Load Balancing and Density Dependent Jump Markov Processes

EXTENDED ABSTRACT

Michael Mitzenmacher *
Department of Computer Science
U.C. Berkeley
Berkeley, CA 94720
mitzen@cs.berkeley.edu

Abstract

We provide a new approach for analyzing both static and dynamic randomized load balancing strategies. We demonstrate the approach by providing the first analysis of the following 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 the expected time a customer spends in the system in equilibrium. The model provides a good abstraction of a simple, efficient load balancing scheme in the setting where jobs arrive at a large system of parallel processors. This model appears more realistic than similar models studied previously, in that it is both dynamic and open: that is, customers arrive over time, and the number of customers is not fixed.*

Our approach consists of two distinct stages: we first develop a limiting, deterministic model representing the behavior as $n \to \infty$, and then show how to translate results from this model to results for large, but finite, values of n. The analysis of the deterministic model is interesting in its own right. This methodology proves effective for studying a number of similar problems, and simulations demonstrate that the method accurately predicts system behavior even for relatively small systems.

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 some constant number d of servers independently and uniformly at random with replacement from the n servers, and waits for service at the server 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.[1] We call this problem the *supermarket model*, or the *supermarket system*. 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 model studied by Azar *et al.* [7], 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. We note that an analysis of the more realistic *open* dynamic model, the supermarket model, has remained open. As described in [7], models of this kind have a number of applications to computing problems, including resource allocation, hashing, and on-line 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 jobs arrive and execute at a large system of parallel processors, and variations of the model are suitable for studying the behavior of dynamic hash tables.

In this paper, we introduce a new approach for analyzing the supermarket model and other related static and dynamic randomized load balancing strategies. Our strategy has two main components:

- We define an idealized process, corresponding to a system of infinite size. This process is cleaner and easier to analyze because its behavior is completely deterministic.

*This work was supported in part by the ONR and in part by NSF Grant CCR-9505448.

[1] Some might object to the requirement that the service time is exponentially distributed. We note that by using our techniques one can also develop approximations for other service distributions, including the case where the service time is constant. The details appear in [27].

- We relate the idealized system to the finite system, carrying over the analysis with bounded error.

For example, in the supermarket model, the intuition is that if we look at the *fraction* of servers containing at least k customers for every k, the system evolves in an almost deterministic way as $n \to \infty$. The analysis of this system is interesting in its own right. Then we bound the deviation between a system of finite size n and the infinite system.

Besides providing the first analysis of the supermarket model, we note that this approach also provides a clean, systematic approach to analyzing several other load balancing models. In particular, the method provides a means of finding accurate numerical estimates of performance.

The following result is typical of our method:

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

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

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

The summation is derived from the infinite system, and the $o(1)$ term arises when we bound the error between the infinite system and the system 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 2.3, for λ close to 1 there is an exponential improvement in the expected time a customer spends in the system.

1.1 Previous work

The power of using two hash functions for load balancing was demonstrated by Karp, Luby, and Meyer auf der Heide [16] in an application to PRAM simulation. They analyzed a static case, where a number of customers are to be permanently distributed among a fixed number of servers. The static problem was further developed and analyzed by Azar *et al.* in [7], and results for other static settings are given by MacKenzie *et al.* [24] and Adler *et al.* [3]. Related dynamic models, where one is concerned with the behavior of a system over an arbitrary time interval, have proven more difficult. A less realistic dynamic model was successfully analyzed by Azar *et al.* [7], and a dynamic edge orientation problem related to load balancing was analyzed by Ajtai *et al.* in [5]. We note that our method can also be applied to these dynamic models and to the static model of [7], providing new insight and results. 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.

Although the supermarket model can be described naturally as a queueing problem, little appears known about it in the queueing theory literature. When $d = 1$, the arrival processes at the queues are independent, and the problem becomes trivial. When $d > 1$, analysis proves difficult because the sizes of the various queues are dependent. A great deal of work has been done to analyze the variant where the customer selects the shortest queue (for example, see [2, 10, 25] for references); for more than two queues, only approximations or asymptotic expressions are known. The supermarket model has been studied by Eager *et al.* [12] and Zhou [31]. In fact, Eager *et al.* also use an approach based on Markov chains for their analysis; however, the authors make the crucial assumption that the state of each queue is stochastically independent of the state of any other queue [12, 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. Our work substantially improves upon their work by avoiding these assumptions, as well as by introducing several new directions in the analysis of these systems. Zhou 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 infinite systems we will use Kurtz's work on *density dependent jump Markov processes* [13, 20, 21, 22, 23], with some extensions specific to our problems. Kurtz's work has previously been applied to matching problems on random graphs [14, 17, 18] as well as some queueing models [28]; here, we apply it for the first time to load balancing problems. Given the increasing use of Markov chains in the analysis of algorithms, we believe that this technique may be more widely applicable than previously expected. Indeed, one goal of this paper is to encourage further use of this type of analysis.

The rest of the paper is structured as follows: in Section 2, we shall analyze the behavior of the infinite version of the supermarket model. In Section 3, we shall briefly explain Kurtz's work and how to adapt it to relate the finite and infinite versions of the supermarket model; more technical details are available in [27]. In Section 4, we shall describe other dynamic problems to which our technique can be applied, as well as show how it can be used to study the static model of Azar *et al.* [7].

2 The analysis of the supermarket model

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 replacement[2] 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 known to be stable (see, for example, [19]). The comparison argument is entirely similar to those in [29, 30], 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 [6], 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 $s_i(t) = \sum_{k=i}^{\infty} p_i(t) = m_i(t)/n$ 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 $t \geq 0$, $\lim_{i\to\infty} s_i(t) = 0$. Under the same conditions, the expected number of customers per queue, or $\sum_{i=1}^{\infty} s_i(t)$, is finite even as $t \to \infty$.

We can represent the state of the system at any given time by an infinite dimensional vector $\vec{s} = (s_0, s_1, s_2, \ldots)$. 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 *infinite system* related to the finite supermarket system. The time evolution of the infinite system is specified by the following set of differential equations:

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

Let us explain the reasoning behind the system (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 $\lambda n(s_{i-1}^d - s_i^d)$. 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 the system (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 (see [1, p.188, Theorem 4.1.5] or [11, Theorem 3.2]). It should be intuitively clear that as $n \to \infty$ 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.

2.1 Finding a fixed point

We will demonstrate that, given a reasonable condition on the initial point $\vec{s}(0)$, the infinite 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 necessary and sufficient condition for $\vec{s}$ to be a fixed point is that for all i, $\frac{ds_i}{dt} = 0$.

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

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

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 the equations (1) over all $i \geq 1$. (Note that we use $\sum_{i=1}^{\infty} s_i < \infty$ here to ensure that the sum converges

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

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. ■

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

Definition 3 *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, while for $d = 1$ the tails decrease only geometrically (or singly exponentially).

2.2 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 $\pi_i = \lambda^{\frac{d^i - 1}{d-1}}$. 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 4 *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 4 holds for any initial state $\vec{s}$ derived from the initial state of a finite system.

Proof: Let $M(t) = \sup_i [s_i(t)/\pi_i]^{1/d^i}$. 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 $n \to \infty$ (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 $s_i(0) = M(0)^{d^i}\pi_i$. 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 non-decreasing in s_j for $j \neq i$. The result then follows from [11, pp. 70-74].

We now show that the s_i decrease doubly exponentially (in the infinite 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 $M(t) \leq M(0)$, $M(0)^{d^i} \geq s_i(t)/\pi_i$ for $t \geq 0$, or

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

Note that $\lambda^{1/(d-1)}M(0)$, 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. ■

To show convergence, we make use of a *potential function* (also called a *Lyapunov function* in the dynamical systems literature) $\Phi(t)$, which in some sense measures the distance of the current location to the fixed point.

Definition 5 *The potential function* Φ *is said to* converge exponentially to 0, *or simply to converge exponentially, if* $\Phi(t) \leq c_0 \mathrm{e}^{-\delta t}$ *for some constant* $\delta > 0$ *and a constant* c_0 *which may depend on the state at* $t = 0$.

We find a potential function Φ that shows that the system converges exponentially quickly to its fixed point. A natural potential function to consider is $\Phi(t) = \sum_{i=1}^{\infty} |s_i(t) - \pi_i|$, 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.

Theorem 6 *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 $\vec{\pi}$ 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

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

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 $\epsilon_{i-1}, \epsilon_i$, and ϵ_{i+1} are positive or negative. Let us consider the case where they are all negative (which, by Theorem 4, 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 $\epsilon_{i-1}, \epsilon_i$, and ϵ_{i+1} are positive, the above term (2) involving ϵ_i is negated, but now ϵ_i is positive. If $\epsilon_{i-1}, \epsilon_i$ and ϵ_{i+1} have mixed signs, this can only decrease the value of the term (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 whenever $\Phi(0) < \infty$.

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, [26, p. 16].) That is, we may define

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

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 $\epsilon_i < 0$. ∎

Theorem 6 yields the following corollary:

Corollary 7 *Under the conditions of Theorem 6, 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 6 are all at least 1 for $i \geq 1$, $\Phi(t) \geq d(t)$ and the corollary is immediate. ∎

Corollary 7 shows that the L_1 distance to the fixed point converges exponentially quickly to 0. Hence, from any suitable starting point, the infinite 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.

2.3 The expected time in the infinite system

Using Theorems 4 and 6, we now examine the expected time a customer spends in the infinite system.

Corollary 8 *The expected time a customer spends in the infinite supermarket system for $d \geq 2$, subject to the condition of Theorem 4, converges as $t \to \infty$ to*

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

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

Proof: An incoming customer that arrives at time t becomes the ith 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 $t \to \infty$, by Corollary 7, the infinite 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 $\sum_{i=0}^{\infty} \pi_i^d = \sum_{i=1}^{\infty} \lambda^{\frac{d^i-d}{d-1}}$ 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 infinite system $s_i(t) \leq \pi_i$ for all t by Theorem 4. ■

Recall that $T_1(\lambda) = 1/(1-\lambda)$ from standard queueing theory. Analysis of the summation in Corollary 8, which we omit in this extended abstract, reveals the following:

Theorem 9 *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,*

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

Choosing from $d > 1$ queues hence yields an exponential improvement in the expected time a customer spends in the infinite system, and as $\lambda \to 1^-$ 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 [7].

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 shortly. 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 *infinite 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 infinite 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 infinite system and the finite system are

- The infinite system is deterministic; the finite system is random.
- The infinite 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.

Theorem 10 *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*

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

where the $o(1)$ *is understood as* $n \to \infty$ *and depends on* T *and* λ. ■

Theorem 10 says the expected time in a finite system is the same as that for the infinite system (Corollary 8) plus an $o(1)$ term. Similarly, we can bound the maximum load:

Theorem 11 *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,*[3] *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 3.3.

3.2 Density dependent Markov chains

We now give a more technical presentation of the background for Kurtz's theorem. We begin with the definition of a density dependent family of Markov chains, as in Kurtz

[3] For this paper, *with high probability* will mean with probability $1 - O(1/n)$ and all logarithms have base e.

[23, Chapter 7], although we extend the definition to countably many dimensions. For convenience we drop the vector notation where it can be understood by context. Let $\mathbf{Z}^*$ be either $\mathbf{Z}^d$ for some dimension d, or $\mathbf{Z}^{\mathbf{N}}$, as appropriate. Given a set of transitions $L \subseteq \mathbf{Z}^*$ and a collection of nonnegative functions β_l for $l \in L$ defined on a subset $E \subset \mathbf{R}^*$, 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 $E_n = E \cap \{n^{-1}k : k \in \mathbf{Z}^*\}$ and the transition rates of X_n are

$$q^{(n)}_{x,y} = n\beta_{n(y-x)}(x), \;\; x, y \in E_n.$$

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 $q^{(n)}_{k,k+l} = n\beta_l(k/n) = n\beta_l(\vec{s})$, where $\beta_{e_i}(\vec{s}) = \lambda(s^2_{i-1} - s^2_i)$, and $\beta_{-e_i}(\vec{s}) = s_i - s_{i+1}$. These rates determined our infinite system (1).

It follows from [23, Chapter 7], that a Markov process $\hat{X}_n$, with intensities $q^{(n)}_{k,k+l} = n\beta_l(k/n)$ 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 $\mathbf{Z}^{\mathbf{N}}$ 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

$$F(x) = \sum_l \beta_l(x), \tag{3}$$

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

$$X_n(t) = X_n(0) + \sum_{l \in L} \frac{l}{n} \tilde{Y}_l \left(n \int_0^t \beta_l(X_n(u)) du \right) + \int_0^t F(X_n(u)) du, \tag{4}$$

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

$$X(t) = x_0 + \int_0^t F(X(u)) du, \; t \geq 0 \, , \tag{5}$$

where $x_0 = \lim_{n \to \infty} X(0)$. An interpretation relating equations (4) and (5) is that as $n \to \infty$, the value of the centered Poisson process $\tilde{Y}_l(x)$ 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 system (1), as can be seen by taking the derivative of equation (5). Also, in the supermarket model we have $x_0 = X_n(0) = (1, 0, 0, \ldots)$ in the case where we begin with the empty system.

We now present Kurtz's theorem (generalized to countably infinite dimensions).

Theorem 12 [Kurtz] *Suppose we have a density dependent family (of possibly countably infinite dimension) satisfying the Lipschitz condition*

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

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

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

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

$$\sum_{l \in L} |l| \sup_{x \in K} \beta_l(x) < \infty. \tag{6}$$

Then

$$\lim_{n \to \infty} \sup_{u \leq t} |X_n(u) - X(u)| = 0 \; a.s. \qquad \blacksquare$$

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 infinite system as well. These bounds generally take the same form as Chernoff-type bounds, up to constant factors.

3.3 Simulation results

We provide the results of some simulations based on the supermarket model. 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, and the first 10,000 steps are ignored in recording data in order to give the

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
	0.99	5.9275	5.4320	9.1227
3	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
	0.95	2.6137	2.5351	3.1002
	0.99	4.4080	3.8578	14.2607
5	0.50	1.0340	1.0312	0.2637
	0.70	1.1766	1.1681	0.7250
	0.80	1.3419	1.3289	0.9789
	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 1. The supermarket model: 100 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

Table 2. The supermarket model: 500 Queues

system time to approach equilibrium. For arrival rates of up to 95% of the service rate (i.e. $\lambda = 0.95$), the predictions are within a few percent of the simulation results. Even at 99% of capacity, the prediction is within 10% 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 give in Table 2.

The simulations clearly demonstrate the impact of having two choices. As previously mentioned, expected time a customer spends in the system in equilibrium given one choice ($d = 1$) is $1/(1-\lambda)$. Hence, as shown in Table 2, when $\lambda = 0.99$ the expected time in the system when $d = 1$ is 100.00; with two choices, this drops to under 6. Allowing additional choices leads to much less significant improvements. When the arrival rate is smaller the effect is less dramatic, but still apparent. The qualitative behaviors that we predicted with our analysis are thus readily observable in our simulations even of relatively small systems. This lends weight to the predictive power of our theoretical results in practical settings.

4 Other load balancing problems

4.1 Dynamic problems

One aspect of our approach that is appealing is that it generalizes quite easily to many similar dynamic models: one need only set up the right differential equations. Furthermore, even if one cannot prove convergence of the infinite system to a fixed point, one can generally calculate the solution to the differential equations numerically, and then use the numerical solution to bound or predict performance of specific systems over a fixed interval of time. We list here some other interesting models to which this method applies. In each case we have found the fixed point and shown that the fixed point is *stable*, in that the L_1 distance to the fixed point is nonincreasing in the infinite system. (Note the stability of the fixed point is different than the stability of the system!) For many cases we have shown exponential convergence to the fixed point. Here we simply briefly describe the models: the details are deferred to the full version of the paper, or can be found in [27].

1. Bounded buffers: Allow at most B customers in any queue; customers that arrive and choose only full queues are turned away.
2. Customer types: Customers proceed directly to one queue with probability p, and the shortest of two queues with probability $1-p$. This system can be used to model priorities, with higher priority customers getting more choices.
3. Threshold systems: Customers choose one server randomly, and stay there if the queue has at most T customers; otherwise, they proceed to another randomly selected queue.
4. Closed systems: Customers, after completing service, are recirculated through the system.
5. Edge orientation problem [5]: This problem is set on a complete graph of n vertices. The weight of a vertex is the difference between its indegree and outdegree. At each time step a random edge arrives and is directed toward the adjacent vertex with smaller weight.

This list is not meant to be exhaustive; the method should apply to many other similar load balancing systems.

4.2 The empty bins problem

We now demonstrate the applicability of the infinite systems approach to static problems by considering the $\text{GREEDY}(d)$ strategy of [7]. In this setting, initially there are n balls and n bins. Balls arrive sequentially. Upon arrival, each ball chooses d bins independently and uniformly at random (with replacement), and is then placed in the least loaded of these bins (ties being broken arbitrarily). With high probability, the maximum load is

$\log \log n / \log d + O(1)$. (Similar results also appeared in [16].)

We first consider the following simple question: how many bins remain empty after the protocol GREEDY(d) terminates? The question can also be seen as a matching problem: given a bipartite graph with n vertices on each side such that each vertex on the left has d edges to vertices chosen independently and uniformly at random on the right, what is the expected size of the greedy matching obtained by sequentially matching vertices on the left to a random unmatched neighbor? This question has been previously solved in the limiting case as $n \to \infty$ by Hajek using similar techniques [14]. We shall begin by briefly repeating his argument with some additional insights. We then extend the argument to the more general load balancing problem.

Theorem 13 *Suppose cn balls are thrown into n bins according to the* GREEDY*(d) protocol for some constant c. Let Y_{cn} be the number of non-empty bins when the process terminates. Then $\lim_{n\to\infty} \mathbf{E}[\frac{Y_{cn}}{n}] = y_c$, where $y_c < 1$ satisfies*

$$c = \sum_{i=0}^{\infty} \frac{y_c^{id+1}}{(id+1)}.$$

Proof: We set up the problem as a density dependent Markov chain. We let t be the time at which exactly $x(t) = nt$ balls have been thrown, and we let $y(t)$ be the fraction of non-empty bins. At time t, the probability that a ball finds at least one empty bin among its d choices is $1 - y^d$, and hence we have $\frac{dy}{dt} = 1 - y^d$. Instead of solving this equation for y in terms of t, we solve for t in terms of y: $\frac{dt}{dy} = \frac{1}{1-y^d} = \sum_{i=0}^{\infty} y^{id}$. We integrate, yielding

$$t_0 = \sum_{i=0}^{\infty} \frac{y(t_0)^{id+1}}{(id+1)}. \qquad (7)$$

From equation (7), given d we can solve for $y(t_0)$ for any value of t_0 using for example binary search.[4] In particular, when $t_0 = c$, all of the balls have been thrown, and the process terminates. Plugging $t_0 = c$ into equation (7) and applying Kurtz's theorem yields the theorem, with $y_c = y(c)$. ■

A marked difference between the static problem and the supermarket model is that in the static case we are only interested in the progress of the process over a fixed time interval, while in the dynamic case we are interested in the behavior of the model over an arbitrary period of time. In this respect, the static problem is easier than the corresponding dynamic problem.

[4] One could also attempt to solve the differential equation for y in terms of t. Standard integral tables [9] give such equations when $d = 2, 3$ and 4.

We may further conclude from the details of the proof of Kurtz's theorem that with high probability, $|Y_{cn}/n - y_c|$ is $O(\sqrt{\log n/n})$. Hence the number of empty bins is sharply concentrated around its expected value. One can also show that Y_{cn} is close to its mean with high probability using standard martingale arguments and the method of bounded differences, which we do here. We assume familiarity with basic martingale theory; see, for example, [4, Chapter 7]. We use the following form of the martingale tail inequality due to Azuma [8]:

Lemma 14 [Azuma] *Let $X_0, X_1, \ldots X_m$ be a martingale sequence such that for each k, $|X_k - X_{k-1}| \le 1$. Then for any $\alpha > 0$,*

$$\mathbf{Pr}(|X_m - X_0| > \alpha\sqrt{m}) < 2e^{-\alpha^2/2}. \qquad ■$$

Theorem 15 $\mathbf{Pr}(|Y_{cn} - \mathsf{E}[Y_{cn}]| > \alpha\sqrt{cn}) < 2e^{-\alpha^2/2}$ *for any $\alpha > 0$.*

Proof: We present an argument similar to that presented in [15, Theorem 2]. For $0 \le j \le cn$, let $\mathcal{F}_j$ be the σ-field of events corresponding to the possible states after j balls have been placed, and $Z_j = \mathsf{E}[Y_{cn}|\mathcal{F}_j]$ be the associated conditional expectation of Y_{cn}. Then the random variables $\{Z_j\}_{j=0}^{cn}$ form a Doob martingale, and it is clear that $|Z_j - Z_{j-1}| \le 1$. The theorem now follows from Lemma 14. ■

Theorem 15 implies that Y_{cn} is within $O(\sqrt{n \log n})$ of its expected value with high probability. Unlike the infinite system method, however, the martingale approach does not immediately lead us to the value to which Y_{cn}/n converges. This appears to be an advantage of the infinite system approach.

4.3 Bins with fixed load

We can extend the previous analysis to find the fraction of bins with load k for any constant k as $n \to \infty$. We first establish the appropriate density dependent Markov chain. Let $s_i(t)$ be the fraction of bins with load at least i at time t, where again at time t exactly nt balls have been thrown; note that $s_0(t) = 1$ for all t. Then the corresponding differential equations regarding the growth of the s_i (for $i \ge 1$) are easily determined:

$$\frac{ds_i}{dt} = s_{i-1}^d - s_i^d.$$

The differential equation (similar to the system (1) for the supermarket model) has the following simple interpretation: for there to be an increase in the number of bins with at least i balls, the d choices must all be from bins with load at least $i - 1$, but not all from bins of load at least i.

	$d=2$	1 million	$d=3$	1 million
s_1	0.7616	0.7616	0.8231	0.8230
s_2	0.2295	0.2295	0.1765	0.1765
s_3	0.0089	0.0089	0.00051	0.00051
s_4	0.000006	0.000007	$< 10^{-11}$	0
s_5	$< 10^{-11}$	0	$< 10^{-11}$	0

Table 3. Predictions vs. simulations for GREEDY(d).

We are not aware of how to determine explicit formulae for $s_i(t)$ in general. However, this system of differential equations can be solved numerically using standard methods; for up to any fixed k, we can accurately determine $s_k(t)$.

Using Kurtz's theorem or martingales one can show that these results will be accurate with high probability. We also demonstrate that our technique accurately predicts the behavior of the GREEDY(d) algorithm by comparing with simulation results. The first and third columns of Table 3 shows the predicted values of s_i for $d = 2$ and $d = 3$. From these results, with $d = 2$, one would not expect to see bins with load five until billions of balls have been thrown. Similarly, choosing $d = 3$ one expects a maximum load of three until billions of balls have been thrown. These results match simulation results presented in [3] and [6]. We also present the averages from one hundred simulations of one million balls, which further demonstrate the accuracy of this technique. This accuracy is a marked advantage of this approach; previous techniques have not provided ways of concretely predicting actual performance.

We also note that we can use Kurtz's theorem to give an alternative proof of of the $\frac{\log\log n}{\log d} + O(1)$ bounds for the GREEDY(d) process. Details appear in [27].

5 Acknowledgments

The author thanks Alistair Sinclair for several helpful discussions and for suggestions improving the presentation of the paper. The author also thanks Andrei Broder, Yossi Azar, Richard Karp, and Michael Luby for helpful discussions.

References

[1] R. Abraham, J. Marsden, and T. Raitu. *Manifolds, Tensor Analysis, and Applications*. Addison-Wesley, 1983.

[2] I. J. B. F. Adan, G. van Houtum, and J. van der Wal. Upper and lower bounds for the waiting time in the symmetric shortest queue system. *Annals of Operations Research*, 48:197–217, 1994.

[3] M. Adler, S. Chakrabarti, M. Mitzenmacher, and L. Rasmussen. Parallel randomized load balancing. In *Proceedings of the 27th ACM Symposium on the Theory of Computing*, pages 238–247, 1995.

[4] N. Alon, J. Spencer, and P. Erdös. *The Probabilistic Method*. John Wiley and Sons, 1992.

[5] M. Atjai, J. Aspnes, M. Naor, Y. Rabani, L. Schulman, and O. Waarts. Fairness in scheduling. In *Proceedings of the Sixth Annual ACM/SIAM Symposium on Discrete Algorithms*, pages 477–485, 1995.

[6] Y. Azar, A. Broder, A. Karlin, and E. Upfal. Balanced allocations. Journal version, preprint.

[7] Y. Azar, A. Broder, A. Karlin, and E. Upfal. Balanced allocations. In *Proceedings of the 26th ACM Symposium on the Theory of Computing*, pages 593–602, 1994.

[8] K. Azuma. Weighted sums of certain dependent random variables. *Tôhoku Mathematical Journal*, 19:357–367, 1967.

[9] W. H. Beyer, editor. *CRC Handbook of Mathematical Sciences: 6th Edition*. CRC Press, 1987.

[10] J. P. C. Blanc. The power-series algorithm applied to the shortest-queue model. *Operations Research*, 40:157–167, 1992.

[11] K. Deimling. *Ordinary Differential Equations in Banach Spaces*. Springer-Verlag, 1977. Lecture Notes in Mathematics. Vol. 96.

[12] D. L. Eager, E. D. Lazokwska, and J. Zahorjan. Adaptive load sharing in homogeneous distributed systems. *IEEE Transactions on Software Engineering*, 12:662–675, 1986.

[13] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence.* John Wiley and Sons, 1986.

[14] B. Hajek. Asymptotic analysis of an assignment problem arising in a distributed communications protocol. In *Proceedings of the 27th Conference on Decision and Control*, pages 1455–1459, 1988.

[15] A. Kamath, R. Motwani, K. Palem, and P. Spirakis. Tail bounds for occupancy and the satisfiability threshold conjecture. In *Proceedings of the 35th IEEE Symposium on Foundations of Computer Science*, 1994.

[16] R. M. Karp, M. Luby, and F. Meyer auf der Heide. Efficient PRAM simulation on a distributed memory machine. In *Proceedings of the 24th ACM Symposium on Theory of Computing*, pages 318–326, 1992.

[17] R. M. Karp and M. Sipser. Maximum matchings in sparse random graphs. In *Proceedings of the 22nd IEEE Symposium on Foundations of Computer Science*, pages 364–375, 1981.

[18] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 27th Conference on Decision and Control*, pages 352–358, 1990.

[19] L. Kleinrock. *Queueing Systems, Volume I: Theory*. John Wiley and Sons, 1976.

[20] T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7:49–58, 1970.

[21] T. G. Kurtz. Limit theorems for sequences of jump Markov processes approximating ordinary differential processes. *Journal of Applied Probability*, 8:344–356, 1971.

[22] T. G. Kurtz. Strong approximation theorems for density dependent Markov chains. *Stochastic Processes and Applications*, 6:223–240, 1978.

[23] T. G. Kurtz. *Approximation of Population Processes. CBMS-NSF Regional Conf. Series in Applied Math*. SIAM, 1981.

[24] P. D. MacKenzie, C. G. Plaxton, and R. Rajaraman. On contention resolution protocols and associated probabilistic phenomena. Department of Computer Science TR-94-06, University of Texas at Austin, April 1994. An extended abstract appears in STOC 1994.

[25] R. Menich and R. Serfozo. Optimality of routing and servicing in dependent parallel processing systems. *Queueing Systems*, 9:403–418, 1991.

[26] A. N. Michel and R. K. Miller. *Qualitative Analysis of Large Scale Dynamical Systems*. Academic Press, Inc., 1977.

[27] M. Mitzenmacher. *The Power of Two Choices in Randomized Load Balancing*. PhD thesis, University of California, Berkeley, September 1996.

[28] A. Shwartz and A. Weiss. *Large Deviations for Performance Analysis*. Chapman & Hall, 1995.

[29] R. R. Weber. On the optimal assignment of customers to parallel servers. *Journal of Applied Probability*, 15:406–413, 1978.

[30] W. Winston. Optimality of the shortest line discipline. *Journal of Applied Probability*, 14:181–189, 1977.

[31] S. Zhou. A trace-driven simulation study of dynamic load balancing. *IEEE Transactions on Software Engineering*, 14:1327–1341, 1988.