

Simple Analytic Modeling of Software Contention

Daniel A. Menascé¹

Abstract

Being able to model contention for software resources (e.g., a critical section or database lock) is paramount to building performance models that capture all aspects of the delay encountered by a process as it executes. Several methods have been offered for dealing with software contention and with message blocking in client-server systems. We present in this paper a simple, straightforward, easy to understand and implement, approach to modeling software contention using queuing networks. The approach consists of a two-level iterative process. Two queuing networks are considered: one represents software resources and the other hardware resources. Multiclass models are allowed and both open and closed queuing networks can be used at the software layer. Any solution technique—exact or approximate—can be used at any of the levels. This technique falls in the general nature of fixed-point approximate models and is similar in nature to other approaches. The main difference lies in its simplicity.

1 Introduction

The response time of a transaction or request submitted to a computer system can be broken down into three components. The first is the total service time, i.e., the total time spent by transactions obtaining service from the various physical resources such as processors, disks, and networks. The second component is the total time spent waiting to use a physical resource. Finally, the third component (*software contention*), is the time spent to access a software resource such as a non-reentrant software module, a software thread, or a database lock. It is important to take into account the effect of software contention when estimating the total response time.

Agrawal and Buzen [1] present an approximate iterative technique to model serialization delays that employs a conventional queuing network (QN) comprised of servers that represent actual processors and I/O devices plus additional aggregate servers that represent serialized processing activity. The technique iteratively determines the values of the service times at all devices by computing adjustment factors used to elongate the original service times. The formulas provided for computing these adjustment factors are neither simple nor intuitive. The formulation is presented for single class models only. The extension to multiple classes is said to follow the same principles but is acknowledged to be computation-

ally difficult [1]. The work in [21] is a two-level model of serialization delays. At the lower level, a QN that represents the physical resources is used to compute transition rates for a higher level Markov Chain model. This formulation does not generalize to multiple classes and does not deal with other types of software contention issues. The approach taken in [7] for analyzing QNs with serialization delays consists in transforming the network to one with population constraints and using an approximation technique to solve the latter problem. Other analytic work on specific software contention situations can be found in [11] and [16].

Multilayer client-server queuing network models were introduced in [15] to handle synchronous and asynchronous exchange of messages between clients and servers. The problem of modeling CORBA based distributed systems was addressed in [9] through the use of QNs with simultaneous resource possession. These networks are solved through decomposition into a set of auxiliary product-form QNs and iteration between these networks.

A very important body of work aimed at estimating software contention issues is the Method of Layers (MOL), also called Layered Queuing Networks (LQN) [17, 23]. Also related and very relevant is the work on stochastic rendez-vous networks [14, 22]. Layered queuing networks were extended to handle parallel tasks in [4]. The LQN approach while, in principle, similar to the technique presented here differs in ways that will be explained in the concluding remarks, after our technique has been presented.

The main contribution of this paper is that it provides a simple analytical modeling technique for estimating software contention delays. The technique consists of a two-layer queuing network model. Software resources are modeled by a software queuing network (QN) and physical resources by a hardware queuing network. The effect of software contention on the physical QN and the effect of physical resource contention on the software QN requires an iterative technique to solve this dependency. The technique relies on existing solution techniques for open, closed, multiclass queuing networks and can be used with both product-form QNs and approximations [2] used to handle situations such as simultaneous resource possession [2, 8, 12], and priorities [2, 3, 18]. Multiservers can be handled by using load-dependent servers [10, 13] or approximations [19]. The basic idea developed in this paper was first proposed in [10] but no algorithm and no validation was presented there. While the other techniques referenced earlier can also handle the situations addressed by our approach, we believe that the strength

¹Department of Computer Science, MS 4A5, George Mason University, Fairfax, VA 22030-4444, menasce@cs.gmu.edu

of our method lies in its simplicity. Anyone who knows how to solve exact or approximated multiclass QNs can use the method to estimate software contention.

The rest of the paper is organized as follows. Section 2 illustrates the main rationale behind the technique through a simple motivating example in which N processes compete for a single critical section and execute non-critical section code after they complete the critical section. A single class algorithm is also presented in that section. The next section presents a more complex example and compares our results with those obtained with other techniques. Section four shows how the technique presented in section 2 can be extended to model other cases. Section 5 presents the multiclass algorithm and section 6 presents some concluding remarks.

2 A Simple Example - Contention for a Critical Section

Consider the case of N processes P_1, \dots, P_N that alternate between executing non-critical section and critical section code. Any number of processes can be concurrently executing their non-critical code. But, only one of them can be executing the critical section code. If a process attempts to enter its critical section code while another process is inside the critical section, the attempting process is put to sleep at the queue associated with the semaphore that controls access to the critical section. Let us assume in this simple example that the only physical device used by all processes during the execution of the critical and non-critical section code is the CPU and that the CPU scheduling discipline is processor-sharing.

If we just consider the software phases of a process execution, we can depict a process by a *software queuing network* (SQN) as in the top part of Fig. 1. The SQN has two resources (software modules). One is a delay-resource (illustrated as a rectangle) that corresponds to the non-critical section code; there is no queuing for a software resource during this phase. The other software resource is a queuing resource, which corresponds to the critical section code.

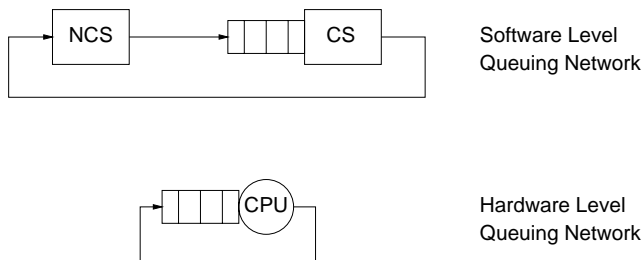


Figure 1: Software and Hardware Queuing Networks for the Critical Section Example

While a process—a customer in the SQN—is using the software resources in the SQN, it is also using or waiting to use physical resources (e.g., CPUs and disks). Delay resources

are used in the SQN to represent software resources for which there is no software contention. The queuing network associated with the physical resources, the *hardware queuing network* (HQN), is shown in the bottom part of Fig. 1. Customers in the HQN are processes using the physical resources as a result of the execution of software modules. The time spent at the NCS and CS resources in the SQN depends on the contention for access to the physical resources, the CPU in this case. Also, the number of processes contending for the CPU, i.e., the number of processes in the HQN, is equal to the number of concurrent processes that are not blocked waiting for entry to the critical section. The blocked processes are sleeping and are therefore not present in any HQN queue. Therefore, the customer population in the HQN is equal to $N - B$ where B is the number of processes blocked for a software resource.

We now define some notation before we proceed with an explanation on how contention for access to software resources can be computed. Let

- $D_{j,i}^{\text{sh}}$: service demand, i.e., total service time of a process executing software module j when using physical resource i in the HQN. The superscript “sh” indicates that these service demands are related to the mapping between the software and hardware QNs. For example, $D_{\text{nscs,cpu}}^{\text{sh}}$ is the total CPU time for the execution of the non-critical section code. This time does not include the time spent waiting to use the CPU while executing non-critical section code.
- D_j^{s} : service demand, i.e., total service time to execute module j in the SQN. The superscript “s” indicates that this service demand relates to the software QN. For example, $D_{\text{nscs}}^{\text{s}}$ is the total service time to execute the non-critical section code. The service demand D_j^{s} is the sum of all service times at all physical devices during the execution of module j . Thus,

$$D_j^{\text{s}} = \sum_{\forall i} D_{j,i}^{\text{sh}}. \quad (1)$$

- D_i^{h} : service demand, i.e., total service time of a process at physical resource i in the hardware QN. For example, $D_{\text{cpu}}^{\text{h}}$ is the total service time of a process at the CPU. This time, is the sum of the service demands due to the execution of all modules of the process. Thus,

$$D_i^{\text{h}} = \sum_{\forall j} D_{j,i}^{\text{sh}}. \quad (2)$$

For example, $D_{\text{cpu}}^{\text{h}} = D_{\text{nscs,cpu}}^{\text{sh}} + D_{\text{cs,cpu}}^{\text{sh}}$.

- $R_i'(n)$: residence time, i.e., total time spent by a process at physical resource i , waiting for or receiving service, when there are n processes at the HQN.

We now show the iterative algorithm used to estimate the software contention time and to compute all performance metrics

(e.g., response time and throughput). The inputs to the algorithm are the service demands $D_{j,i}^{sh}$ and the number N of processes. The algorithm iterates until successive values of the number, B , of processes blocked for software contention are sufficiently close.

- Step 1 - Initialization:

$$D_j^s \leftarrow \sum_{\forall i} D_{j,i}^{sh} \quad (3)$$

$$D_i^h \leftarrow \sum_{\forall j} D_{j,i}^{sh} \quad (4)$$

$$B^0 \leftarrow 0 \quad \text{initial value for } B \quad (5)$$

$$k \leftarrow 1 \quad \text{iteration counter} \quad (6)$$

- Step 2 - Solve the SQN with D_j^s as service demands and N as customer population.
- Step 3 - Compute the average number of blocked processes B^k . In the case of our example, this is equal to the average number of processes waiting in the queue for the CS resource. So, $B^k = \bar{n}_{cs} - U_{cs}$, where \bar{n}_{cs} is the average number of processes at resource CS in the SQN and U_{cs} is the utilization of resource CS. In general,

$$B^k = \sum_{\forall j} L_j \quad (7)$$

where L_j is the average number of processes in the waiting line for software resource j .

- Step 4 - Solve the HQN with D_i^h as service demands and $N^h = N - B^k$ as customer population. Note that the solution to a QN with a non-integer customer population can be obtained using Schweitzer's approximation [20].
- Step 5 - Adjust the service demands at the SQN to account for contention at the physical resources. So,

$$D_j^s \leftarrow \sum_{\forall i} \frac{D_{j,i}^{sh}}{D_i^h} \times R'_i(N^h). \quad (8)$$

So, for our example we have

$$D_{ncs}^s \leftarrow \frac{D_{ncs,cpu}^{sh}}{D_{cpu}^h} \times R'_{cpu}(N^h) \quad (9)$$

$$D_{cs}^s \leftarrow \frac{D_{cs,cpu}^{sh}}{D_{cpu}^h} \times R'_{cpu}(N^h). \quad (10)$$

- Step 6 (Convergence Check Step):
If $|(B^k - B^{k-1})/B^k| > \xi$ then $k \leftarrow k + 1$ and go to step 2.

We now give a justification for Step 5 of the above algorithm using our critical section example. The residence time equation for MVA applied to the HQN is

$$R'_{cpu}(N) = D_{cpu}^h [1 + \bar{n}_{cpu}(N - 1)]. \quad (11)$$

But,

$$D_{cpu}^h = D_{ncs,cpu}^{sh} + D_{cs,cpu}^{sh}. \quad (12)$$

So, using Eqs. (11) and (12), we get

$$R'_{cpu}(N) = D_{ncs,cpu}^{sh} [1 + \bar{n}_{cpu}(N - 1)] + D_{cs,cpu}^{sh} [1 + \bar{n}_{cpu}(N - 1)]. \quad (13)$$

The first term of the right-hand side of Eq. (13) is the total time (waiting + service) spent at the CPU by a process while executing non-critical section code and the second term is the total time spent at the CPU by a process while executing the critical section code. So, for example, using Eqs. (11) and (13) we can write the total time spent at the CPU while executing the non-critical section code as

$$\frac{D_{ncs,cpu}^{sh}}{D_{cpu}^h} \times R'_{cpu}(N). \quad (14)$$

Figure 2 shows the percentage of the total response time spent waiting to enter the critical section as a function of the number of concurrent processes. The service demands used to obtain the graph are: $D_{ncs,cpu}^{sh} = 0.2$ sec and $D_{cs,cpu}^{sh} = 0.1$ sec. As it can be seen, as the concurrency level increases, contention for software resources dominates the response time of a process.

3 A More Complex Example

In this section we apply the algorithm presented in the previous section to an example that appears in [1]. We use this example for two reasons. First, because the example is more complex than the one presented before. It has two critical sections at the software level (see top portion of Fig. 3). The hardware QN is composed of one CPU and three disks (see bottom part of Fig. 3). Processes use both CPU and disk while in the critical section. The second reason for using the example in [1] is that global balance solutions are provided in that paper. So, we can compare the results provided by our technique with exact results.

The service demands (i.e., the values of $D_{j,i}^{sh}$) are given in Table 1 and are the same as in [1]. The last row shows the values of D_j^s and the last column contains the values of D_i^h .

Table 2 shows the results (i.e., throughputs) obtained with our technique (called SQN-HQN) and with global balance equations, as reported in [1], for eight values of the number of concurrent processes N . The last three columns of the table show the absolute % relative error obtained with our technique and with those of [1] and [7] relative to the global balance equations solution. The error of our approach is very small and does not exceed 3.05% in this example; for $N > 1$, it is always smaller than the one reported in [1]. The error obtained with our method is higher than the one for [7] for $N > 4$ but our method is much simpler to implement. It should also be noticed that the results obtained with the SQN-HQN technique are consistently pessimistic.

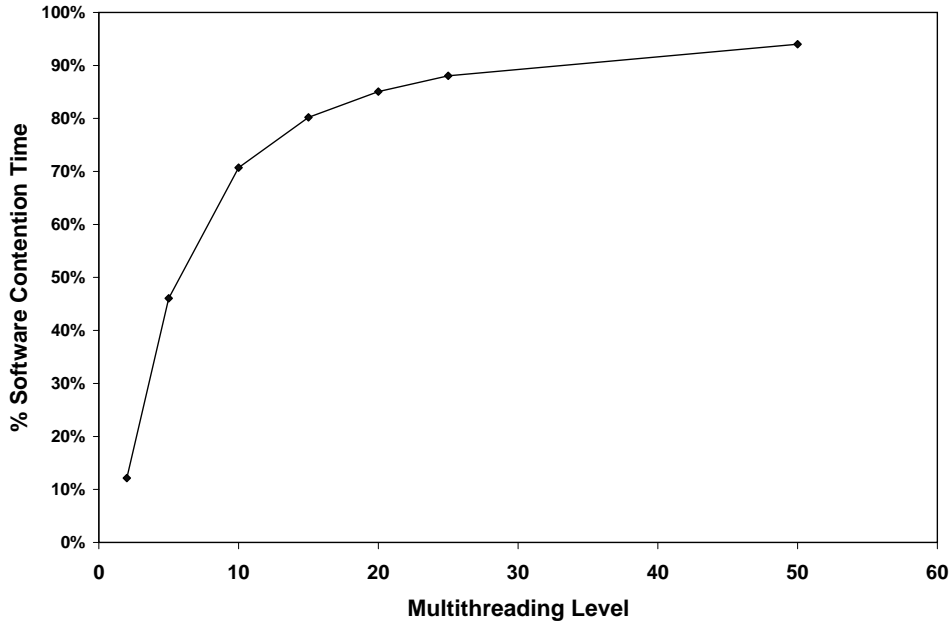


Figure 2: Percent of Software Contention Time vs. Multithreading Level in the Critical Section Example

Table 1: Service Demands (in sec) for the Example in [1]

Device	Software Module			Hardware Demands
	NCS	CS 1	CS 2	
CPU	0.2000	0.0600	0.0808	0.3408
Disk 1	0.0560	0.0576	0.000	0.1136
Disk 2	0.0360	0.0000	0.1212	0.1572
Disk 3	0.0360	0.0000	0.0000	0.0360
Software Demands	0.3280	0.1176	0.2020	

Table 2: Throughput(processes/sec) for Example in [1] (GBS = Global Balance Solution, SQN-HQN= technique presented in this paper, ASM = algorithm in [1], and ASPA = algorithm in [7])

N	SQN-HQN	GBS	Absolute % Relative Error		
			SQN-HQN	ASM	ASPA
1	1.544	1.54	0.27	0.00	-
2	2.088	2.11	1.06	4.60	-
3	2.317	2.37	2.22	5.89	4.2
4	2.428	2.49	2.49	5.99	2.8
5	2.487	2.56	2.86	5.87	2.0
6	2.521	2.60	3.05	5.78	1.5
7	2.541	2.62	3.00	5.75	1.5
8	2.555	2.63	2.86	5.77	1.1

4 Other Cases

In this section we explain how the algorithm presented in Section 2 can be extended to cover other situations. The first is the case in which a customer in the SQN spends time not

directly associated with the execution of a software module (e.g., think times or network time). The second case discusses how open QNs can be used to represent SQNs.

4.1 Modeling Non-Software Resources

There are cases in which there is no correspondence between a resource at the hardware and the software layers. For example, Fig. 4 shows the case in which user think time as well as network waiting and transmission time need to be modeled. In these cases, we need to incorporate these two resources—a delay device for the think time and a queuing device for the network—in both the SQN and HQN networks. We call them *non-software* resources. Note that we depict software resources by rectangles in the SQN and all other resources by circles. The computation of B^k in Step 3 of the algorithm presented in the previous section only includes processes blocked for software resources.

4.2 Open QN at the Software Level

The technique presented in this paper can be extended to include the case in which the SQN is modeled as an open queuing network as illustrated in Fig. 5. In that case, processes

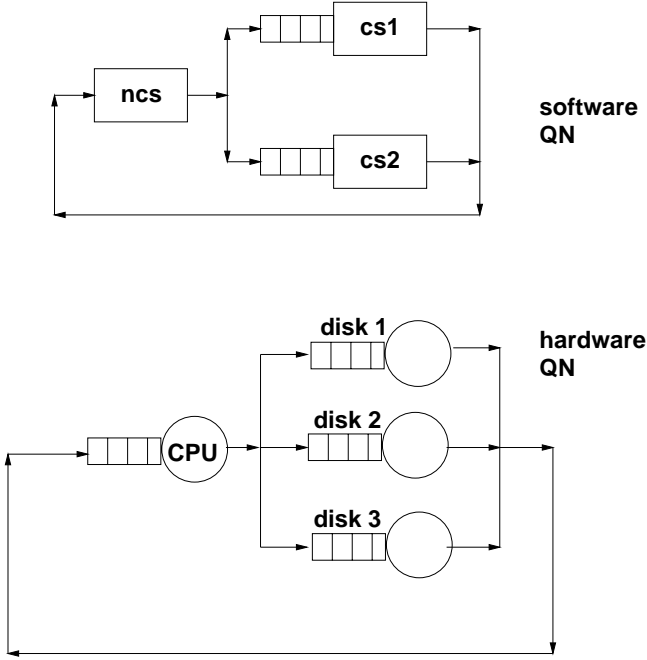


Figure 3: Software and Hardware Queuing Networks for the Example in [1]

arrive at a rate of λ requests/sec.

Steps 2 and 4 of the algorithm described in Section 2 should be replaced by the following to contemplate the case of open SQNs.

- new Step 2 - Solve the SQN with D_j^s as service demands and λ as arrival rate and obtain \bar{N} the average number of customers in the SQN.
- new Step 4 - Solve the HQN with D_i^h as service demands and $N^h = \bar{N} - B^k$ as customer population. Note that the solution to a QN with a non-integer customer population can be obtained using Schweitzer's approximation [20].

5 The Multiclass Algorithm

The algorithm presented in Section 2 generalizes in a straightforward manner to multiple classes. We define here the generalized notation and present the multiclass algorithm. Let

- R : number of classes in the SQN and in the HQN.
- $\vec{N}^s = (N_1, \dots, N_R)$: population vector for the SQN. N_r is the multithreading level for class r .
- $\vec{N}^h = (N_1, \dots, N_R)$: population vector for the HQN. N_r is the number of customers in class r .

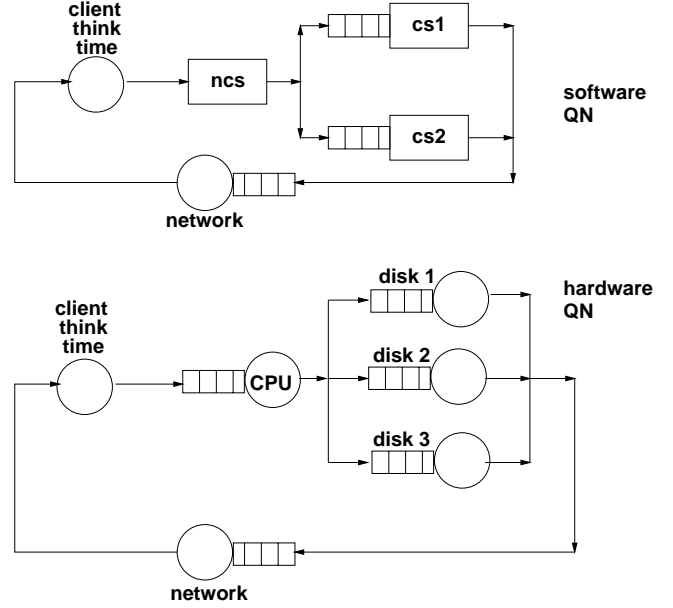


Figure 4: Example with Client Think Time and Network

- $\vec{B} = (B_1, \dots, B_R)$: vector of number of processes blocked, i.e., waiting for a software resource. B_r is the number of class r processes blocked for a software resource.
- $D_{j,i;r}^{sh}$: service demand, i.e., total service time at physical resource i of a class r software module j .
- $D_{j;r}^s$: service demand, i.e., total service time to execute module j of class r in the SQN. The service demand $D_{j;r}^s$ is the sum of all service times at all physical devices during the execution of module j . Thus,

$$D_{j;r}^s = \sum_{\forall i} D_{j,i;r}^{sh}. \quad (15)$$

- $D_{i;r}^h$: service demand, i.e., total service time at physical resource i for class r in the hardware QN. This time, is the sum of the service demands due to the execution of all modules of a process. Thus,

$$D_{i;r}^h = \sum_{\forall j} D_{j,i;r}^{sh}. \quad (16)$$

- $R'_{i;r}(\vec{N}^h)$: residence time, i.e., total time spent by a class r process at resource i , waiting for or receiving service, when the customer population at the HQN is given by the vector \vec{N}^h .

We now show the iterative algorithm used to estimate the software contention time and to compute all performance metrics (e.g., response time and throughput). The inputs to the algorithms are the service demands $D_{j,i;r}^{sh}$ and the vector \vec{N}^s . The algorithm iterates until successive values of \vec{B} are sufficiently close.

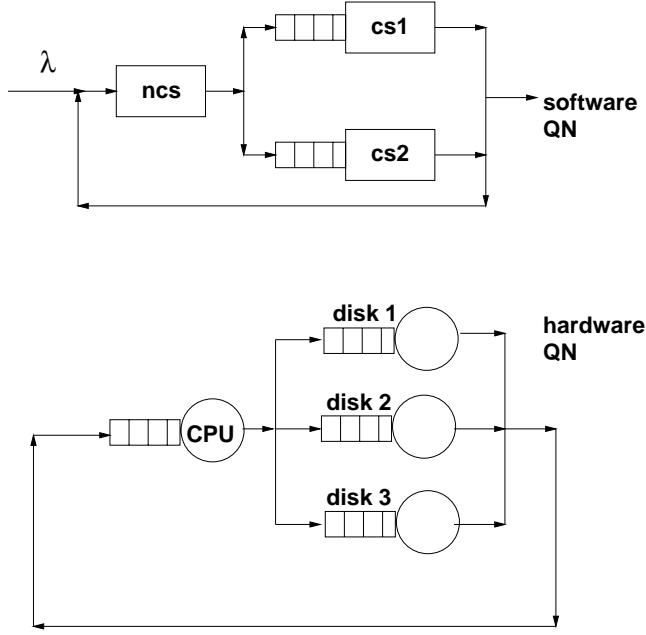


Figure 5: Open QN at the Software Level

- Step 1 - Initialization:

$$D_{j;r}^s \leftarrow \sum_{\forall i} D_{j,i;r}^{\text{sh}} \quad r = 1, \dots, R \quad (17)$$

$$D_{i;r}^h \leftarrow \sum_{\forall j} D_{j,i;r}^{\text{sh}} \quad r = 1, \dots, R \quad (18)$$

$$\vec{B} \leftarrow \vec{0} \quad (19)$$

$$k \leftarrow 1 \quad (20)$$

- Step 2 - Solve the SQN with $D_{j;r}^s$ as service demands and \vec{N}^s as customer population.
- Step 3 - Compute the average number of blocked processes per class as

$$\vec{B}^k = \left(\sum_{\forall j} L_{j,1}, \dots, \sum_{\forall j} L_{j,r}, \dots, \sum_{\forall j} L_{j,R} \right) \quad (21)$$

where $L_{j,r}$, $r = 1, \dots, R$, is the average number of class r processes in the waiting line for software resource j .

- Step 4 - Solve the HQN with $D_{i;r}^h$ as service demands and population vector $\vec{N}^h = \vec{N}^s - \vec{B}^k$ as customer population. Note that the solution to a QN with a non-integer customer population can be obtained using Schweitzer's approximation [20]. The solution to the HQN provides the residence time values $R'_{i;r}(\vec{N}^h)$ for all devices i and classes r .

- Step 5 - Adjust the service demands at the SQN to account for contention at the physical resources. So,

$$D_{j;r}^s \leftarrow \sum_{\forall i} \frac{D_{j,i;r}^{\text{sh}}}{D_{i;r}^h} \times R'_{i;r}(\vec{N}^h). \quad (22)$$

- Step 6 (Convergence Check Step): If $\max_r | (B_r^k - B_r^{k-1}) / B_r^k | > \xi$ then $k \leftarrow k + 1$ and go to step 2.

The same considerations discussed in section 4 can be applied to the multiclass case. Situations that require modeling a queue for a multithreaded software server can be considered by including a multiserver queue in the SQN and using approximation techniques [19] to solve QNs with multiservers or QNs with load-dependent servers [10, 13].

6 Concluding Remarks

This paper presented a simple technique to estimate software contention using analytic models. The approach uses well-known hierarchical modeling methods solved through iteration. The technique consists of two queuing networks. The top level—the software queuing network (SQN)—has servers associated with software resources. A class of customers represents processes with similar behavior in their use of software resources and use of the underlying physical resources while using software resources. Multiple classes of customers are allowed. Software resources that do not generate any software contention are modeled as delay resources in the SQN and contention for software resources is represented by queuing resources. These may have a single server (as in the case of a critical section or a serialization delay) or multiple servers as in the case of a multithreaded process. The SQN may be open or closed. Any known technique, exact or approximate, may be used to solve the SQN. Situations that involve processes that request service more than once (multi-phased processes) from the same software resource and have significantly different service times at each phase, can be modeled using class-switching QNs as done in [5, 6].

The bottom-level QN represents the hardware resources. Service demands at the software level are iteratively adjusted to account for contention at the physical level. Again, any exact or approximate technique can be used to solve the hardware QN, which is modeled as a closed QN.

Our approach differs from the LQN method [17] in that i) only two layers are considered, ii) any QN solution technique can be used at the SQN and HQN without any need to adapt specific residence time equations as done in [17], and iii) open and closed QNs are allowed at the SQN. The hope is that by providing a straightforward and easy to understand method for modeling software contention, those who are familiar with solution methods for multiclass QNs will be able to readily model software contention without the need to use new types of QNs. While simple, the technique presented here was validated and is general enough to be applied to multiclass situations.

References

- [1] S. C. Agrawal and J. P. Buzen, "The Aggregate Server Method for Analyzing Serialization Delays in Computer Systems," *ACM Transactions on Computer Systems*, Vol. 1, No. 2, May 1983, pp. 116–143.
- [2] S. C. Agrawal, "Metamodeling: A study of Approximations in Queuing Models," MIT Press, 1985.
- [3] V.A.F. Almeida and D. A. Menascé, "Approximate Modeling of CPU Preemptive Resume Priority Scheduling Using Operational Analysis," Proc. 10th European Computer Manufactures Association (ECOMA) Conference on Computer Measurement, Munich, Germany, October 12-15, 1982.
- [4] G. Franks and C. M. Woodside, "Performance of Multi-Level Client-Server Systems with Parallel Service Operations," Proc. First Int. Workshop on Software and Performance (WOSP'98), Santa Fe, NM, Oct. 1998, pp. 120–130.
- [5] A. Harbitter and D. A. Menascé, "The Performance of Public Key Enabled Kerberos Authentication in Mobile Computing Applications," Eighth ACM Conference on Computer and Communications Security (CCS-8), Philadelphia, PA, November 5-8, 2001.
- [6] A. Harbitter and D. A. Menascé, "Performance of Public Key-Enabled Kerberos Authentication in Large Networks," Proc. 2001 IEEE Symposium on Security and Privacy, Oakland, California, May 13-16, 2001.
- [7] P. A. Jacobson and E. D. Lazowska, A Reduction Technique for Evaluating Queuing Networks with Serialization Delays," *Performance '83*, eds. A. K. Agrawal and S. K. Tripathi, North-Holland Publishing Company, 1983, pp. 45–59.
- [8] P. A. Jacobson and E. D. Lazowska, "Analyzing Queuing Networks with Simultaneous Resource Possession," *Comm. ACM* 25, 2, Feb. 1982, pp. 142–151.
- [9] P. Kahkipuro, "Performance Modeling Framework for CORBA Based Distributed Systems," Ph.D. Dissertation, Technical Report A-2000-3, Department of Computer Science, University of Helsinki, Finland, 2000.
- [10] E. D. Lazowska, J. Zahorjan, G. S. Graham, K. C. Sevcik, "Quantitative System Performance," Prentice Hall, 1984.
- [11] C. M. Lladó and P. G. Harrison, "Performance Evaluation of an Enterprise JavaBean Server Implementation," Proc. Second International Workshop on Software and Performance, Ottawa, Canada, Sept. 17-2-, 2000, pp. 180–188.
- [12] D. A. Menascé, O. Pentakalos and Y. Yesha, "An Analytic Model of Hierarchical Mass Storage Systems with Network-Attached Storage Devices," Proc. of the 1996 ACM Sigmetrics Conference Philadelphia, PA, May 1996.
- [13] D. A. Menascé and V. A. F. Almeida, "Capacity Planning and Performance Modeling: from mainframes to client-server systems," Prentice Hall, Upper Saddle River, NJ, 1994.
- [14] D. Petriu, H. Amer, S. Majumdar, I. Abdull-Fatah, "Using Analytic Models for Predicting Middleware Performance," Proc. Second International Workshop on Software and Performance, Ottawa, Canada, Sept. 17-2-, 2000, pp. 189–194.
- [15] S. Ramesh and H. G. Perros, "A Multilayer Client-Server Queuing Network Model with Synchronous and Asynchronous Messages," *IEEE Tr. Software Eng.*, Vol. 26, no. 11, Nov. 2000, pp. 1086–1100.
- [16] P. Reeser and R. Hariharan, "Analytic Model of Web Servers in Distributed Environments," Proc. Second International Workshop on Software and Performance, Ottawa, Canada, Sept. 17-2-, 2000, pp. 158–167.
- [17] J. A. Rolia and K. C. Sevcik, "The Method of Layers," *IEEE Tr. Software Eng.*, vol. 21, no. 8, 1995, pp. 689–700.
- [18] K. C. Sevcik, "Priority Scheduling Disciplines in Queuing Network Models of Computer Systems," Proc. IFIP Congress 77, North-Holland Publishing Co., Amsterdam, 1977, pp. 565–570.
- [19] A. Seidmann, P. Schweitzer, and S. Shalev-Oren, "Computerized Closed Queueing Network Models of Flexible Manufacturing Systems," *Large Scale Syst. J.*, North Holland, vol. 12, pp. 91–107, 1987.
- [20] P. Schweitzer, Approximate analysis of multiclass closed network of queues, in *International Conference on Stochastic Control and Optimization*, Amsterdam, 1979.
- [21] Thomasian, A., "Queuing Network Models to Estimate Serialization Delays in Computer Systems," *Performance '83*, eds. A. K. Agrawal and S. K. Tripathi, North-Holland Publishing Company, 1983, pp. 61–81.
- [22] C. M. Woodside, J. E. Neilson, D. C. Petriu, and S. Majumdar, "The Stochastic Rendez-vous Network Model for Performance of Synchronous Client-Server-like Distributed Software," *IEEE Tr. Computers*, vol. 44, no. 1, 1995.
- [23] C. M. Woodside, "An Active-Server Model for the Performance of Parallel Programs Written Using Rendezvous," *J. Systems and Software*, 1986, pp. 125–131.