Performance Evaluation of a Single Node with General Arrivals and Service

Alexandre Brandwajn[1] and Thomas Begin[2]

[1] Baskin School of Engineering, University of California Santa Cruz, USA
[2] LIP, Université Lyon 1, France
alexb@soe.ucsc.edu, thomas.begin@ens-lyon.fr

Abstract. Queueing delays experienced by packets buffered at a node are among the most difficult to predict when considering the performance of a flow in a network. The arrivals of packets at a node tend to be highly variable so that a finite-buffer single-server queue with general arrivals and service emerges as a natural model of a network link. In this paper we propose an approach to the solution of such a queue when the times between arrivals and service times are represented as acyclic phase-type distributions. The proposed solution approach, based on the use of conditional probabilities, is conceptually simple, easy to implement in a standard computer language, numerically robust and reasonably fast. In addition to standard steady-state probabilities and queue size averages, the proposed approach produces the probabilities of the state of the queue found by an arriving packet, in particular, the packet loss probability, directly linked to the QoS perceived by the user.

Keywords: Capacity planning; buffer sizing; quality of service; modeling; single-server queue; *Ph*/*Ph*/*1*/*N* queue.

1 Introduction and state of the art

When considering the performance of a flow at a network node, queueing delays experienced by packets (or frames) buffered at the node are among the most difficult to predict [23]. Given that a network interface can only transmit one packet at a time, a single-server queue seems a natural model of a network link. There is general agreement that the arrivals of packets at a node are highly variable and a Poisson process is not a valid representation [17, 10, 30]. With variable size packets, the transmission time (service time) is also variable and, generally, not exponentially distributed. Additionally, buffers have finite capacity, sometimes quite small [40], so that the buffer overflow probability and resulting packet loss are of major interest (e.g., because of its effect on TCP congestion window [3]). Hence, the underlying queueing model is a so-called *G*/*G*/*1*/*N* queue [2].

While the equilibrium solution of the *M/M/1/N* queue (with Poisson arrivals and exponential service) is straightforward [2], there are no readily usable analytical results for the more realistic *G/G/1/N* queue. Thus, there are essentially three possible approaches to the solution of such a queue.

The first one is to solve the state equations of the queueing system numerically. A commonly adopted approach is to represent the times between arrivals, as well as the service times, as phase-type distributions [28, 27], and then attempt to solve the resulting balance equations. Here, besides general methods for systems of linear equations (e.g., [35, 12]), which include direct iterative methods (e.g., Gauss-Seidel, over-relaxation [33]), specific approaches for Quasi Birth and Death processes could be used (e.g., Matrix Geometric [27, 26, 24], Spectral Expansion [25]). The latter methods may require a non-trivial degree of mathematical sophistication [15].

The second possible approach is to use discrete-event simulation. While simple in principle, this solution turns out to be fraught with problems. To achieve a reasonable accuracy, serious problems appear when simulating heavy-tailed distributions (such as the Pareto distribution for the time between arrivals) involving rare but important events (cf. [4, 13, 5]). Additionally, the simulation times tend to elongate considerably when the link utilization approaches saturation. As a result, simulation times may become long, making difficult the systematic exploration of a large number of parameter values.

When there is no easy exact solution and the simulation becomes exceedingly long, a natural alternative is to look for simple approximations that would be accurate enough to provide meaningful estimates for delays and loss probabilities. There are several bounds and approximations for the *G/G/1* queue with an infinite buffer [21, 22, 34, 37, 2, 19, 9, 31], which may be used to estimate delays for lightly loaded links but not to assess loss probabilities or dimension buffers. There are fewer results for the *G/G/1/N* queue [20, 38, 39], either limited to the first two moments of the inter-arrival and service times distributions or based on heavy-traffic assumption, which may not be applicable in a realistic network with moderate levels of losses. To the best of our knowledge there are no readily applicable good approximate results for the loss probability.

This paper falls into the first category of approaches. Within the framework of acyclic phase-type representation of general distributions, we propose to use conditional probabilities to obtain a conceptually simple numerical solution that is easy to implement, generally fast, and stable in practice. Our approach allows us to obtain the state probabilities at instants of packet arrival, and hence, the packet loss probability, as well as such customary measures as the expected time in the queue. We have implemented our method in standard C language, and, in the very many test we ran, it performed well, including with phase representations of heavy-tailed distributions.

Our paper is organized as follows. In Section 2 we describe our model and the proposed solution method. Section 3 discusses the practical performance of the method. We use two non-trivial examples to compare the performance of the proposed approach with that of a discrete-event simulation. Section 4 concludes this paper.

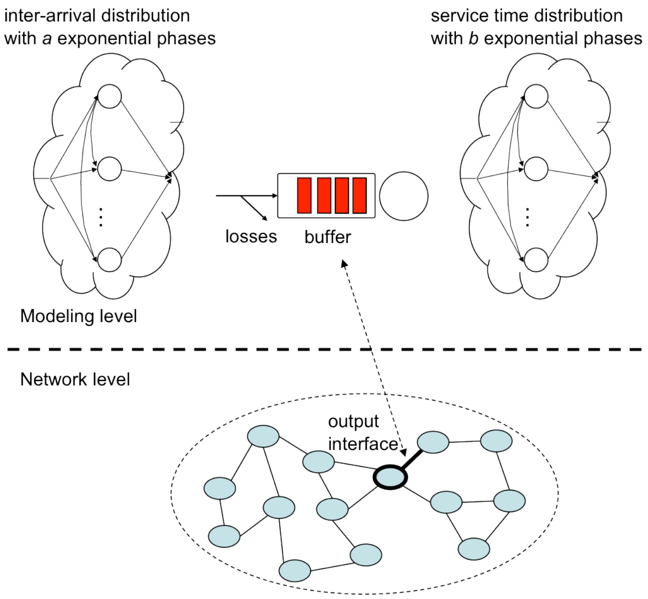
inter-arrival distribution
with *a* exponential phases
service time distribution
with *b* exponential phases
losses
buffer
Modeling level
Network level
output
interface

Fig. 1. Modeling a network link or an output interface as a *Ph*/*Ph*/*1*/*N* queue

2 Model and its Solution

We focus our attention on a single node viewed as a *G*/*G*/*1*/*N* queue with a FIFO discipline. We adopt the acyclic phase-type representation of the time between arrivals of packets and of their service times. We use the terms requests and packets (or frames) interchangeably. As is well known, any distribution can be approximated arbitrarily closely by a finite number of exponential phases [28, 27], and algorithms exist to represent a given distribution by a phase-type distribution (e.g., [16, 6, 29, 18, 36, 11]). Several authors address more specifically the representation of heavy-tailed distributions (for instance, [16, 18, 11]). Thus, our *G*/*G*/*1*/*N* model becomes what is known as the *Ph*/*Ph*/*1*/*N* queue. The advantage of the phase-type representation is that we are able to easily obtain the state equations for our model in the form of the regular balance equations.

We denote by a the number of phases in the arrival process, and by b the number of phases in the service process (see Fig. 1). We let λ_j be the rate (intensity) of phase

$j, j = 1,...,a$, r_{jl} the probability that phase j is followed by phase l ($l > j$), $\hat{r}_j$ the probability that the arrival process ends after phase j, and τ_j the probability that the arrival process starts in phase j. Similarly, for the service process, we denote by μ_i the rate (intensity) of phase $i, i = 1,...,b$, q_{il} the probability that phase i is followed by phase l ($l > i$), $\hat{q}_i$ the probability that the service process end after phase i, and σ_i the probability that the service process starts in phase $i, i = 1,...,b$. As an example, with a hyper-exponential service time distribution, we have $q_{il} = 0$ for all values of i and l, and $\hat{q}_i = 1$ for all phases i. Clearly, for any distribution, we must have $\sum_{j=1}^{a} \tau_j = 1$, $\sum_{j=1}^{a} \hat{r}_j + \sum_{l=j+1}^{a} r_{jl} = 1$, and, similarly, $\sum_{i=1}^{b} \sigma_i = 1$, $\sum_{i=1}^{b} \hat{q}_i + \sum_{l=i+1}^{b} q_{il} = 1$.

A possible state description for the *Ph/Ph/1/N* queue is $p(j,i,n)$ which is the probability that the arrival process is in its phase $j, j = 1,...,a$, the service process is in its phase $i, i = 1,...,b$, and there is a total of $n, n = 0,...,N$ requests in the system (queued and in service). The state probabilities are subject to the normalizing condition $\sum_{j=1}^{a}\sum_{i=1}^{b}\sum_{n=0}^{N} p(j,i,n) = 1$. The balance equations together with this normalizing condition define a system of linear equations for the probabilities $p(j,i,n)$, and classical numerical solution methods for this type of queueing system seek to solve this system of linear equations (e.g.,Chapter 3 [7], Chapter 8 [14]) as discussed in the Introduction.

We adopt the same basic state description but we use the fact that, from the definition of conditional probability, we have

$$p(j,i,n) = p(j,i \mid n)\, p(n) \tag{1}$$

where $p(n)$ is the probability that there are $n, n = 0,...,N$ requests in the system, and $p(j,i \mid n)$ is the conditional probability of the current phases of arrival and service given the number of requests. Note that, for $n = 0$, the index of the current phase of the service process becomes meaningless, and we write simply $p(j \mid 0)$ to denote the state of the arrival process given that the system is idle.

It is not difficult to show that the probability $p(n)$ can be expressed as

$$p(n) = \frac{1}{G}\prod_{k=1}^{n} \frac{\alpha(k-1)}{u(k)}, \tag{2}$$

where

$$\alpha(n) = \begin{cases} \sum_{j=1}^{a} \lambda_j \hat{r}_j \sum_{i=1}^{b} p(j,i \mid n), & \text{if } n > 0 \\ \sum_{j=1}^{a} \lambda_j \hat{r}_j\, p(j \mid 0), & \text{if } n = 0 \end{cases} \tag{3}$$

is the conditional rate of request arrivals with n requests in the system, and

$$u(n) = \sum_{i=1}^{b} \mu_i \hat{q}_i \sum_{j=1}^{a} p(j,i \mid n), \qquad n \geq 1 \tag{4}$$

denotes the conditional rate of request service given n. G is a normalizing constant chosen so that $\sum_{n=0}^{N} p(n) = 1$.

The probability of a lost request (packet loss probability) can be obtained as

$$p_{loss} = \frac{\alpha(N)p(N)}{\sum_{k=0}^{N} \alpha(k)p(k)}, \tag{5}$$

the probability that there is no wait is given by

$$p_{no\,wait} = \frac{\alpha(0)\,p(0)}{\sum_{k=0}^{N} \alpha(k)\,p(k)}. \tag{6}$$

More generally, we have for the probability that an arriving request finds n requests already present in the system

$$P_A(n) = \frac{\alpha(n)\,p(n)}{\sum_{k=0}^{N} \alpha(k)\,p(k)}, \quad n = 0, \ldots, N, \tag{7}$$

Thus, to obtain the quantities of interest (such as the loss probability), it suffices to obtain the conditional rates $\alpha(n)$ and $u(n)$, i.e., to solve for the conditional probabilities $p(j,i \mid n)$. Using the basic identity (1) in the balance equations, we obtain the following system of equations for the conditional probabilities $p(j,i \mid n)$:

for $n = 2, \ldots, N-1$

$$\begin{aligned} p(j,i \mid n)(\lambda_j + \mu_i) &= \sum_{l=1}^{j-1} p(l,i \mid n)\lambda_l r_{lj} + \sum_{l=1}^{i-1} p(j,l \mid n)\mu_l q_{li} \\ &+ u(n)\tau_j \sum_{l=1}^{a} p(l,i \mid n-1)\lambda_l \hat{r}_l / \alpha(n-1) + \alpha(n)\sigma_i \sum_{l=1}^{b} p(j,l \mid n+1)\mu_l \hat{q}_l / u(n+1) \end{aligned} \tag{8}$$

for $n = N$

$$\begin{aligned} p(j,i \mid N)(\lambda_j + \mu_i) &= \sum_{l=1}^{j-1} p(l,i \mid N)\lambda_l r_{lj} + \sum_{l=1}^{i-1} p(j,l \mid N)\mu_l q_{li} \\ &+ u(N)\tau_j \sum_{l=1}^{a} p(l,i \mid N-1)\lambda_l \hat{r}_l / \alpha(N-1) + \tau_j \sum_{l=1}^{a} p(l,i \mid N)\lambda_l \hat{r}_l \end{aligned} \tag{9}$$

for $n = 1$

$$\begin{aligned} p(j,i \mid 1)(\lambda_j + \mu_i) &= \sum_{l=1}^{j-1} p(l,i \mid 1)\lambda_l r_{lj} + \sum_{l=1}^{i-1} p(j,l \mid 1)\mu_l q_{li} \\ &+ u(1)\tau_j \sigma_i \sum_{l=1}^{a} p(l \mid 0)\lambda_l \hat{r}_l / \alpha(0) + \alpha(1)\sigma_i \sum_{l=1}^{b} p(j,l \mid 2)\mu_l \hat{q}_l / u(2) \end{aligned} \tag{10}$$

for $n = 0$

$$p(j \mid 0)\lambda_j = \sum_{l=1}^{j-1} p(l \mid 0)\lambda_l r_{lj} + \alpha(1)\sum_{i=1}^{b} p(j,i \mid 1)\mu_i \hat{q}_i / u(2) \tag{11}$$

Note that we have $\sum_{j=0}^{a}\sum_{i=0}^{b} p(j,i\,|\,n) = 1$ for $n > 0$ and $\sum_{j=0}^{a} p(j\,|\,0) = 1$, i.e., we have a separate normalizing condition for each value of the number of requests in the system.

We propose a simple iteration akin to a recent iterative method for $C_m/C_k/c$-like queues [8] to solve this system of non-linear equations. Details of this easy-to-implement iteration are given in the Appendix. In the next section, we apply our approach to study the effect of buffer size on loss probability in a network node. Additionally, we compare the performance of our method with discrete-event simulation when the times between arrivals are represented by a Pareto-like distribution.

3 Performance of the Proposed Approach

In our first example, we study the performance of our approach using a model of a node with a finite buffer of 50 requests. The request arrival process is represented by a Pareto-like distribution. The latter was obtained using the PhFit package [16], and has a total of 16 exponential phases, 10 of which are used for the heavy-tailed part of the distribution. The parameters of this phase distribution are given in Table 1 in the Appendix. The service times in our *Ph/Ph/1/50* model are represented by a three-phase distribution chosen so as to match the first three moments of a reported packet mix in real networks [1] on a link operating at about 2.5 Gb/s. We used the method proposed by Bobbio et al. [6] to perform the matching of the first three moments. We give in Table 2 in the Appendix the parameters of the resulting phase distribution.

Fig. 2 and Fig. 3 illustrate the results obtained. In all examples we set the convergence stringency of our algorithm at $\varepsilon = 10^{-7}$ (increasing the convergence stringency does not change our results). In Fig. 2 we show the number of iterations in our method needed to solve the *Ph/Ph/1/50* model for server utilization ranging from 0.1 to 0.9. Fig. 3 displays the execution times of a C implementation on an Intel processor running at 2.2 GHz. For comparison, we include the execution times of a discrete-event simulation of the same model on a processor of the same speed. The simulation has been run long enough to attain a confidence interval of no more than 5% around the exact value of the average number of requests in the queue at 95% confidence level. We used 7 independent replications in our runs.

We observe in Fig. 2 that, for the server utilization levels considered, the proposed method requires only a few tens of iterations (below 100 in all cases). The simulation performance is displayed in terms of the number of packets generated to achieve the desired confidence intervals, and ranges from a total of about 350,000 to 8,000,000 packets. In Fig. 3 we compare the CPU times for both approaches. In the example considered, the CPU time for the proposed solution remains below 0.25 s in all cases. By contrast, the CPU time for the simulation tends to increase significantly as the server utilization increases, exceeding 3 s even for moderate loads. Note that we ran the simulation only for as long as was needed to achieve the desired accuracy, knowing the exact value from our solution method. In practice, one does not know in advance the correct value, so that the simulation needs to be run until the estimated confidence interval at the desired confidence level is sufficiently narrow. This might result in longer simulation times.

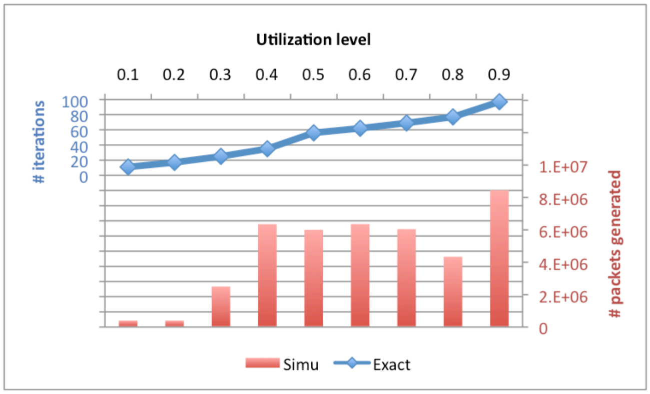

Fig. 2. Number of iterations of proposed approach and number of packets generated in simulation vs. link utilization

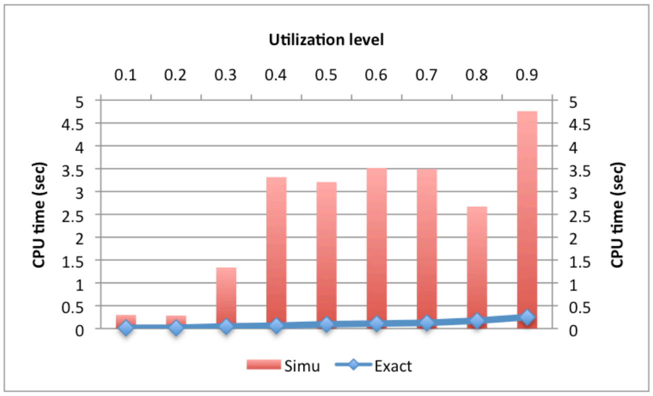

Fig. 3. CPU time for proposed approach and simulation vs. link utilization

In our second example we apply our approach to study the effect of the buffer size on the packet loss probability in a hypothetical configuration. Obviously, the issue of buffer sizing is not new and has received considerable attention (e.g., [40, 3]). In our hypothetical configuration, the rate of packet arrivals is set to a value such that with unrestricted buffer, i.e., with no packets lost, the resulting link utilization would be 0.8. We keep the same type of Pareto-like arrival process as in our first example. For the service process, we assume a single packet size of 1500 bytes and a link speed of 2.5 Gb/s. We represent the resulting service times as a 10-phase Erlang distribution.

Fig. 4 shows the loss probability as a function of the buffer size obtained using the proposed exact solution of our *Ph/Ph/1/N* queue. We observe the rapid decrease of losses as the buffer size increases. The ability to predict accurately the loss probability is important to properly size buffers, in particular in the context of TCP [40].

The loss probability estimated by simulation follows the same general pattern. However, as shown in Fig. 5 in the Appendix, the relative error versus the exact value for the loss probability varies somewhat erratically. The simulation appears to have trouble estimating accurately smaller loss probabilities. We think that this is due to the heavy-tailed nature of the Pareto-like arrival process and the general difficulty of estimating the probability of rare events in a straightforward simulation (similar issues have been pointed out by other authors [4, 13, 5]). If desired, the simulation accuracy could be improved at the expense of non-negligible additional sophistication (e.g., Chapter 5 [32]).

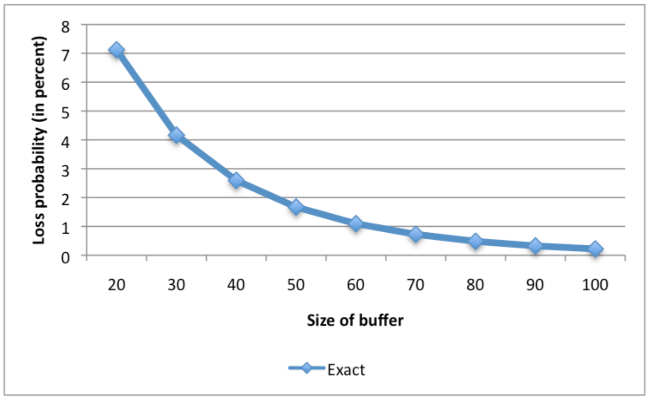

Fig. 4. Predicted loss probability as a function of buffer size

We believe that, given its reasonable speed and conceptual simplicity, the proposed approach offers an attractive alternative to simulation for buffer sizing under non-Poisson assumptions. In that respect, it is worthwhile noting that, for instance, for a buffer size of 80 ($N = 80$) the proposed solution approach computes some 14000 unknowns (the probabilities $p(j,i\,|\,n)$) in fewer than 100 iterations. Additionally, unlike simulation, analytical or numerical solutions produce directly the desired result and not an estimated confidence interval.

The next section concludes our paper.

4 Conclusion

In this paper we propose an approach to the solution of a single server queue with a finite buffer and general (but independent) arrivals and service represented through phase-type distributions. Such a *Ph*/*Ph*/*1*/*N* queue can be viewed as a model of a link or a node in a network. The proposed solution approach is based on the use of conditional probabilities. It is conceptually simple, easy to implement in a standard computer language, numerically robust and reasonably fast.

To deploy our method in practice, one needs to represent the distribution of the times between arrivals and the service time distribution as acyclic phase-type distributions, and then apply our iterative procedure.

In addition to standard steady-state probabilities and queue size averages, the proposed approach produces the probabilities of the state of the queue found by an arriving packet, in particular, the packet loss probability. Note that the state of the queue found by an arriving packet is directly linked to the QoS perceived by the user.

We illustrate the performance of the proposed method using an assumed Pareto-like distribution of the times between arrivals and non-exponential service time distributions. We apply our method to the sizing of a buffer in a hypothetical configuration. Additionally, we compare the performance of the proposed approach with that of a discrete-event simulation. Besides being an order of magnitude faster than the simulation, our exact solution has the advantage of avoiding non-trivial issues that may appear in simulation when dealing with heavy-tailed distribution and rare events.

We believe that the proposed iterative solution for a *Ph*/*Ph*/*1*/*N* queue is simple enough to be of practical interest for researchers and practitioners who are not necessarily specialists in queueing theory or advanced numerical methods.

References

1. CAIDA: IP packet length distribution [Online document]. Available at: http://www.caida.org/analysis/AIX/plen hist (2000).
2. Allen, A.O.: Probability, Statistics and Queueing Theory. Second Edition, Academic Press, London (1990)
3. Appenzeller, G., Keslassy, I., McKeown, N.: Sizing router buffers. In Proceedings of SIGCOMM Comput. Commun. Rev. vol. 34, 4, pp. 281--292 (2004)
4. Asmussen, K., Binswanger, K., Hojgaard, B.: Rare events simulation for heavy-tailed distributions. In Bernoulli. vol. 6, 2, pp. 303--322 (2000)
5. Begin, T., Brandwajn, A.: Note sur la simulation d'une file *M*/*G*/*1* selon la distribution du temps de service. In Proceedings of ALGOTEL. (2010)
6. Bobbio, A., Horváth, A., Telek, M.: Matching three moments with minimal acyclic phase type distributions. In Stochastic Models, vol. 21, 2-3, pp. 303--326 (2005)
7. Bolch, G., Greiner, S., Meer, H., Trivedi, K.: Queueing Networks and Markov Chains. Second Edition, Wiley-Interscience (2005)
8. Brandwajn, A., Begin, T.: Preliminary Results on a Simple Approach to *G*/*G*/*c*-like Queues. In Proceedings of the 16th ASMTA (2009).

9. Buzacott, J.A., Shanthikumar, J.G.: Stochastic models of manufacturing systems. Prentice Hall, Englewood Cliffs (1993)

10. Crovella, M.E., Bestavros, A.: Self-similarity in World Wide Web traffic: evidence and possible causes. In IEEE/ACM Trans. Netw. vol. 5, 6, pp. 835--846 (1997)

11. Feldmann, A., Whitt, W.: Fitting Mixtures of Exponentials to Long-Tail Distributions to Analyze Network Performance Models. In Proceedings of INFOCOM (1997)

12. Grassmann, W.K., Taksar, M.I., Heyman, D.P.: Regenerative analysis and steady state distribution for Markov chains. In Oper. Res. vol. 33, pp. 1107--1116 (1985)

13. Gross, D., Shortle, J.F., Fischer, M.J., Masi, D.M.: Simulation input analysis: difficulties in simulating queues with Pareto service. In Proceedings of the 34th Winter Simulation Conference, pp. 407--415 (2002)

14. Gross, D., Shortie, J.F., Thompson, J.M., Harris, C.M.: Fundamentals of Queueing Theory (4th Ed.) John Wiley & Sons, Inc (2008)

15. Haverkort, B.R., van Moorsel, A.P., Dijkstra, A.: MGMtool: A performance modelling tool based on matrix geometric techniques. In Computer performance evaluation, modelling techniques and tools, pp. 397--401 (1992)

16. Horváth, A., Telek, M.: PhFit: A General Phase-Type Fitting Tool. In International Conference on Computer Performance Evaluation, Modelling Techniques and Tools, pp. 14--17 (2002)

17. Jiang, H., Dovrolis, C.: Why is the internet traffic bursty in short time scales? In ACM SIGMETRICS, pp. 241--252 (2005)

18. Khayari, R.A., Sadre, R., Haverkort, B.R.: Fitting world-wide web request traces with the EM-algorithm. In Performance Evaluation. vol. 52, 2-3, pp. 175--191 (2003)

19. Kimura, T.: Approximating the Mean Waiting Time in the *GI/G/s* Queue. In Journal of the Operational Research Society. vol. 42, 11, pp. 959--970 (1991)

20. Kouvatsos, D.D.: A maximum entropy queue length distribution for the *G/G/1* finite capacity queue. In Proceedings of SIGMETRICS Perform. Eval. Rev. vol. 14, 1, 1986, pp. 224--236.

21. Kraemer, W., Langenbach-Belz, M.: Approximate formulae for the delay in the queueing system *G/G/1*. In Proceedings of International Teletraffic Congress (1976)

22. Kuehn, P.J.: Approximate analysis of general queueing networks by decomposition. In IEEE Transactions on Communications. vol. 27, pp. 113--126 (1979)

23. Kurose, J.F., Ross, K.W.: Computer Networking: a Top-Down Approach. 5th. Addison-Wesley Publishing Company (2009)

24. Latouche, G., Ramaswami, V.: A logarithmic reduction algorithm for quasi-birth-and-death processes. In Journal of Applied Probability. vol. 30, pp. 650--674 (1993)

25. Mitrani, I., Chakka, R.: Spectral expansion solution for a class of Markov models: application and comparison with the matrix-geometric method. In Performance Evaluation. vol. 23, 3, pp. 241--260 (1995)

26. Naoumov, V., Krieger, U.R., Wagner, D.: Analysis of a Multi-Server Delay-Loss System with a General Markovian Arrival Process. In Matrix-analytic Methods in Stochastic Models, pp. 43--66 (1996)

27. Neuts, M.F.: Matrix-geometric solutions in stochastic models. An algorithmic approach. Dover books on advanced mathematics (1994)

28. O'Cinneide, C.A.: Characterization of phase-type distributions. In Communications in Statistics: Stochastic Models. vol. 6, 1, pp. 1--57 (1990)

29. Osogami, T., Harchol-Balter, M.: Closed form solutions for mapping general distributions to quasi-minimal PH distributions. In Perform. Eval. vol. 63, 6, pp. 524--552 (2006)

30. Paxson, V., Floyd, S.: Wide area traffic: the failure of Poisson modeling. In IEEE/ACM Trans. Netw. vol. 3, 3, pp. 226--244 (1995)

31. Rao, B.V., Feldman, R.M.: Numerical approximations for the steady-state waiting times in a *GI/G/1* queue. In Queueing Syst. Theory Appl. vol. 31, 1-2, pp. 25--42 (1999)

32. Rubino, G., Tuffin, B.: Rare Event Simulation Using Monte Carlo Methods. Wiley Publishing (2009)

33. Saad, Y., van der Vorst, H.A.: Iterative solution of linear systems in the 20th century. In J. Comput. Appl. Math. vol. 123, 1-2, pp. 1--33 (2000)

34. Shanthikumar, J.G., Buzacott, J.A.: On the approximations of the single-server queue. In International Journal Production Research. vol. 18, pp. 761--773 (1980)

35. Stewart, W.J.: Introduction to the numerical solution of Markov chains. Princeton University Press (1994)

36. Thummler, A., Buchholz, P., Telek, M.: A Novel Approach for Phase-Type Fitting with the EM Algorithm. In IEEE Trans. Dependable Secur. Comput. vol. 3, 3, pp. 245--258 (2006)
37. Whitt, W.: An Interpolation Approximation for the Mean Workload in a *GI/G/1* Queue. In Operations Research. vol. 37, 6, pp. 936--952 (1989)
38. Whitt, W.: Approximations for the *GI/G/m* queue. In Production and Operations Management. vol. 2, 2, pp. 114--161 (1993)
39. Whitt, W.: Heavy-Traffic Limits for Loss Proportions in Single-Server Queues. In Queueing Syst. Theory Appl. vol. 46, 3/4, pp. 507--536 (2004)
40. Wischik, D., McKeown, N.: Part I: buffer sizes for core routers. In Proceedings of SIGCOMM Comput. Commun. Rev. vol. 35, 3, pp. 75--78 (2005)

Appendix

A. Details of the proposed iterative solution

Our goal is to solve the system of non-linear equations (8) through (11). We use a superscript to denote the iteration number, and for sake of simplicity we introduce the following notation for $n \geq 1$:

$$A^{(k)}(j,i,n) = \begin{cases} \dfrac{1}{(\lambda_j + \mu_i)} \sum_{l=1}^{j-1} p^{(k)}(l,i \mid n)\lambda_l r_{lj}, \text{ if } \quad n < N \\ \dfrac{1}{(\lambda_j + \mu_i)} [\sum_{l=1}^{j-1} p^{(k)}(l,i \mid n)\lambda_l r_{lj} + \sum_{l=1}^{a} p^{(k)}(l,i \mid N)\lambda_l \hat{r}_l], \text{ if } n = N \end{cases}$$

$$B^{(k)}(j,i,n) = \frac{1}{(\lambda_j + \mu_i)} \sum_{l=1}^{i-1} p^{(k)}(j,l \mid n)\mu_l q_{li}$$

$$C^{(k)}(j,i,n) = \begin{cases} \dfrac{\tau_j}{(\lambda_j + \mu_i)} \sum_{j=1}^{a} p^{(k)}(l,i \mid n-1)\lambda_l \hat{r}_l / \alpha^{(k)}(n-1), \text{ if } n > 1 \\ \dfrac{\sigma_i \tau_j}{(\lambda_j + \mu_i)} \sum_{j=1}^{a} p^{(k)}(l \mid 0)\lambda_l \hat{r}_l / \alpha^{(k)}(0), \text{ if } n = 1 \end{cases}$$

$$D^{(k)}(j,i,n) = \begin{cases} \dfrac{\sigma_i}{(\lambda_j + \mu_i)} \sum_{l=1}^{b} p^{(k)}(j,l \mid n+1)\mu_l \hat{q}_l / u^{(k)}(n+1), \text{ if } n < N \\ 0, \text{ if } n = N \end{cases}$$

Our solution proceeds as follows

1) Select initial distributions $p^{(0)}(j\,|\,0)$ and $p^{(0)}(j,i\,|\,n)$ for $n \geq 1$. We must have $\sum_{j=0}^{a}\sum_{i=0}^{b} p^{(0)}(j,i\,|\,n) = 1$ for $n \geq 1$, and $\sum_{j=0}^{a} p^{(0)}(l\,|\,0) = 1$. A simple equally likely distribution will work just fine ($p^{(0)}(j\,|\,0) = 1/a$ and $p^{(0)}(j,i\,|\,n) = 1/(ab)$ for $n \geq 1$). Compute also the corresponding values $\alpha^{(0)}(n)$ and $u^{(0)}(n)$.

2) At iteration $k = 1,...,$

- Consider consecutive increasing values of the number of requests in the system $n = 0,1,...,N$
- Enumerate phases of the arrival process from j equal 1 to a
- Enumerate phases of the service process from i equal 1 to b
- Compute new non-normalized values of $p^{(k)}(j,i\,|\,n)$

for $n \geq 1$

$$p^{(k)}(j,i\,|\,n) \leftarrow A^{(k)}(j,i,n) + B^{(k)}(j,i,n) + u^{(k-1)}(n)C^{(k)}(j,i,n) + \alpha^{(k-1)}(n)D^{(k-1)}(j,i,n)$$

for $n = 0$

$$p^{(k)}(j\,|\,0) \leftarrow \frac{1}{\lambda_j}[\sum_{l=1}^{j-1} p^{(k)}(l\,|\,0)\lambda_l r_{lj} + \alpha^{(k-1)}(0)\sum_{i=1}^{b} p^{(k-1)}(j,i\,|\,1)\mu_i \hat{q}_i / u^{(k-1)}(1)]$$

- Normalize newly computed values so that $\sum_{j=0}^{a}\sum_{i=0}^{b} p^{(k)}(j,i\,|\,n) = 1$ for $n \geq 1$, and $\sum_{j=0}^{a} p^{(k)}(l\,|\,0) = 1$
- Compute new values for the conditional arrival and service rates

$$\alpha^{(k)}(n) = \begin{cases} \sum_{j=1}^{a} \lambda_j \hat{r}_j \sum_{i=1}^{b} p^{(k)}(j,i\,|\,n), \text{ if } n > 0 \\ \sum_{j=1}^{a} \lambda_j \hat{r}_j p^{(k)}(j\,|\,0), \text{ if } n = 0 \end{cases}$$

$$u^{(k)}(n) = \sum_{i=1}^{b} \mu_i \hat{q}_i \sum_{j=1}^{a} p^{(k)}(j,i\,|\,n) \text{ for } n \geq 1.$$

- If $\left\|1 - \alpha^{(k-1)}(n) / \alpha^{(k)}(n)\right\| < \varepsilon$ and $\left\|1 - u^{(k-1)}(n) / u^{(k)}(n)\right\| < \varepsilon$ where ε is the desired convergence stringency (e.g., $\varepsilon = 10^{-7}$), finish the iteration and proceed to Step 3. Otherwise, perform another iteration step, i.e., go back to Step 2.

3) Compute $p(n)$ using formula (2) and any other quantities of interest using formulas (5) to (7).

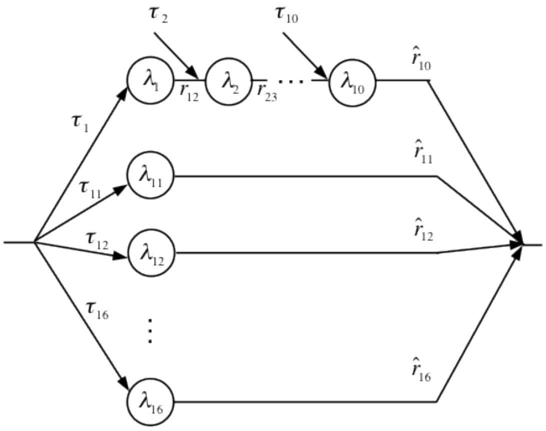

Fig. 6. Structure of the Pareto-like distribution of the time between arrivals

B. Parameters of the Pareto-Like distribution of the time between arrivals

The Pareto-like distribution used in our examples to represent the times between arrivals has the structure shown below in Fig. 6. The phase rate values given in the following table correspond to a mean time between arrivals of 7.797. For a different time between arrivals, these phase rates need to be scaled up or down by multiplying each rate by the ratio $7.797 / T_a$ where T_a is the desired mean time between arrivals, i.e., the inverse of the desired rate of request arrivals.

Table 1. Pareto-like distribution for the time between arrivals used in Section 3.

Probabilities		**Phase rate**	
τ_1	1.39201642e-004	λ_1	1.49162788e+001
τ_2	1.77752368e-004	λ_2	1.34254874e+001
τ_3	1.91426701e-005	λ_3	1.06568964e+001
τ_4	3.45978796e-005	λ_4	9.00035659e+000
τ_5	4.75808142e-005	λ_5	7.61426145e+000
τ_6	1.99877132e-005	λ_6	4.77130864e+000
τ_7	7.26038667e-005	λ_7	3.75562889e+000
τ_8	1.39327654e-004	λ_8	2.37959901e+000
τ_9	6.32355379e-002	λ_9	1.97633786e+000
τ_{10}	2.83641569e-001	λ_{10}	1.76580558e+000

τ_{11}	2.00491113e-006	λ_{11}	2.23396046e-005
τ_{12}	3.08716391e-005	λ_{12}	1.69138281e-004
τ_{13}	4.11347558e-004	λ_{13}	9.82050775e-004
τ_{14}	5.33726475e-003	λ_{14}	5.48273618e-003
τ_{15}	6.52276871e-002	λ_{15}	3.01135332e-002
τ_{16}	5.81463522e-001	λ_{16}	1.55607459e-001
$r_{12}, r_{23}, ... = \hat{r}_{10}, \hat{r}_{11}, ... = 1$			

C. Parameters of the three-phase distribution of example 1

The parameters of this distribution (see Fig. 7) have been selected [6] so that its three first moments (or, equivalently, its mean, variance and skewness) match those of a mix of three packet sizes (50 bytes, 500 bytes and 1500 bytes) transmitted on a 2.5 Gb/s link. The resulting distribution has a mean of 1, coefficient of variation of 1.41, and skewness of 1.84.

Table 2. Distribution for the service time used in Section 3.

Probabilities		Phase rate	
σ_1	4.7162e-001	μ_1	9.8864504e-001
σ_3	$1-\sigma_1$	μ_2	9.8864504e-001
$q_{12}, q_{23} = \hat{q}_3 = 1$		μ_3	2.2820616e+001

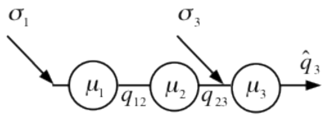

Fig. 7. Structure of the service time distribution in first example

D. Equations for the case of blocked arrivals

In cases where the source of requests is regulated by acknowledgments from the receiver, one can view the source as becoming blocked when the buffer at the receiver is at capacity. In such a case, the state equations given in Section 2 need to be modified for the values of $n = N-1$ and $n = N$. These modified equations lead to the following changes in the iterative scheme

- for $n = N-1$

$$p^{(k)}(j,i\,|\,N-1) \leftarrow A^{(k)}(j,i,N-1) + B^{(k)}(j,i,N-1) + u^{(k-1)}(N-1)C^{(k)}(j,i,N-1)$$
$$+\,\alpha^{(k-1)}(N-1)E^{(k-1)}(j,i,N-1)$$

where $E^{(k)}(j,i,N-1) = \dfrac{\sigma_i \tau_j}{(\lambda_j + \mu_i)} \sum_{l=1}^{b} p^{(k)}(l\,|\,N)\mu_l \hat{q}_l / u^{(k)}(N)$. Note that $p(i\,|\,N)$ denotes the probability of the service process being in its stage $i, i = 1,\ldots,b$ given that the buffer is at capacity and the source is blocked.

- for $n = N$

$$p^{(k)}(i\,|\,N) \leftarrow \frac{1}{\mu_i}[\sum_{l=1}^{i-1} p^{(k)}(l\,|\,N)\mu_l q_{li} + u^{(k-1)}(N)\sum_{j=1}^{a} p^{(k)}(j,i\,|\,N-1)\lambda_j \hat{r}_j / \alpha^{(k)}(N-1)]$$

For all other values of n, the iteration proceeds as outlined in Appendix A.

Note that, with blocked arrivals, we have $\alpha(N) = 0$, i.e., there are no packets lost. The fraction of time during which the arrivals are blocked is given by $p(N)$.

E. Accuracy of discrete-event simulation with heavy-tailed distributions

In Fig. 5 we show the simulation results for 3 different simulation runs whose lengths range from 700 thousand to 70 million packets generated. Interestingly, increasing the length of the simulation does not necessarily reduce the relative error. In our example, for a larger buffer size, the latter can exceed 20% with 7 million packets generated.

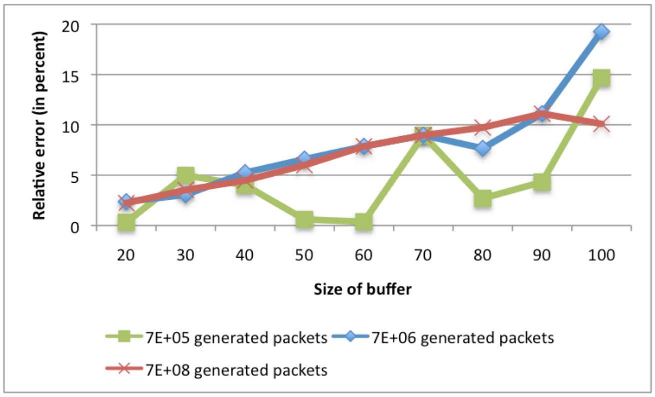

Fig. 5. Relative error for loss probability of simulation vs. exact results as a function of buffer size