On a Generalization of Join-the-Shortest-Queue Scheduling with a Bias

Essia H. Elhafsi and Mart Molle

Technical Report UCR-CS-2005-09001
September 8, 2005

Department of Computer Science and Engineering
University of California, Riverside CA 92521
{essia, mart}@cs.ucr.edu

Abstract

We consider a system of two parallel queues sharing a common arrival process, where the arriving customers select which queue to join based on a generalized Join the Shortest Queue (JSQ) policy. We assume that both servers run at the same speed. However, we assume that departing customers face some additional delay after leaving the system, and that this downstream delay is different between the two queues. Therefore, we introduce a bias into the routing policy, such that customers prefer the path with lower downstream delay unless the current queue length for the fast path exceeds the current queue length for the slow path by at least b customers. We model the number of customers in the two-queue system as a two-dimensional Markov chain, and present a closed-form solution for its steady-state distribution for the case of heavy traffic. When $b = 0$, our results for the $JSQ + b$ policy can be applied to JSQ, for which no closed-form solution is currently available.

Keywords:- Parallel servers; Threshold routing; Asymmetric networks; Queueing systems; State-dependent routing; Closed-form solution.

I. Introduction

The Join the Shortest Queue (JSQ) scheduling policy, in which a newly arriving job is dispatched to the server with the fewest waiting jobs, is an easy to implement and highly effective load balancing scheme. For systems with a single server per queue, the JSQ policy is optimal in the sense of maximizing the discounted number of service completions in any specified time interval [20]. It was also shown [7] that JSQ minimizes the expected total time required to serve all jobs that arrived before some fixed time limit. Unfortunately, even though JSQ has been studied since at least 1958, explicit simple formulae for its stationary probabilities have not yet been obtained [10].

In this paper we present a closed-form solution to the steady state probabilities for two parallel queues under the JSQ scheduling policy. We model the system as a two-dimensional Markov chain and solve for the stationary probabilities. The advantage of this analysis is that we deal directly with the equations of the steady state probabilities without using generating functions or diffusion approximation [11].

In addition, our results cover a more general scheduling policy we call Join the Shortest Queue with Bias ($JSQ + b$). In this case, a newly arriving job will join Q_2 if no more than b additional jobs are waiting at Q_2 than at Q_1; otherwise the new job joins Q_1. We refer to b as the routing threshold. This generalization is motivated by a desire to improve the scheduling decision by taking a more global view of the system into account. For example, suppose Q_1 and Q_2 are the first stages of two alternate paths for completing a multi-stage job. If the costs and/or delays for the remaining stages differed from one path to another, then a job might be willing to tolerate some (local) extra delay at the first stage in order to enter the (globally) faster path (see Figure 1). Alternatively, under the cost model from [4] Q_1 might have a higher admission cost than Q_2 so jobs will seek a balance between higher queueing delays versus lower admission prices in deciding which queue they should join.

Given the difficulty of the analysis of the basic JSQ system, for which the solution is known to be an infinite sum of product and very complicated to explore, we do not anticipate finding an easy or simple solution to the more general $JSQ + b$ system. Hence we focus on the problem of finding exact solution in the case of heavy traffic.

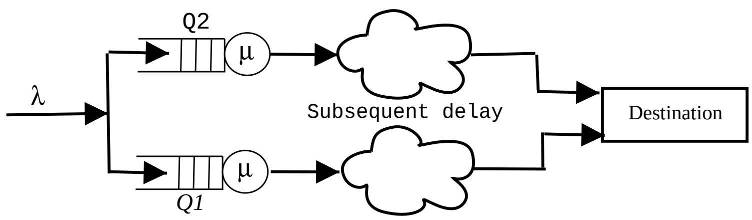

Fig. 1. Two parallel queues

The rest of the paper is organized as follows. In Section 2, we present a brief summary of previous work on the analysis of JSQ scheduling. In Section 3, we present our analysis of $JSQ + b$ scheduling for two parallel queues, including the state-transition diagram, global balance equations, and its steady-state probabilities in closed-form. In Section 4, illustrate the importance of the routing bias in $JSQ + b$ scheduling when jobs experience some extra delay after exiting Q_1. Finally, in Section 5 we state our conclusions.

II. Background

The theoretical analysis of the JSQ scheduling policy is known to be a very challenging problem [16]. The main difficulty is that the job assignments depend on the state of the system, and the state space is multidimensional. Thus, a number of numerical results and approximations have been obtained for the case of N queues and a single server per queue.

Conolly [3] and Rao [19] truncated one or both queues to finite size in order to obtain numerical results. Gertsbakh [10] and Rao [19] used the matrix geometric technique, developed by Neuts [18], to calculate the state occupancy probabilities approximately for two queue systems with equal service rates. In [15], the authors developed an asymptotic approach to obtain approximation to the steady state joint distribution of a two parallel symmetric and asymmetric queueing system. Blanc [2] considered parallel queues with servers of unequal service rates and reported numerical data concerning the means and standard deviations of the waiting time distributions for the case of N queues and a single server. The method was based on power series expansion and recursion and requires substantial computational effort. Lin [16] developed an iterative procedure, for single and multi-servers queues, to estimate the average service rates for different states to obtain the average response time. Kang and Lipsky [12] developed a recursive method to calculate the mean response time of a system consisting of two homogeneous exponential servers and extended it for finite population for general server system in [13]. Nelson and Philips in [17] derived an approximation for the mean response time of a multiple queue system. They used a closed-form equation to calculate the approximate mean response time for systems with general inter-arrival and service time distribution. They report an error of less than 0.5%.

Some attempts have been done in finding closed-form solutions. Flatto [9] obtained complex closed forms for an open network with two identical servers and exponential times. He studied the problem via transform methods and found the mean number of jobs in the systems as an infinite sum that simplifies only under heavy traffic. The transform provided by Flatto cannot be easily inverted to determine the equilibrium probabilities. The waiting time distribution and mean number of jobs in the system are also derived as infinite sums and simplified under heavy traffic assumption in [9] and [14]. Based on the results in [9], Zhao and Grassmann [21] developed an algorithm to compute the probability that there are exactly k jobs in each queue and the joint distribution of the queue lengths in the system. Adan [1] used the compensation approach to determine a closed-form solution to a two $M/M/1$ queueing system with Erlang servers. The solution was an infinite sum of products.

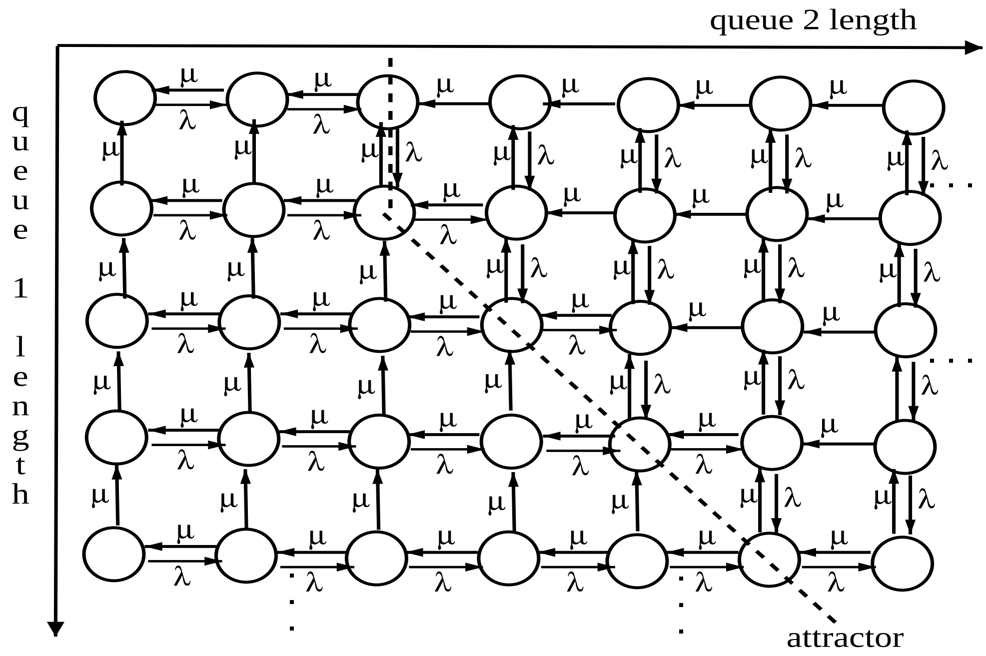

Fig. 2. The transition diagram for the CTMC $\mathbf{Q}(t)$. The *attractor line* is shown for $b = 1$

III. Analysis

1. Model Description

The system considered in this paper consists of two identical queues that we denote by Q_1 and Q_2, as shown in Figure 1. Each queue has infinite storage capacity and the same service rate μ. Arrivals to the system are Poisson with rate λ.

We assume that Q_1 and Q_2 are the first stages of the two alternate paths for completing a multi-stage job. The costs or delays for the remaining stages may differ from one path to another; the delay of the path behind Q_1 is larger than the delay of the path behind Q_2 by a factor of d ($d \geq 0$).

Let q_1 and q_2, respectively, be the total number of jobs currently at each queue, including the job in service (if any). At each arrival instant, the new job joins one of the queues based on the following $JSQ + b$ scheduling policy:

- The first b arrivals to the system are always sent to Q_2.
- Thereafter, if $(q_2 - q_1) \leq b$, then the arrival will be routed to Q_2.
- Otherwise, it will be routed to Q_1.

Note that when $b = 0$, this policy reduces to the basic JSQ policy, which is known to be optimal for the unbiased dynamic scheduling problem.

2. System Model

We model the system as a two-dimensional Markov chain, whose the state may be represented as the tuple (q_1, q_2). The state transition diagram of the system is shown in Figure 2. The routing policies described in the previous section dictate the transition behavior between states.

If we assume that the system is stable (i.e., $\rho \equiv \frac{\lambda}{2\mu} < 1$), then the Markov chain will be ergodic and its equilibrium distribution π_{q_1,q_2} can be found as the (unique) solution to the set of global flow-balance equations derived from Figure 2. For completeness, we provide expressions for the global balance equations in all states (q_1, q_2) in the state space.

We define the *attractor line* in Figure 2 as the set of states where $q_2 - q_1 = b$. Notice that for all states that are *not* on the attractor line, an arrival (i.e., an outgoing transition with label λ) always points towards the attractor line. Thus, for all states below (to the left of) the attractor line, an arrival moves the current state towards the right, and the following global balance equation holds for an interior state within this region:

$$(\lambda + 2\mu)\pi_{q_1,q_2} = \mu\pi_{q_1,q_2+1} + \lambda\pi_{q_1,q_2-1} + \mu\pi_{q_1+1,q_2} \qquad 0 < q_1,\ 0 < q_2 < q_1 + b \tag{1}$$

Using I_C to represent the indicator function whose value is 1 if the condition C is true and 0 otherwise, we can easily generalize Equation (1) so that it also covers the boundary states below the *attractor line*:

$$(\lambda + \mu \cdot (I_{q_1>0} + I_{q_2>0}))\pi_{q_1,q_2} = \mu\pi_{q_1,q_2+1} + I_{q_2>0} \cdot \lambda\pi_{q_1,q_2-1} + \mu\pi_{q_1+1,q_2} \qquad q_2 < q_1 + b \tag{2}$$

Conversely, for all states above (to the right of) the *attractor line*, an arrival moves the current state downwards. The only ambiguity is what to do when the current state is directly on the attractor line. We will always break "ties" in favor of Q_2. Thus, we need to define separate global balance equations for states on the *attractor line*:

$$(\lambda + \mu \cdot (I_{q_1>0} + 1))\pi_{q_1,q_2} = \mu\pi_{q_1,q_2+1} + \lambda\pi_{q_1,q_2-1} + \mu\pi_{q_1+1,q_2} + I_{q_1>0} \cdot \lambda\pi_{q_1-1,q_2} \qquad q_2 = q_1 + b \tag{3}$$

for the first line of states above the *attractor line*:

$$(\lambda + \mu \cdot (I_{q_1>0} + 1))\pi_{q_1,q_2} = \lambda\pi_{q_1,q_2-1} + I_{q_1>0} \cdot \lambda\pi_{q_1-1,q_2} + \mu\pi_{q_1,q_2+1} + \mu\pi_{q_1+1,q_2} \qquad q_2 = q_1 + b + 1 \tag{4}$$

and for all other states above the *attractor line*:

$$(\lambda + \mu \cdot (I_{q_1>0} + 1))\pi_{q_1,q_2} = I_{q_1>0} \cdot \lambda\pi_{q_1-1,q_2} + \mu\pi_{q_1,q_2+1} + \mu\pi_{q_1+1,q_2} \qquad q_2 > q_1 + b + 1 \tag{5}$$

By numerically solving the system of Equations (2)–(5), it is easy to see that in equilibrium the system spends almost all of its time in states on or very close to the *attractor line*.

3. Analytical Solution

For well-structured transition matrices, the literature is full of methods to solve the global balance equations, such as matrix geometric methods [18] and tensor algebra [8]. In our case however, such standard techniques are difficult to apply because the state space is infinite and the transition matrix does not have a very nice structure. Thus, we adopted a different method to solve the balance equations, following an approach commonly used in the product-form queueing network literature. Here we *guess* the form of a general solution to the problem, and then *check* the validity of the result by substitution into the global balance equations.

In this case, our candidate solution has the form:

$$\pi_{q_1,q_2} = f_1(q_1)f_2(q_2) \qquad \text{for all } q_1, q_2 \text{ away from the boundaries} \tag{6}$$

where $f_1(q_1) = c_1\varphi^{q_1}$ and $f_2(q_2) = c_2\psi^{q_2}$. The parameters c_1 and c_2 have been added to create sufficient freedom in determining the model. Substituting the candidate solution given in Equation (6) into Equations (2)-(5), and then solving for parameter values which satisfy the global balance conditions leads us to the solution of the stationary probabilities.

We only solve for the stationary probabilities $\{\pi_{q_1,q_2}\}$ when the queues are full ($q_1 \geq 1$ and $q_2 \geq 1$ and more precisely when condition C is true). $\{\pi_{q_1,q_2}\}$ are given by the Equations (7)–(9) where $\alpha = (2 + \rho)$ and $\beta = \frac{\rho^2}{\alpha}$.

For all states on either side of the *attractor line*, we have:

$$\pi_{q_1 q_2} = \begin{cases} \alpha^{q_2-b-1} \cdot \beta^{q_1-1}\ \pi_{1,b+1} & q_1 \geq 1,\ 1 \leq q_2 \leq b + q_1 \\ \beta^{q_2-b-1}\ \alpha^{q_1}\ \frac{1}{\rho}\ \pi_{1,b+1} & 1 \leq q_1 \leq q_2 - b - 1 \end{cases} \tag{7}$$

For those states along the *attractor line*, these equations reduce to the following:

$$\pi_{q_1,b+q_1} = (\alpha\beta)^{q_1-1}\,\pi_{1,b+1} \quad q_1 \geq 1 \tag{8}$$

The final unknown $\pi_{1,b+1}$ is determined by solving the following normalizing equation:

$$\sum_{q_1=0}^{\infty}\sum_{q_2=0}^{\infty}\pi_{q_1,q_2} = 1 \tag{9}$$

IV. Significance of the Routing Bias

The routing bias in $JSQ + b$ scheduling provides a significant performance improvement over pure JSQ scheduling when our system model includes an additional source of state-independent delay behind Q_1, as shown in Figure 1.

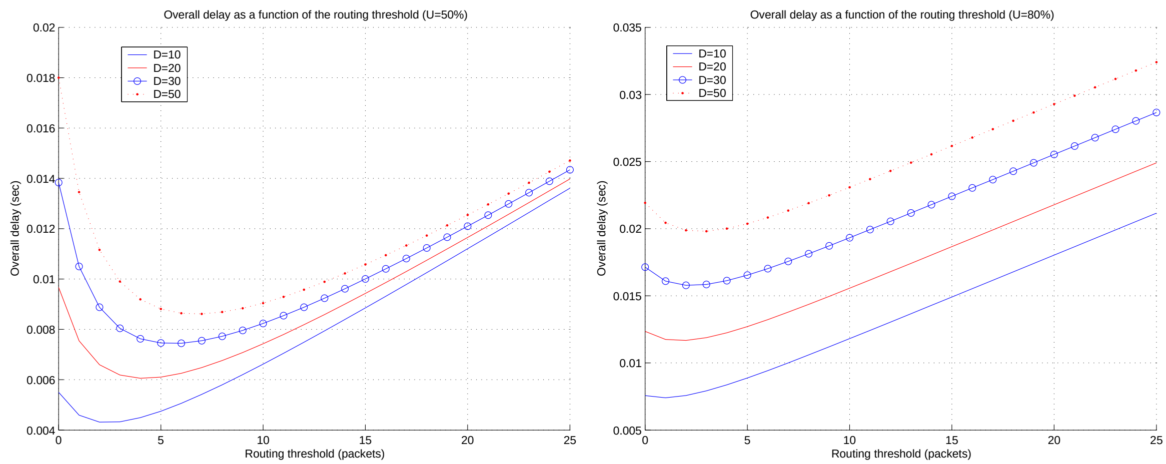

Fig. 3. Reducing the mean delay via threshold routing when there is some extra delay behind Q_1.

This improvement is clearly visible in Figure 3, which shows the mean delay as a function of b at $\rho = 0.5$ (left side) and $\rho = 0.8$ (right side), for several values of the extra delay (i.e., $d = 10,\ 20,\ 30,\ 50$ times the mean service time, $1/\mu$). Notice that in all cases, the minimum delay occurs for some $b > 0$. Moreover, the optimum value of b is always much smaller than d, and becomes even smaller as ρ increases.

It is interesting to note that the optimal solution to the standard parallel queue scheduling problem ($d = 0$) is a "greedy" policy, i.e., JSQ. However, when we generalize the problem by adding some extra delay behind one queue ($d > 0$) then the "greedy" policy, i.e., $JSQ + b$ with $b/\mu = d$, is clearly *not* optimal in this case.

V. Conclusion

In this paper we present a closed-form solution to the joint distribution of the queue lengths in steady-state when a generalization of the JSQ scheduling is applied to two parallel queues in a network with heavy traffic. Our scheduling policy is equivalent to the JSQ policy when the routing threshold is zero. Thus, our results also provide a closed-form solution to the steady-state probabilities for JSQ scheduling as a special case. These results provide a third option for evaluating the performance of such systems, which is both less computationally intensive than "brute-force" numerical solution of the global balance equations, i.e., $\pi \cdot Q = 0$, and more accurate than matrix geometric approximations we describe in [6].

We also provide a counterexample to show that neither JSQ nor the "greedy" version of $JSQ + b$ are optimal policies when jobs face some additional delay behind one of the queues. This observation is further extended in [5] where we present a closed-form relation between the routing threshold and the extra delay behind the queues. This work is being extended to systems where each queue feeds multiple servers, and to systems with more than two queues.

References

[1] I. Adan. Explicit solutions for queueing problems. *Memoradum COSOR 97-05, Eindhoven university of technology*, April 1997.

[2] J. Blanc. The power series algorithm applied to the shortest queue model. *Operations Research*, 40(1):157–167, Jan-Feb 1992.

[3] B. W. Conolly. The autostrada queueing problem. *J. Appl. Prob.*, 21:394–403, 1993.

[4] P. Dube, V. S. Borkar, and D. Manjunath. Differential join prices for parallel queues: Social optimality, dynamic pricing algorithms and application to internet pricing. In *INFOCOM*, 2002.

[5] E. H. Elhafsi, M. Molle, and D. Manjunath. On the application of forking nodes in queueing networks (*submitted*).

[6] E. H. Elhafsi, M. Molle, and D. Manjunath. Can we use product-form solution techniques in networks with asymmetric paths? In *Proceedings of the International symposium on Performance Evaluation of computer and telecommuniaction systems (SPECTS 2005)*, pages 1–2, 2005.
[7] A. Ephremides, P. Varaiya, and J.Walrand. A simple dynamic routing problem. *IEEE Trans. Automatic Control*, 25:690–693, August 1980.
[8] P. Fernandes, B. Plateau, and W. Stewart. Efficient descriptor-vector multiplications in stochastic automata networks. *Journal of the ACM*, 45(3):381–414, May 1998.
[9] L. Flatto and Hahn S. Two parallel queues created by arrivals with two demands I. *SIAM Journal of Applied Mathematics*, 44(5):1041–1053, October 1984.
[10] I. Gertsbakh. The shorter queue problem: A numerical study using the matrix geometric solution. *European J. Operation Research*, 15:374–381, 1984.
[11] D. Iglhart and W. Witt. Multiple channel queues in heavy traffic-I. *Adv. Appl. Prob.*, 2:150–177, 1970.
[12] S. Kang. Load sharing in the dynamic environments. *Ph.D thesis, University of Connecticut, Storrs*, December 1995.
[13] S. Kang and L. R. Lipsky. Steady state solution for join-the-shortest-queue policy in general server system. *http://uconn.edu/ lester/papers/new.ps.gz*, February 1996.
[14] J.F.C Kingman. Two similar queues in parallel. *Ann. Math. Stat.*, 32:1314–1323, 1961.
[15] C. Knessl, B. Matkowsky, Z. Schuss, and C. Tier. Two parallel queues with dynamic routing. *IEEE Transactions on Communications*, 34(12):1170–1175, December 1986.
[16] H. Lin and C. Raghavendra. An approximation analysis of the join the shortest queue (JSQ) policy. *IEEE Transaction on Parallel and Distributed Systems*, 7(3):301–307, March 1996.
[17] R. D. Nelson and T.K. Philis. An approximation to the response time for the shortest queue routing. *Performance Evaluation Review*, 17(1), May 1989.
[18] M. F. Neuts. *Matrix geometric Solutions in Stochastic Models*. John Hopkins University Press, Baltimore, 1980.
[19] B. M. Rao and M.J.M. Posner. Algorithmic and approximation analyses of the shorter queue model. *Naval Research Logistics*, 34:381–398, 1977.
[20] J. Winston. Optimality of the shortest line discipline. *J. Appl. Prob.*, 14:181–198, 1997.
[21] Y. Zhao and W.K. Grassman. A numerically stable algorithm for two server queue model. *Queueing Systems*, 8:59–79, 1991.