The Asymptotics of Selecting the Shortest of Two, Improved

Michael Mitzenmacher* Berhold Vöcking †

We investigate variations of a novel, recently proposed load balancing scheme based on small amounts of choice. To motivate this work, we first provide the relevant background. When n balls are thrown into n bins, the maximum load, or balls in a bin, is $\frac{\ln n}{\ln \ln n}(1+o(1))$ with high probability. Suppose instead that n balls are sequentially placed into n bins so that for each ball, $d \geq 2$ bins are chosen independently and uniformly at random from the n bins, and the ball is placed in the bin with the fewest balls, ties being broken arbitrarily. Then in this case the maximum load is only $\frac{\ln \ln n}{\ln d} \pm \Theta(1)$ with high probability [1].

This result was generalized to natural queueing models independently in [6] and [2, 3]. Suppose that tasks arrive at a bank of n First In First Out processors as a Poisson process of rate λn, where $\lambda < 1$; tasks require an exponentially distributed amount of service with mean 1. If each task queues at a random processor, then in the stationary distribution the probability that a server has at least k tasks is simply λ^k. If instead each task chooses two processors at random and queues at the shorter, then in the fluid limit process representing the limiting behavior as n grows, the probability that a server has at least k tasks converges to $\lambda^{\frac{d^k-1}{d-1}}$. That is, the tails of the processor queue lengths decrease *doubly exponentially*, rather than exponentially, when $d > 1$. The effect is clear in simulations even for relatively small values of n.

In this paper, we consider a variation based on the work by Vöcking [5]. In his basic model, n bins are split into d groups of size n/d. We think of these groups as being laid out linearly, with the first group (bins 1 to n/d) being the *leftmost*, and so on. A ball is placed by choosing a bin independently and uniformly at random from each of the d groups and placing the ball into the least loaded of the d bins; ties are broken by placing the ball in the leftmost (tied) bin. Vöcking showed that the maximum number of balls in a bin in this case is $\frac{\ln \ln n}{d \ln \phi_d} \pm \Theta(1)$ with high probability, where here ϕ_d corresponds to the exponent of growth for a generalized Fibonacci sequence. (For reference, $\phi_2 = 1.61 < \phi_3 < \phi_4 \ldots < 2$.) Surprisingly, coordinating ties improves performance! These results also hold for certain dynamic models; see [5] for more details. To differentiate the approaches, we call Vöcking's the *d-left* scheme and the original the *d-random* scheme.

Here, following the work of [2], we examine the d-left scheme using fluid limit models corresponding to a family of differential equations. Because of limited space we consider only the queueing variation. For more details and results, see the extended draft of this paper [4].

Suppose arriving tasks choose a processor according to the d-random scheme. We let w_i represent the fraction of the bins with at least i tasks in the queue (including the one being served). Then the differential equations describing the fluid limit are

$$\frac{dw_i}{dt} = \lambda(w_{i-1}^d - w_i^d) - (w_i - w_{i+1}). \quad (1)$$

Of course $w_0(t) = 1$ always.

The key to understanding these systems generally revolves around their *fixed points*, where all the derivatives are zero. Moreover, as we expect the system to be stable, we seek fixed points where the w_i fall to 0 as i grows. It can be checked [3, 6] that the unique such fixed point for the d-random queueing system, which we denote by w_i^*; is given by $w_i^* = \lambda^{\frac{d^i-1}{d-1}}$. Hence, we find that even for $d = 2$, the tails of the queue lengths decrease doubly exponentially.

We now examine the d-left scheme. We first consider the case $d = 2$, for which the fluid limit equations are

$$\frac{dy_i}{dt} = 4\lambda\left(y_{i-1} - y_i\right) z_{i-1} - \left(y_i - y_{i+1}\right) \; ; \; (2)$$

*Harvard University, Computer Science Department. 29 Oxford St., Cambridge, MA 02138. e-mail:michaelm@eecs.harvard.edu.

†International Computer Science Institute, Berkeley, CA 94704. e-mail: voecking@icsi.berkeley.edu. Supported by a grant of the "Gemeinsames Hochschulsonderprogramm III von Bund und Ländern" through the DAAD.

$$\frac{dz_i}{dt} = 4\lambda\left(z_{i-1} - z_i\right) y_i - \left(z_i - z_{i+1}\right), \qquad (3)$$

where y_i and z_i represent the fraction of servers with queue length at least i on the left and right, respectively. We may simplify for general d by using x_{jd+k} to represent the fraction of servers that have queue length at least j and lie in the kth group from the left. Then the fluid limit model yields the following family of differential equations:

$$\frac{dx_i}{dt} = d^d\lambda\left(x_{i-d} - x_i\right)\left(\prod_{j=i-d+1}^{i-1} x_j\right) - \left(x_i - x_{i+d}\right).$$

This system appears substantially more complex than (1), even in the simplest case where $d = 2$. Indeed, we have not been able to prove that there is a unique fixed point (satisfying $x_i \to 0$ as $i \to \infty$), although that appears to be the case based on simulating the process determined by the differential equations. We will from here on assume that there is indeed such a unique fixed point for this system. With this assumption, we can prove that the tails for this system at the fixed point are term by term no greater than the tails at the fixed point of the d-random system. Let $u_i = \sum_{k=0}^{d-1} x_{id+k}$. That is, u_i is the total fraction of servers with load at least i. Clearly $u_0 = 1$. We represent the values at the fixed point for the d-left scheme by u_i^* and x_i^*. Since all derivatives are zero at the fixed point, we have

$$d^d\lambda\left(x_{i-d}^* - x_i^*\right)\left(\prod_{j=i-d+1}^{i-1} x_j^*\right) - \left(x_i^* - x_{i+d}^*\right) = 0.$$

Summing these expressions from $i = jd$ to infinity yields $u_j^* = d^d\lambda \prod_{i=(j-1)d}^{jd-1} x_i^*$. Hence, for instance, it follows that $u_1^* = \lambda$, verifying that the arrival rate equals the departure rate at the fixed point. But now by the arithmetic-geometric mean inequality

$$u_j^* = d^d\lambda \prod_{i=(j-1)d}^{jd-1} x_i^* \le \lambda\left(\sum_{i=(j-1)d}^{jd-1} x_i^*\right)^d = \lambda\left(u_{j-1}^*\right)^d.$$

A simple direct induction now yields that $u_j^* \le \lambda^{\frac{d^j-1}{d-1}}$, and hence in equilibrium in the fluid limit the d-random system has larger tails than a d-left system. In fact, as explained in [4], the tails decrease *Fibonacci exponentially*, which is faster than the double exponetial decrease of the d-random system.

	$d=2$		$d=3$	
i	w_i^*	u_i^*	w_i^*	u_i^*
1	9.0e-1	9.0e-1	9.0e-1	9.0e-1
2	7.3e-1	7.3e-1	6.5e-1	6.5e-1
3	4.8e-1	4.8e-1	2.5e-1	2.5e-1
4	2.1e-1	2.0e-1	1.4e-2	9.1e-3
5	3.8e-2	3.2e-2	2.9e-6	5.8e-10
6	1.3e-3	5.8e-4	2.9e-17	1.4e-54
7	1.5e-6	2.5e-8	9.7e-51	

Table 1: Results for d-left (u^*) vs. d-random (w^*) in the fluid limit model when $\lambda = 0.9$.

We present results obtained from finding the fixed point from the fluid limit model through simulating the differential equations. We demonstrate the behavior in Table 1 by focusing on the case $\lambda = 0.9$ and comparing fixed point values for the d-left and d-random schemes. The improvement from using the d-left scheme on measures such as the expected time is small, but real. The improvement is more pronounced if we consider a different measure of performance. As we have shown, the d-left scheme leads to a faster decrease in the tails of the loads. Hence, the main benefit of the d-left over the d-random scheme is that the probability of joining a long queue, and thereby spending a long time in the system, is much smaller under the d-left scheme.

References

[1] Y. Azar, A. Broder, A. Karlin, and E. Upfal. Balanced Allocations. In *Proc. of the 26th ACM Symp. on the Theory of Computing*, 1994, pp. 593–602.

[2] M. Mitzenmacher. The Power of Two Choices in Randomized Load Balancing. Ph.D. thesis, U. C. Berkeley, September 1996.

[3] M. Mitzenmacher. Load Balancing and Density Dependent Jump Markov Processes. In *Proc. of the 37^{th} IEEE Symp. on Foundations of Computer Science*, 1996, pp. 213–222.

[4] M. Mitzenmacher and B. Vöcking. The Asymptotics of Selecting the Shortest of Two, Improved. Extended abstract available at www.eecs.harvard.edu/~michaelm/NEWWORK/papers.html.

[5] B. Vöcking, How Asymmetry Helps Load Balancing. Extended abstract, to appear in FOCS '99.

[6] N.D. Vvedenskaya, R.L. Dobrushin, and F.I. Karpelevich. Queueing System with Selection of the Shortest of Two Queues: an Asymptotic Approach. *Problems of Information Transmission*, Vol 32, 1996, pp. 15–27.