

The Power of Two Choices in Randomized Load Balancing

Michael Mitzenmacher, *Member, IEEE*

Abstract—We consider the following natural model: Customers arrive as a Poisson stream of rate λn , $\lambda < 1$, at a collection of n servers. Each customer chooses some constant d servers independently and uniformly at random from the n servers and waits for service at the one with the fewest customers. Customers are served according to the first-in first-out (FIFO) protocol and the service time for a customer is exponentially distributed with mean 1. We call this problem the *supermarket model*. We wish to know how the system behaves and in particular we are interested in the effect that the parameter d has on expected time a customer spends in the system in equilibrium. Our approach uses a limiting, deterministic model representing the behavior as $n \rightarrow \infty$ to approximate the behavior of finite systems. The analysis of the deterministic model is interesting in its own right. Along with a theoretical justification of this approach, we provide simulations that demonstrate that the method accurately predicts system behavior, even for relatively small systems. Our analysis provides surprising implications: Having $d = 2$ choices leads to exponential improvements in the expected time a customer spends in the system over $d = 1$, whereas having $d = 3$ choices is only a constant factor better than $d = 2$. We discuss the possible implications for system design.

Index Terms—Load balancing, queuing theory, distributed systems, limiting systems, choices.

1 INTRODUCTION

CONSIDER the following natural dynamic model: Customers arrive as a Poisson stream of rate λn , $\lambda < 1$, at a collection of n servers. Each customer chooses d servers independently and uniformly at random with replacement from the n servers for some fixed constant d . The customer waits for service at the server from these d choices currently containing the fewest customers (ties being broken arbitrarily). Customers are served according to the first-in first-out (FIFO) protocol and the service time for a customer is exponentially distributed with mean 1. We call this model the *supermarket model*, or the *supermarket system* (Fig. 1). We wish to know how the system behaves in equilibrium and, in particular, we are interested in the expected time a customer spends in the system in equilibrium. The supermarket model proves difficult to analyze because of dependencies: Knowing the length of one queue affects the distribution of the length of all the other queues.

The supermarket model can be seen as a generalization of the static load balancing model studied by Azar et al. [1], in which there are a fixed number of customers to be distributed who never leave the system. They also analyze a different *closed* dynamic model, where the number of customers remains fixed over all time and a customer who completes service is recirculated in the system. They do not analyze the *open* dynamic model in their work. As described in [1], models of this kind have a number of applications to computing problems, including resource

allocation, hashing, and online load balancing. Our results apply to dynamic variations of these applications. For example, the supermarket model provides a good abstraction of a simple, efficient load balancing scheme in the setting where tasks arrive and execute at a large system of parallel processors.

In this paper, we analyze the supermarket model and introduce techniques that prove useful for studying other randomized load balancing strategies. Our strategy is to define an idealized process, corresponding to a system of infinite size, where the number of servers n goes to infinity. This idealized process is given by a family of differential equations, whose behavior is cleaner and easier to analyze because its behavior is completely deterministic. This idealized system can be related to systems with finite n using the appropriate mathematical tools. In practice, we find the method provides a means of finding accurate numerical estimates of performance, as we demonstrate with simulations. Besides providing an analysis of the supermarket model, this approach also provides a clean, systematic approach to analyzing several other load balancing models.

The following theorem is representative of the results we obtain:

Theorem 1. *For any fixed T and $d \geq 2$, the expected time a customer spends in an initially empty supermarket system over the first T units of time is bounded above by*

$$\sum_{i=1}^{\infty} \lambda^{\frac{i}{d-1}} + o(1),$$

where the $o(1)$ term is understood as $n \rightarrow \infty$ (and may depend on T and λ).

• The author is with Harvard University, 33 Oxford St., Cambridge, MA 02138. E-mail: michaelm@eecs.harvard.edu.

Manuscript received 11 Sept. 1997; revised 25 Jan. 2000; accepted 16 Feb. 2001.

For information on obtaining reprints of this article, please send e-mail to: tpd@computer.org, and reference IEEECS Log Number 105653.

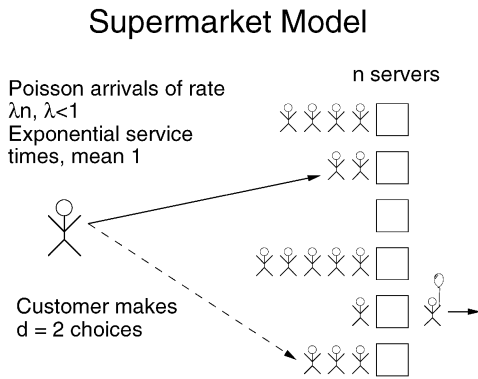


Fig. 1. The supermarket model, with $d = 2$.

The summation is derived from the limiting system of differential equations and the $o(1)$ term arises when we bound the error between this system and the random process for a fixed n . The combination of the two analyses yields the theorem. This result should be compared to the case of $d = 1$, where in equilibrium the expected time is $1/(1 - \lambda)$. As we describe in Section 2.4, for λ close to 1, there is an exponential improvement in the expected time a customer spends in the system.

The exponential improvement from two choices suggests an excellent rule of thumb in the design of distributed load balancing systems: Systems where items have two (or a small number of) choices can perform almost as well as a perfect load balancing system with global load knowledge. Indeed, because a system based on two choices can have significantly lower overhead, it is possible it may perform better than apparently better but more complicated load balancing algorithms. Although our model only demonstrates this rule of thumb in a very simple setting, this effect has been noted in distributed systems with much greater structure [4], [5], [30]. We believe that this rule of thumb will prove useful in many systems to come.

We note that at approximately the same time as the initial publication of this work [24], [25], similar results for the supermarket model were derived independently by Vvedenskaya et al. [33]. In particular, they study the same model and use an approach similar to ours in that they also derive the same family of differential equations that describe the limiting behavior as the number of servers n goes to infinity. As a consequence, some of the results of this paper duplicate (or are in some respects weaker than) similar results in [33].

There are, however, noticeable differences. We summarize some of them here; the interested reader is encouraged to read [33] for a clearer picture. We apply simpler techniques, similar in spirit to potential function arguments, that should be more accessible. Our approach allows us to prove *exponential convergence* of the limiting differential equations, a stronger result than the convergence result of [33]. Our simpler techniques, however, yield results that depend on how long the system runs; the results of [33] invoke more powerful theory and can therefore state results in terms of the equilibrium distribution of the system. Although the consequences are similar in practice, in this respect the results of [33] are clearly stronger than ours.

Further differences include that we study the expected time a customer spends in the system in substantially more detail and test the performance on systems of reasonable size through simulation. We find that having d choices in the supermarket model provides an exponential improvement in the average time a customer spends in the system. Our detailed argument yields a surprising connection to the previous results of Azar et al. [1]. Our simulation experiments also demonstrate that despite the theory holding essentially only in the limit as the number of servers n goes to infinity, in practice one sees dramatic improvements even in small systems for reasonable arrival rates such as $\lambda < 0.95$.

1.1 Related Work

The power of having two choices has been noted before, although primarily in the static setting. An analysis using two hash functions for load balancing was provided by Karp et al. [11] in an application to PRAM simulation. They analyzed a static case, where a number of memory items are to be permanently distributed among a fixed number of servers, and demonstrated an exponential improvement in the maximum load. The static problem was further developed and analyzed by Azar et al. in [1]. Results for other static settings are given by MacKenzie et al. [19], Adler et al. [2], and Stemmann [32]. Related dynamic models, where one is concerned with the behavior of a system over an arbitrary time interval, have proven more difficult. A different (and arguably less realistic) dynamic model was successfully analyzed by Azar et al. [1]. We note that our method can also be applied to these dynamic models and to the static model of [1], providing new insight and results. (See [24], [25], [26] for details.) Previously, these problems have all been attacked by applying complicated arguments based on Chernoff-type bounds. Our approach has several advantages: It is extremely general, it is simple to apply, and it provides more detailed and accurate numerical information about these systems.

The supermarket model, and variations thereof, were studied in previous queueing theory work by Eager et al. [8]. This work and later related papers [9], [22], [23] also use an approach based on Markov chains for their analysis. However, the authors base their work on the assumption that the state of each queue is stochastically independent of the state of any other queue [8, p. 665]. The authors also claim (without justification) that this approach is exact in the asymptotic limit as the number of queues grows to infinity. Besides introducing several new directions in the analysis of these systems, we explain how to justify these assumptions. Zhou [37] examines the load balancing strategies proposed by Eager et al. as well as others using a trace-driven simulation. Both Eager et al. and Zhou suggest that simple randomized load balancing schemes, based on choosing from a small subset of processors, appear quite effective in practice.

To bound the error between the finite and limiting systems we will use Kurtz's work on *density dependent jump Markov processes* [7], [14], [15], [16], [17]. Kurtz's work has previously been applied to matching problems on random graphs [10], [12], [13]. More recently, the technique of using differential equations has been used in queueing theory [31]

and several works related to random graph structures [3], [18], [29], [36].

Although in this paper we focus solely on the supermarket model, the techniques developed here can be used to examine a wide variety of related models. In particular, one can develop models for other service distributions [26] or for cases where the available load information may be inexact [26], [27].

Finally, a great deal of additional work on the general theme of the power of two choices has appeared since this paper was written. As a starting point, we suggest the recent survey [28].

The rest of the paper is structured as follows: In Section 2, we analyze the behavior of the infinite version of the supermarket model. In Section 3, we briefly explain Kurtz's work and how to adapt it to relate the finite and limiting versions of the supermarket model; more technical details are available in the Appendix. In Section 4, we provide simulation results demonstrating the accuracy of this methodology, even for a small number of servers. We conclude with some thoughts on the implications of these results.

2 THE ANALYSIS OF THE SUPERMARKET MODEL

2.1 Preliminaries

Recall the definition of the supermarket model: Customers arrive as a Poisson stream of rate λn , $\lambda < 1$, at a collection of n FIFO servers. Each customer chooses some constant d servers independently and uniformly at random with replacement¹ and queues at the server currently containing the fewest customers with ties being broken arbitrarily. The service time for a customer is exponentially distributed with mean 1. The following lemma, which we state without proof, will be useful:

Lemma 1. *The supermarket system is stable for every $\lambda < 1$; that is, the expected number of customers in the system remains finite for all time.*

Remark. Lemma 1 can be proven by a simple comparison argument against the system in which each customer queues at a random server (that is, where $d = 1$); in this system, each server acts like an M/M/1 server with arrival rate λ , which is well known to be stable (see any introductory queueing theory text). The comparison argument is entirely similar to those in [34], [35], which show that choosing the shortest queue is optimal subject to certain assumptions on the service process; alternatively, an argument based on majorization, such as that in [1], is possible. A similar argument also shows that the size of the longest queue in a supermarket system of size n is stochastically dominated by the size of the longest queue in a set of n independent M/M/1 servers.

We now introduce a representation of the system that will be convenient throughout our analysis. We define $n_i(t)$ to be the number of queues with i customers at time t , $m_i(t)$ to be the number of queues with at least i customers at time

t , $p_i(t) = n_i(t)/n$ to be the fraction of queues of size i , and $s_i(t) = \sum_{k=i}^{\infty} p_k(t) = m_i(t)/n$ to be the tails of the $p_i(t)$. We drop the reference to t in the notation where the meaning is clear. As we shall see, the s_i prove much more convenient to work with than the p_i . In an *empty system*, which corresponds to one with no customers, $s_0 = 1$ and $s_i = 0$ for $i \geq 1$. By comparing this system with a system of M/M/1 queues as in the remark after Lemma 1, we have that if $s_i(0) = 0$ for some i , then for all $t \geq 0$, $\lim_{i \rightarrow \infty} s_i(t) = 0$. Under the same conditions, the expected number of customers per queue, or $\sum_{i=1}^{\infty} s_i(t)$, is finite even as $t \rightarrow \infty$.

We can represent the state of the system at any given time by an infinite dimensional vector $\vec{s} = (s_0, s_1, s_2, \dots)$. Note that our state only includes information regarding the number of queues of each size; this is all the information we require. It is clear that for each value of n , the supermarket model can be considered as a Markov chain on the above state space.

We now introduce a deterministic limiting system related to the finite supermarket system. The time evolution of the limiting system is specified by the following set of differential equations:

$$\begin{cases} \frac{ds_i}{dt} = \lambda(s_{i-1}^d - s_i^d) - (s_i - s_{i+1}) & \text{for } i \geq 1, \\ s_0 = 1. \end{cases} \quad (1)$$

Let us explain the reasoning behind (1). Consider a supermarket system with n queues and determine the expected change in the number of servers with at least i customers over a small period of time of length dt . The probability a customer arrives during this period is $\lambda n dt$ and the probability an arriving customer joins a queue of size $i - 1$ is $s_{i-1}^d - s_i^d$. (This is the probability that all d servers chosen by the new customer are of size at least $i - 1$, but not all are of size at least i .) Thus, the expected change in m_i due to arrivals is exactly $\lambda n (s_{i-1}^d - s_i^d) dt$. Similarly, the probability a customer leaves a server of size i in this period is $n_i dt = n(s_i - s_{i+1}) dt$. Hence, if the system behaved according to these expectations, we would have

$$\frac{dm_i}{dt} = \lambda n (s_{i-1}^d - s_i^d) - n(s_i - s_{i+1}).$$

Removing the factor of n permeating the equations yields (1). That this infinite set of differential equations has a unique solution given appropriate initial conditions is not immediately obvious; however, it follows from standard results in analysis (the Picard approximation method; see also [20, p.188, Theorem 4.1.5] or [6, Theorem 3.2]). It should be intuitively clear that as $n \rightarrow \infty$ the behavior of the supermarket system approaches that of this deterministic system; this is justified by Kurtz's theorem, which is explained in Section 3. For now, we simply take this set of differential equations to be the appropriate limiting process.

2.2 Finding a Fixed Point

We will demonstrate that, given a reasonable condition on the initial point $\vec{s}(0)$, the limiting process converges to a *fixed point*. A fixed point (also called an *equilibrium point* or a *critical point*) is a point \vec{p} such that if $\vec{s}(t) = \vec{p}$ then $\vec{s}(t') = \vec{p}$ for all $t' \geq t$. It is clear that for the supermarket model a

1. We note that our results also hold with minor variations if the d queues are chosen without replacement.

necessary and sufficient condition for \vec{s} to be a fixed point is that for all i , $\frac{ds_i}{dt} = 0$.

Lemma 2. *System (1), with $d \geq 2$, has a unique fixed point with $\sum_{i=1}^{\infty} s_i < \infty$ given by*

$$s_i = \lambda^{\frac{i-1}{d-1}}.$$

Proof. It is easy to check that the proposed fixed point satisfies $\frac{ds_i}{dt} = 0$ for all $i \geq 1$. Conversely, from the assumption $\frac{ds_i}{dt} = 0$ for all i , we can derive that $s_1 = \lambda$ by summing (1) over all $i \geq 1$. (Note that we use $\sum_{i=1}^{\infty} s_i < \infty$ here to ensure that the sum converges absolutely. That $s_1 = \lambda$ at the fixed point also follows intuitively from the fact that at the fixed point, the rate at which customers enter and leave the system must be equal.) The result then follows from (1) by induction. \square

The condition, $\sum_{i=1}^{\infty} s_i < \infty$, which corresponds to the average number of customers per queue being finite, is necessary; $(1, 1, \dots)$ is also a fixed point, which corresponds to the number of customers at each queue going to infinity.

Definition 1. *A sequence $(x_i)_{i=0}^{\infty}$ is said to decrease doubly exponentially if and only if there exist positive constants $N, \alpha < 1, \beta > 1$, and γ such that for $i \geq N$, $x_i \leq \gamma\alpha^{\beta^i}$.*

It is worth comparing the result of Lemma 2 to the case where $d = 1$ (i.e., all servers are M/M/1 queues), for which the fixed point is given by $s_i = \lambda^i$. The key feature of the supermarket system is that for $d \geq 2$ the tails s_i decrease doubly exponentially at the fixed point while for $d = 1$ the tails decrease only geometrically (or singly exponentially).

2.3 Convergence to the Fixed Point

We now show that every trajectory of the supermarket system converges to the fixed point of Lemma 2 in an appropriate metric. Denote the above fixed point by $\vec{\pi} = (\pi_i)$, where $\pi_i = \lambda^{\frac{i-1}{d-1}}$. We shall assume that $d \geq 2$ in what follows unless otherwise specified.

We begin with a result that shows the system has an invariant, which restricts in some sense how far any s_i can be from the corresponding value π_i .

Theorem 2. *Suppose there exists some j such that $s_j(0) = 0$.*

Then, the sequence $(s_i(t))_{i=0}^{\infty}$ decreases doubly exponentially for all $t \geq 0$, where the associated constants are independent of t . In particular, if the system begins empty, then $s_i(t) \leq \pi_i$ for all $t \geq 0$.

Note that the hypothesis of Theorem 2 holds for any initial state \vec{s} derived from the initial state of a finite system. (However, as shown in [33], even this limitation can be removed.)

Proof. Let $M(t) = \sup_i [s_i(t)/\pi_i]^{1/d}$. We first show that $M(t) \leq M(0)$ for all $t \geq 0$. We will then use this fact to show that the s_i decrease doubly exponentially.

A natural, intuitive proof proceeds as follows: In the case where there are a finite number of queues, an inductive coupling argument can be used to prove that, if we increase some $s_i(0)$, thereby increasing the number of customers in the system, the expected value of all s_j after any time t increases as well. Extending this to the

limiting case as the number of queues $n \rightarrow \infty$ (so that the s_j behave according to their expectations), we have that increasing $s_i(0)$ can only increase all the $s_j(t)$ and, hence, $M(t)$ for all t .

So, to begin, let us increase all $s_i(0)$ (including $s_0(0)$!) so that $s_i(0) = M(0)^d \pi_i$. But then it is easy to check that the initial point is a fixed point (albeit possibly with $s_0 > 1$) and, hence, $M(t) = M(0)$ in the raised system. We conclude that in the original system $M(t) \leq M(0)$ for all $t \geq 0$.

A more formal proof that increasing $s_i(0)$ only increases all $s_j(t)$ relies on the fact that the ds_i/dt are *quasimonotone*: that is, ds_i/dt is nondecreasing in s_j for $j \neq i$. The result then follows from [6, pp. 70-74].

We now show that the s_i decrease doubly exponentially (in the limiting model). Let j be the smallest value such that $s_j(0) = 0$, which exists by the hypothesis of the theorem. Then,

$$M(0) \leq [1/\pi_{j-1}]^{1/d^{j-1}} < 1/\lambda^{1/(d-1)}.$$

Since $M(t) \leq M(0)$, $M(0)^d \geq s_i(t)/\pi_i$ for $t \geq 0$ or

$$s_i(t) \leq \pi_i M(0)^d = \lambda^{-1/(d-1)} (\lambda^{1/(d-1)} M(0))^d.$$

Note that $\lambda^{1/(d-1)} M(0) < 1$ since

$$M(0) < 1/\lambda^{1/(d-1)}.$$

Hence, the s_i decrease doubly exponentially, with

$$\alpha = \lambda^{1/(d-1)} M(0)$$

and $\beta = d$. In particular, if the system begins empty, then $s_i(t) \leq \pi_i$ for all t and i . \square

To show convergence, we find a *potential function* (also called a *Lyapunov function* in the dynamical systems literature) $\Phi(t)$ with the following properties:

1. The potential function is related to the distance between the current point on the trajectory and the fixed point.
2. The potential function is strictly decreasing, except at the fixed point.

The intuition is that the potential function shows that the system heads toward the fixed point. By finding a suitable potential function, we will also be able to say how fast the system approaches the fixed point. A natural potential function to consider is $D(t) = \sum_{i=1}^{\infty} |s_i(t) - \pi_i|$, which measures the L_1 -distance between the two points. Our potential function will actually be a weighted variant of this, namely $\Phi(t) = \sum_{i=1}^{\infty} w_i |s_i(t) - \pi_i|$ for suitably chosen weights w_i .

Definition 2. *The potential function Φ is said to converge exponentially to 0 or simply to converge exponentially if $\Phi(t) \leq c_0 e^{-\delta t}$ for some constant $\delta > 0$ and a constant c_0 which may depend on the state at $t = 0$.*

We note that Vvedenskaya et al. [33] provide an alternative convergence proof that does not yield exponential convergence. Their proof requires transforming the vector of s_i values into a different infinite dimensional vector, on which they show convergence coordinate by coordinate.

Theorem 3. Let $\Phi(t) = \sum_{i=1}^{\infty} w_i |s_i(t) - \pi_i|$, where, for $i \geq 1$, the w_i are appropriately chosen constants to be determined, satisfying $w_i \geq 1$. If $\Phi(0) < \infty$, then Φ converges exponentially to 0. In particular, if there exists a j such that $s_j(0) = 0$, then Φ converges exponentially to 0.

Proof. Define $\epsilon_i(t) = s_i(t) - \pi_i$. As usual, we drop the explicit dependence on t when the meaning is clear. For convenience, we assume that $d = 2$; the proof is easily modified for general d .

As $d\epsilon_i/dt = ds_i/dt$, we have from (1)

$$\begin{aligned} \frac{d\epsilon_i}{dt} &= \lambda[(\pi_{i-1} + \epsilon_{i-1})^2 - (\pi_i + \epsilon_i)^2] - (\pi_i + \epsilon_i - \pi_{i+1} - \epsilon_{i+1}) \\ &= \lambda(2\pi_{i-1}\epsilon_{i-1} + \epsilon_{i-1}^2 - 2\pi_i\epsilon_i - \epsilon_i^2) - (\epsilon_i - \epsilon_{i+1}), \end{aligned}$$

where the last equality follows from the fact that $\bar{\pi}$ is a fixed point.

As $\Phi(t) = \sum_{i=1}^{\infty} w_i |\epsilon_i(t)|$, the derivative of Φ with respect to t , $d\Phi/dt$, is not well defined if $\epsilon_i(t) = 0$. We shall explain how to cope with this problem at the end of the proof and we suggest the reader proceed by temporarily assuming $\epsilon_i(t) \neq 0$.

Now,

$$\begin{aligned} \frac{d\Phi}{dt} &= \sum_{i:\epsilon_i>0} w_i [\lambda(2\pi_{i-1}\epsilon_{i-1} + \epsilon_{i-1}^2 - 2\pi_i\epsilon_i - \epsilon_i^2) - (\epsilon_i - \epsilon_{i+1})] - \\ &\quad \sum_{i:\epsilon_i<0} w_i [\lambda(2\pi_{i-1}\epsilon_{i-1} + \epsilon_{i-1}^2 - 2\pi_i\epsilon_i - \epsilon_i^2) - (\epsilon_i - \epsilon_{i+1})]. \end{aligned}$$

Let us look at the terms involving ϵ_i in this summation. (Note: ϵ_1 terms are a special case, which can be included in the following if we take $w_0 = 0$.) There are several cases, depending on whether ϵ_{i-1} , ϵ_i and ϵ_{i+1} are positive or negative. Let us consider the case where they are all negative (which, by Theorem 2, is always the case when the system is initially empty). Then, the term involving ϵ_i is

$$-w_{i-1}\epsilon_i + w_i(2\lambda\pi_i\epsilon_i + \lambda\epsilon_i^2 + \epsilon_i) - w_{i+1}(2\lambda\pi_i\epsilon_i + \lambda\epsilon_i^2). \quad (2)$$

We wish to choose w_{i-1} , w_i , and w_{i+1} so that this term is at most $\delta w_i \epsilon_i$ for some constant $\delta > 0$. It is sufficient to choose them so that

$$(w_i - w_{i-1}) + (2\lambda\pi_i + \lambda\epsilon_i)(w_i - w_{i+1}) \geq \delta w_i$$

or, using the fact that $|\epsilon_i| \leq 1$,

$$w_{i+1} \leq w_i + \frac{w_i(1 - \delta) - w_{i-1}}{\lambda(2\pi_i + 1)}.$$

We note that the same inequality would be sufficient in the other cases as well: For example, if all of ϵ_{i-1} , ϵ_i and ϵ_{i+1} are positive, the above term (2) involving ϵ_i is negated, but now ϵ_i is positive. If ϵ_{i-1} , ϵ_i and ϵ_{i+1} have mixed signs, this can only decrease the value of (2).

It is simple to check inductively that one can choose an increasing sequence of w_i (starting with $w_0 = 0, w_1 = 1$) and a δ such that the w_i satisfy the above restriction. For example, we break the terms up into two subsequences. The first subsequence consists of all w_i such that π_i

satisfies $\lambda(2\pi_i + 1) \geq \frac{1+\lambda}{2}$. For these i we can choose $w_{i+1} = w_i + \frac{w_i(1-\delta) - w_{i-1}}{3}$. Because this subsequence has only finitely many terms, we can choose a suitably small δ so that this sequence is increasing. For sufficiently large i , we must have $\lambda(2\pi_i + 1) < \frac{1+\lambda}{2} < 1$, and for these i we may set $w_{i+1} = w_i + \frac{2w_i(1-\delta) - 2w_{i-1}}{1+\lambda}$. This subsequence of w_i will be increasing for suitably small δ and, hence, $w_i \geq 1$ for all $i \geq 1$. Further, this sequence of w_i is dominated by a geometrically increasing sequence and, hence, if $s_j(0) = 0$ for some j , then $\Phi(0) < \infty$.

Comparing terms involving ϵ_i in Φ and $d\Phi/dt$ yields that $d\Phi/dt \leq -\delta\Phi$. Hence, $\Phi(t) \leq \Phi(0)e^{-\delta t}$ and, thus, Φ converges exponentially.

We now consider the technical problem of defining $d\Phi/dt$ when $\epsilon_i(t) = 0$ for some i . Since we are interested in the forward progress of the system, it is sufficient to consider the upper right-hand derivatives of ϵ_i . (See, for instance, [21, p. 16].) That is, we may define

$$\left. \frac{d|\epsilon_i|}{dt} \right|_{t=t_0} \equiv \lim_{t \rightarrow t_0^+} \frac{|\epsilon_i(t)|}{t - t_0},$$

and similarly for $d\Phi/dt$. Note that this choice has the following property: If $\epsilon_i(t_0) = 0$, then $\left. \frac{d|\epsilon_i|}{dt} \right|_{t=t_0} \geq 0$, as it intuitively should be. The above proof applies unchanged with this definition of $d\Phi/dt$ with the understanding that the case $\epsilon_i > 0$ includes the case where $\epsilon_i = 0$ and $d\epsilon_i/dt \geq 0$ and similarly for the case $\epsilon_i < 0$. \square

Theorem 3 yields the following corollary:

Corollary 1. Under the conditions of Theorem 3, the L_1 distance from the fixed point $d(t) = \sum_{i=1}^{\infty} |s_i(t) - \pi_i|$ converges exponentially to 0.

Proof. As the w_i of Theorem 3 are all at least 1 for $i \geq 1$, $\Phi(t) \geq d(t)$ and the corollary is immediate. \square

Corollary 1 shows that the L_1 distance to the fixed point converges exponentially quickly to 0. Hence, from any suitable starting point, the limiting system quickly becomes extremely close to the fixed point. Although it seems somewhat unusual that we first had to prove exponential convergence for a weighted variation of the L_1 distance in order to prove exponential convergence of the L_1 distance, it appears that this approach was necessary.

Remark. Note that in the proof of Theorem 3, if we set all the w_i equal to 1 and let $\delta = 0$, the appropriate inequalities hold. This yields a simple proof that the L_1 distance to the fixed point is nonincreasing over time on the trajectory given by the differential equations. Although such a result is not as strong as the exponential convergence or even as strong as the convergence result of [33], this technique proves useful for other problems where exponential convergence may not be possible (see [25], [26]).

2.4 The Expected Time in the Limiting System

Using Theorems 2 and 3, we now examine the expected time a customer spends in the limiting system. We emphasize that the expected time in the system is easily determined by the fixed point and, hence, this corollary is also easily determined by the results in [33].

Corollary 2. *The expected time a customer spends in the limiting supermarket system for $d \geq 2$, subject to the condition of Theorem 2, converges as $t \rightarrow \infty$ to*

$$T_d(\lambda) \equiv \sum_{i=1}^{\infty} \lambda^{\frac{d-i}{d-1}}.$$

Furthermore, $T_d(\lambda)$ is an upper bound on the expected time in the limiting system for all t when the system is initially empty.

Proof. An incoming customer that arrives at time t becomes the i th customer in the queue with probability $s_{i-1}(t)^d - s_i(t)^d$. Hence, the expected time a customer that arrives at time t spends in the system is

$$\sum_{i=1}^{\infty} i(s_{i-1}(t)^d - s_i(t)^d) = \sum_{i=0}^{\infty} s_i(t)^d.$$

As $t \rightarrow \infty$, by Corollary 1, the limiting system converges to the fixed point in the L_1 metric. Hence, the expected time a customer spends in the system can be made arbitrarily close to

$$\sum_{i=0}^{\infty} \pi_i^d = \sum_{i=1}^{\infty} \lambda^{\frac{d-i}{d-1}}$$

for all customers arriving at time $t \geq t_0$ for some sufficiently large t_0 and the result follows. The second result follows since we know that in an initially empty limiting system $s_i(t) \leq \pi_i$ for all t by Theorem 2. \square

Recall that $T_1(\lambda) = 1/(1-\lambda)$ from standard queueing theory. Intuitively, one would expect an exponential improvement in the expected time in the system going from one choice to some constant $d \geq 2$ choices since the tails of the queue lengths decrease doubly exponentially instead of exponentially. Analysis of the summation in Corollary 2 in fact reveals the following:

Theorem 4. *For $\lambda \in [0, 1]$ and $d \geq 2$, $T_d(\lambda) \leq c_d(\log T_1(\lambda))$ for some constant c_d dependent only on d . Furthermore,*

$$\lim_{\lambda \rightarrow 1^-} \frac{T_d(\lambda)}{\log T_1(\lambda)} = \frac{1}{\log d}.$$

Proof. We prove only the limiting statement as $\lambda \rightarrow 1^-$; the other statement is proved similarly. Let $\lambda' = \lambda^{1/(d-1)}$. Then,

$$T_d(\lambda) = \sum_{i=1}^{\infty} \lambda^{\frac{d-i}{d-1}} = \frac{\sum_{i=1}^{\infty} (\lambda')^{d-i}}{\lambda^{d/(d-1)}}.$$

Hence,

$$\begin{aligned} \lim_{\lambda \rightarrow 1^-} \frac{T_d(\lambda)}{\log T_1(\lambda)} &= \lim_{\lambda \rightarrow 1^-} \frac{\sum_{i=1}^{\infty} (\lambda')^{d-i}}{-\log(1-\lambda)\lambda^{d/(d-1)}} \\ &= \lim_{\lambda \rightarrow 1^-} \frac{\sum_{i=1}^{\infty} \lambda'^{d-i} \log(1-\lambda')}{-\log(1-\lambda') \log(1-\lambda) \lambda^{d/(d-1)}}. \end{aligned}$$

In the final expression on the right, the last two terms go to 1 as $\lambda \rightarrow 1^-$. The result then follows from the following lemma. \square

Lemma 3. *Let*

$$F_d(\lambda) = \frac{\sum_{i=0}^{\infty} \lambda^{d^i}}{\log \frac{1}{1-\lambda}}.$$

Then, $\lim_{\lambda \rightarrow 1^-} F_d(\lambda) = 1/\log d$.

Proof. We show that, for any small enough $\epsilon > 0$, there is a corresponding δ such that for $\lambda > 1 - \delta$,

$$\frac{1}{(\log d) + \epsilon} \leq F_d(\lambda) \leq \frac{1}{(\log d) - \epsilon}.$$

We prove only the left inequality; the right inequality is entirely similar. We use the following identity:

$$\prod_{i=0}^{\infty} (1 + \lambda^{d^i} + \lambda^{2d^i} + \dots + \lambda^{(d-1)d^i}) = \frac{1}{1-\lambda}.$$

From this identity, it follows that

$$\sum_{i=0}^{\infty} \log(1 + \lambda^{d^i} + \lambda^{2d^i} + \dots + \lambda^{(d-1)d^i}) = \log \frac{1}{1-\lambda}.$$

For a given ϵ , let $\epsilon' = \epsilon/2$, and let

$$z = \sup \left[\{0\} \cup \left\{ x : 0 < x \leq 1, \frac{\log(1+x+x^2+\dots+x^{d-1})}{x} > \log d + \epsilon' \right\} \right].$$

Note that $z < 1$. For any fixed λ , we split up the summation in the previous equation to obtain

$$\begin{aligned} &\sum_{i: \lambda^{d^i} \leq z} \log(1 + \lambda^{d^i} + \dots + \lambda^{(d-1)d^i}) \\ &+ \sum_{i: \lambda^{d^i} > z} \log(1 + \lambda^{d^i} + \dots + \lambda^{(d-1)d^i}) \quad (3) \\ &= \log \frac{1}{1-\lambda}. \end{aligned}$$

The leftmost term of (3) is bounded by a constant, dependent only on z and independent of λ . Hence,

$$\begin{aligned} &\sum_{i: \lambda^{d^i} \leq z} \log(1 + \lambda^{d^i} + \dots + \lambda^{(d-1)d^i}) \\ &+ \sum_{i: \lambda^{d^i} > z} \log(1 + \lambda^{d^i} + \dots + \lambda^{(d-1)d^i}) \quad (4) \\ &\leq c_z + (\log d + \epsilon') \sum_{i: \lambda^{d^i} \leq z} \lambda^{d^i} + (\log d + \epsilon') \sum_{i: \lambda^{d^i} > z} \lambda^{d^i}, \end{aligned}$$

where c_z is a constant dependent only on z and is independent of λ . Combining (3) and (4) yields

$$(\log d + \epsilon') \sum_{i=0}^{\infty} \lambda^{di} + c_z \geq \log \frac{1}{1 - \lambda}$$

or

$$F_d(\lambda) + \frac{c_z}{(\log d + \epsilon')(\log \frac{1}{1-\lambda})} \geq \frac{1}{(\log d) + \epsilon'}.$$

We now choose δ small enough so that for $\lambda > 1 - \delta$,

$$\frac{1}{(\log d) + \epsilon'} - \frac{c_z}{(\log d + \epsilon')(\log \frac{1}{1-\lambda})} \geq \frac{1}{(\log d) + \epsilon'},$$

and the lemma follows. \square

Hence, choosing from $d > 1$ queues yields an exponential improvement in the expected time a customer spends in the limiting system and as $\lambda \rightarrow 1^-$ the choice of d affects the time only by the constant factor $\log d$. These results are remarkably similar to those for the static case studied in [1], where again the choice of d only affects the bound on the maximum load by a $\log d$ factor in the denominator. This correspondence is perhaps not surprising; the static system studied in [1] can also be related to a similar set of differential equations [24]. A more direct connection between these results, however, remains lacking and might be enlightening.

3 FROM INFINITE TO FINITE: KURTZ'S THEOREM

3.1 An Overview of Kurtz's Theorem

The supermarket model is an example of a *density dependent family of jump Markov processes*, the formal definition of which we shall give in the Appendix. Informally, such a family is a one parameter family of Markov processes, where the parameter n corresponds to the total population size (or, in some cases, area or volume). The states can be normalized and interpreted as measuring population densities so that the transition rates depend only on these densities. As we have seen in the supermarket model, the transition rates between states depend only upon the densities s_i . Hence, the supermarket model fits our informal definition of a density dependent family. The limiting system corresponding to a density dependent family is the limiting model as the population size grows arbitrarily large.

Kurtz's work provides a basis for relating the limiting system for a density dependent family to the corresponding finite systems. Essentially, Kurtz's theorem provides a law of large numbers and Chernoff-like bounds for density dependent families. We provide some intuition for this result. The primary differences between the limiting system and the finite system are

- The limiting system is deterministic; the finite system is random.
- The limiting system is continuous; the finite system has jump sizes that are discrete values.

Imagine starting both systems from the same point for a small period of time. Since the jump rates for both processes are initially the same, they will have nearly the same behavior. Now, suppose that, if two points are close in the state space, then their transition rates are also close: This is

called the *Lipschitz condition* and it is a precondition for Kurtz's theorem. Then, even after the two processes separate, if they remain close, they will still have nearly the same behavior. Continuing this process inductively over time, we can bound how far the processes separate over any interval $[0, T]$.

We can apply Kurtz's results to the supermarket model to obtain bounds on the expected time a customer spends in the system and the maximum queue length. The proofs, however, are somewhat technical; these technicalities are due to the fact that the states (s_0, s_1, \dots) are infinite-dimensional. The interested reader may examine [25].

Theorem 5. *For any fixed T , the expected time a customer spends in an initially empty supermarket system with $d \geq 2$ over the interval $[0, T]$ is bounded above by*

$$\sum_{i=1}^{\infty} \lambda^{\frac{i-d}{d-1}} + o(1),$$

where $o(1)$ is understood as $n \rightarrow \infty$ and depends on T and λ .

Theorem 5 says the expected time in a finite system is the same as that for the limiting system (Corollary 2) plus a $o(1)$ term. Similarly, we can bound the maximum load:

Theorem 6. *For any fixed T , the length of the longest queue in an initially empty supermarket system with $d \geq 2$ over the interval $[0, T]$ is $\frac{\log \log n}{\log d} + O(1)$ with high probability,² where the $O(1)$ term depends on T and λ .*

Hence, in comparing the systems where customers have one choice and customers have $d \geq 2$ choices, we see that the second yields an exponential improvement in both the expected time in the system and in the maximum observed load for sufficiently large n . In practice, simulations reveal that this behavior is apparent even for relatively small n over long periods of time, as shown in Section 4.

Also, we note that a more sophisticated use of the underlying theory, as shown in [33], can yield stronger results, in the following sense. It can be shown that as the number of servers n grows to infinity, the equilibrium distribution of the system becomes concentrated at the fixed point. Hence, after running the system a sufficiently long time, the expected time a customer spends in the system and the maximum load will be as in Theorems 5 and 6, with the improvement that the low order terms will not depend on T .

3.2 Further Implications for Finite Systems

Intuitively, Kurtz's theorem says that the random, finite system tends to travel along the same trajectory as the path defined by the differential equations. This fact, combined with the exponential convergence we have proven in Section 2.3, suggests that the system reaches its equilibrium quickly and that its equilibrium distribution is sharply concentrated around the fixed point. Moreover, even when the random process deviates from the expected path, its tendency is to quickly return on a trajectory toward the fixed point. One would therefore expect that simulations of

2. Here, *with high probability* will mean with probability $1 - O(1/n)$ and all logarithms have base e .

TABLE 1
Average Time in the Supermarket Model: 100 Queues

Choices	λ	Simulation	Prediction	Rel. Error (%)
2	0.50	1.2673	1.2657	0.1289
	0.70	1.6202	1.6145	0.3571
	0.80	1.9585	1.9475	0.5742
	0.90	2.6454	2.6141	1.1981
	0.95	3.4610	3.3830	2.3028
3	0.99	5.9275	5.4320	9.1227
	0.50	1.1277	1.1252	0.2146
	0.70	1.3634	1.3568	0.4858
	0.80	1.5940	1.5809	0.8314
	0.90	2.0614	2.0279	1.6533
5	0.95	2.6137	2.5351	3.1002
	0.99	4.4080	3.8578	14.2607
	0.50	1.0340	1.0312	0.2637
	0.70	1.1766	1.1681	0.7250
	0.80	1.3419	1.3289	0.9789
5	0.90	1.6714	1.6329	2.3564
	0.95	2.0730	1.9888	4.2363
	0.99	3.4728	2.9017	19.6825

TABLE 2
Average Time in the Supermarket Model: 500 Queues

Choices	λ	Simulation	Prediction	Rel. Error (%)
1	0.99		100.00	
2	0.99	5.5413	5.4320	2.0121
3	0.99	3.9518	3.8578	2.4366
5	0.99	3.0012	2.9017	3.4305

and the first 10,000 steps are ignored in recording data in order to give the system time to approach equilibrium. For arrival rates of up to 95 percent of the service rate (i.e., $\lambda = 0.95$), the predictions are within a few percent of the simulation results. Even at 99 percent of capacity, the prediction is within 10 percent when two queues are selected. It is not surprising that the error increases as the arrival rate or the number of choices available to a customer increases as these parameters affect the error term in Kurtz’s theorem. As one would expect, however, the approximation does improve if the number of queues is increased, as can be seen by the results for 500 queues given in Table 2.

The simulations clearly demonstrate the impact of having two choices as well as the accuracy of the differential equations. The effect is made apparent by Fig. 2, which compares the expected time in the system given by simulations of 100 queues and two choices to the prediction given by the differential equations and *the logarithm* of the expected time in equilibrium when just one choice ($d = 1$) is used. The points for 100 queues and the prediction coincide nearly exactly. Fig. 2 also shows similar results for a much smaller system of only eight queues. For systems this small, one would expect the limiting model to be less accurate, and of course it is, especially under high arrival rates. However, at moderate arrival rates it still proves a good approximation to actual behavior and it demonstrates the appropriate trend as the arrival rate increases.

The qualitative behaviors that we predicted with our analysis are thus readily observable in our simulations, even of relatively small systems on the order of 100 queues.

the supermarket system would tend to closely match the results predicted by the limiting system. In the next section, we demonstrate that this in fact holds, even for a very small number of queues.

4 SIMULATION RESULTS

We provide the results of some simulations based on the supermarket model, focusing on the average time a customer spends in the system. In these simulations, choices were made without replacement as this method is more likely to be used in practice. (The limiting system is the same regardless of whether the choices are made with or without replacement; we have described the model with replacement to simplify the exposition.) The results of Table 1 are based on a system of $n = 100$ queues at various arrival rates. The results are based on the average of 10 runs, where each run consists of a simulation of 100,000 time steps




Fig. 2. Comparing simulations of $d = 2$ to the limiting system and $d = 1$.

TABLE 3
Variance in the Time in Supermarket Model: 100 Queues

Choices	λ	Simulation	Prediction	Rel. Error (%)
2	0.50	1.4966	1.4923	0.2881
	0.70	2.1304	2.1139	0.7805
	0.80	2.7415	2.7148	0.9835
	0.90	3.9689	3.8669	2.6378
	0.95	5.4326	5.1442	5.6063
	0.99	10.5856	8.1444	29.974

This lends weight to the predictive power of our theoretical results in practical settings.

While we have focused on the average time a customer spends in the system in this section, because the fixed point of the limiting system gives a complete distribution over queue lengths, it would also be possible to use the limiting system to predict other measures of the system as well. For example, in many systems the variance of the time in the system may also be important. The limiting system proves useful for determining this quantity as well. Table 3 compares the predicted variance with the average of the sample variance from the runs used for Table 1. We note that, as with the average time in the system, the accuracy improves with the number of queues simulated and worsens as the arrival rate increases.

5 CONCLUSIONS

We have demonstrated that in a simple dynamic load balancing model allowing customers to choose between the shortest of two servers yields an exponential improvement over distributing customers uniformly at random. Although this model is arguably unrealistic for many applications, it can be extended to other similar domains [24], [25], [26] and we believe that, more generally, as a rule of thumb, it will prove useful in the design of distributed systems.

Perhaps the most interesting open question is to include *locality* in this model. For example, suppose the servers are arranged in a ring and each customer chooses to queue either at a server chosen uniformly at random or its right neighbor. The locality complicates the state space sufficiently that it appears the techniques used here will not apply. Finding other scenarios where the power of two choices proves evident, both in theory and in practice, also remains open for further exploration.

APPENDIX

DENSITY DEPENDENT MARKOV CHAINS

For the interested reader, we now give a more technical presentation of the background for Kurtz's theorem. The reader is recommended to [17] or [31] for a more complete discussion and derivation. We begin with the definition of a density dependent family of Markov chains, as in Kurtz [17, chapter 7]. For convenience, we drop the vector notation where it can be understood by context. Given a set of transitions $L \subseteq \mathbf{Z}^D$ for some fixed D and a collection of nonnegative functions β_l for $l \in L$ defined on a subset $E \subset \mathbf{R}^D$, a *density dependent family of Markov chains*

X_n is a sequence $\{X_n\}$ of jump Markov processes such that the state space of X_n is $E_n = E \cap \{n^{-1}k : k \in \mathbf{Z}^D\}$ and the transition rates of X_n are

$$q_{x,y}^{(n)} = n\beta_{n(y-x)}(x), \quad x, y \in E_n.$$

As an example of this definition, consider the supermarket model for $d = 2$ with n queues. The state of the system $\vec{s} = k/n$, where \vec{s} represents the state by the fraction of servers of size at least i and k represents the state by the number of servers of size at least i . Note that we may think of the state of the system either as \vec{s} or k , as they are the same except for a scale factor. The possible transitions from k is given by the set $L = \{\pm e_i : i \geq 1\}$, where the e_i are standard unit vectors; these transitions occur either when a customer enters or departs. The transition rates are given by $q_{k,k+l}^{(n)} = n\beta_l(k/n) = n\beta_l(\vec{s})$, where $\beta_{e_i}(\vec{s}) = \lambda(s_{i-1}^2 - s_i^2)$ and $\beta_{-e_i}(\vec{s}) = s_i - s_{i+1}$. These rates determined our limiting system (1). Note that, in this case, our system is not in \mathbf{Z}^D for any fixed D but in \mathbf{Z}^N and, hence, technically does not fit Kurtz's definition of a density dependent Markov chain. This technical problem is, in this case, minor; see [25] for explicit details.

It follows from [17, Chapter 7] that a Markov process \hat{X}_n with intensities $q_{k,k+l}^{(n)} = n\beta_l(k/n)$ satisfies

$$\hat{X}_n(t) = \hat{X}_n(0) + \sum_{l \in L} l Y_l \left(n \int_0^t \beta_l \left(\frac{\hat{X}_n(u)}{n} \right) du \right),$$

where the $Y_l(x)$ are independent standard Poisson processes. This equation has a natural interpretation: The process at time t is determined by the starting point and the rate of each transition integrated over the history of the process. In the supermarket system, \hat{X}_n is the unscaled process with state space \mathbf{Z}^N that records the number of servers with at least i customers for all i and $\hat{X}_n(0)$ is the initial state, which we usually take to be the empty system.

We set

$$F(x) = \sum_l \beta_l(x), \quad (5)$$

and by setting $X_n = n^{-1}\hat{X}_n$ to be the appropriate scaled process, we have from the above:

$$X_n(t) = X_n(0) + \sum_{l \in L} \frac{l}{n} \tilde{Y}_l \left(n \int_0^t \beta_l(X_n(u)) du \right) + \int_0^t F(X_n(u)) du, \quad (6)$$

where $\tilde{Y}_l(x) = Y_l(x) - x$ is the Poisson process centered at its expectation.

Kurtz's theorem shows that the deterministic limiting process is given by

$$X(t) = x_0 + \int_0^t F(X(u)) du, \quad t \geq 0, \quad (7)$$

where $x_0 = \lim_{n \rightarrow \infty} X(0)$. An interpretation relating (6) and (7) is that, as $n \rightarrow \infty$, the value of the centered Poisson process $\tilde{Y}_i(x)$ will go to 0 by the law of large numbers. In the supermarket model, the deterministic process corresponds exactly to the differential equations we have in (1), as can be seen by taking the derivative of (7). Also, in the supermarket model we have $x_0 = X_n(0) = (1, 0, 0, \dots)$ in the case where we begin with the empty system.

We now present Kurtz's theorem.

Theorem 7 [Kurtz]. *Suppose we have a density dependent family satisfying the Lipschitz condition*

$$|F(x) - F(y)| \leq M|x - y|$$

for some constant M . Further, suppose $\lim_{n \rightarrow \infty} X(0) = x_0$ and let X be the deterministic process:

$$X(t) = x_0 + \int_0^t F(X(u))du, \quad t \geq 0.$$

Consider the path $\{X(u) : u \leq t\}$ for some fixed $t \geq 0$ and assume that there exists a neighborhood K around this path satisfying

$$\sum_{l \in L} |l| \sup_{x \in K} \beta_l(x) < \infty. \quad (8)$$

Then,

$$\lim_{n \rightarrow \infty} \sup_{u \leq t} |X_n(u) - X(u)| = 0 \quad a.s.$$

Kurtz's theorem says that the limiting process is indeed the deterministic process given by the appropriate differential equations. Although we do not show it here, one can use the proof of Kurtz's theorem to bound the deviation between the finite and the limiting system as well. These bounds generally take the same form as Chernoff-type bounds, up to constant factors.

We note that the most important condition for Kurtz's theorem is that the underlying density dependent Markov chain is Lipschitz. Hence, we show here that the limiting supermarket model is Lipschitz.

Lemma 4. *The supermarket model satisfies the Lipschitz condition.*

Proof. For the supermarket model, we have

$$F(x) = \sum_{i=1}^{\infty} \lambda(x_{i-1}^d - x_i^d) - (x_i - x_{i+1}).$$

Let $x = (x_i)$ and $y = (y_i)$ be two states of the supermarket model. Then,

$$\begin{aligned} |F(x) - F(y)| &\leq \sum_{i=1}^{\infty} |\lambda(x_{i-1}^d - x_i^d) - (x_i - x_{i+1}) - \lambda(y_{i-1}^d - y_i^d) \\ &\quad + (y_i - y_{i+1})| \\ &\leq 2 \sum_{i=0}^{\infty} |x_i - y_i| + 2\lambda \sum_{i=0}^{\infty} |x_i^d - y_i^d| \\ &\leq \sum_{i=0}^{\infty} (2 + 2d\lambda) |x_i - y_i|, \end{aligned}$$

where we have used the fact that $0 \leq x_i, y_i \leq 1$ for all i . \square

ACKNOWLEDGMENTS

The author thanks Alistair Sinclair for his guidance while this work was developed. The author also thanks Andrei Broder, Yossi Azar, Richard Karp, and Michael Luby for many helpful discussions. Most of this work was done while the author was at the U.C. Berkeley Department of Computer Science and at Compaq Systems Research Center, Palo Alto, California. This work was supported in part by an Office of Naval Research graduate research fellowship, US National Science Foundation Grant CCR-9505448, US National Science Foundation CAREER Grant CCR-9983832, and an Alfred P. Sloan Research Fellowship. Part of this work has previously appeared as part of the author's thesis and in *Proceedings for the 37th Annual Symposium on Foundations of Computer Science*.

REFERENCES

- [1] Y. Azar, A. Broder, A. Karlin, and E. Upfal, "Balanced Allocations," *Proc. 26th ACM Symp. Theory of Computing*, pp. 593-602, 1994.
- [2] M. Adler, S. Chakrabarti, M. Mitzenmacher, and L. Rasmussen, "Parallel Randomized Load Balancing," *Proc. 27th ACM Symp. Theory of Computing*, pp. 238-247, 1995.
- [3] D. Achlioptas and M. Molloy, "The Analysis of a List-Coloring Algorithm on a Random Graph," *Proc. 38th IEEE Symp. Foundations of Computer Science*, 1997.
- [4] A. Bestavros, "Load Profiling: A Methodology for Scheduling Real-Time Tasks in a Distributed System," *Proc. ICDCS '97: IEEE Int'l Conf. Distributed Computing Systems*, May 1997.
- [5] M. Dahlin, "Serverless Network File Systems," PhD thesis, Univ. of California, Berkeley, 1995.
- [6] K. Deimling, "Ordinary Differential Equations in Banach Spaces," *Lecture Notes in Math.*, vol. 96, 1977.
- [7] S.N. Ethier and T.G. Kurtz, *Markov Processes: Characterization and Convergence*. John Wiley and Sons, 1986.
- [8] D.L. Eager, E.D. Lazowska, and J. Zahorjan, "Adaptive Load Sharing in Homogeneous Distributed Systems," *IEEE Trans. Software Eng.*, vol. 12, pp. 662-675, 1986.
- [9] D.L. Eager, E.D. Lazowska, and J. Zahorjan, "A Comparison of Receiver-Initiated and Sender-Initiated Adaptive Load Sharing," *Performance Evaluation Rev.*, vol. 16, pp. 53-68, 1986.
- [10] B. Hajek, "Asymptotic Analysis of an Assignment Problem Arising in a Distributed Communications Protocol," *Proc. 27th Conf. Decision and Control*, pp. 1455-1459, 1988.
- [11] R.M. Karp, M. Luby, and F. Meyer auf der Heide, "Efficient PRAM Simulation on a Distributed Memory Machine," *Proc. 24th ACM Symp. Theory of Computing*, pp. 318-326, 1992.
- [12] R.M. Karp and M. Sipser, "Maximum Matchings in Sparse Random Graphs," *Proc. 22nd IEEE Symp. Foundations of Computer Science*, pp. 364-375, 1981.
- [13] R.M. Karp, U.V. Vazirani, and V.V. Vazirani, "An Optimal Algorithm for On-Line Bipartite Matching," *Proc. 22nd ACM Symp. Theory of Computing*, pp. 352-358, 1990.
- [14] T.G. Kurtz, "Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes," *J. Applied Probability*, vol. 7, pp. 49-58, 1970.
- [15] T.G. Kurtz, "Limit Theorems for Sequences of Jump Markov Processes Approximating Ordinary Differential Processes," *J. Applied Probability*, vol. 9, pp. 344-356, 1971.
- [16] T.G. Kurtz, "Strong Approximation Theorems for Density Dependent Markov Chains," *Stochastic Processes and Applications*, vol. 6, pp. 223-240, 1978.
- [17] T.G. Kurtz, *Approximation of Population Processes*. SIAM Press, 1981.
- [18] M. Luby, M. Mitzenmacher, M.A. Shokrollahi, D. Spielman, and V. Stemann "Practical Loss-Resilient Codes," *Proc. 29th ACM Symp. Theory of Computing*, pp. 150-159, 1997.
- [19] P.D. MacKenzie, C.G. Plaxton, and R. Rajaraman, "On Contention Resolution Protocols and Associated Probabilistic Phenomena," *J. ACM*, vol. 45, no. 2, pp. 324-378, 1998.

- [20] R. Abraham, J. Marsden, and T. Raitu, *Manifolds, Tensor Analysis, and Applications*. Addison-Wesley, 1983.
- [21] A.N. Michel and R.K. Miller, "Qualitative Analysis of Large Scale Dynamical Systems," *Math. Science and Eng.*, vol. 134, 1977.
- [22] R. Mirchandaney, D. Towsley, and J.A. Stankovic, "Analysis of the Effects of Delays on Load Sharing," *IEEE Trans. Computers*, vol. 38, pp. 1513-1525, 1989.
- [23] R. Mirchandaney, D. Towsley, and J.A. Stankovic, "Adaptive Load Sharing in Heterogeneous Distributed Systems," *J. Parallel and Distributed Computing*, vol. 9, pp. 331-346, 1990.
- [24] M. Mitzenmacher, "Load Balancing and Density Dependent Jump Markov Processes," *Proc. 37th IEEE Symp. Foundations of Computer Science*, pp. 213-222, 1996.
- [25] M. Mitzenmacher, "The Power of Two Choices in Randomized Load Balancing," PhD thesis, Univ. of California, Berkeley, Sept. 1996.
- [26] M. Mitzenmacher, "On the Analysis of Randomized Load Balancing Schemes," *Theory of Computing Systems*, vol. 32, pp. 361-386, 1999.
- [27] M. Mitzenmacher, "How Useful is Old Information?" *IEEE Trans. Parallel and Distributed Systems*, vol. 11, no. 1, pp. 6-20, Jan. 2000.
- [28] M. Mitzenmacher, A. Richa, and R. Sitaraman, "The Power of Two Random Choices: A Survey of Techniques and Results," *Handbook of Randomized Computing*, vol. 1, June 2001.
- [29] B. Pittel, J. Spencer, and N. Wormald, "Sudden Emergence of a Giant k-Core in a Random Graph," *J. Combinatorial Theory, Series B*, vol. 67, pp. 111-151, 1996.
- [30] J.R. Santos and R. Muntz, "Design of the RIO (Randomized I/O) Storage Server," technical report, Computer Science Dept., Univ. of Calif., Los Angeles, May 1997.
- [31] A. Shwartz and A. Weiss, *Large Deviations for Performance Analysis*. Chapman and Hall, 1995.
- [32] V. Stemann, "Parallel Balanced Allocations," *Proc. Eighth Ann. ACM Symp. Parallel Algorithms and Architectures*, pp. 261-269, 1996.
- [33] 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, pp. 15-27, 1996.
- [34] R. Weber, "On the Optimal Assignment of Customers to Parallel Servers," *J. Applied Probability*, vol. 15, pp. 406-413, 1978.
- [35] W. Winston, "Optimality of the Shortest Line Discipline" *J. Applied Probability*, vol. 14, pp. 181-189, 1977.
- [36] N.C. Wormald, "Differential Equations for Random Processes and Random Graphs," *Annals of Applied Probability*, vol. 5, pp. 1217-1235, 1995.
- [37] S. Zhou, "A Trace-Driven Simulation Study of Dynamic Load Balancing," *IEEE Trans. Software Eng.*, vol. 14, pp. 1327-1341, 1988.



Michael Mitzenmacher received the PhD degree in computer science from the University of California, Berkeley, in 1996. From 1996 to 1998, he worked as a researcher at Digital Systems Research Center, now owned by Compaq. In January 1999, he joined Harvard University as an assistant professor of computer science. His research interests focus on random processes and algorithms, including load balancing, algorithms for the Web, network protocols, stochastic bin packing, coding theory, and compression. He is a member of the IEEE.

► **For more information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.**