

BALANCED ALLOCATIONS*

YOSSI AZAR[†], ANDREI Z. BRODER[‡], ANNA R. KARLIN[§], AND ELI UPFAL[¶]

Abstract. Suppose that we sequentially place n balls into n boxes by putting each ball into a randomly chosen box. It is well known that when we are done, the fullest box has with high probability $(1 + o(1)) \ln n / \ln \ln n$ balls in it. Suppose instead that for each ball we choose two boxes at random and place the ball into the one which is less full at the time of placement. We show that with high probability, the fullest box contains only $\ln \ln n / \ln 2 + O(1)$ balls—exponentially less than before. Furthermore, we show that a similar gap exists in the infinite process, where at each step one ball, chosen uniformly at random, is deleted, and one ball is added in the manner above. We discuss consequences of this and related theorems for dynamic resource allocation, hashing, and on-line load balancing.

Key words. urn models, occupancy problems, on-line algorithms, resource allocation, hashing, load balancing

AMS subject classifications. 68Q25, 68M20, 68P10, 60G99, 60G10, 90C40

PII. S0097539795288490

1. Introduction. Suppose that we sequentially place n balls into n boxes by putting each ball into a randomly chosen box. Properties of this random allocation process have been extensively studied in the probability and statistics literature. (See, e.g., [20, 17].) One of the classical results in this area is that, asymptotically, when the process has terminated, with high probability (that is, with probability $1 - o(1)$) the fullest box contains $(1 + o(1)) \ln n / \ln \ln n$ balls. (G. Gonnet [16] has proven a more accurate result, $\Gamma^{-1}(n) - 3/2 + o(1)$.)

Consider a variant of the process above whereby each ball comes with d possible destinations, chosen independently and uniformly at random. (Hence the d destinations are not necessarily distinct.) The ball is placed in the least full box among the d possible locations. Surprisingly, even for $d = 2$, when the process terminates the fullest box has only $\ln \ln n / \ln 2 + O(1)$ balls in it. Thus, this apparently minor change in the random allocation process results in an exponential decrease in the maximum occupancy per location. The analysis of this process is summarized as follows

THEOREM 1.1. *Suppose that m balls are sequentially placed into n boxes. Each ball is placed in the least full box, at the time of the placement, among d boxes, $d \geq 2$, chosen independently and uniformly at random. Then after all the balls are placed,*

- *with high probability, as $n \rightarrow \infty$ and $m \geq n$, the number of balls in the fullest box is $(1 + o(1)) \ln \ln n / \ln d + \Theta(m/n)$;*

*Received by the editors July 5, 1995; accepted for publication (in revised form) August 4, 1997; published electronically September 14, 1999.

<http://www.siam.org/journals/sicomp/29-1/28849.html>

[†]Department of Computer Science, Tel Aviv University, Tel-Aviv 69978, Israel (azar@math.tau.ac.il). This research was supported by the Israel Academy of Sciences.

[‡]Digital Systems Research Center, 130 Lytton Avenue, Palo Alto, CA 94301 (broder@src.dec.com).

[§]Department of Computer Science, University of Washington, Seattle, WA 98195 (karlin@cs.washington.edu).

[¶]IBM Almaden Research Center, San Jose, CA 95120 and Department of Applied Mathematics, The Weizmann Institute of Science, Rehovot, Israel (eli@wisdom.weizmann.ac.il). Work at the Weizmann Institute was supported in part by the Norman D. Cohen Professorial Chair of Computer Science.

- *in particular, with high probability, as $n \rightarrow \infty$ and $m = n$, the number of balls in the fullest box is $\ln \ln n / \ln d + \Theta(1)$;*
- *any other on-line strategy that places each ball into one of d randomly chosen boxes results in stochastically more balls¹ in the fullest box.*

It is also interesting to study the infinite version of the random allocation process. There, at each step a ball is chosen uniformly at random and removed from the system, and a new ball appears. The new ball comes with d possible destinations, chosen independently at random, and it is placed into the least full box among these d possible destinations.

The analysis of the case $d = 1$ in this infinite stochastic process is simple since the location of any ball does not depend on the locations of other balls in the system. Thus, for $d = 1$, in the stationary distribution, with high probability the fullest box has $\Theta(\log n / \log \log n)$ balls. The analysis of the case $d \geq 2$ is significantly harder, since the locations of the current n balls might depend on the locations of balls that are no longer in the system. We prove that when $d \geq 2$, in the stationary distribution, the fullest box has $\ln \ln n / \ln d + O(1)$ balls, with high probability. Thus, the same exponential gap holds in the infinite process. Theorem 1.2 is proven in section 4.

THEOREM 1.2. *Consider the infinite process with $d \geq 2$, starting at time 0 in an arbitrary state. There is a constant c such that for any fixed $T > cn^2 \log \log n$, the fullest box at time T contains, with high probability, less than $\ln \ln n / \ln d + O(1)$ balls. Thus, in the stationary distribution, with high probability, no box contains more than $\ln \ln n / \ln d + O(1)$ balls.*

Karp, Luby, and Meyer auf der Heide [18] were the first to notice a dramatic improvement when switching from one hash function to two in the context of PRAM simulations. In fact, it is possible to use a result from [18] to derive a weaker form of our static upper bound. (For details see [7].)

A preliminary version of this paper has appeared in [7]. Subsequently, Adler et al. [1] analyzed parallel implementation of the balanced allocation mechanism and obtained interesting communication vs. load tradeoffs.

A related question was considered by Broder et al. [10]. In their model the set of choices is such that there is a placement that results in maximum load equal to one. The question they analyze is, what is the expected maximum load under a random order of insertion under the greedy strategy?

More recent results, based on the balanced allocation paradigm, have appeared in [23, 24, 25, 12].

1.1. Applications. Our results have a number of interesting applications to computing problems. We elaborate here on three of them.

1.1.1. Dynamic resource allocation. Consider a scenario in which a user or a process has to choose between a number of identical resources on-line (choosing a server to use among the servers in a network, choosing a disk to store a directory, etc.). To find the least loaded resource, users may check the load on all resources before placing their requests. This process is expensive, since it requires sending an interrupt to each of the resources. A second approach is to send the task to a random resource. This approach has minimum overhead, but if all users follow it, the difference in load between different servers will vary by up to a logarithmic factor. Our analysis suggests

¹By this we mean that for any other strategy and any k the probability that the number of balls in the fullest box is greater than k is at least the probability that the number of balls in the fullest box is greater than k under the greedy strategy. See Corollary 3.6.

a more efficient solution. If each user samples the load of two resources and sends his request to the least loaded, the total overhead is small, and the load on the n resources varies by only a $O(\log \log n)$ factor.

1.1.2. Hashing. The efficiency of a hashing technique is measured by two parameters: the expected and the maximum access time. Our approach suggests a simple hashing technique, similar to hashing with chaining. We call it *2-way chaining*. It has $O(1)$ expected and $O(\log \log n)$ maximum access time. We use two random hash functions. The two hash functions define two possible entries in the table for each key. The key is inserted to the least full location, at the time of the insertion. Keys in each entry of the table are stored in a linked list. Assume that n keys are sequentially inserted by this process into a table of size n . As shown in section 5, the expected insertion and look-up time is $O(1)$, and our analysis summarized above immediately implies that with high probability the maximum access time is $\ln \ln n / \ln 2 + O(1)$, vs. the $\Theta(\log n / \log \log n)$ time when only one random hash function is used.

An advantage of our scheme over some other known techniques for reducing worst-case behavior of hashing (e.g., [14, 13, 11]) is, that it uses only two hash functions, it is easy to parallelize, and it does not involve rehashing of data. Other commonly used schemes partition the available memory into multiple tables, and use a different hash function in each table. For example, the Fredman, Komlos, Szemerédi scheme for perfect hashing [14] uses up to n different hash functions to get $O(1)$ worst-case access time (not on-line however), and the algorithm of Broder and Karlin [11] uses $O(\log \log n)$ hash functions to achieve $O(\log \log n)$ maximum access time on-line but using rehashings.

Karp, Luby, and Meyer auf der Heide [18] studied the use of two hash functions in the context of PRAM simulations. Other PRAM simulations using multiple hash functions were developed and analyzed in [21].

1.1.3. Competitive on-line load balancing. Consider the following on-line load balancing problem: We are given a set of n servers and a sequence of arrivals and departures of tasks. Each task comes with a list of servers on which it can be executed. The load balancing algorithm has to assign each task to a server on-line, with no information on future arrivals and departures of tasks. The goal of the algorithm is to minimize the maximum load on any server. The quality of an on-line algorithm is measured by the *competitive ratio*: the ratio between the maximum load it achieves and the maximum load achieved by the optimal off-line algorithm that knows the whole sequence in advance. This load balancing problem models, for example, communication in heterogeneous networks containing workstations, I/O devices, etc. Servers correspond to communication channels and tasks correspond to requests for communication links between devices. A network controller must coordinate the channels so that no channel is too heavily loaded.

On-line load balancing has been studied extensively against worst-case adversaries [9, 6, 5, 3, 8, 4]. For permanent tasks (tasks that arrive but never depart), Azar, Naor and Rom [9] showed that the competitive ratio of the greedy algorithm is $\log n$ and that no algorithm can do better. For temporary tasks (tasks that depart at unpredictable times), the works of Azar, Broder, and Karlin [6] and Azar et al. [8] show that there is an algorithm with competitive ratio $\Theta(\sqrt{n})$ and that no algorithm can do better.

It is interesting to compare these high competitive ratios, obtained from inputs generated by an adversary, to the competitive ratio against randomly generated inputs. Our results show that under reasonable probabilistic assumptions the compet-

itive ratios for both permanent and temporary tasks are significantly better. In the case of permanent tasks, if the set of servers on which a task can be executed is a small set (that is, constant size ≥ 2) chosen at random, the competitive ratio decreases from $\Theta(\log n)$ to $\Theta(\log \log n)$. In the case of temporary tasks, if we further assume that at each time step a randomly chosen existent task is replaced by a new task, then at any fixed time the ratio between the maximum on-line load and the maximum off-line load is $\Theta(\log \log n)$ with high probability. Further details are presented in section 6.

2. Definitions and notation. We consider two stochastic processes: the *finite process* and the *infinite process*.

The finite process. There are n boxes, initially empty, and m balls. Each ball is allowed to go into $d \geq 1$ boxes chosen independently and uniformly at random. The balls arrive one by one, and a *placement algorithm* must decide on-line (that is, without knowing what choices are available to future balls) in which box to put each ball as it comes. Decisions are irrevocable. We will subsequently refer to this setup as a (m, n, d) -problem.

The infinite process. There are n boxes, initially containing n balls in an arbitrary state. (For example, all the balls could be in one box.) At each step, one random ball is removed, and one new ball is added; the new ball is allowed to go into $d \geq 1$ boxes chosen independently and uniformly at random. Once again, a placement algorithm must decide on-line (that is, without knowing what choices are available to future balls and without knowing which ball will be removed at any future time) in which box to put each arriving ball. Decisions are irrevocable.

We use the following notations for the random variables associated with a placement algorithm A . Note that the state at time t refers to the state immediately after the placement of the t th ball.

$\lambda_j^A(t)$ called the *load* of box j , is the number of balls in box j at time t , resulting from algorithm A .

$\nu_k^A(t)$ is the number of boxes that have load k at time t .

$\nu_{\geq k}^A(t)$ is the number of boxes that have load $\geq k$ at time t , that is, $\nu_{\geq k}^A(t) = \sum_{i \geq k} \nu_i^A(t)$.

h_t^A called the *height* of ball t (= the ball that arrives at time t), is the number of balls at time t in the box where ball t is placed. In other words, the first ball to be placed in a particular box has height 1, the second ball has height 2, etc.

$\mu_k^A(t)$ is the number of balls that have height k at time t .

$\mu_{\geq k}^A(t)$ is the number of balls that have height $\geq k$ at time t , that is, $\mu_{\geq k}^A(t) = \sum_{i \geq k} \mu_i^A(t)$.

We omit the superscript A when it is clear which algorithm we are considering. Constants were chosen for convenience, and we made no attempts to optimize them.

Algorithm GREEDY assigns ball j to the box that has the lowest load among the d random choices that j has. We use the superscript G for GREEDY.

The basic intuition behind the proofs that follow is simple: Let $p_i = \mu_{\geq i}/n$. Since the available choices for each ball are independent and $\nu_{\geq i} \leq \mu_{\geq i}$, we roughly have (“on average” and disregarding conditioning) $p_{i+1} \leq p_i^d$, which implies a doubly exponential decrease in p_i , once $\mu_{\geq i} < n/2$. Of course the truth is that $\mu_{\geq i+1}$ is strongly dependent on $\mu_{\geq i}$ and a rather complex machinery is required to construct a correct proof.

3. The finite process. We use the notation $B(n, p)$ to denote a binomially distributed random variable with parameters n and p , and start with the following standard lemma, whose proof is omitted.

LEMMA 3.1. *Let X_1, X_2, \dots, X_n be a sequence of random variables with values in an arbitrary domain, and let Y_1, Y_2, \dots, Y_n be a sequence of binary random variables, with the property that $Y_i = Y_i(X_1, \dots, X_i)$. If*

$$\Pr(Y_i = 1 \mid X_1, \dots, X_{i-1}) \leq p,$$

then

$$\Pr\left(\sum Y_i \geq k\right) \leq \Pr(B(n, p) \geq k),$$

and similarly if

$$\Pr(Y_i = 1 \mid X_1, \dots, X_{i-1}) \geq p,$$

then

$$\Pr\left(\sum Y_i \leq k\right) \leq \Pr(B(n, p) \leq k). \quad \square$$

We now turn to the analysis of the finite process. In what follows, we omit the argument t when $t = m$, that is, when the process terminates. In the interest of a clearer exposition, we start with the case $m = n$, although the general case (Theorem 3.7) subsumes it.

THEOREM 3.2. *The maximum load achieved by the GREEDY algorithm on a random (n, n, d) -problem is less than $\ln \ln n / \ln d + O(1)$ with high probability.*

Proof. Since the d choices for a ball are independent, we have

$$\Pr(h_t \geq i + 1 \mid \nu_{\geq i}(t-1)) = \frac{(\nu_{\geq i}(t-1))^d}{n^d}.$$

Let \mathcal{E}_i be the event that $\nu_{\geq i}(n) \leq \beta_i$ where β_i will be exposed later. (Clearly, \mathcal{E}_i implies that $\nu_{\geq i}(t) \leq \beta_i$ for $t = 1, \dots, n$.) Now fix $i \geq 1$ and consider a series of binary random variables Y_t for $t = 2, \dots, n$, where

$$Y_t = 1 \text{ iff } h_t \geq i + 1 \text{ and } \nu_{\geq i}(t-1) \leq \beta_i.$$

(Y_t is 1 if the height of the ball t is $\geq i + 1$ despite the fact that the number of boxes that have load $\geq i$ is less than β_i .)

Let ω_j represent the choices available to the j th ball. Clearly,

$$\Pr(Y_t = 1 \mid \omega_1, \dots, \omega_{t-1}) \leq \frac{\beta_i^d}{n^d} \stackrel{\text{def}}{=} p_i.$$

Thus we can apply Lemma 3.1 to conclude that

$$(3.1) \quad \Pr\left(\sum Y_t \geq k\right) \leq \Pr(B(n, p_i) \geq k).$$

Observe that conditioned on \mathcal{E}_i , we have $\mu_{\geq i+1} = \sum Y_t$. Therefore

$$(3.2) \quad \Pr(\mu_{\geq i+1} \geq k \mid \mathcal{E}_i) = \Pr\left(\sum Y_t \geq k \mid \mathcal{E}_i\right) \leq \frac{\Pr(\sum Y_t \geq k)}{\Pr(\mathcal{E}_i)}.$$

Combining (3.1) and (3.2) we obtain that

$$(3.3) \quad \Pr(\nu_{\geq i+1} \geq k \mid \mathcal{E}_i) \leq \Pr(\mu_{\geq i+1} \geq k \mid \mathcal{E}_i) \leq \frac{\Pr(B(n, p_i) \geq k)}{\Pr(\mathcal{E}_i)}.$$

We can bound large deviations in the binomial distribution with the formula (see, for instance, [2, Appendix A])

$$(3.4) \quad \Pr(B(n, p_i) \geq ep_i n) \leq e^{-p_i n},$$

which inspires us to set

$$\beta_i = \begin{cases} n, & i = 1, 2, \dots, 5; \\ \frac{n}{2e}, & i = 6; \\ \frac{e\beta_{i-1}^d}{n^{d-1}}, & i > 6. \end{cases}$$

With these choices $\mathcal{E}_{\geq 6} = \{\nu_6 \leq n/(2e)\}$ holds with certainty, and from (3.3) and (3.4), for $i \geq 6$

$$\Pr(\neg \mathcal{E}_{i+1} \mid \mathcal{E}_i) \leq \frac{1}{n^2 \Pr(\mathcal{E}_i)},$$

provided that $p_i n \geq 2 \ln n$. Since

$$\Pr(\neg \mathcal{E}_{i+1}) \leq \Pr(\neg \mathcal{E}_{i+1} \mid \mathcal{E}_i) \Pr(\mathcal{E}_i) + \Pr(\neg \mathcal{E}_i),$$

it follows that for $p_i n \geq 2 \ln n$

$$(3.5) \quad \Pr(\neg \mathcal{E}_{i+1}) \leq \frac{1}{n^2} + \Pr(\neg \mathcal{E}_i).$$

To finish the proof let i^* be the smallest i such that $\beta_{i^*}^d/n^d \leq 2 \ln n/n$. Notice that $i^* \leq \ln \ln n / \ln d + O(1)$ since

$$\beta_{i+6} = \frac{ne^{(d^i-1)/(d-1)}}{(2e)^{d^i}} \leq \frac{n}{2^{d^i}}.$$

As before,

$$\Pr(\nu_{\geq i^*+1} \geq 6 \ln n \mid \mathcal{E}_{i^*}) \leq \frac{\Pr(B(n, 2 \ln n/n) \geq 6 \ln n)}{\Pr(\mathcal{E}_{i^*})} \leq \frac{1}{n^2 \Pr(\mathcal{E}_{i^*})},$$

and thus

$$(3.6) \quad \Pr(\nu_{\geq i^*+1} \geq 6 \ln n) \leq \frac{1}{n^2} + \Pr(\neg \mathcal{E}_{i^*}).$$

Finally,

$$\begin{aligned} \Pr(\mu_{\geq i^*+2} \geq 1 \mid \nu_{\geq i^*+1} \leq 6 \ln n) &\leq \frac{\Pr(B(n, (6 \ln n/n)^d) \geq 1)}{\Pr(\nu_{\geq i^*+1} \leq 6 \ln n)} \\ &\leq \frac{n(6 \ln n/n)^d}{\Pr(\nu_{\geq i^*+1} \leq 6 \ln n)} \end{aligned}$$

by the Markov inequality, and thus

$$(3.7) \quad \Pr(\mu_{\geq i^*+2} \geq 1) \leq \frac{(6 \ln n)^d}{n^{d-1}} + \Pr(\nu_{\geq i^*+1} \geq 6 \ln n).$$

Combining (3.7), (3.6), and (3.5), we obtain that

$$\Pr(\nu_{\geq i^*+2} \geq 1) \leq \frac{(6 \ln n)^d}{n^{d-1}} + \frac{i^* + 1}{n^2} = o(1),$$

which implies that with high probability the maximum load achieved by GREEDY is less than $i^* + 2 = \ln \ln n / \ln d + O(1)$. \square

We now prove a matching lower bound.

THEOREM 3.3. *The maximum load achieved by the GREEDY algorithm on a random (n, n, d) -problem is at least $\ln \ln n / \ln d - O(1)$ with high probability.*

Proof. Let \mathcal{F}_i be the event that $\nu_{\geq i}(n(1-1/2^i)) \geq \gamma_i$ where γ_i will be exposed later. For the time being, it suffices to say that $\gamma_{i+1} < \gamma_i/2$. We want to compute $\Pr(\neg \mathcal{F}_{i+1} \mid \mathcal{F}_i)$. To this aim, for t in the range $R = \{n(1-1/2^i)+1, \dots, n(1-1/2^{i+1})\}$, let Z_t be defined by

$$Z_t = 1 \text{ iff } h_t = i + 1 \text{ or } \nu_{\geq i+1}(t-1) \geq \gamma_{i+1},$$

and observe that while $\nu_{\geq i+1}(t-1) < \gamma_{i+1}$, if $Z_t = 1$, then the box where the t th ball is placed had load exactly i at time $t-1$. This means that all the d choices that ball t had pointed to boxes with load $\geq i$ and at least one choice pointed to a box with load exactly i .

Now let ω_j represent the choices available to the j th ball. Using

$$\begin{aligned} \Pr(A \vee B \mid C) &= \Pr(A \wedge \bar{B} \mid C) + \Pr(B \mid C) \\ &= \Pr(A \mid \bar{B} \wedge C) \Pr(\bar{B} \mid C) + \Pr(B \mid C) \geq \Pr(A \mid \bar{B} \wedge C), \end{aligned}$$

and in view of the observation above, we derive that

$$(3.8) \quad \Pr(Z_t = 1 \mid \omega_1, \dots, \omega_{t-1}, \mathcal{F}_i) \geq \left(\frac{\gamma_i}{n}\right)^d - \left(\frac{\gamma_{i+1}}{n}\right)^d \geq \frac{1}{2} \left(\frac{\gamma_i}{n}\right)^d \stackrel{\text{def}}{=} p_i.$$

Applying Lemma 3.1 we get

$$\Pr\left(\sum_{t \in R} Z_t \leq k \mid \mathcal{F}_i\right) \leq \Pr(B(n/2^{i+1}, p_i) \leq k).$$

We now choose

$$\begin{aligned} \gamma_0 &= n; \\ \gamma_{i+1} &= \frac{\gamma_i^d}{2^{i+3} n^{d-1}} = \frac{n}{2^{i+3}} \left(\frac{\gamma_i}{n}\right)^d = \frac{1}{2} \frac{n}{2^{i+1}} p_i. \end{aligned}$$

Since $\Pr(B(N, p) < Np/2) < e^{-Np/8}$ (see, for instance, [2, Appendix A]), it follows that

$$(3.9) \quad \Pr(B(n/2^{i+1}, p_i) \leq \gamma_{i+1}) = o(1/n^2),$$

provided that $p_i n / 2^{i+1} \geq 17 \ln n$. Let i^* be the largest integer for which this holds. Clearly $i^* = \ln \ln n / \ln d - O(1)$.

Now observe that by the definition of \mathcal{F} and Z_t , the event $\{\sum_{t \in R} Z_t \geq \gamma_{i+1}\}$ implies \mathcal{F}_{i+1} . Thus in view of (3.8) and (3.9)

$$\Pr(\neg \mathcal{F}_{i+1} \mid \mathcal{F}_i) \leq \Pr\left(\sum_{t \in R} Z_t \leq \gamma_{i+1} \mid \mathcal{F}_i\right) = o(1/n^2),$$

and therefore

$$\begin{aligned} \Pr(\mathcal{F}_{i^*}) &\geq \Pr(\mathcal{F}_{i^*} \mid \mathcal{F}_{i^*-1}) \times \Pr(\mathcal{F}_{i^*-1} \mid \mathcal{F}_{i^*-2}) \times \cdots \times \Pr(\mathcal{F}_1 \mid \mathcal{F}_0) \times \Pr(\mathcal{F}_0) \\ &\geq (1 - 1/n^2)^{i^*} = 1 - o(1/n), \end{aligned}$$

which completes the proof. \square

We now turn to showing that the GREEDY algorithm is stochastically optimal under our model, that is, we assume that each ball has d destinations chosen uniformly at random and that all balls have equal weight. (The optimality is not preserved if either condition is violated.) It suffices to consider only deterministic algorithms since randomized algorithms can be considered as a distribution over deterministic algorithms.

We say that a vector $\bar{v} = (v_1, v_2, \dots, v_n)$ majorizes a vector \bar{u} , written $\bar{v} \succeq \bar{u}$, if for $1 \leq i \leq n$, we have $\sum_{1 \leq j \leq i} v_{\pi(j)} \geq \sum_{1 \leq j \leq i} u_{\sigma(j)}$, where π and σ are permutations of $1, \dots, n$ such that $v_{\pi(1)} \geq v_{\pi(2)} \geq \cdots \geq v_{\pi(n)}$ and $u_{\sigma(1)} \geq u_{\sigma(2)} \geq \cdots \geq u_{\sigma(n)}$.

LEMMA 3.4. *Let \bar{v} and \bar{u} be two positive integer vectors such that $v_1 \geq v_2 \geq \cdots \geq v_n$ and $u_1 \geq u_2 \geq \cdots \geq u_n$. If $\bar{v} \succeq \bar{u}$ then also $\bar{v} + \bar{e}_i \succeq \bar{u} + \bar{e}_i$, where \bar{e}_i is the i th unit vector, that is $\bar{e}_{i,j} = \delta_{i,j}$.*

Proof. Let $S_j(\bar{x})$ be the sum of the j largest components of the vector \bar{x} . Notice first that for all j

$$(3.10) \quad S_j(\bar{x}) \leq S_j(\bar{x} + \bar{e}_i) \leq S_j(\bar{x}) + 1.$$

By hypothesis, for all j , we have $S_j(\bar{v}) \geq S_j(\bar{u})$. To prove the lemma we show that for all j , we also have $S_j(\bar{v} + \bar{e}_i) \geq S_j(\bar{u} + \bar{e}_i)$. Fix j . By (3.10) if $S_j(\bar{v}) > S_j(\bar{u})$, then $S_j(\bar{v} + \bar{e}_i) \geq S_j(\bar{u} + \bar{e}_i)$. Now assume $S_j(\bar{v}) = S_j(\bar{u})$. There are three cases to consider:

Case 1. $i \leq j$. Then

$$S_j(\bar{v} + \bar{e}_i) = S_j(\bar{v}) + 1 = S_j(\bar{u}) + 1 = S_j(\bar{u} + \bar{e}_i).$$

Case 2. $i > j$ and $u_j > u_i$. Since $u_j \geq u_i + 1$, it follows that $S_j(\bar{u}) = S_j(\bar{u} + \bar{e}_i)$ and therefore

$$S_j(\bar{v} + \bar{e}_i) \geq S_j(\bar{v}) = S_j(\bar{u}) = S_j(\bar{u} + \bar{e}_i).$$

Case 3. $i > j$ and $u_j = u_{j+1} = \cdots = u_i$. Observe first that since $S_{j-1}(\bar{v}) \geq S_{j-1}(\bar{u})$, $S_j(\bar{v}) = S_j(\bar{u})$, and $S_{j+1}(\bar{v}) \geq S_{j+1}(\bar{u})$, we have

$$v_j \leq u_j \quad \text{and} \quad v_{j+1} \geq u_{j+1}.$$

Hence

$$v_j \geq v_{j+1} \geq u_{j+1} = u_j \geq v_j.$$

We conclude that $v_j = u_j = v_{j+1} = u_{j+1}$, and thus $S_{j+1}(\bar{v}) = S_{j+1}(\bar{u})$. Repeating the argument, we obtain that

$$v_j = u_j = v_{j+1} = u_{j+1} = \cdots = v_i = u_i,$$

and therefore

$$S_j(\bar{v} + \bar{e}_i) = S_j(\bar{v}) + 1 = S_j(\bar{u}) + 1 = S_j(\bar{u} + \bar{e}_i). \quad \square$$

Let Ω be the set of all possible n^d choices for each ball and Ω^t be the set of sequences of choices for the first t balls.

THEOREM 3.5. *For any on-line deterministic algorithm A , and $t \geq 0$, there is 1-1 correspondence $f : \Omega^t \rightarrow \Omega^t$ such that for any $\omega_t \in \Omega^t$ the vector of box loads associated with GREEDY acting on ω_t , written*

$$\bar{\lambda}^G(\omega_t) = (\lambda_1^G(\omega_t), \lambda_2^G(\omega_t), \dots, \lambda_n^G(\omega_t)),$$

is majorized by the vector of box loads associated with A acting on $f(\omega_t)$, that is

$$\bar{\lambda}^G(\omega_t) \preceq \bar{\lambda}^A(f(\omega_t)).$$

Proof. To simplify notation we assume $d = 2$. The proof for larger d is analogous. The proof proceeds by induction on t , the length of the sequence. The base case ($t = 0$) is obvious. Assume the theorem valid for t and let f_t be the mapping on Ω^t . Fix a sequence $\omega_t \in \Omega^t$. It suffices to show that we can refine f_t to obtain a 1-1 correspondence for all possible 1-step extensions of ω_t . Without loss of generality, renumber the boxes such that

$$\lambda_1^G(\omega_t) \geq \lambda_2^G(\omega_t) \geq \cdots \geq \lambda_n^G(\omega_t),$$

and let π be a permutation of $1, \dots, n$ such that

$$\lambda_{\pi(1)}^A(f_t(\omega_t)) \geq \lambda_{\pi(2)}^A(f_t(\omega_t)) \geq \cdots \geq \lambda_{\pi(n)}^A(f_t(\omega_t)).$$

Let (i, j) be two choices for the $t + 1$ ball. For every i, j we define

$$f_{t+1}(\omega_t \diamond (i, j)) = f_t(\omega_t) \diamond (\pi(i), \pi(j)),$$

where “ \diamond ” represents extension of sequences.

Clearly f_{t+1} is 1-1. We need to show that

$$\bar{\lambda}^G(\omega_t \diamond (i, j)) \preceq \bar{\lambda}^A(f_t(\omega_t) \diamond (\pi(i), \pi(j))).$$

Notice that when the sequence ω_t is extended by the step (i, j) for any algorithm, exactly one component of the vector $\bar{\lambda}(\omega_t)$ changes, namely either $\lambda_i(\omega_t)$ or $\lambda_j(\omega_t)$ increases by one. Assume that $i \geq j$; then

$$\bar{\lambda}^G(\omega_t \diamond (i, j)) = \bar{\lambda}^G(\omega_t) + \bar{e}_i \preceq \bar{\lambda}^A(f_t(\omega_t)) + \bar{e}_{\pi(i)} \preceq \bar{\lambda}^A(f_t(\omega_t) \diamond (\pi(i), \pi(j))),$$

where the first inequality follows from Lemma 3.4 and the second is due to the fact that

$$\bar{\lambda}^A(f_t(\omega_t)) + \bar{e}_{\pi(i)} \preceq \bar{\lambda}^A(f_t(\omega_t)) + \bar{e}_{\pi(j)}. \quad \square$$

COROLLARY 3.6. *For any fixed k and any t*

$$\Pr(\max_i \lambda_i^A(t) > k) \geq \Pr(\max_i \lambda_i^G(t) > k).$$

We are now ready to discuss the general case of the finite process.

THEOREM 3.7. *The maximum load achieved by the GREEDY algorithm on a random (m, n, d) -problem, with $d \geq 2$ and $m \geq n$, is, with high probability, less than $(1 + o(1)) \ln \ln n / \ln d + O(m/n)$.*

Proof. We start by replaying the proof of Theorem 3.2, taking into account the fact that there are now m balls. So let \mathcal{E}_i be the event that $\nu_{\geq i}(m) \leq \beta_i$, and define $p_i = \beta_i^d/n^d$. Following the proof of Theorem 3.2 we derive that

$$\Pr(\nu_{\geq i+1} \geq k \mid \mathcal{E}_i) \leq \frac{\Pr(B(m, p_i) \geq k)}{\Pr(\mathcal{E}_i)}.$$

Suppose that for some value x we set $\beta_x = n^2/(2em)$ and show that \mathcal{E}_x holds with high probability, that is,

$$(3.11) \quad \Pr\left(\nu_x \geq \frac{n^2}{2em}\right) = o(1).$$

Then

$$\beta_{i+x} = \frac{n}{2^{d^i}} \left(\frac{me}{n}\right)^{(d^i-1)/(d-1)-d^i} \leq \frac{n}{2^{d^i}},$$

and continuing as before, we obtain that

$$\Pr(\mu \geq x + \ln \ln n / \ln d + 2) = o(1).$$

It remains to be shown that x can be taken to be $O(m/n) + o(\ln \ln n / \ln d)$. First assume that $m/n \geq w(n)$ where $w(n)$ is an increasing function of n , but $w(n) = o(\ln \ln n / \ln d)$. Then we claim that we can take $x = \lceil em/n \rceil$.

Consider a placement algorithm, denoted R , that always puts a ball in the box corresponding to the first choice offered. This is entirely equivalent with the case $d = 1$, the classical occupancy problem. The load within a box under this process is a binomial random variable $B(m, 1/n)$, and therefore (via (3.4)), the probability that the load within a box exceeds em/n is bounded by $e^{-m/n}$. Now consider the height of the t th ball, denoted h_t^R . The probability that the box into which the t th ball is placed has load greater than em/n is less than $e^{-m/n}$, and therefore the expected number of balls of height $\geq em/n$ satisfies

$$\mathbf{E}(\mu_{\geq em/n}^R) \leq me^{-m/n}.$$

Hence by Markov's inequality

$$\Pr\left(\mu_{\geq em/n}^R \geq \frac{n^2}{2em}\right) \leq \frac{2em^2}{n^2} e^{-m/n} = o(1),$$

since $m/n \rightarrow \infty$.

We claim that Theorem 3.5 implies

$$(3.12) \quad \Pr(\mu_{\geq k}^G \geq r) \leq \Pr(\mu_{\geq k}^R \geq r).$$

Indeed, suppose that there is an outcome ω_t for which GREEDY has exactly i boxes with load greater than or equal to k . As in the proof of Theorem 3.5, renumber the boxes such that

$$\lambda_1^G(\omega_t) \geq \lambda_2^G(\omega_t) \geq \cdots \geq \lambda_n^G(\omega_t).$$

Let $f_t(\omega_t)$ be the corresponding outcome for algorithm R and let π be a permutation of $1, \dots, n$ such that

$$\lambda_{\pi(1)}^R(f_t(\omega_t)) \geq \lambda_{\pi(2)}^R(f_t(\omega_t)) \geq \cdots \geq \lambda_{\pi(n)}^R(f_t(\omega_t)).$$

Then

$$\mu_{\geq k}^G(\omega_t) = \sum_{1 \leq j \leq i} (\lambda_j^G(\omega_t) - (k-1))$$

and

$$\mu_{\geq k}^R(f_t(\omega_t)) \geq \sum_{1 \leq j \leq i} (\lambda_{\pi(j)}^R(f_t(\omega_t)) - (k-1)).$$

But Theorem 3.5 implies that

$$\sum_{1 \leq j \leq i} \lambda_j^R(f_t(\omega_t)) \geq \sum_{1 \leq j \leq i} \lambda_j^G(\omega_t),$$

and by considering all outcomes we obtain (3.12). Therefore,

$$\Pr\left(\mu_{\geq em/n}^G \geq \frac{n^2}{2em}\right) \leq \Pr\left(\mu_{\geq em/n}^R \geq \frac{n^2}{2em}\right),$$

and since

$$\Pr\left(\nu_{\geq em/n}^G \geq \frac{n^2}{2em}\right) \leq \Pr\left(\mu_{\geq em/n}^G \geq \frac{n^2}{2em}\right),$$

we have that for $x = \lceil em/n \rceil$ (3.11) is satisfied.

To remove the assumption $m/n \geq w(n)$, we can simply imagine that the number of balls is increased to $\max(m, nw(n))$. Then the corresponding value of x becomes $O(\max(m, nw(n))/n) = O(m/n) + o(\ln \ln n / \ln d)$. \square

4. The infinite process. In this section we consider the infinite process. Analogously to Theorem 3.5 it is possible to show that the GREEDY algorithm minimizes the expected maximum load on any box. We analyze its performance below. The main theorem of this section is the following.

THEOREM 4.1. *Assume that the infinite process starts in an arbitrary state. Under GREEDY, with $d \geq 2$, there is a constant c such that for any fixed $T \geq cn^2 \log \log n$,*

$$\Pr(\exists j, \lambda_j(T) \geq \ln \ln n / \ln d + O(1)) = o(1).$$

Thus in the stationary distribution the maximum load is $\ln \ln n / \ln d + O(1)$ with high probability.

Proof. For simplicity of presentation we state and prove the results only for $d = 2$. The proof assumes that at time $T - cn^2 \log \log n$ the process is in an arbitrary state and therefore we can let $T = cn^2 \log \log n$ with no loss of generality.

By the definition of the process, the number of balls of height at least i cannot change by more than 1 in a time step, that is $|\mu_{\geq i}(t+1) - \mu_{\geq i}(t)| \leq 1$. The random variable $\mu_{\geq i}(t)$ can be viewed as a random walk on the integers l , $0 \leq l \leq n$. The proof is based on bounding the maximum values taken by the variables $\mu_{\geq i}(t)$ by studying the underlying process.

We define an integer i^* and a decreasing sequence α_i , for $200 \leq i \leq i^* + 1$ as follows:

$$\begin{aligned} \alpha_{200} &= \frac{n}{200}, \\ \alpha_i &= \frac{100\alpha_{i-1}^2}{n} \quad \text{for } i > 200 \text{ and } \alpha_{i-1} \geq \sqrt{n \log_2 n}, \\ \alpha_{i^*} &= 100 \log_2 n, \quad i^* = \text{the smallest } i \text{ for which} \\ &\quad \alpha_{i-1} < \sqrt{n \log_2 n}, \\ \alpha_{i^*+1} &= 100. \end{aligned}$$

Clearly $i^* \leq \ln \ln n / \ln 2 + O(1)$. For future reference, observe also that for $200 < i \leq i^* + 1$

$$(4.1) \quad \alpha_i \geq \frac{100\alpha_{i-1}^2}{n}.$$

We also define an increasing sequence of times: $t_{200} = 0$ and $t_i = t_{i-1} + n^2$ for $i > 200$. Thus $t_{i^*+1} = O(n^2 \log \log n) = O(T)$.

Let $\{\mu_{\geq i}[t^-, t^+] \leq \alpha\}$ denote the event that $\mu_{\geq i}(t) \leq \alpha$ for all t , such that $t^- < t \leq t^+$, and similarly, let $\{\mu_{\geq i}[t^-, t^+] > \alpha\}$ denote the event that $\mu_{\geq i}(t) > \alpha$ for all t , such that $t^- < t \leq t^+$. We define the events C_i as follows:

$$\begin{aligned} C_{200} &= \{\nu_{\geq 200}[t_{200}, T] \leq 2\alpha_{200}\} \equiv \{\nu_{\geq 200}[0, T] \leq n/100\}; \\ C_i &= \{\mu_{\geq i}[t_i, T] \leq 2\alpha_i\} \quad \text{for } i > 200. \end{aligned}$$

Note that C_{200} always holds, and for $i > 200$, the event $C - i$ implies that $\nu_{\geq i}[t_i, T] \leq 2\alpha_i$.

We shall prove inductively that for all $i = 200, \dots, i^* + 1$

$$(4.2) \quad \Pr(-C_i) \leq \frac{2i}{n^2}.$$

This implies that the event $\{\mu_{\geq i^*+1}[t_{i^*+1}, T] \leq 200\}$ occurs with probability $1 - o(1)$, and therefore with high probability, for every j , $\lambda_j(T) \leq i^* + 201 = \ln \ln n / \ln 2 + O(1)$, which completes the proof of the main part of the theorem.

Finally, we show that in the stationary distribution

$$\Pr(\forall j, \lambda_j \leq \log \log n + O(1)) = 1 - o(1).$$

Indeed, let S be the set of states such that for all j , $\lambda_j(t) \leq \log \log n + O(1)$. Let $s(t)$ be the state of the chain at time t . Then the previous observation implies that

$$\Pr(s(t+T) \notin S \mid s(t)) = o(1).$$

Let π be the stationary distribution; then

$$\sum_{i \notin S} \pi_i = \sum_j \Pr(s(t+T) \notin S \mid s(t) = j) \cdot \pi_j = \sum_j \pi_j o(1) = o(1),$$

which completes the proof of the theorem assuming (4.2). To prove it, we show that conditioned on C_{i-1} :

- a. With high probability $\mu_{\geq i}(t)$ becomes less than α_i before time t_i . (This is shown in Lemma 4.2.)
- b. If $\mu_{\geq i}(t)$ becomes less than α_i at any time before T , from then until T , with high probability, it does not become larger than $2\alpha_i$. (This is shown in Lemma 4.3.)

The two facts above imply that if C_{i-1} holds, then with high probability $\mu_{\geq i}[t_i, T] \leq 2\alpha_i$, that is, C_i holds as well.

Base case. The base case is straightforward since $\Pr(\neg C_{200}) = \Pr(\neg\{\nu_{\geq 200}[0, T] \leq n/100\}) = 0$.

Induction. Suppose that

$$(4.3) \quad \Pr(\neg C_{i-1}) \leq \frac{2(i-1)}{n^2},$$

where $200 < i \leq i^* + 1$.

Let $s(t)$ be the state at time t . It is easy to verify the following bounds on the underlying transition probabilities. For any t ,

$$(4.4) \quad \Pr(\mu_{\geq i}(t+1) > \mu_{\geq i}(t) \mid s(t)) \leq \left(\frac{\nu_{\geq(i-1)}(t)}{n}\right)^2 \leq \left(\frac{\mu_{\geq(i-1)}(t)}{n}\right)^2$$

and

$$(4.5) \quad \Pr(\mu_{\geq i}(t+1) < \mu_{\geq i}(t) \mid s(t)) \geq \frac{\mu_{\geq i}(t)}{n} \left(1 - \left(\frac{\nu_{\geq(i-1)}(t)}{n}\right)^2\right) \geq \frac{\mu_{\geq i}(t)}{2n}.$$

From (4.4) and (4.5) we obtain that the transition probabilities satisfy

$$\Pr(\mu_{\geq i}(t+1) > \mu_{\geq i}(t) \mid \mu_{\geq i-1}(t) \leq 2\alpha_{i-1}) \leq \left(\frac{2\alpha_{i-1}}{n}\right)^2 \stackrel{def}{=} q_i^+$$

and

$$\Pr(\mu_{\geq i}(t+1) < \mu_{\geq i}(t) \mid \mu_{\geq i}(t) \geq \alpha_i) \geq \frac{\alpha_i}{2n} \stackrel{def}{=} q_i^-.$$

Thus in view of (4.1)

$$q_i^+ \leq \frac{\alpha_i}{25n}.$$

We define two new binary random variables for $0 < t \leq T$ as follows:

$$X_t = 1 \text{ iff } \mu_{\geq i}(t) > \mu_{\geq i}(t-1) \text{ and } \mu_{\geq i-1}(t-1) \leq 2\alpha_{i-1},$$

and

$$Y_t = 1 \text{ iff } \mu_{\geq i}(t) < \mu_{\geq i}(t-1) \text{ or } \mu_{\geq i}(t-1) \leq \alpha_i.$$

Clearly

$$(4.6) \quad \Pr(X_t = 1) \leq q_i^+ \quad \text{and} \quad \Pr(Y_t = 1) \geq q_i^-.$$

We also define F_i to be the event

$$F_i \stackrel{def}{=} \{\exists t^* \in [t_{i-1}, t_i] \text{ s.t. } \mu_{\geq i}(t^*) \leq \alpha_i\};$$

thus $\neg F_i$ is the event

$$\neg F_i = \{\mu_{\geq i}[t_{i-1}, t_i] > \alpha_i\}.$$

Two lemmas are necessary in order to conclude that $\Pr(\neg C_i) \leq 2i/n^2$.

LEMMA 4.2. *Under the inductive hypothesis*

$$\Pr(\neg F_i \mid C_{i-1}) \leq \frac{1}{n^2}.$$

Proof. Notice that conditioned on C_{i-1} , the sum $\sum_{t \in [t_{i-1}, t_i]} X_t$ is the number of times $\mu_{\geq i}(t)$ increased in the interval $[t_{i-1}, t_i]$; similarly, if within this interval $\mu_{\geq i}$ did not become less than α_i , then $\sum_{t \in [t_{i-1}, t_i]} Y_t$ equals the number of times $\mu_{\geq i}(t)$ decreased in this interval. We conclude that

$$\begin{aligned} \Pr(\neg F_i \mid C_{i-1}) &\leq \Pr\left(\sum_{t \in [t_{i-1}, t_i]} Y_t - \sum_{t \in [t_{i-1}, t_i]} X_t \leq n \mid C_{i-1}\right) \\ &\leq \frac{1}{\Pr(C_{i-1})} \Pr\left(\sum_{t \in [t_{i-1}, t_i]} Y_t - \sum_{t \in [t_{i-1}, t_i]} X_t \leq n\right). \end{aligned}$$

In view of (4.6) and Lemma 3.1, Chernoff-type bounds imply that for every $i \leq i^* + 1$

$$\Pr\left(\sum_{t \in [t_{i-1}, t_i]} X_t > 2n^2 q_i^+\right) \leq \Pr(B(n^2, q_i^+) \geq 2n^2 q_i^+) \leq e^{-\Omega(n^2 q_i^+)} = o(1/n^c)$$

and

$$\Pr\left(\sum_{t \in [t_{i-1}, t_i]} Y_t < \frac{1}{2}n^2 q_i^-\right) \leq \Pr(B(n^2, q_i^-) \leq \frac{1}{2}n^2 q_i^-) \leq e^{-\Omega(n^2 q_i^-)} = o(1/n^c)$$

for any constant c . On the other hand, in view of (4.1),

$$\frac{1}{2}n^2 q_i^- - 2n^2 q_i^+ \geq \frac{1}{4}n\alpha_i - \frac{2}{25}n\alpha_i \geq \frac{n\alpha_i}{10} \geq n,$$

and therefore we conclude that

$$\Pr(\neg F_i \mid C_{i-1}) \leq \frac{1}{n^c \Pr(C_{i-1})}$$

for any constant c . Taking $c = 3$ and using the inductive hypothesis on C_{i-1} (4.3) completes the proof. \square

LEMMA 4.3. *Under the inductive hypothesis*

$$\Pr(\neg C_i \mid C_{i-1}, F_i) \leq \frac{1}{n^2}.$$

Proof. Since $\Pr(A \mid B \wedge C) \leq \Pr(A \wedge B \mid C)$ we get that

$$\begin{aligned}
& \Pr(-C_i \mid C_{i-1}, F_i) \\
& \leq \Pr(-C_i \wedge F_i \mid C_{i-1}) \\
& \leq \Pr(\exists t_1, t_2 \in [t_{i-1}, T] \\
& \quad \text{s.t. } \mu_{\geq i}(t_1) = \alpha_i, \mu_{\geq i}(t_2) = 2\alpha_i, \mu_{\geq i}[t_1, t_2] \geq \alpha_i \mid C_{i-1}) \\
& \leq \sum_{t_{i-1} \leq t_1 < t_2 \leq T} \Pr(\mu_{\geq i}(t_1) = \alpha_i, \mu_{\geq i}(t_2) = 2\alpha_i, \mu_{\geq i}[t_1, t_2] \geq \alpha_i \mid C_{i-1}) \\
& \leq \sum_{t_{i-1} \leq t_1 < t_2 \leq T} \Pr\left(\sum_{t \in [t_1, t_2]} X_t - \sum_{t \in [t_1, t_2]} Y_t \geq \alpha_i \mid C_{i-1}\right) \\
& \leq \frac{1}{\Pr(C_{i-1})} \sum_{t_{i-1} \leq t_1 < t_2 \leq T} \Pr\left(\sum_{t \in [t_1, t_2]} X_t - \sum_{t \in [t_1, t_2]} Y_t \geq \alpha_i\right).
\end{aligned}$$

Fix t_1 and t_2 and let $\Delta = t_2 - t_1$. We now consider four cases.

A. $\Delta \leq n$ and $i \leq i^*$:

$$\Pr\left(\sum_{t \in [t_1, t_2]} X_t \geq \alpha_i\right) \leq \binom{\Delta}{\alpha_i} (q^+)^{\alpha_i} \leq \left(\frac{e\Delta}{\alpha_i} \cdot \frac{\alpha_i}{25n}\right)^{\alpha_i} \leq n^{-100}.$$

B. $\Delta \leq n \log n$ and $i = i^* + 1$:

$$\begin{aligned}
\Pr\left(\sum_{t \in [t_1, t_2]} X_t \geq \alpha_{i^*+1}\right) & \leq \binom{\Delta}{\alpha_{i^*+1}} (q^+)^{\alpha_{i^*+1}} \leq \left(\frac{e\Delta}{\alpha_{i^*+1}} \cdot \frac{4\alpha_{i^*}^2}{n^2}\right)^{\alpha_{i^*+1}} \\
& \leq \left(\frac{en \log n}{100} \cdot \frac{4 \cdot 100^2 \log^2 n}{n^2}\right)^{100} \leq n^{-100}.
\end{aligned}$$

C. $\Delta \geq n$ and $i \leq i^*$: Again using large deviation bounds and the fact that $\alpha_i \Delta \geq 100 \log n$ we obtain that

$$\Pr\left(\sum_{t \in [t_1, t_2]} Y_t \leq \frac{1}{2} q^- \Delta\right) \leq e^{-q^- \Delta/8} = e^{-\alpha_i \Delta / (16n)} \leq n^{-6.1}$$

and that

$$\Pr\left(\sum_{t \in [t_1, t_2]} X_t \geq \frac{1}{2} q^- \Delta\right) \leq \left(\frac{2eq^+ \Delta}{q^- \Delta}\right)^{q^- \Delta/2} \leq \left(\frac{4e}{25}\right)^{\alpha_i \Delta / (4n)} \leq n^{-25}.$$

D. $\Delta \geq n \log n$ and $i = i^* + 1$: We use the same proof as case **C** using the fact that $\alpha_{i^*+1} \Delta \geq 100 \log n$.

Thus in all four cases,

$$\Pr\left(\sum_{t \in [t_1, t_2]} X_t - \sum_{t \in [t_1, t_2]} Y_t \geq \alpha_i\right) \leq \frac{1}{n^{6.1}},$$

therefore, under the induction hypothesis,

$$\frac{1}{\Pr(C_{i-1})} \sum_{t_{i-1} \leq t_1 < t_2 \leq T} \Pr\left(\sum_{t \in [t_1, t_2]} X_t - \sum_{t \in [t_1, t_2]} Y_t \geq \alpha_i\right) \leq \frac{2T^2}{n^{6.1}} \leq \frac{1}{n^2}. \quad \square$$

Returning to the proof of (4.2), by using the induction hypothesis, Lemmas 4.2 and 4.3, and the law of total probability, we can complete the induction as follows:

$$\begin{aligned} \Pr(\neg C_i) &= \Pr(\neg C_i \mid C_{i-1}) \cdot \Pr(C_{i-1}) \\ &\quad + \Pr(\neg C_i \mid \neg C_{i-1}) \cdot \Pr(\neg C_{i-1}) \quad \text{Now apply IH} \\ &\leq \Pr(\neg C_i \mid C_{i-1}) + 2(i-1)/n^2 \\ &= \Pr(\neg C_i \mid C_{i-1}, F_i) \cdot \Pr(F_i \mid C_{i-1}) \quad \text{Now apply Lemma 4.3} \\ &\quad + \Pr(\neg C_i \mid C_{i-1}, \neg F_i) \cdot \Pr(\neg F_i \mid C_{i-1}) \quad \text{Now apply Lemma 4.2} \\ &\quad + 2(i-1)/n^2 \\ &\leq 1/n^2 + 1/n^2 + 2(i-1)/n^2 = 2i/n^2. \quad \square \end{aligned}$$

5. Hashing. We define a simple hashing algorithm, called *2-way chaining*, by analogy with the popular direct chaining method. We use two random hash functions. For each key, the two hash functions define two indices in a table. Each table location contains a pointer to a linked list. When a new key arrives, we compare the current length of the two lists associated to the key, and the key is inserted at the end of the shortest list. (The direct chaining method corresponds to having only one associated random index.)

For searching, the two hash values are computed, and two linked lists are searched in alternate order. (That is, after checking the i th element of the first list, we check the i th element of the second list, then element $i+1$ of the first list, and so on.) When the shorter list is exhausted, we continue searching the longer list until it is exhausted as well. (In fact, if no deletions are allowed, we can stop after checking only one more element in the longer list. For the analysis below, this is immaterial.)

Assume that n keys are sequentially inserted by this process to a table of size n . Theorem 1.1 analysis implies that with high probability the maximum access time, which is bounded by twice the length of the longest list, is $2 \ln n \ln n / \ln 2 + O(1)$, versus the $\Theta(\log n / \log \log n)$ time when one random hash function is used. More generally, if m keys are stored in the table with d hash functions, then the maximum access time under this scheme is $2(1 + o(1)) \ln \ln n / \ln d + \Theta(m/n)$.

Next we show that the average access time of 2-way chaining is no more than twice the average access time of the standard direct chaining method. As customary, we discuss the average access time separately for successful searches and unsuccessful searches. The latter, denoted C'_G , is bounded by twice the expected cost of checking a list chosen uniformly at random. Therefore

$$C'_G(m, n) \leq 2 + \frac{2m}{n}.$$

For successful searches, the cost C_G , is given by

$$C_G(m, n) \leq \frac{2}{m} \sum_{1 \leq i \leq m} h_i = \frac{2}{m} \sum_{1 \leq j \leq n} \binom{\lambda_j + 1}{2},$$

where all the notations are as in section 2. Since we know that ν_k eventually decreases doubly exponentially, we can bound C_G via the inequality

$$C_G(m, n) \leq \frac{2}{m} \sum_{k>0} k\nu_{\geq k}.$$

However, we can achieve better bounds, using the majorization Theorem 3.5. We start from the following.

LEMMA 5.1. *Let $\bar{v} = (v_1, v_2, \dots, v_n)$ and $\bar{u} = (u_1, v_2, \dots, u_n)$ be two positive integer vectors. If $\bar{v} \succeq \bar{u}$, then*

$$\sum_{1 \leq i \leq n} v_i^2 \geq \sum_{1 \leq i \leq n} u_i^2.$$

This lemma is a special case of a well-known theorem from majorization (see, e.g., [22]), but for completeness we present a proof.

Proof. Let \bar{x} be a n -vector and let (\bar{x}, \bar{u}) denote the inner product of \bar{x} and \bar{u} . Consider the linear program

$$\text{Maximize } (\bar{x}, \bar{u}) \text{ subject to } \bar{x} \leq \bar{v} \text{ and } \bar{x} \geq 0.$$

It is easy to check that $\bar{x} = \bar{u}$ is a feasible point and that the optimal solution is $\bar{x} = \bar{v}$. Hence $(\bar{u}, \bar{u}) \leq (\bar{v}, \bar{u})$. Now consider the same program with the objective function (\bar{x}, \bar{v}) . Then again $\bar{x} = \bar{u}$ is a feasible point and the optimal solution is $\bar{x} = \bar{v}$. Hence $(\bar{u}, \bar{u}) \leq (\bar{u}, \bar{v}) \leq (\bar{v}, \bar{v})$. \square

Consider now the standard direct chaining method. In our terminology it corresponds to the random placement algorithm R and it therefore majorizes G . It is well known that the cost for successful search for direct chaining is [19, Ex. 6.4.34]

$$C_R(m, n) = \frac{1}{m} \sum_{1 \leq j \leq n} \binom{\lambda_j^R + 1}{2} = 1 + \frac{m-1}{2n}.$$

Applying the lemma above we obtain that the cost of successful search in 2-way chaining satisfies

$$C_G(m, n) \leq 2 + \frac{m-1}{n}.$$

6. Competitive on-line load balancing.

6.1. Preliminaries. The on-line load balancing problem is defined as follows. Let M be a set of servers (or machines) that is supposed to run a set of tasks that arrive and depart in time. Each task j has associated with it a weight, or load, $w(j) \geq 0$, an arrival time $\tau(j)$, and a set $M(j) \subset M$ of servers capable of running it. We distinguish among two variants of this problem: the case of *permanent tasks*, tasks that arrive but never depart, and the case of *temporary tasks*, tasks that depart the system at a time unknown in advance.

As soon as each task arrives, it must be assigned to exactly one of the servers capable of running it, and once assigned, it can not be transferred to a different server. The assigned server starts to run the task immediately, and continues to run it until the task departs.

When task j arrives, an *assignment algorithm* must select a server $i \in M(j)$, and assign task j to it.

The *load* on server i at time t , denoted $L_i^A(t)$, is the sum of the weights of all the tasks running on server i at time t under assignment algorithm A .

Let σ be a sequence of task arrivals and departures, and let $|\sigma|$ be the time of the last arrival. Then the cost, $C_A(\sigma)$, of an assignment algorithm A on the sequence σ is defined as

$$C_A(\sigma) = \max_{0 \leq t \leq |\sigma|; i \in M} L_i^A(t).$$

An *on-line assignment algorithm* must assign an arriving task j at time $\tau(j)$ to a server in $M(j)$ knowing only $w(j)$, $M(j)$, and the past and current state of the servers—the decision is made without any knowledge about future arrivals or departures. The *optimal off-line assignment algorithm*, denoted OPT , assigns arriving tasks knowing the entire sequence of task arrivals and departures and does so in a way that minimizes cost.

The worst-case behavior of an on-line algorithm A is characterized by the *competitive ratio*, defined as the supremum over all sequences σ of the ratio $C_A(\sigma)/C_{\text{OPT}}(\sigma)$.

To characterize the average behavior of A , let $C_A(\mathcal{P})$ (resp., $C_{\text{OPT}}(\mathcal{P})$) be the expected cost of algorithm A (resp., OPT) on sequences σ generated by the distribution \mathcal{P} . The competitive ratio of an on-line algorithm, A *against distribution* \mathcal{P} , is defined as the ratio $C_A(\mathcal{P})/C_{\text{OPT}}(\mathcal{P})$.

Finally, the GREEDY algorithm is formally defined as follows.

ALGORITHM GREEDY. Upon arrival of a task j , assign it to the server in $M(j)$ with the current minimum load (ties are broken arbitrarily).

6.1.1. Permanent tasks. For permanent tasks, Azar, Naor, and Rom [9] have shown that the competitive ratio of the GREEDY algorithm is $\Theta(\log n)$ and that no algorithm can do better.

To bring this problem into our framework, we present our results for the case where for each task j the set of servers that can run it, $M(j)$, consists of $d \geq 2$ servers chosen uniformly at random (with replacement), the number of requests $|\sigma|$ equals n , and all weights are equal. Let \mathcal{P}_d be the associated probability distribution on request sequences.

LEMMA 6.1. *With probability $1 - O(1/n)$, $C_{\text{OPT}}(\mathcal{P}_d) = O(1)$.*

Proof. We show that with high probability there is an assignment with cost 10 for the case $d = 2$. A fortiori the result is true for $d > 2$.

The problem can be reduced to showing that in a random n -by- n bipartite graph (U, V, E) where each node in U has two random edges to V , there is an assignment of value 10. Arbitrarily break U into 10 pieces of size $n/10$ each. We show that each of these pieces contains a perfect matching. By Hall's theorem, the probability that there is no such assignment is bounded by the probability that there is a set of size k in one of the pieces of U whose neighborhood has size less than k . Ipso facto, this probability is at most

$$10 \sum_{k \leq n/10} \binom{n}{k-1} \binom{n/10}{k} \left(\left(\frac{k-1}{n} \right)^2 \right)^k.$$

Using standard approximations to the binomial coefficients, this is at most

$$10 \sum_{k \leq n/10} \left(\frac{en}{k-1} \right)^{k-1} \left(\frac{en}{10k} \right)^k \left(\left(\frac{k-1}{n} \right)^2 \right)^k.$$

Finally, rewriting and simplifying yield

$$\frac{10}{n} \sum_{k \leq n/10} \frac{k-1}{e} \left(\frac{e^2 n^2 (k-1)^2}{10k(k-1)n^2} \right)^k = \frac{10}{n} \sum_{k \leq n/10} \frac{k-1}{e} \left(\frac{e^2}{10} \right)^k = O\left(\frac{1}{n}\right). \quad \square$$

(A more delicate analysis [15] shows that the maximum load achieved by the off-line case is 2 with high probability, for $d \geq 2$ and $m \leq 1.6n$.)

LEMMA 6.2. *With high probability, $C_{\text{GREEDY}}(\mathcal{P}_d) = O(\log \log n / \log d)$*

Proof. The proof follows immediately from Theorem 3.2. \square

Thus, we obtain the following theorem.

THEOREM 6.3. *The competitive ratio of the GREEDY algorithm against the distribution \mathcal{P}_d is $O(\log \log n / \log d)$, and no algorithm can do better.*

Proof. The proof follows from Lemmas 6.1 and 6.2 and Corollary 3.6. \square

6.1.2. Temporary tasks. For temporary tasks, the results of Azar, Broder, and Karlin [6] and Azar et al. [8] showed that there is an algorithm with competitive ratio $\Theta(\sqrt{n})$ and that no algorithm can do better.

It is difficult to construct a natural distribution of task arrivals and departures. As an approximation, we consider the following stochastic process \mathcal{S} : First, n tasks arrive; for each task, the set of servers that can run it consists of $d \geq 2$ servers chosen uniformly at random (with replacement). Then the following repeats forever: a random task among those present departs and a random task arrives, which again may be served by any one of d random servers. Clearly, in such an infinite sequence, eventually there will be n tasks which can only be served by one server, and so for trivial reasons the competitive ratio for long enough sequences is 1. However, we can state a competitiveness result in the following way:

THEOREM 6.4. *Let $L_A[t]$ be the maximum load on any server at time t , for tasks arriving according to the stochastic process \mathcal{S} , and assigned using algorithm A , that is, $L_A[t] = \max_{i \in M} L_i^A(t)$. Then for any fixed $t > 0$, with high probability,*

$$\frac{L_{\text{GREEDY}}[t]}{L_{\text{OPT}}[t]} = O(\log \log n).$$

Proof. The proof follows from Lemma 6.1 and Theorem 4.1. \square

7. Experimental results. The bound proven in Theorem 3.2 for the $O(1)$ term in the formula for the upper bound on the maximum load is rather weak (≈ 8), so it might be the case that for practical values of n , the constant term dominates the $\ln \ln n / \ln d$ term. However, experiments seem to indicate that this is not the case, and in fact even for small values of n , the maximum load achieved with $d = 2$ is substantially smaller than the maximum load achieved with $d = 1$. For the values we considered, $256 \leq n \leq 16777216$, increasing d beyond 2 has only limited further effect. For each value of n and d we ran 100 experiments. The results are summarized in Table 7.1.

TABLE 7.1
Experimental maximum load ($m = n$).

n	$d = 1$	$d = 2$	$d = 3$	$d = 4$
2^8	3 1% 4 40% 5 41% 6 15% 7 3%	2 10% 3 90%	2 84% 3 16%	2 99% 3 1%
2^{12}	5 12% 6 66% 7 17% 8 4% 9 1%	3 99% 4 1%	2 12% 3 88%	2 91% 3 9%
2^{16}	7 48% 8 43% 9 9%	3 64% 4 36%	3 100%	2 23% 3 77%
2^{20}	8 28% 9 61% 10 10% 13 1%	4 100%	3 100%	3 100%
2^{24}	9 12% 10 73% 11 13% 12 2%	4 100%	3 100%	3 100%

Acknowledgments. We wish to thank Martin Dyer, Alan Frieze, and Greg Nelson for several very useful discussions.

REFERENCES

- [1] M. ADLER, S. CHAKRABARTI, M. MITZENMACHER, AND L. RASMUSSEN, *Parallel randomized load balancing*, in Proc. 27th Annual ACM Symposium on Theory of Computing, 1995, pp. 238–247.
- [2] N. ALON AND J. H. SPENCER, *The Probabilistic Method*, John Wiley & Sons, New York, 1992.
- [3] B. AWERBUCH, Y. AZAR, AND S. PLOTKIN, *Throughput-competitive on-line routing*, in 34th Annual Symposium on Foundations of Computer Science, Palo Alto, CA, 1993, IEEE, Piscataway, NJ, pp. 32–40.
- [4] B. AWERBUCH, Y. AZAR, S. PLOTKIN, AND O. WAARTS, *Competitive routing of virtual circuits with unknown duration*, in Proc. ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 1994, pp. 321–327.
- [5] Y. AZAR, J. ASPNES, A. FIAT, S. PLOTKIN, AND O. WAARTS, *On-line routing of virtual circuits with applications to load balancing and machine scheduling*, J. Assoc. Comput. Mach., 44 (1997), pp. 486–504.
- [6] Y. AZAR, A. Z. BRODER, AND A. R. KARLIN, *On-line load balancing*, Theoret. Comput. Sci., 130 (1994), pp. 73–84.
- [7] Y. AZAR, A. Z. BRODER, A. R. KARLIN, AND E. UPFAL, *Balanced allocations*, in Proc. 26th Annual ACM Symposium on the Theory of Computing, Montreal, Quebec, Canada, 1994, pp. 593–602.
- [8] Y. AZAR, B. KALYANASUNDARAM, S. PLOTKIN, K. PRUHS, AND O. WAARTS, *On-line load balancing of temporary tasks*, J. Algorithms, 22 (1997), pp. 93–110.
- [9] Y. AZAR, J. NAOR, AND R. ROM, *The competitiveness of on-line assignments*, J. Algorithms, 18 (1995), pp. 221–237.
- [10] A. Z. BRODER, A. M. FRIEZE, C. LUND, S. PHILLIPS, AND N. REINGOLD, *Balanced allocations*

- for tree-like inputs, Inform. Process. Lett., 55 (1995), pp. 329–332.
- [11] A. Z. BRODER AND A. R. KARLIN, *Multilevel adaptive hashing*, in Proc. 1st Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 1990, pp. 43–53.
 - [12] A. CZUMAJ AND V. STEMANN, *Randomized allocation processes*, in Proc. 38th Annual IEEE Symposium on Foundations of Computer Science, 1997, Miami Beach, FL, 1997, pp. 194–203.
 - [13] M. DIETZFELBINGER, A. R. KARLIN, K. MEHLHORN, F. MEYER AUF DER HEIDE, H. ROHNERT, AND R. E. TARJAN, *Dynamic perfect hashing: Upper and lower bounds*, SIAM J. Comput., 23 (1994), pp. 738–761.
 - [14] M. L. FREDMAN, J. KOMLÓS, AND E. SZEMERÉDI, *Storing a sparse table with $O(1)$ worst case access time*, J. Assoc. Comput. Mach., 31 (1984), pp. 538–544.
 - [15] A. M. FRIEZE, *Personal communication*. 1994.
 - [16] G. H. GONNET, *Expected length of the longest probe sequence in hash code searching*, J. Assoc. Comput. Mach., 28 (1981), pp. 289–304.
 - [17] N. L. JOHNSON AND S. KOTZ, *Urn Models and Their Application*, John Wiley & Sons, New York, 1977.
 - [18] R. M. KARP, M. LUBY, AND F. MEYER AUF DER HEIDE, *Efficient PRAM simulation on a distributed memory machine*, in Proc. 24th Annual ACM Symposium on the Theory of Computing, Victoria, British Columbia, Canada, 1992, pp. 318–326.
 - [19] D. E. KNUTH, *The Art of Computer Programming, Vol. III: Sorting and Searching*, Addison-Wesley, Reading, MA, 1973.
 - [20] V. F. KOLCHIN, B. A. SEVASTYANOV, AND V. P. CHISTYAKOV, *Random Allocations*, John Wiley & Sons, New York, 1978.
 - [21] P. D. MACKENZIE, C. G. PLAXTON, AND R. RAJARAMAN, *On contention resolution protocols and associated probabilistic phenomena*, in Proc. 26th Annual ACM Symposium on the Theory of Computing, Montreal, Quebec, Canada, 1994, pp. 153–162.
 - [22] A. W. MARSHALL AND I. OLKIN, *Inequalities: Theory of Majorization and Its Applications*, Academic Press, New York, 1979.
 - [23] M. MITZENMACHER, *Load balancing and density dependent jump Markov processes*, in Proc. 37th Annual Symposium on Foundations of Computer Science, Burlington, VT, 1996, pp. 213–222.
 - [24] M. MITZENMACHER, *The Power of Two Choices in Randomized Load Balancing*, Ph.D. thesis, University of California, Berkeley, 1996.
 - [25] V. STEMANN, *Parallel balanced allocations*, in Proc. 8th Annual ACM Symposium on Parallel Algorithms and Architectures, Padua, Italy, 1996, pp. 261–269.