

PAM — A Noniterative Approximate Solution Method for Closed Multichain Queueing Networks*

Ching-Tarn Hsieh
AT&T Bell Laboratories
Naperville, IL 60566

Simon S. Lam
Department of Computer Sciences
University of Texas at Austin
Austin, TX 78712-1188

ABSTRACT

Approximate MVA algorithms for separable queueing networks are based upon an iterative solution of a set of modified MVA formulas. Although each iteration has a computational time requirement of $O(MK^2)$ or less, many iterations are typically needed for convergence to a solution. (M denotes the number of queues and K the number of closed chains or customer classes.) We present some faster approximate solution algorithms that are *noniterative*. They are suitable for the analysis and design of communication networks which may require tens to hundreds, perhaps thousands, of closed chains to model flow-controlled virtual channels. Three PAM algorithms of increasing accuracy are presented. Two of them have time and space requirements of $O(MK)$. The third algorithm has a time requirement of $O(MK^2)$ and a space requirement of $O(MK)$.

1. Introduction

We are interested in fast solution algorithms for separable queueing networks with a large number of closed chains (also called customer classes), such as models of communication networks where each closed chain represents a flow-controlled virtual channel [11, 12]. Tens to hundreds, perhaps thousands, of such flow-controlled virtual channels between source-destination node pairs may be active at a time in a practical network. Exact solution algorithms, such as the convolution and MVA algorithms, are obviously not applicable since their

computational time and space requirements grow exponentially with the number of chains [13, 15]. In fact, they cannot be used for most networks with more than 6 or 7 chains. The tree convolution algorithm [10] as well as the tree MVA algorithms [6, 20] are more computationally efficient for the solution of networks with chains whose routes are sparse, by exploiting routing information. Nevertheless, the largest networks solved by the tree convolution algorithm have 32-50 chains, which are not sufficient for modeling many real networks. Furthermore, most algorithms for network design problems (e.g., optimal routing, topology optimization, etc.) involve a heuristic search for solutions that are optimum according to various network performance measures. A very fast algorithm is needed to evaluate these network performance measures at each of the numerous intermediate steps of such a heuristic search.

For the same reason, the approximate solution methods which have been shown empirically to have good accuracy [1, 3, 4, 14, 16, 17, 22] are deemed to be still too slow for communication network design problems. All of these methods are based upon an iterative solution of a set of modified MVA formulas. A single iteration of the Schweitzer algorithm has a computation time of $O(MK)$ while a single iteration of Linearizer or AQL has a computation time of $O(MK^2)$ where M denotes the number of queues and K the number of closed chains in the network [4, 16, 22]. (It was shown recently by de Souza e Silva and Muntz [18] that a single iteration of Linearizer can be implemented with a computation time of $O(MK^2)$ instead of $O(MK^3)$ as indicated by Chandy and Neuse [4].) Typically many iterations are needed for convergence to a solution, and it has been shown that for some networks convergence occurs extremely slowly [22]. Lastly, the accuracy of these methods has been examined only for networks with a small number of chains, i.e., those that can be solved exactly by the MVA algorithm.

Various methods for computing performance bounds are available [5, 7, 21]. However, for networks with a large number of chains, these bounds are generally too loose to be useful for communication network design.

The above considerations led us to investigate a class of faster approximation solution algorithms for closed multichain queueing networks. In particular, all of our algorithms are *noniterative*. We consider separable queue-

*This work was performed at the University of Texas at Austin under National Science Foundation grant no. ECS-8304734 and grant no. NCR-8613338.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

ing networks of fixed-rate servers (also called queue-independent servers) and delay servers (also called infinite-server centers). Like all of the approximate solution methods referenced, our method is also based upon the MVA recursion formula. In our study of proportional throughput upper bounds for single-chain queueing networks, we found that they are very accurate for networks with small to medium population sizes [7]. These throughput upper bounds are obtained by distributing the population of a chain over the servers it visits proportional to server loads. In a multichain network, such a proportional distribution leads to approximations rather than upper bounds of chain throughputs. We found that these approximate solutions provide chain throughputs, mean end-to-end delays, and server utilizations that are sufficiently accurate for the analysis and design of communication networks [8, 11] and possibly various other distributed systems that have a large number of customer classes (see Sections 3 and 5); an iterative solution to improve accuracy is not necessary.

We shall refer to our method as the *Proportional Approximation Method (PAM)*. We present three algorithms of increasing accuracy that are based upon the distribution of a chain's population proportional to loads to get initial estimates of mean queue lengths: PAM_BASIC, PAM_IMPROVED, and PAM_TWO. The accuracy improvement of PAM_IMPROVED over PAM_BASIC is obtained by a simple scaling operation to ensure that the utilization of each server does not exceed one. The additional accuracy of PAM_TWO is obtained by executing the final two steps of the MVA recursion instead of just the last step. The computational time requirements are $O(MK)$ for PAM_BASIC and PAM_IMPROVED, and $O(MK^2)$ for PAM_TWO. Since PAM algorithms do not iterate, these are their total time requirements. All three PAM algorithms have space requirements of $O(MK)$.

The rest of this paper is organized as follows. Algorithms PAM_BASIC and PAM_IMPROVED are presented in Section 2. In Section 3, we study their accuracy by comparing their predictions with exact solutions given by the tree convolution algorithm (TCA) for a network example and also for 100 randomly generated networks. These networks have fixed-rate servers only and possess characteristics of models of communication networks. We also show a correlation between the approximation errors for a chain and the utilization of any bottleneck server the chain visits. (It is important to understand such a correlation because most network optimization algorithms are concerned with, in one way or another, eliminating bottlenecks). Algorithm PAM_TWO is presented in Section 4. In Section 5, we study the accuracy of PAM_TWO and PAM_IMPROVED by comparing their predictions with exact solutions given by the MVA algorithm for 500 networks generated randomly as specified by Zahorjan et al. [22]. These networks have both fixed-rate and delay servers and possess characteristics of models of computer systems.

2. Two PAM Algorithms with $O(MK)$ Computation Time

Consider a closed queueing network with K routing chains (also called customer classes by other authors). Let N_k be the customer population of the k^{th} chain. We refer to the vector $\underline{N}=(N_1, N_2, \dots, N_K)$ as the *population vector* of the queueing network. If N_k in \underline{N} is greater than or equal to 1, then $\underline{N}-\underline{1}_k$ refers to the population vector \underline{N} with a chain k customer removed. Let τ_{mk} denote the load, or traffic intensity, of chain k customers at server m . (In models of communication networks, we also refer to τ_{mk} as the mean service time of chain k customers at server m with the assumption that visit ratios are the same for all servers visited by chain k customers.) We use $D_{mk}(\underline{n})$ to denote the approximate mean delay of chain k customers at server m , $q'_{mk}(\underline{n})$ to denote the approximate mean queue length of chain k customers at server m , and $T_k(\underline{n})$ to denote the approximate throughput of chain k customers, in a network whose population vector is \underline{n} .

PAM algorithms are based upon MVA formulas [15]. However, the iterations from population vector $\underline{0}$ to population vector $\underline{N}-\underline{1}_k$ in the MVA algorithm are skipped. Instead, mean queue length approximations $q'_{mh}(\underline{N}-\underline{1}_k)$ at population vector $\underline{N}-\underline{1}_k$ are obtained by distributing the population N_h of chain h over the servers it visits, proportional to τ_{mh} . This approach was the basis of proportional throughput upper bounds for single-chain networks in [7]. In multichain networks, this approach leads to approximations rather than upper bounds, as we shall see.

We present below two algorithms with time and space requirements of $O(MK)$. The first algorithm, PAM_BASIC, calculates throughputs of individual chains. The second algorithm, PAM_IMPROVED, calculates throughputs of chains and the utilizations of servers. The algorithm then checks server utilizations to see if any utilization exceeds 1. (This is possible because the chain throughputs are approximations.) The throughputs of those chains that visit a server whose utilization exceeds 1 are then scaled down.

Algorithm PAM_BASIC

Step 1: Calculate proportional approximations of mean queue lengths from

$$\gamma_{mk} = \tau_{mk} / \sum_{i=1}^M \tau_{ik}$$

for $m = 1, 2, \dots, M, k = 1, 2, \dots, K$

$$q'_{mh}(\underline{N}) = \gamma_{mh} N_h$$

$$q'_{mh}(\underline{N}-\underline{1}_k) = \begin{cases} q'_{mh}(\underline{N}) & \text{if } h \neq k \\ q'_{mh}(\underline{N}) - \gamma_{mh} & \text{if } h = k \end{cases}$$

for $m = 1, 2, \dots, M, h = 1, 2, \dots, K$, and $k = 1, 2, \dots, K$.

Step 2: Calculate approximate mean delay of chain k at server m and approximate throughput of chain k from the following MVA formulas:

$$D_{mk}(N) = \begin{cases} \tau_{mk} \left(1 + \sum_{h=1}^K q'_{mh}(N-1_k) \right) & \text{if } m \text{ is a fixed-rate server} \\ \tau_{mk} & \text{if } m \text{ is a delay server} \end{cases}$$

for $m = 1, 2, \dots, M$ and $k = 1, 2, \dots, K$, and

$$T_k(N) = \frac{N_k}{\sum_{m=1}^M D_{mk}(N)} \quad \text{for } k = 1, 2, \dots, K.$$

Total throughput of the network, if needed, is equal to the summation of the chain throughputs.

Note that PAM_BASIC requires $(3M+1)K$ multiplications and divisions. The number of additions and subtractions is also $O(MK)$ if the sums $\sum_{i=1}^M \tau_{ik}$ in Step 1 and the sums $\sum_{h=1}^K q'_{mh}(N-1_k)$ in Step 2 are computed prior to looping over $m = 1, 2, \dots, M$ and $k = 1, 2, \dots, K$.

Algorithm PAM_IMPROVED

The first two steps of this algorithm are the same as those of Algorithm PAM_BASIC.

Step 3: Calculate server utilizations from the following formula:

$$U_m(N) = \sum_{k=1}^K \tau_{mk} T_k(N)$$

for $m = 1, 2, \dots, M$, where $U_m(N)$ is the utilization of server m at population vector N .

Step 4: Find the largest utilization S_k among the fixed-rate servers visited by chain k ,

$$S_k = \max_{\substack{m \text{ in} \\ \text{chain } k \text{ and} \\ m \text{ is a fixed-rate server}}} U_m(N)$$

for $k=1, 2, \dots, K$.

Step 5: (Scale down throughputs of individual chains if necessary.)

$$\text{If } S_k > 1 \text{ then } T_k(N) := T_k(N) / S_k.$$

Step 6: Calculate total throughput, and recalculate server utilizations if the throughput of any chain has been scaled down,

$$T(N) = \sum_{k=1}^K T_k(N)$$

and

$$U_m(N) = \sum_{k=1}^K \tau_{mk} T_k(N), \quad \text{for } m = 1, 2, \dots, M.$$

3. Experimental Results

We study in this section the accuracy of PAM_BASIC and PAM_IMPROVED using queueing networks that have characteristics of models of communication networks. Errors in the approximate chain throughputs, end-to-end delays, and server utilizations predicted by the PAM algorithms are obtained by comparing them with exact results calculated by TCA.

A communication network is specified by a set of nodes interconnected by full-duplex communication links. In our queueing network model, each link is modeled by two fixed-rate servers, one for each direction of the link; nodes are not modeled. Each flow-controlled virtual channel is modeled by a closed chain; the chain population corresponds to the flow-control window size [8, 11, 12, 14]. The route of a virtual channel is specified by the sequence of nodes it visits, which uniquely determines the sequence of fixed-rate servers visited by the corresponding chain. An additional fixed-rate server is inserted between the destination node and the source node of the route to form a closed chain. This is referred to as the *source server* and its service rate represents the packet generation rate of the user process that provides input to the virtual channel.

We examine the accuracy of the PAM algorithms by first studying a network example. We then present statistical results from a 100-network experiment. (The networks considered in this section have fixed-rate servers only.) We also examine carefully a correlation between the approximation errors for a chain and the maximum server utilization among those servers visited by the chain. Understanding such a correlation is important in applying PAM algorithms to communication network design problems because most network optimization algorithms are concerned with rerouting or reconfiguring a network to alleviate congestion at bottlenecks.

Network example

The network has 12 nodes, 15 links, and 20 chains. The set of links specified by node pairs and their capacities are shown in Table 1. The average packet length is 240 bits. The mean service time at the source server of each chain is 0.3 second. The populations and routes for all chains are shown in Table 2.

The throughputs calculated by TCA, PAM_BASIC, and PAM_IMPROVED, and percentage errors of the two approximation algorithms, relative to TCA's exact results, are listed in Table 3. Note that the throughputs calculated by PAM_IMPROVED and PAM_BASIC are the same for those chains that do not visit a "bottleneck" queue, i.e., those that do not require execution of Step 5 in

PAM_IMPROVED. For those chains that do visit bottleneck queues, PAM_IMPROVED has smaller approximation errors than PAM_BASIC. The throughputs of individual chains calculated by PAM_BASIC and PAM_IMPROVED are also plotted together with exact values in Figures 1 and 2 respectively. The server loads are fairly heavy and unbalanced. The maximum server utilization, calculated by TCA, is 0.990. The average utilization over servers with nonzero utilizations is 0.398.

Link	Capacity (bits/second)
(1, 4)	1200
(8, 10)	9600
(1, 11)	9600
(8, 5)	4800
(1, 5)	2400
(2, 6)	2400
(3, 10)	1200
(7, 8)	4800
(9, 11)	2400
(12, 2)	4800
(1, 8)	1200
(3, 6)	9600
(4, 11)	9600
(7, 10)	1200
(9, 12)	1200

Table 1. Links and their capacities in the network example.

Chain	Population	Route (in node sequence)
1	2	5 8 7
2	3	6 2
3	2	6 3 10
4	3	6 2 12 9 11 4
5	4	9 12 2
6	3	8 5 1 11
7	2	6 3 10 8 1
8	2	9 11 1 8
9	2	7 8 5
10	2	10 8 1 11 4
11	3	12 2
12	3	12 2 6
13	2	1 11 9 12 2
14	2	6 3 10 8 7
15	3	8 5 1 11
16	3	9 11 1 5 8 10 3 6
17	2	8 10
18	2	2 6 3 10 8 7
19	2	8 10
20	3	6 2 12 9 11 4

Table 2. Populations and routes of chains in the network example.

Chain	TCA	PAM_BASIC		PAM_IMPROVED	
	Value	Value	Error(%)	Value	Error(%)
1	3.0946	3.021	2	3.021	2
2	3.0995	3.069	1	3.069	1
3	1.3954	1.764	26	1.437	3
4	1.8674	1.923	3	1.923	3
5	2.7153	2.784	3	2.784	3
6	3.0215	2.974	2	2.974	2
7	1.0678	1.307	22	1.065	0
8	2.0149	2.07	3	2.070	3
9	3.0655	2.989	2	2.989	2
10	2.2768	2.148	6	2.148	6
11	3.2901	3.258	1	3.258	1
12	3.1181	3.058	2	3.058	2
13	1.6333	1.607	2	1.607	2
14	1.3055	1.63	25	1.327	2
15	3.0215	2.974	2	2.974	2
16	2.197	2.254	3	2.254	3
17	3.3054	3.279	1	3.279	1
18	1.1794	1.431	21	1.166	1
19	3.3054	3.279	1	3.279	1
20	1.8674	1.923	3	1.923	3
Total throughput	47.8417	48.744	2	47.606	0

Table 3. Exact throughputs, approximate throughputs, and approximation errors for the network example.

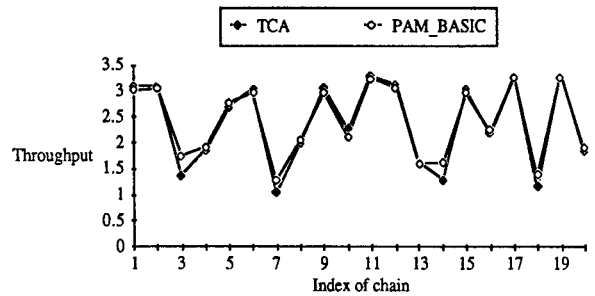


Figure 1. Throughputs of chains calculated by TCA and PAM_BASIC.

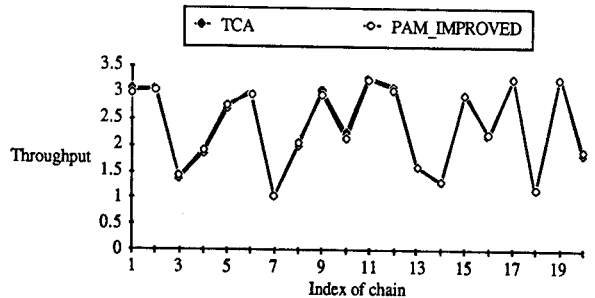


Figure 2. Throughputs of chains calculated by TCA and PAM_IMPROVED.

The approximation errors of PAM_BASIC for chains 3, 7, 14, and 18 are significantly larger than errors for other chains. Table 1 shows that those four chains are the ones that visit the server (communication channel) from node 3 to node 10. This server has the maximum utilization, calculated by TCA to be 0.990, in the network. An explanation of these large errors is the following: Proportional approximation, as the basis for throughput upper bounds in [7] and PAM algorithms herein, underestimates the mean queue lengths of a chain at servers that are highly utilized relative to other servers visited by the chain. While PAM_BASIC has large errors for such chains, Table 3 and Figure 2 show that PAM_IMPROVED do not have large errors in its throughputs for chains 3, 7, 14, and 18, suggesting that its throughput scaling operation in Step 5 is quite effective.

Despite the relatively large errors of PAM_BASIC in predicting the throughputs of chains visiting a bottleneck sever, it is still quite useful in network design algorithms. For example, it is used in the optimal routing algorithm presented in [11] because it is the fastest of the PAM algorithms. Also, in choosing the best route for a chain among various candidates, a route that visits a highly utilized server will not likely be chosen; thus the accuracy of PAM_BASIC's throughput prediction for such a route is irrelevant.

Most network optimization algorithms are concerned, in one way or another, with the alleviation of traffic at highly utilized servers. In such applications, the accuracy of PAM_BASIC at an intermediate step of the optimization algorithm is not very important. For example, in designing the link capacities of a network, identifying the communication channel from node 3 to node 10 to be a bottleneck will probably lead to an increase in its capacity. Suppose the communication channel capacity is increased from 1200 bits/second to 4800 bits/second and PAM_BASIC is again applied to calculate chain throughputs. Figure 3 shows the approximation errors of PAM_BASIC for the two cases of the network example, where "No bottleneck" denotes the network with a 4800 bits/second communication channel from node 3 to node 10. Note that after such a capacity increase, approximation errors of PAM_BASIC for all chains fall below 5%. The largest server utilization in the so-called "No bottleneck" network is 0.870, which is still quite high. The correlation between errors of PAM_BASIC and maximum server utilizations is illustrated again in Figure 7 for the 100 randomly generated networks to be presented below. Additional experimental results are presented in [9].

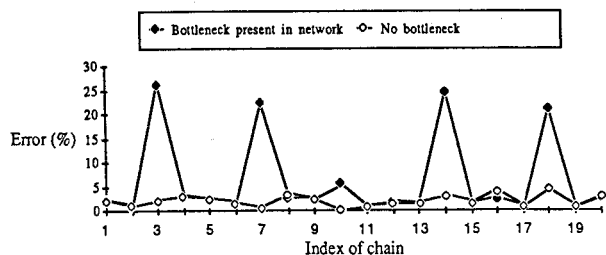


Figure 3. Approximation errors of PAM_BASIC with and without a bottleneck in the network example.

Statistical results from 100 networks

We next present some statistical results from applying PAM_IMPROVED and TCA to 100 models of communication networks generated randomly as described in [8, 9, 19]. Statistics of parameters of the 100 networks generated are summarized in Table 4, where #packets is the summation of all chain populations in a network, and the average queue utilization is computed over those servers with nonzero utilizations.

Percentage errors of total throughputs, chain throughputs, and server utilizations of PAM_IMPROVED, relative to TCA exact results, are shown in Figures 4, 5, and 6. The maximum percentage error in the total throughput calculated by PAM_IMPROVED is 3.55% (see Figure 4). Figures 5 and 6 show that the chain throughputs and server utilizations calculated by PAM_IMPROVED are also very accurate with only a small number of exceptions. The maxima, means, and variances of the percentage errors of PAM_IMPROVED for the three performance measures are shown in Table 5.

	Max.	Min.	Ave.
1. #nodes	25	7	15
2. #links	37	9	20
3. #queues	74	18	40
4. #chains	43	8	23
5. #packets	107	19	57
6. Ave. queue util.	0.546	0.291	0.419

Table 4. Statistics of 100 test networks.

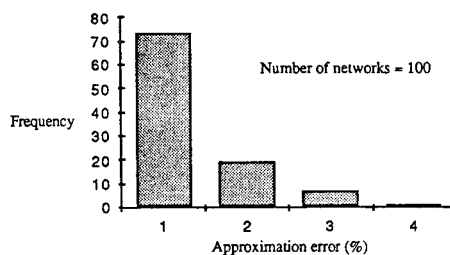


Figure 4. Distribution of approximation errors of total throughputs calculated by PAM_IMPROVED for the 100 networks.

Figure 7 shows a strong correlation between the approximation error of PAM_BASIC in predicting a chain's throughput and the maximum utilization among servers visited by the chain. The points plotted in Figure 7 were obtained by selecting one chain from each of the 100 randomly generated networks and applying PAM_BASIC and TCA to calculate its approximate and exact throughputs. Note that PAM_BASIC is quite accurate if the maximum utilization is less than about 0.9.

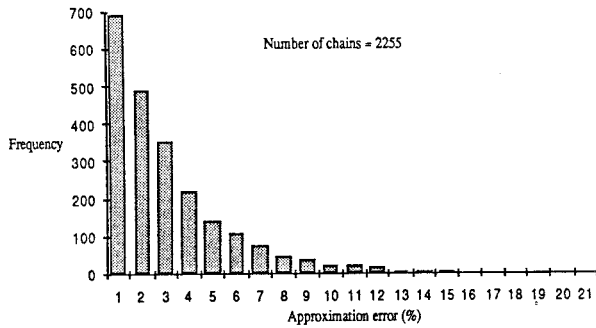


Figure 5. Distribution of approximation errors of throughputs of individual chains calculated by PAM_IMPROVED for the 100 networks.

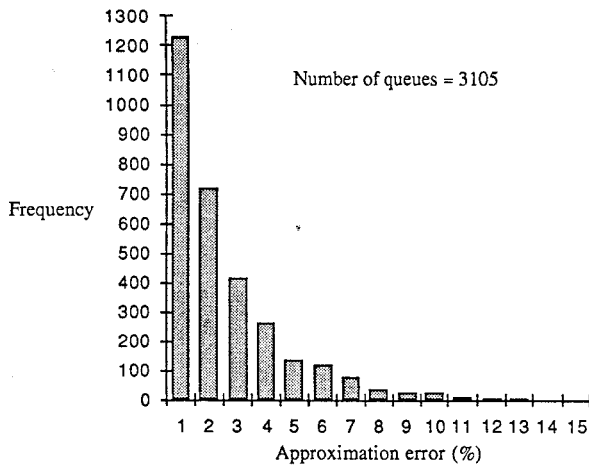


Figure 6. Distribution of approximation errors of server utilizations calculated by PAM_IMPROVED for the 100 networks.

	Statistics of percentage errors				
	#samples	Minimum	Maximum	Mean	Variance
Total throughput	100	0.009	3.554	0.824	0.002
Throughput of chain	2255	0.000	20.503	2.667	7.167
Utilization of queue	3104	0.000	14.769	2.064	4.263

Table 5. Statistics of approximation errors of total throughput, chain throughputs, and server utilizations calculated by PAM_IMPROVED for the 100 networks.

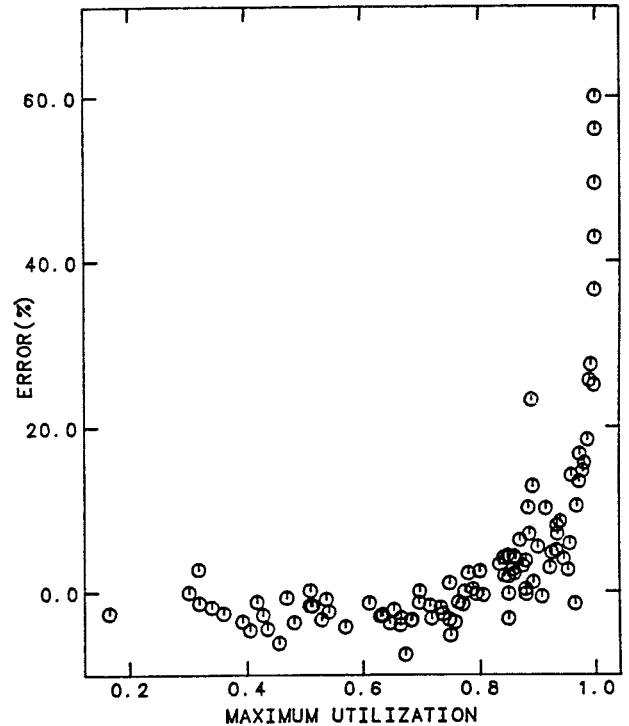


Figure 7. Correlation between percentage error of chain throughput calculated by PAM_BASIC and the maximum utilization of the servers visited by the chain.

4. PAM Algorithm with $O(MK^2)$ Computation Time

A simple approach to improve the accuracy of the PAM algorithms presented in the previous section, without resorting to an iterative solution, is to execute the last two steps of the MVA recursion instead of just the last step, again using the proportional approximation to get initial mean queue length estimates. Our algorithm, presented below, will be referred to as PAM_TWO. Such an approach to trade computation time for accuracy is not unlike the approach of Linearizer which improves the accuracy of Schweitzer's algorithm [4, 16] and the approach of bound hierarchies [5, 7]. The time requirement of PAM_TWO is $O(MK^2)$. The space requirement is $O(MK)$.

Algorithm PAM_TWO

Step 1: Calculate proportional approximations of mean queue lengths from

$$\gamma_{mk} = \tau_{mk} / \sum_{i=1}^M \tau_{ik}$$

for $m = 1, 2, \dots, M, k = 1, 2, \dots, K$

$$q'_{ml}(N) = \gamma_{ml} N_l$$

$$q'_{ml}(N-1_k-1_h) = \begin{cases} q'_{ml}(N) & \text{if } l \neq k \text{ and } l \neq h \\ q'_{ml}(N) - \gamma_{ml} & \text{if } (l=k \text{ or } l=h) \text{ and } h \neq k \\ q'_{ml}(N) - 2\gamma_{ml} & \text{if } l=k=h \end{cases}$$

for $m = 1, 2, \dots, M$, $l = 1, 2, \dots, K$, $h = 1, 2, \dots, K$, and $k = 1, 2, \dots, K$. Note that $q'_{ml}(N-1_k-1_h)$ may have a negative value if $N_l < 2$. In this case, $q'_{ml}(N-1_k-1_h)$ is assigned the value of zero.

Step 2: For $k = 1, 2, \dots, K$, repeat Step 2.1 and Step 2.2.

Step 2.1: For $h=1, 2, \dots, K$, repeat the following calculations:

$$D_{mh}(N-1_k) = \begin{cases} \tau_{mh} \left(1 + \sum_{l=1}^K q'_{ml}(N-1_k-1_h) \right) & \text{if } m \text{ is a fixed-rate server} \\ \tau_{mh} & \text{if } m \text{ is a delay server} \end{cases}$$

for $m = 1, 2, \dots, M$,

$$T_h(N-1_k) = \begin{cases} \frac{N_h}{\sum_{m=1}^M D_{mh}(N-1_k)} & \text{if } h \neq k \\ \frac{N_h - 1}{\sum_{m=1}^M D_{mh}(N-1_k)} & \text{if } h = k, \end{cases}$$

$$q'_{mh}(N-1_k) = D_{mh}(N-1_k) \times T_h(N-1_k)$$

for $m = 1, 2, \dots, M$.

Step 2.2: Calculate the following,

$$D_{mk}(N) = \begin{cases} \tau_{mk} \left(1 + \sum_{h=1}^K q'_{mh}(N-1_k) \right) & \text{if } m \text{ is a fixed-rate server} \\ \tau_{mk} & \text{if } m \text{ is a delay server} \end{cases}$$

for $m = 1, 2, \dots, M$, and

$$T_k(N) = \frac{N_k}{\sum_{m=1}^M D_{mk}(N)}$$

Step 3: Execute Steps 3-6 of PAM_IMPROVED.

5. More Experimental Results

The 100 networks used in Section 3 have characteristics of communication networks: large number of chains with sparse routes and small chain populations. In this section, we examine the accuracy of PAM_TWO and PAM_IMPROVED using 500 networks generated randomly according to the specification used in the study of errors of iterative AMVA algorithms by Zahorjan et al. [22, Table 2]. (There is one difference: We set the minimum number of servers at 5 instead of 2.) The network generation parameters are shown in Table 6, where U indicates the uniform distribution. Networks generated by these parameters possess characteristics of computer systems rather than communication networks. Note that every network generated has four chains and each chain visits every server in the network. Compared to the 100 networks in Section 3, these 500 networks have larger chain populations and much larger service time variations; there are also delay servers in addition to fixed-rate servers.

Scheduling Discipline	Prob[Infinite Server] = 0.05 Prob[Load Independent] = 0.95
Population Size	Class 1 : U(1, 10) Class 2 : U(1, 5) Class 3 : U(1, 5) Class 4 : U(1, 5)
Loadings	U(0.1, 50.0)
Number of Centers	U(5, 50)

Table 6. Network generation parameters.

We used PAM_TWO and PAM_IMPROVED to calculate approximate chain throughputs and the MVA algorithm to calculate exact chain throughputs for the 500 networks generated. The average and maximum percentage errors in the approximate chain throughputs, relative to exact MVA solutions, are shown in Table 7 for PAM_IMPROVED and PAM_TWO. Both PAM algorithms have small average errors but fairly large maximum errors. Table 7 also shows that PAM_TWO is more accurate than PAM_IMPROVED.

Although the maximum percentage errors are large, Figure 8 shows that only a very small fraction of the chains have percentage errors larger than 10% while more than three-fourths of the chains have less than 1% error for PAM_TWO and less than 3% error for PAM_IMPROVED.

In Section 3, we found a strong correlation between the approximation error of a chain's throughput calculated by PAM_BASIC and the maximum utilization among servers visited by the chain. Table 8 shows that such a correlation

also exists for PAM_IMPROVED and PAM_TWO. For those chains which do not visit servers with utilizations between 0.95 and 1, both the average and maximum approximation errors become substantially smaller for both PAM algorithms (see row 2 of Table 8).

Technique	Measure	Error(%)	#Networks
PAM_IMPROVED	Average	2.3	500
	Maximum	40.3	
PAM_TWO	Average	0.8	500
	Maximum	30.8	

Table 7. Relative errors in chain throughputs calculated by PAM_IMPROVED and PAM_TWO for 500 networks.

Max. util. ≤	Approximation errors of chain throughputs					
	PAM_IMPROVED			PAM_TWO		
	Max.	Ave.	#Chains	Max.	Ave.	#Chains
1.0	40.3	2.3	2000	30.8	0.8	2000
0.95	11.4	1.6	1784	9.6	0.5	1860
0.9	9.4	1.5	1732	5.4	0.5	1784
0.8	8.4	1.2	1556	3.1	0.3	1612
0.7	4.8	1.0	1392	1.7	0.2	1416
0.6	3.8	0.8	1136	1.4	0.2	1164
0.5	2.1	0.6	744	1.4	0.1	788

Table 8. Correlation between percentage errors in approximate chain throughputs and maximum server utilizations for PAM_IMPROVED and PAM_TWO.

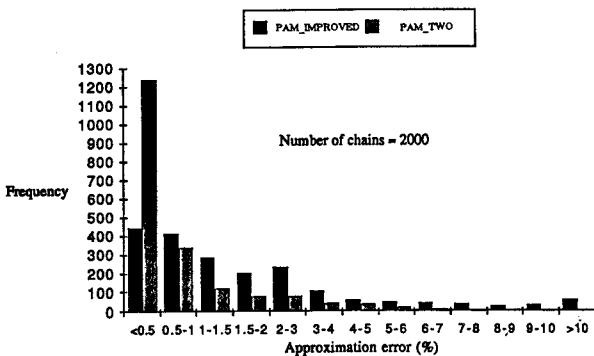


Figure 8. Distribution of approximation errors of throughputs of individual chains calculated by PAM_IMPROVED and PAM_TWO for 500 networks.

Although we generated the 500 networks similar to what Zahorjan et al. did in their study of three iterative AMVA algorithms, a direct comparison of the accuracy of PAM and the iterative AMVA algorithms cannot be made. There are two reasons. First, we evaluated approximation errors in chain throughputs and server utilizations (we are also interested in mean end-to-end delays of virtual channels as a performance measure but these can be obtained from chain throughputs and Little's formula). The iterative AMVA algorithms were evaluated by the maximum approximation error in mean queue lengths $q_{mk}(N)$ in [22] and also mean delays $D_{mk}(N)$ and server utilizations $U_{mk}(N)$ in [4] for all m and k . The performance measures of AMVA have finer granularity than ours. On the other hand, the AMVA studies used a specially defined measure of error called "tolerance error" instead of the usual relative error which we use. The tolerance error used by Zahorjan et al. [22] is the following:

$$\max_{m,k} [q_{mk}(N) - q_{mk}^*(N)] / N_k$$

where $q_{mk}(N)$ is the approximate value and $q_{mk}^*(N)$ is the exact value. Note that N_k is used in the denominator instead of $q_{mk}^*(N)$. It was argued that tolerance error was used in place of relative error because the latter measure is very sensitive to small (absolute) errors in small values. (That is exactly how we got most of the large percentage errors for PAM algorithms since we use relative error as our measure.)

The tolerance errors in mean queue lengths of the three iterative AMVA algorithms from the last column of Table 4 in [22] are reproduced in Table 9. However, it is not possible to make a direct comparison between the relative errors in Table 7 for PAM and the tolerance errors in Table 9 for AMVA. On the one hand, for the same numerical values of approximate and exact solutions, tolerance errors are much smaller than relative errors. On the other hand, $q_{mk}(N)$ is a measure having finer granularity than chain throughput (which is obtained from a summation of mean delays) and is expected to have larger approximation errors.

Technique	Measure	Tolerance Error (%)	#Networks
Schweitzer's method	Average	2.5	2000
	Maximum	30.5	
Linearizer	Average	0.5	2000
	Maximum	2.4	
AQL	Average	0.3	2000
	Maximum	3.3	

Table 9. Tolerance errors in mean queue lengths calculated by 3 iterative AMVA algorithms from Zahorjan et al. [22].

To carry out a direct comparison of AMVA and PAM, we will have to implement the AMVA algorithms or modify our PAM algorithms (which we have not done). In fact, the comparison can only be made for networks with a small number of closed chains, which is the domain of applications of iterative AMVA algorithms. For such networks, we conjecture that Linearizer and AQL are more accurate than PAM algorithms (at the expense of substantially more computation time for both Linearizer and AQL and more space for Linearizer). For networks with a large number of closed chains, however, the convergence behavior of iterative AMVA algorithms is still unknown. In fact, Zahorjan et al. [20] show that the accuracy of iterative AMVA algorithms gets worse as the number of closed chains increases from 1 to 4.

6. Conclusion

PAM algorithms have been designed for the approximate solution of queueing networks with a large number of closed chains and relatively small chain populations. For this class of networks, PAM is the only available technique whose accuracy has been studied and whose time and space requirements can be met. Because they are noniterative, they are suitable for many communication network design and optimization problems that are typically based upon a heuristic search for an optimum; a very fast evaluation of network performance is needed at each step of such a heuristic search.

REFERENCES

- [1] Y. Bard, "Some Extensions to Multiclass Queueing Network Analysis," *Performance of Computer Systems*, M. Arato et al. (ed.), North-Holland, Amsterdam, 1979.
- [2] F. Baskett, K. M. Chandy, R. R. Muntz, and F. Palacios, "Open, Closed and Mixed Networks of Queues with Different Class of Customers," *JACM*, Vol. 22, 1975, pp. 248-260.
- [3] W.-M. Chow, "Approximations for Large Scale Closed Queueing Networks," *Performance Evaluation*, Vol. 3, No. 1, Feb. 1983, pp. 1-12.
- [4] K. M. Chandy and D. Neuse, "Linearizer: A Heuristic Algorithm for Queueing Network Models of Computer Systems," *Comm. ACM*, Vol. 25, No. 2, Feb. 1982, pp. 126-134.
- [5] D. L. Eager and K. C. Sevcik, "Bound Hierarchies for Multiple-Class Queueing Networks," *Journal of ACM*, Vol. 33, No. 1, Jan. 1986, pp. 179-206.
- [6] K. P. Hoyme, S. C. Bruell, P. V. Afshari, and R. Y. Kain, "A Tree-Structured Mean Value Analysis Algorithm," *ACM Trans. on Computer Systems*, Vol. 4, No. 2, May 1986, pp. 178-185.
- [7] C.-T. Hsieh and S. S. Lam, "Two Classes of Performance Bounds for Closed Queueing Networks," *Performance Evaluation*, Vol. 7, No. 1, 1987, pp. 3-30.
- [8] C.-T. Hsieh, "Models and Algorithms for the Design of Store-and-Forward Communication Networks," Ph.D. Thesis, Department of Computer Sciences, University of Texas at Austin, 1987.
- [9] C.-T. Hsieh and S. S. Lam, "A Noniterative Approximate Solution Method for Closed Multichain Queueing Networks," Technical Report TR-87-28, Dept. of Computer Sciences, University of Texas at Austin, July 1987.
- [10] S. S. Lam and Y. L. Lien, "A Tree Convolution Algorithm for the Solution of Queueing Networks," *Comm. ACM*, Vol. 26, No. 3, March 1983, pp. 203-215.
- [11] S. S. Lam and Ching-Tarn Hsieh, "Modeling, Analysis, and Optimal Routing of Flow-Controlled Communication Networks," Technical Report TR-87-24, Dept. of Computer Sciences, Univ. of Texas at Austin, June 1987.
- [12] S. S. Lam and Y. L. Lien, "Modeling and Analysis of Flow-Controlled Packet Switching Networks," *Proc. 7th Data Communications Symposium*, Mexico City, October 1981.
- [13] M. Reiser and H. Kobayashi, "Queueing Networks with Multiple Closed Chains: Theory and Computational Algorithms," *IBM J. Res. Develop.*, Vol. 21, 1975, pp. 283-294.
- [14] M. Reiser, "A Queueing Network Analysis of Computer Communication Networks with Window Flow Control," *IEEE Trans. on Communication*, Vol. COM-27, 1979, pp. 1199-1209.
- [15] M. Reiser and S. Lavenberg, "Mean Value Analysis of Closed Multichain Queueing Networks," *Journal of ACM*, Vol. 27, No. 2, April 1980, pp. 313-322.
- [16] P. Schweitzer, "Approximate Analysis of Multiclass Closed Networks of Queues," *Proc. Int. Conf. Stochastic Control and Optimization*, Amsterdam, 1979.
- [17] E. de Souza e Silva, S. S. Lavenberg, and R. R. Muntz, "A Clustering Approximation Technique for Queueing Network Models with a Large Number of Chains," *IEEE Trans. on Computers*, Vol. C-35, No. 5, May 1986, pp. 419-430.
- [18] E. de Souza e Silva and R. R. Muntz, "A Note on the Computational Cost of the Linearizer Algorithm for Queueing Networks," Technical Report CSD-870025, Dept. of Computer Science, UCLA, June 1987.
- [19] K. Steiglitz, P. Weiner, and D. J. Kleitman, "The Design of Minimum Cost Survivable Networks," *IEEE Transactions on Circuit Theory*, Vol. CT-16, 1969, pp. 455-460.
- [20] S. Tucci and C. H. Sauer, "The Tree MVA Algorithm," *Performance Evaluation*, Vol. 5, No. 3, August 1985, pp. 187-196.
- [21] J. Zahorjan, K. C. Sevcik, D. L. Eager, and B. Galler, "Balanced Job Bound Analysis of Queueing Networks," *Comm. ACM*, Vol. 25, No. 2, 1982, pp. 132-141.
- [22] J. Zahorjan, D. L. Eager, and H. Sweillam, "Accuracy, Speed, and Convergence of Approximate Mean Value Analysis," Technical Report, Dept. of Computer Science, Univ. of Washington, August 1986; to appear in *Performance Evaluation*.