde{\pi}_{1,1}^1 > \tilde{\pi}_{1,1}$.
 $\mathcal{P}_{2,1}^1 > \mathcal{P}_{2,1} \rightarrow \tilde{\pi}_{2,2}^1 < \tilde{\pi}_{2,2} = \mathcal{P}_{1,1}$ and $\tilde{\pi}_{2,1}^1 > \tilde{\pi}_{2,1}$.
- **Step 2:** We set $\mathcal{P}_{1,1}^2 = \tilde{\pi}_{2,2}^1$ and $\mathcal{P}_{2,1}^2 = \tilde{\pi}_{1,2}^1$.
 $\mathcal{P}_{1,1}^2 < \mathcal{P}_{1,1} \rightarrow \tilde{\pi}_{1,2}^2 > \tilde{\pi}_{1,2} = \mathcal{P}_{2,1}$ and $\tilde{\pi}_{1,1}^2 < \tilde{\pi}_{1,1}$.
 $\mathcal{P}_{2,1}^2 < \mathcal{P}_{2,1} \rightarrow \tilde{\pi}_{2,2}^2 > \tilde{\pi}_{2,2} = \mathcal{P}_{1,1}$ and $\tilde{\pi}_{2,1}^2 < \tilde{\pi}_{2,1}$.
- **Step 3:** We set $\mathcal{P}_{1,1}^3 = \tilde{\pi}_{2,2}^2$ and $\mathcal{P}_{2,1}^3 = \tilde{\pi}_{1,2}^2$.
 $\mathcal{P}_{1,1}^3 > \mathcal{P}_{1,1} \rightarrow \tilde{\pi}_{1,2}^3 < \tilde{\pi}_{1,2} = \mathcal{P}_{2,1}$ and $\tilde{\pi}_{1,1}^3 > \tilde{\pi}_{1,1}$.
 $\mathcal{P}_{2,1}^3 > \mathcal{P}_{2,1} \rightarrow \tilde{\pi}_{2,2}^3 < \tilde{\pi}_{2,2} = \mathcal{P}_{1,1}$ and $\tilde{\pi}_{2,1}^3 > \tilde{\pi}_{2,1}$.
- **Step 2k:** We set $\mathcal{P}_{1,1}^{2k} = \tilde{\pi}_{2,2}^{2k-1}$ and $\mathcal{P}_{2,1}^{2k} = \tilde{\pi}_{1,2}^{2k-1}$.
 $\mathcal{P}_{1,1}^{2k} < \mathcal{P}_{1,1} \rightarrow \tilde{\pi}_{1,2}^{2k} > \tilde{\pi}_{1,2} = \mathcal{P}_{2,1}$ and $\tilde{\pi}_{1,1}^{2k} < \tilde{\pi}_{1,1}$.
 $\mathcal{P}_{2,1}^{2k} < \mathcal{P}_{2,1} \rightarrow \tilde{\pi}_{2,2}^{2k} > \tilde{\pi}_{2,2} = \mathcal{P}_{1,1}$ and $\tilde{\pi}_{2,1}^{2k} < \tilde{\pi}_{2,1}$.

- **Step $2k + 1$:** We set $\mathcal{P}_{1,1}^{2k+1} = \tilde{\pi}_{2,2}^{2k}$ and $\mathcal{P}_{2,1}^{2k+1} = \tilde{\pi}_{1,2}^{2k}$.
 $\mathcal{P}_{1,1}^{2k+1} > \mathcal{P}_{1,1} \rightarrow \tilde{\pi}_{1,2}^{2k+1} < \tilde{\pi}_{1,2} = \mathcal{P}_{2,1}$ and $\tilde{\pi}_{1,1}^{2k+1} > \tilde{\pi}_{1,1}$.
 $\mathcal{P}_{2,1}^{2k+1} > \mathcal{P}_{2,1} \rightarrow \tilde{\pi}_{2,2}^{2k+1} < \tilde{\pi}_{2,2} = \mathcal{P}_{1,1}$ and $\tilde{\pi}_{2,1}^{2k+1} > \tilde{\pi}_{2,1}$.

Next, we would like to prove the following inequality:

$$\mathcal{P}_{i,1}^{2k} < \mathcal{P}_{i,1}^{2k+2} < \mathcal{P}_{i,1} < \mathcal{P}_{i,1}^{2k+3} < \mathcal{P}_{i,1}^{2k+1}, \quad (26)$$

where $i = 1$ or 2 and $k \geq 0$.

One can use the mathematical induction to prove the above inequality.

1. When $k = 0$, it is trivial to verify the inequality is correct for class 1 or class 2.
2. Assume $k = m$, the inequality holds for class 1 and class 2.
3. As $k = m + 1$, for class 1, we have

$$\mathcal{P}_{2,1}^{2m-1} > \mathcal{P}_{2,1}^{2m+1} \rightarrow \mathcal{P}_{1,1}^{2m} = \tilde{\pi}_{2,2}^{2m-1} < \tilde{\pi}_{2,2}^{2m+1} = \mathcal{P}_{1,1}^{2m+2},$$

that is, we use the higher blocking probability for class 2, we get lower state probability $\tilde{\pi}_{2,2}^{2m-1}$, which is the blocking probability of class 1 for the next round.

Similarly, we can prove another side of the inequality for class 1: $\mathcal{P}_{2,1}^{2m} < \mathcal{P}_{2,1}^{2m+2} \rightarrow \mathcal{P}_{1,1}^{2m+3} < \mathcal{P}_{1,1}^{2m+1}$. Of course, the same approach can be used to prove the inequality for class 2.

Once we show that (26) holds, it is not difficult to show that $\mathcal{A}_{i,b}^k$ approaches $\mathcal{A}_{i,b}$ as k increases, i.e., e_i^k is close to 0.

In the following, we give a proof of the convergence of the iterative method for a 2-class multiserver system.

Lemma 1. Consider a 2-class, K -server threshold-based system as follows: $K, K_1, K_2, F_1, R_1, F_2, a$, and R_2 . Assume that there exists a stationary state distribution for the system, the proposed iterative method for solving this system will converge to the steady state distribution.

Proof. Let $\tilde{\pi}_{i,j}$ be the steady state probability for class i with allocated j servers, we have $\tilde{\pi}_1 = (\tilde{\pi}_{1,1}, \tilde{\pi}_{1,2}, \dots, \tilde{\pi}_{1,K_1})$, and $\tilde{\pi}_2 = (\tilde{\pi}_{2,1}, \tilde{\pi}_{2,2}, \dots, \tilde{\pi}_{2,K_2})$. According to our method to compute the blocking probabilities, for class 1, we have $\mathcal{P}_{1,1} = \dots = \mathcal{P}_{1,K-K_2-1} = 0$ and

$$\mathcal{P}_{1,x} = \frac{\tilde{\pi}_{2,K-x}}{\sum_{m=1}^{K-x} \tilde{\pi}_{2,m}},$$

where $K - K_2 \leq x \leq K - 2$. Similarly, for class 2 we get $\mathcal{P}_{2,1} = \dots = \mathcal{P}_{2,K-K_1-1} = 0$ and

$$\mathcal{P}_{2,y} = \frac{\tilde{\pi}_{1,K-y}}{\sum_{m=1}^{K-y} \tilde{\pi}_{1,m}},$$

where $K - K_1 \leq y \leq K - 2$. Therefore, our goal is to solve two groups of nonhomogeneous of linear equations, i.e., $\tilde{\pi}_1 \mathcal{A}_1 = b_1$ and $\tilde{\pi}_2 \mathcal{A}_2 = b_2$, where $\mathcal{P}_{1,x}$ and $\mathcal{P}_{2,y}$ are defined above.

One can use the mathematical induction to prove the following inequalities for both classes:

$$\mathcal{P}_{1,x}^{2k} < \mathcal{P}_{1,x}^{2k+2} < \mathcal{P}_{1,x} < \mathcal{P}_{1,x}^{2k+3} < \mathcal{P}_{1,x}^{2k+1}, \quad (27)$$

where $K - K_2 \leq x \leq K - 2$ and $k \geq 0$,

$$\mathcal{P}_{2,y}^{2k} < \mathcal{P}_{2,y}^{2k+2} < \mathcal{P}_{2,y} < \mathcal{P}_{2,y}^{2k+3} < \mathcal{P}_{2,y}^{2k+1}, \quad (28)$$

where $K - K_1 \leq y \leq K - 2$ and $k \geq 0$.

When the inequalities of (27) and (28) hold, it is straightforward to show that $\mathcal{A}_{i,b}^k$ will approach to $\mathcal{A}_{i,b}$ as k increases, i.e., e_i^k is close to 0. $\square$

For the general N -class K -server system, we note that the following inequality still holds for any blocking probability $\mathcal{P}_{i,j}$, i.e., $\mathcal{P}_{i,j}^{2k} < \mathcal{P}_{i,j}^{2k+2} < \mathcal{P}_{i,j} < \mathcal{P}_{i,j}^{2k+3} < \mathcal{P}_{i,j}^{2k+1}$, where $k \geq 0$. In other words, the above equation still holds under the operation of the convolution operator $\otimes$ and the function g defined in Section 3. Due to the lack of space, we do not give the detailed steps of the proof here. We note that the convergence rate of our iterative method is still an open question, which needs further study and is beyond the scope of this work.

5 NUMERICAL EXAMPLES AND VALIDATION OF APPROXIMATION

In this section, we present numerical results to illustrate

1. the accuracy of our iterative methodology as compared with simulation,
2. the effect of server activation rate on performance measures,
3. the benefits of resource sharing among heterogeneous workload classes, and
4. the effects of threshold values on performance measures.

Since checking the accuracy of our iterative methods is done by comparing them with simulations, all simulation results are given with at least $95\% \pm 5\%$ confidence. In all experiments presented here, our iterative approach uses $\epsilon = 0.0000001$ (refer to Section 3 for details). Moreover, the computation time of our iterative methods is more than two orders of magnitude faster than that of simulations. This empirical evidence indicates our iterative method is fast and fairly accurate.

5.1 Experiment A: Accuracy of the Iterative Method

To compare the accuracy of our iterative method in computing performance measures of the Markovian system $\mathcal{M}$, we also simulate $\mathcal{M}$, i.e., for the purpose of validating the proposed solution technique. We use mean response time for class i , $1 \leq i \leq N$, as our main performance metric of interest. Parameter settings for all test cases in Experiment A are listed in Table 3.

For each test case under Experiment A, we assume all classes have the same service rate $\mu = 1$. The server activation rate is the same for all classes, and we vary this activation rate. In particular, we define $\mathcal{R} = \beta/\mu$ (where we drop the class notation for simplicity of exposition). Fig. 3 illustrates the mean response time under test case 1 for $\mathcal{R} = 1$ and 10.0 . It is easy to observe that the difference between the two results is small (e.g., in the case of Fig. 3, the largest difference is $\approx 7\%$). Since such small differences are difficult to visualize using graphs, we present the remainder of the accuracy related experiments using tables.

Tables 4, 5, 6, and 7 to illustrate several other experiments of validating the accuracy of our technique. Due to the large size of the tables, we only give the iterative result and the percentage error. In all cases, the percentage error (%E) is defined as:

TABLE 3
Parameter Settings for Experiments

Case 1 (for M^n) Exp. A	$K = 10, K_1=K_2=K_3=4, \alpha_1 : \alpha_2 : \alpha_3 = 1 : 3 : 6,$ $F_1 = [4,8,12], R_1 = [2,6,10], F_2 = [8,12,16], R_2 = [5,9,13], F_3 = [6,10,14], R_3 = [3,7,11],$ $\mu_1=\mu_2=\mu_3=1.0,$
Case 2 (for M^n) Exp. A	$K = 12, K_1=K_2=K_3=3, K_4=5; \alpha_1 : \alpha_2 : \alpha_3 : \alpha_4 = 1 : 1 : 1 : 2,$ $F_1 = [6,10], R_1 = [4,7], F_2 = [4,8], R_2 = [2,4], F_3 = [8,12], R_3 = [6,9], F_4 = [5,9,13,17], R_4 = [3,6,9,12]$ $\mu_1=\mu_2=\mu_3=\mu_4=1.0,$
Case 3 (for M^n) Exp. A	$K = 12, K_1=K_2=K_3=3, K_4=5; \alpha_1 : \alpha_2 : \alpha_3 : \alpha_4 = 1 : 2 : 3 : 4,$ $F_1 = [6,10], R_1 = [4,7], F_2 = [4,8], R_2 = [2,4], F_3 = [8,12], R_3 = [6,9], F_4 = [5,9,13,17], R_4 = [3,6,9,12]$ $\mu_1=\mu_2=\mu_3=\mu_4=1.0,$
Case 4 (for M^a) Exp. A	$K = 12, K_1=K_2=K_3=3, K_4=5; \alpha_1 : \alpha_2 : \alpha_3 : \alpha_4 = 1 : 1 : 1 : 2,$ $F_1 = [6,10], R_1 = [4,7], F_2 = [4,8], R_2 = [2,4], F_3 = [8,12], R_3 = [6,9], F_4 = [5,9,13,17], R_4 = [3,6,9,12]$ $a_1 = [3,3], a_2 = [3,3], a_3 = [3,3], a_4 = [3,3,3,3] \mu_1=\mu_2=\mu_3=\mu_4=1.0,$
Case 5 (for M^a) Exp. A & C	$K = 12, K_1=K_2=K_3=3, K_4=5; \alpha_1 : \alpha_2 : \alpha_3 : \alpha_4 = 1 : 2 : 3 : 4,$ $F_1 = [6,10], R_1 = [4,7], F_2 = [4,8], R_2 = [2,4], F_3 = [8,12], R_3 = [6,9], F_4 = [5,9,13,17], R_4 = [3,6,9,12]$ $a_1 = [3,3], a_2 = [3,3], a_3 = [3,3], a_4 = [3,3,3,3] \mu_1=\mu_2=\mu_3=\mu_4=1.0,$
Case 6 (for M^a) Exp. B	$K = 8, K_1=K_2=3, K_3=4, \alpha_1 : \alpha_2 : \alpha_3 = 1 : 1 : 1,$ $F_1 = [6,10], R_1 = [4,7], F_2 = [6,10], R_2 = [4,7], F_3 = [6,10,14], R_3 = [4,7,10],$ $a_1 = [3,3], a_2 = [3,3], a_3 = [3,3,3], \mu_1=\mu_2=\mu_3=1.0,$
Case 7 (for M^a) Exp. D	$K = 4, K_1=K_2=3, \alpha_1 : \alpha_2 = 2 : 3; a_1 = [3,3], a_2 = [3,3]; \mu_1=\mu_2=1.0; \mathcal{R} = 1;$ Configuration A: $F_1 = [4,8], R_1 = [2,6], F_2 = [10,20], R_2 = [5,10];$ Configuration B: $F_1 = [4,8], R_1 = [2,6], F_2 = [20,40], R_2 = [15,35];$ Configuration C: $F_1 = [4,8], R_1 = [2,6], F_2 = [40,80], R_2 = [30,70];$
Case 8 (for M^a) Exp. D	$K = 4, K_1=K_2=3, \alpha_1 : \alpha_2 = 2 : 3; a_1 = [3,3], a_2 = [3,3]; \mu_1=\mu_2=1.0; \mathcal{R} = 1;$ Configuration A: $F_1 = [4,8], R_1 = [2,6], F_2 = [6,9], R_2 = [4,7];$ Configuration B: $F_1 = [20,40], R_1 = [15,35], F_2 = [30,45], R_2 = [25,40];$ Configuration C: $F_1 = [40,80], R_1 = [30,70], F_2 = [60,90], R_2 = [50,80];$
Case 9 (for M^a) Exp. E	$K = 8, K_1=K_2=3, K_3=4, \alpha_1 : \alpha_2 : \alpha_3 = 2 : 3 : 4,$ $F_1 = [6,10], R_1 = [4,7], F_2 = [6,10], R_2 = [4,7], F_3 = [6,10,14], R_3 = [4,7,10],$ $a_1 = [3,3], a_2 = [3,3], a_3 = [3,3,3], \mu_1=\mu_2=\mu_3=1.0,$

$$\%E = \frac{|\text{simulation result} - \text{iterative result}|}{\text{simulation result}} \times 100\%.$$

As can be observed, in these experiments, we track the performance metrics closely for all classes and the maximum error is around 10 percent.

As is probably expected, in our experiments, the higher error cases corresponded to fairly high contention cases. These are also the cases that likely corresponded to “poor” designs where a reduction in contention for resources between classes is needed in order to obtain a system with good performance characteristics. In most of our experiments (some of which are presented below), the performance improvements that could be obtained, for instance, through better threshold settings, were significantly higher (percentage-wise) than the loss in accuracy due to our approximation. Hence, this is a good indication that our

iterative technique is a useful tool for fast and fairly accurate assessment of threshold-based designs that can be used, for instance, for searching for good threshold settings.

Note that, we have performed extensive experiments and those results are similar to the ones included here. In general, the percentage error in most testing cases was within 10 percent, but around 14 percent in a few cases. Next, we illustrate some of the performance tradeoffs and designs that can be studied using our technique.

5.2 Experiment B: Effect of Service Activation Rate on System Performance

In this experiment, we consider how the server activation rate may affect the average response time of different classes of customers. Let us consider test case 6 in Table 3. We vary the server activation ratio $\mathcal{R}$ from 0.1 to 100. In

Fig. 3. Experiment A—accuracy testing (Case 1): Average response time of three classes under different server activation ratios, $\mathcal{R}$. (a) Server activation ratio $\mathcal{R} = 1.0$ (b) Server activation ratio $\mathcal{R} = 10.0$.

TABLE 4
Experiment A—Accuracy Testing (Case 2): Average Response Time for Server Activation Ratio $\mathcal{R} = 1$

λ	$E[T_1]$ (iterative)	$E[T_2]$ (iterative)	$E[T_3]$ (iterative)	$E[T_4]$ (iterative)	% error (class 1)	% error (class 2)	% error (class 3)	% error (class 4)
1.000	1.249403209	1.240894331	1.249966817	1.596548157	0.188442	0.102801	0.091916	1.082684
2.00	1.632332041	1.530185205	1.659140114	2.346862510	0.589992	1.655815	0.628414	6.931096
3.000	2.137886898	1.800536338	2.327797341	2.877125939	2.607489	4.531356	1.872697	8.398753
4.000	2.655828539	2.024466859	3.203617359	3.258738327	5.961660	7.101779	4.766918	4.467737
5.000	3.086978739	2.210009611	4.010064726	3.549132931	7.808542	8.701947	6.984125	0.872970
6.000	3.414312029	2.371654188	4.556764880	3.748743137	7.871845	8.730359	6.793586	5.205084
7.000	3.666371001	2.520343773	4.875718116	3.873368054	6.090529	8.135594	5.009943	7.823094
8.000	3.876706835	2.667619678	5.069148468	3.957965747	3.398553	6.806763	2.094356	8.864857
9.000	4.076410481	2.833209795	5.216584527	4.054224035	0.768549	5.588292	0.612172	8.653662
10.000	4.294252513	3.044849701	5.367470840	4.258024598	0.163318	6.478889	0.981928	6.684466

TABLE 5
Experiment A—Accuracy Testing (Case 3): Average Response Time for Server Activation Ratio $\mathcal{R} = 1.0$

λ	$E[T_1]$ (iterative)	$E[T_2]$ (iterative)	$E[T_3]$ (iterative)	$E[T_4]$ (iterative)	% error (class 1)	% error (class 2)	% error (class 3)	% error (class 4)
1.000	1.111101334	1.240894331	1.427769462	1.596548157	0.107716	0.098113	0.223379	1.196270
2.000	1.249403209	1.530185205	2.327797341	2.346862510	0.139604	1.440456	1.633202	6.963889
3.000	1.422120871	1.800536338	3.633427552	2.877125937	0.155003	4.476466	6.016418	8.367091
4.000	1.632332041	2.024466858	4.556723717	3.258737881	0.753543	6.891397	6.901301	4.475396
5.000	1.875214383	2.210009348	4.979858567	3.549114921	1.383353	8.523428	3.532388	0.853224
6.000	2.137893813	2.371635231	5.191372533	3.748505544	3.590440	8.978094	0.609548	5.259355
7.000	2.403485514	2.519945274	5.373455187	3.871855859	5.422789	7.976693	3.374543	7.634374
8.000	2.657202108	2.664366352	5.735811514	3.952018413	7.275292	6.907009	4.165632	8.751759
9.000	2.889913738	2.822985811	7.140900686	4.037086312	8.396802	5.234919	3.636787	8.237973

TABLE 6
Experiment A—Accuracy Testing (Case 4): Average Response Time for Server Activation Ratio $\mathcal{R} = 1.0$

λ	$E[T_1]$ (iterative)	$E[T_2]$ (iterative)	$E[T_3]$ (iterative)	$E[T_4]$ (iterative)	% error (class 1)	% error (class 2)	% error (class 3)	% error (class 4)
1.000	1.249403209	1.240894331	1.249966817	1.596548157	0.188442	0.102801	0.091916	1.082684
2.000	1.632332041	1.530185205	1.659140114	2.346862510	0.589992	1.655815	0.628414	6.931096
3.000	2.137886898	1.800536338	2.327797341	2.877125937	2.607489	4.531356	1.872697	8.398753
4.000	2.655828538	2.024466858	3.203617358	3.258737880	5.961660	7.101779	4.766918	4.467751
5.000	3.086978271	2.210009329	4.010064114	3.549114866	7.807762	8.699709	6.980227	0.879640
6.000	3.414279404	2.371633439	4.556723717	3.748504405	7.872210	8.735477	6.847014	5.201449
7.000	3.665701028	2.519889412	4.874913397	3.871844972	6.090198	8.137421	4.974583	7.772737
8.000	3.871141342	2.663543765	5.062810736	3.951957924	3.526046	7.062987	2.360242	8.736353
9.000	4.053114585	2.814672095	5.191371673	4.036856602	1.660845	6.500396	0.292736	8.240439
10.000	4.235582423	2.994503482	5.306591592	4.216504326	1.340852	7.132456	0.282897	6.970984
11.000	4.459602796	3.243621554	5.457264738	4.763969135	2.877606	8.946037	1.377628	5.284929
12.000	4.815342733	3.651148297	5.735634028	7.980613244	5.087811	9.514688	4.097968	3.596882

TABLE 7
Experiment A—Accuracy Testing (Case 5): Average Response Time for Server Activation Ratio $\mathcal{R} = 1.0$

λ	$E[T_1]$ (iterative)	$E[T_2]$ (iterative)	$E[T_3]$ (iterative)	$E[T_4]$ (iterative)	% error (class 1)	% error (class 2)	% error (class 3)	% error (class 4)
1.000	1.111101334	1.240894331	1.427769462	1.596548157	0.174672	0.138768	0.014092	1.301958
2.000	1.249403209	1.530185205	2.327797341	2.346862510	0.190320	1.628630	1.589170	6.970281
3.000	1.422120871	1.800536338	3.633427552	2.877125937	0.017206	4.421477	6.187979	4.456612
4.000	1.632332041	2.024466858	4.556723717	3.258737880	0.582615	7.084183	6.795926	0.910167
5.000	1.875214307	2.210009329	4.979858554	3.549114866	1.270779	8.620992	3.584426	5.289322
6.000	2.137886898	2.371633439	5.191371673	3.748504405	3.064727	8.844618	0.579377	7.780205
7.000	2.403313249	2.519889412	5.373437048	3.871844972	4.904982	8.076828	3.464238	8.703671
8.000	2.655828538	2.663543765	5.735634028	3.951957924	6.529355	6.896232	3.719736	8.307669
9.000	2.885087046	2.814672095	7.139840165	4.036856602	8.310797	5.611460	1.103102	7.699883
9.500	2.989496161	2.899102696	10.324209470	4.105645593	9.692825	5.219340	2.013128	7.569769

other words, server activation rate can be 10 times slower than the average service rate or up to 100 times faster than the average service rate. Table 8 illustrates the average response time for all classes of customers under different traffic loadings. As we can observe, it is important to

consider server activation issues in the performance analysis of dynamic resource management systems using threshold-based techniques. In particular, when the system is operating at a low server activation ratio (i.e., when $\mathcal{R}$ is small), the average response time for all classes is higher

TABLE 8

Experiment B—Service Activation Rate Experiment (Case 6): Average Response Time for all Classes under Different Server Activation Ratios and Traffic Loadings

λ	$\mathcal{R}$ activation ratio	$E[T_1]$ (iterative)	$E[T_2]$ (iterative)	$E[T_3]$ (iterative)
0.900	0.100	1.425080947	1.425080947	1.425071989
0.900	1.000	1.422120871	1.422120871	1.422119765
0.900	10.00	1.421043244	1.421043244	1.421042825
0.900	100.0	1.420915968	1.420915968	1.420915582
3.600	0.100	5.379141763	5.379141763	5.393195375
3.600	1.000	3.414267196	3.414267196	3.387999226
3.600	10.00	3.145905856	3.145905856	3.128791883
3.600	100.0	3.118374248	3.118374248	3.101954561
5.400	0.100	7.299335678	7.299335678	7.223295397
5.400	1.000	4.050733764	4.050733764	3.927137459
5.400	10.00	3.587663185	3.587663185	3.465490149
5.400	100.0	3.539729932	3.539729932	3.418567750
8.100	0.100	10.616787909	10.616787909	7.852223350
8.100	1.000	6.283156018	6.283156018	4.297965780
8.100	10.00	5.764497387	5.764497387	3.741218408
8.100	100.0	5.709881527	5.709881527	3.682093400

than in a situation where the system is operating at high server activation ratios. The performance difference is more prominent when we have moderate to high traffic loadings.

5.3 Experiment C: Benefits of Resource Sharing

In this experiment, we illustrate the benefits of resource sharing. In particular, we consider the case wherein the number of common resources, K , is less than the peak resource demand $\sum_{i=1}^N K_i$. The reasons for using fewer resources are to reduce the overall system operational cost and to take advantage of the fact that the probability of all classes being at their high workload demands simultaneously might be quite low. Of course, one needs to quantify the above advantages, especially when hysteresis control is used in threshold-based system. We consider test case 5 in Table 3. In this test case, the total peak resource demand is $\sum K_i = 14$. We consider two cases, where we set the values of K , the total amount of system resources, to be 11 and 14, respectively. Note that when $K = 14$, we will not have resource contention among classes. For this experiment, we assume a homogeneous server activation ratio $\mathcal{R} = \beta/\mu = 1.0$. Fig. 4 illustrates the average response time of three classes of workload, under both cases. As we can

Fig. 4. Experiment C—Resource sharing experiment (Case 5): Average response time for four classes of workload with $K = 11$ and 14.

observe, the *difference* in the average response time is *not* significant, i.e., the maximum percentage error between cases $K = 11$ and $K = 14$ is 6 percent. This implies that, in this experiment, one can use fewer resources, with resource saving of around 21 percent, while at the same time providing *comparable* performance (i.e., comparable to the case where the amount of resources is equal to the peak workload demands of all classes).

5.4 Experiment D: Effects of Threshold Settings

As stated in Section 1, we believe that our solution technique for this multiclass, multiserver threshold-based queueing system with hysteresis behavior can be used as an efficient tool for searching for good threshold values. In this experiment, we illustrate that this is an important issue by showing that different threshold values can result in vastly different performance measures. Since our solution can quickly generate corresponding performance results, it should be a good tool for exploring proper threshold settings. In particular, we consider a four servers systems with two classes whose parameter settings are listed in Table 3. Note that, in each test case, we have three configurations A, B, and C, where each configuration has different threshold settings. Fig. 5 depicts the corresponding expected response time results, which illustrate that changes in threshold values result in significant changes in the expected response time of the different classes. Therefore, one can use the proposed solution technique to search for the proper threshold values such that the

(a)

(b)

Fig. 5. Experiment D—Threshold setting experiment (Case 7 and Case 8): Average response time for two workload classes under different thresholds F_i and R_i . (a) Case 7: Mean response time versus different F_i and R_i . (b) Case 8: Mean response time versus different F_i and R_i .

Fig. 6. Experiment E—Instantaneous model versus noninstantaneous model (Case 9). (a) $R = 0.1$. (b) $R = 100$.

expected response time of the different workloads is satisfactory, within certain levels of system loadings.

5.5 Experiment E: Instantaneous Model versus Noninstantaneous Model

The motivation for this experiment is to study the difference between the instantaneous model and the noninstantaneous model and when is important to take the server activation ratio into consideration for design an efficient multiclass multiserver system.

In Fig. 6a, the results show that as the activation ratio is smaller than the service rate, the system performance is mostly dominated by the server activation time. This is why the curve of the instantaneous model deviates from the simulation curve a lot (e.g., for class 1, the largest difference is $\approx 100\%$) while the curve of the noninstantaneous model is close to the simulation curve. That is, it is important to consider the server activation ratio in order to construct an precisely analytic model for the system. On the other hand, as illustrated in Fig. 6b, when the service rate is much faster than the server activation rate, both the results of instantaneous and noninstantaneous models are similar with that of the simulation. In other words, the instantaneous model is a good approximation one as the server activation ratio is large.

6 CONCLUSIONS

In this paper, we considered efficient and accurate computation of a multiclass multiserver threshold-based system with *noninstantaneous* server activation, which is of use in studying performance characteristics of many important applications. We presented an efficient iterative approximation technique for solving such models and illustrated through numerical results that this technique is reasonably accurate (in the presented experiments the deviation from the exact solution is within $\approx 10\%$) and fast (with more than two orders of magnitude improvement in computation time compared to simulations). Moreover, our numerical results also illustrated that 1) server activation characteristics have a significant effect on the system's performance, 2) dynamic resource sharing, through the use of threshold-based techniques, can result in significant cost savings (i.e., through sharing of a pool of resources among heterogeneous workload classes)

without detrimental effects on the class' performance, and 3) proper threshold settings can have a significant effect on the class' performance characteristics. Consequently, all these results indicate that efficient solution techniques for models of such dynamic resource management systems are critical for proper design and performance studies of these systems. And, we believe that the technique presented in this paper is one such approach which will lead to better system designs.

ACKNOWLEDGMENTS

This work has been funded in part by the US National Science Foundation (NSF) ANI-0070016 grant (part of the joint NSF/CNPq program), by the RGC and the CUHK Mainline Research Grant, by the National Science Council and the Ministry of Education of ROC (under the contract No. NSC92-2622-E-002-002, 94-2213-E-002-093, and 89E-FA06-2-4-8), as well as by the Integrated Media Systems Center, an NSF Engineering Research Center, Cooperative Agreement No. EEC-9529152. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the NSF.

REFERENCES

- [1] P.J. Courtois, *Decomposability: Queueing and Computer System Applications*, ACM monograph series. Academic Press, 1977.
- [2] E. de Souza e Silva, S.S. Lavenberg, and R.R. Muntz, "A Perspective on Iterative Methods for the Approximate Analysis of Closed Queueing Networks," *Math. Computer Performance and Reliability*, G. Iazeola, P.J. Courtois, and A. Hordijk, eds., pp. 225-244, North Holland, 1984.
- [3] L. Golubchik and J.C.S. Lui, "Bounding of Performance Measures for a Threshold-Based Queueing System with Hysteresis," *Proc. 1997 ACM SIGMETRICS Conf.*, June 1997.
- [4] L. Golubchik and J.C.S. Lui, "A Fast and Accurate Iterative Solution of a Multi-Class Threshold-Based Queueing System with Hysteresis," *Proc. ACM SIGMETRICS Conf.*, pp. 196-206, June 2000.
- [5] L. Golubchik and J.C.S. Lui, "Bounding of Performance Measures for Threshold-Based Queueing Systems: Theory and Application to Dynamic Resource Management in Video-on-Demand Servers," *IEEE Trans. Computers*, vol. 51, no. 4, pp. 353-372, Apr. 2002.
- [6] S.C. Graves and J. Keilson, "The Compensation Method Applied to a One-Product Production/Inventory Problem," *J. Math. Operational Research*, vol. 6, pp. 246-262, 1981.

- [7] O.C. Ibe, "An Approximate Analysis of a Multi-Server Queueing System with a Fixed Order of Access," Technical Report RC9346, IBM Research, 1982.
- [8] O.C. Ibe and J. Keilson, "Multi-Server Threshold Queues with Hysteresis," *Performance Evaluation*, vol. 21, pp. 185-212, 1995.
- [9] O.C. Ibe and K. Maruyama, "An Approximation Method for a Class of Queueing Systems," *Performance Evaluation*, vol. 5, pp. 15-27, 1985.
- [10] J. Keilson, *Green's Function Methods in Probability Theory*. Charles Griffin, 1965.
- [11] J. Keilson, *Markov Chain Models: Rarity and Exponentiality*. Springer, 1979.
- [12] F.P. Kelly, *Reversibility and Stochastic Networks*. John Wiley and Sons, 1979.
- [13] R.L. Larsen and A.K. Agrawala, "Control of a Heterogeneous Two-Server Exponential Queueing System," *IEEE Trans. Software Eng.*, vol. 9, pp. 552-526, 1983.
- [14] W. Lin and P.R. Kumar, "Optimal Control of a Queueing System with Two Heterogeneous Servers," *IEEE Trans. Automatic Control*, vol. 29, pp. 696-703, 1984.
- [15] J.D.C. Little, "A Proof of the Queueing Formula $L = \lambda W$," *Operations Research*, vol. 9, pp. 383-387, May 1961.
- [16] J.C.S. Lui and L. Golubchik, "Stochastic Complement Analysis of Multi-Server Threshold Queues with Hysteresis," *Performance Evaluation*, vol. 35, nos. 1-2, pp. 19-48, Mar. 1999.
- [17] C.D. Meyer, "Stochastic Complementation, Uncoupling Markov Chains and the Theory of Nearly Reducible Systems," *SIAM Rev.*, vol. 31, no. 2, pp. 240-272, 1989.
- [18] D. Mitra and I. Ziedins, "Virtual Partitioning by Dynamic Priorities: Fair and Efficient Resource-Sharing by Several Services," *Int'l Zurich Seminar on Digital Comm., Lecture Notes in Computer Science, Broadband Comm.*, B. Plattner ed., pp. 173-185, Springer, 1996.
- [19] D. Mitra and I. Ziedins, "Hierarchical Virtual Partitioning: Algorithms for Virtual Private Networking," *Proc. IEEE GLOBE-COM*, pp. 1784-1791, 1997.
- [20] J.A. Morrison, "Two-Server Queue with One Server Idle below a Threshold," *Queueing Systems*, vol. 7, pp. 325-336, 1990.
- [21] R. Nelson and D. Towsley, "Approximating the Mean Time in System in a Multiple-Server Queue that Uses Threshold Scheduling," *Operations Research*, vol. 35, pp. 419-427, 1987.

Cheng-Fu Chou received the MS and PhD degrees from the University of Maryland, College Park, in 1999 and 2002, respectively. After his graduation, he joined the Computer Science and Information Engineering Department at the National Taiwan University. He has been a visiting scholar in the Computer Science Department at the University of Southern California. His current research interests are in data network,

distributed multimedia systems and wide-area network applications, wireless networks, and their performance evaluation. He is a member of the IEEE Computer Society.

Leana Golubchik received the PhD degree from the University of California at Los Angeles in 1995. She is an associate professor in the Computer Science Department at the University of Southern California. Prior to that, she was an assistant professor and then an associate professor at the University of Maryland at College Park, and an assistant professor at Columbia University. She is a past chair (2003-2005) and vice chair (2001-2003) of ACM SIGMETRICS (2003-2005) and a member of its Board of Directors (1999-2001, 2005-2007). She was a guest coeditor for special issues of *IEEE Transactions on Knowledge and Data Engineering*, the *Parallel Computing Journal*, and the *International Journal of Intelligent Systems*, a program cochair of the 2001 Joint ACM SIGMETRICS/Performance Conference and MIS '99, and a program committee member of numerous conferences. She received several awards, including the US National Science Foundation (NSF) CAREER Award, the Okawa Foundation award, and the IBM & NSF Doctoral Fellowships. She is a member of the IFIP WG 7.3 and Tau Beta Pi.

John Chi-Shing Lui received the PhD degree in computer science from the University of California at Los Angeles (UCLA). After his graduation, he joined the IBM Almaden Research Laboratory/San Jose Laboratory and participated in various R&D projects on file systems and parallel I/O architectures. He later joined the Department of Computer Science and Engineering at the Chinese University of Hong Kong. He has been a visiting professor in Computer Science Departments at UCLA, Columbia University, the University of Maryland at College Park, Purdue University, the University of Massachusetts at Amherst, and Universit degli Studi di Torino in Italy. His research interests span both in system and in theory/mathematics, with current research interests in theoretic/applied topics in data networks, distributed multimedia systems, network security, OS design issues and mathematical optimization, and performance evaluation theory. He received various departmental teaching awards and the CUHK Vice-Chancellor's Exemplary Teaching Award, as well as the corecipient of the IFIP WG 7.3 Performance 2005 Best Student Paper Award. He is an associate editor of the *Performance Evaluation Journal*, a member of the ACM, a senior member of the IEEE, and an elected member in the IFIP WG 7.3. He was the TPC cochair of ACM Sigmetrics 2005.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.