Balanced Allocations: The Heavily Loaded Case

Petra Berenbrink* Artur Czumaj† Angelika Steger‡ Berthold Vöcking§

ABSTRACT

We investigate load balancing processes based on the multiple-choice paradigm. In these randomized processes m balls are inserted into n bins. In the classical single-choice variant each ball is placed simply into a randomly selected bin. In a multiple-choice process each ball can be placed into one out of $d \geq 2$ randomly selected bins. It is well known that having more than one choice for each ball can improve the load balance significantly. In contrast to previous work on multiple-choice processes, we investigate the heavily loaded case, that is, we assume $m \gg n$ rather than $m \approx n$. The best previously known results for the multiple-choice processes in the heavily loaded case were obtained by majorization from the single-choice process. This yields an upper bound of $m/n + O(\sqrt{m \ln n \,/\, n})$. We show, however, that the multiple-choice processes are fundamentally different from the single-choice variant in that they have "short memory". The great consequence of this property is that the deviation of the multiple-choice processes from the optimal allocation (i.e., at most $\lceil m/n \rceil$ balls in every bin) does not increase with the number of balls as in case of the single-choice process.

In particular, we investigate the allocation obtained by two different multiple-choice allocation schemes, the original greedy scheme and the recently presented always-go-left scheme. We show that these schemes result in a maximum load of only $m/n + O(\ln \ln n)$. We point out that our detailed bounds are tight up to additive constants.

Furthermore, we investigate the two multiple-choice algorithms in a comparative study. We present a majorization result showing that the always-go-left scheme obtains a better load balancing than the greedy scheme for any choice of n, m, and d.

*Department of Mathematics and Computer Science, University of Paderborn, Germany. pebe@upb.de. Supported by the DFG Sonderforschungsbereich 376 "Massive Parallelität: Algorithmen, Entwurfsmethoden, Anwendungen".

†Department of Computer and Information Science, New Jersey Institute of Technology. czumaj@cis.njit.edu. Work partly done while the author was with Heinz Nixdorf Institute and Department of Mathematics and Computer Science at the University of Paderborn, Germany. Research partly supported by SFB-DFG 376.

‡Department of Computer Science, Technische Universität München, Germany. steger@informatik.tu-muenchen.de. Supported in parts by the DFG Sonderforschungsbereich 342 "Werkzeuge und Methoden für die Nutzung paralleler Rechnerarchitekturen".

§University of Massachusetts, Amherst. This research was conducted in part while he was staying at the International Computer Science Institute, Berkeley, USA, supported by a stipend of the "Gemeinsames Hochschulsonderprogramm II von Bund und Ländern" through the DAAD.

STOC 2000 Portland Oregon USA

1. INTRODUCTION

The study of *balls-and-bins games* or *occupancy problems* has a very long history. These very common models were used to derive several results in the area of probability theory with many applications to computer science, e.g, hashing or randomized rounding. In particular, balls-and-bins games can be used in order to translate realistic problems into mathematical ones in a natural way. Examples are load balancing and resource allocation in parallel and distributed systems. In general, the goal of a balls-and-bins algorithm is to allocate a set of independent objects (tasks, jobs, balls) to a set of resources (servers, bins, urns) so that the load is distributed among the bins as evenly as possible.

In the classical *single-choice game*, each ball is placed into a bin chosen independently and uniformly at random (i.u.r.). For the case of n bins and $m \geq n \ln n$ balls it is well known that there exists a bin receiving $m/n + \Theta(\sqrt{m \ln n \,/\, n})$ balls (e.g. see [9]). This result holds not only on expectation but with high probability[1]. Let the *max height above average* denote the difference between the number of balls in the fullest bin and the average number of balls per bin. Then the max height above average of the single choice algorithm is $\Theta(\sqrt{m \ln n \,/\, n})$. In other words, the deviation between the randomized single-choice allocation and the optimal allocation increases with the number of balls.

We investigate randomized multiple-choice allocation schemes. The idea of *multiple-choice algorithms* is to reduce the maximum load by choosing a small subset of the bins for each ball at random and placing the ball into one of these bins. Usually, the ball is placed simply into a bin with a minimum number of balls among the d alternatives. It is well known that having more than one choice for each ball can improve the load balancing significantly. Previous analysis, however, are only able to deal with the *lightly loaded case*, i.e., $m = O(n)$. We present the first tight analysis for the *heavily loaded case*, i.e., $m = \omega(n)$. In particular, we investigate two different kinds of well known multiple-choice algorithms, the greedy scheme and the always-go-left scheme.

- Algorithm Greedy$[d]$ chooses $d \geq 2$ locations for each ball i.u.r. from the set of bins. This process has been introduced

[1]We say an event A to occur with high probability (w.h.p.) if $\mathbf{Pr}[A] \geq 1 - n^{-\alpha}$ for an arbitrarily chosen constant $\alpha \geq 1$.

by Azar et al. in [1]. It is assumed that the m balls are inserted one by one, and each ball is placed into the least loaded among its d locations. (If several locations have the same minimum load, then the ball is placed into an arbitrary one among them.) Azar et al. show that the max height above average of Left[d] is only $(\ln\ln n)/\ln d + \Theta(m/n)$, w.h.p.

- Algorithm Left[d] has been introduced and analyzed by Vöcking [11]. This algorithm partitions the set of bins into d groups of equal size. These groups are ordered from left to right. For each ball, we choose the ith location for each ball from the ith group i.u.r. The ball is placed in one of the least loaded bins among these locations. If there are several locations having the same minimum load, the ball is always placed into the leftmost group containing one of these locations. Surprisingly, the use of this unfair tie breaking mechanism leads to a better load balancing than a fair mechanism that solves ties at random. In particular, the max height above average produced by Left[d] is only $\ln\ln n/d^* + \Theta(m/n)$ with $d^* \approx d\ln 2$.

In the lightly loaded case, the bounds above are tight up to additive constants. In the heavily loaded case, however, these bounds are even not as good as the bounds known for the classical single-choice process. In fact, the best known bound for the multiple-choice algorithms in the heavily loaded case are obtained using majorization from the single-choice process showing only that the multiple-choice algorithms does not behave worse than the single-choice process.

Unfortunately, the known methods for analyzing the multiple-choice algorithms do not allow to obtain better results for the heavily loaded case. Both the techniques used in [1] ("layered induction") and [11] ("witness trees") inherently assume a load of $2m/n$ already in their base case. Alternative proof techniques using differential equations as suggested in [7; 8; 10; 12] fail for the heavily loaded case, too, because the concentration results obtained by Kurtz's theorem hold only for a limited number of balls. Therefore, the analysis of the heavily loaded case requires new ideas. Before we proceed with the detailed statement of our results we first provide some terminology.

1.1 Basic definitions and notations

We model the state of the system by *load vectors*. A load vector $u = (u_1, \ldots, u_n)$ specifies that the number of balls in the ith bin is u_i. If u is *normalized* then the entries in the vector are sorted in decreasing order so that u_i describes the number of balls in the ith fullest bin. In case of Greedy[d], the order among the bins does not matter so that we can restrict the state space to normalized vectors. In case of Left[d], however, we need to consider general vectors.

Suppose X_t denotes the load vector at time t, i.e., after inserting t balls using Greedy[d] or Left[d], respectively. Then the stochastic process $(X_t)_{t\in\mathbb{N}}$ corresponds to a Markov Chain $\mathfrak{M} = (\mathbf{X}_t)_{t\in\mathbb{N}}$ whose transition probabilities are defined by the respective allocation process. In particular, X_t is a random variable obeying some probability distribution $\mathcal{L}$ defined by the allocation scheme. We use a standard measure of discrepancy between two probability distributions ϑ and ν on a space Ω. The *variation distance*, defined as

$$\|\vartheta - \nu\| = \frac{1}{2}\sum_{\omega\in\Omega} |\vartheta(\omega) - \nu(\omega)| = \max_{A\subseteq\Omega}(\vartheta(A) - \nu(A)) \ .$$

A basic technique which we apply is *coupling* (cf., e.g., [3]). A coupling for two (possibly the same) Markov chains $\mathfrak{M}_X = (\mathbf{X}_t)_{t\in\mathbb{N}}$ with state space Ω_X and $\mathfrak{M}_Y = (\mathbf{Y}_t)_{t\in\mathbb{N}}$ with state space Ω_Y is a stochastic process $(X_t, Y_t)_{t\in\mathbb{N}}$ on $\Omega_X \times \Omega_Y$ such that each of $(X_t)_{t\in\mathbb{N}}$ and $(Y_t)_{t\in\mathbb{N}}$ is a faithfull copy of $\mathfrak{M}_X$ and $\mathfrak{M}_Y$, respectively.

Another basic concept that we use frequently is *majorization* (cf., e.g., [2]). We say that a vector $u = (u_1, \ldots, u_n)$ is *majorized* by a vector v, written $u \leq v$, if for $1 \leq i \leq n$, if

$$\sum_{1\leq j\leq i} u_{\pi(j)} \leq \sum_{1\leq j\leq i} v_{\sigma(j)},$$

where π and σ are permutations of $1, \ldots, n$ such that $u_{\pi(1)} \geq u_{\pi(2)} \geq \cdots \geq u_{\pi(n)}$ and $v_{\sigma(1)} \geq v_{\sigma(2)} \geq \cdots \geq v_{\sigma(n)}$. Given an allocation scheme $\mathcal{X}$ defining a Markov Chain $\mathfrak{M}_X = (\mathbf{X}_t)_{t\in\mathbb{N}}$ and an allocation scheme $\mathcal{Y}$ defining a Markov Chain $\mathfrak{M}_Y = (\mathbf{Y}_t)_{t\in\mathbb{N}}$, we say that $\mathcal{X}$ is *majorized* by $\mathcal{Y}$ if there is a coupling between the two Markov chains $\mathfrak{M}_X$ and $\mathfrak{M}_Y$ such that $X_t \leq Y_t$, for all $t \in \mathbb{N}$.

In order to express our results of the always-go-left scheme we use Fibonacci numbers. Define $F_d(k) = 0$ for $k \leq 0$, $F_d(1) = 1$, and $F_d(k) = \sum_{i=1}^{d} F_d(k-i)$ for $k \geq 2$. Let $\phi_d = \lim_{k\to\infty} \sqrt[k]{F_d(k)}$, so that $F_d(k) = \Theta(\phi_d^k)$. Notice that ϕ_2 corresponds to the golden ratio. In general $1.61 < \phi_2 < \phi_3 < \cdots < 2$.

1.2 New Results

We present the first tight analysis for multiple-choice algorithms assuming an arbitrary number of balls. In particular, we show that the multiple-choice games are fundamentally different from the classical single-choice game in that they have "short memory".

THEOREM 1. *Let $\varepsilon > 0$. Let $d \geq 2$ be any integer. Let X and Y be any two load vectors describing the allocation of M balls to n bins. Let X_t (Y_t) be the random variable that describes the load vector after allocating t further balls on top of X (Y, respectively) using protocol* Greedy[d]. *Then there is a $\tau = \mathcal{O}(n^2 M + n^4 \ln(M\varepsilon^{-1}))$ such that $\|\mathcal{L}(X_\tau) - \mathcal{L}(Y_\tau)\| \leq \varepsilon$.*

In other words, given any configuration with a maximum load difference Δ between any pair of bins, the Greedy[d] process forgets this inbalance in $\Delta \cdot \mathrm{poly}(n)$ steps. The allocation after inserting further $\Delta \cdot \mathrm{poly}(n)$ balls is undistinguishable from an allocation obtained by starting from a totally balanced system. This is in contrast to the single-choice game requiring $\Delta^2 \cdot \mathrm{poly}(n)$ steps in order to forget a load difference of Δ.

We show that this property yields a fundamental difference between the allocation obtained by the multiple- and the single-choice algorithms. While the allocation of the single-choice algorithm deviates more and more from the optimal allocation with an increasing number of balls, the deviation between the multiple-choice and the optimal allocation is independent from the number of balls.

THEOREM 2. *Suppose we allocate m balls to n bins using* Greedy[d] *with $d \geq 2$. Then the number of bins with load at least $\frac{m}{n} + i + \gamma$ is bounded above by $n \cdot \exp(-d^i)$, w.h.p., where γ denotes a suitable constant.*

This result is tight up to additive constants in the sense that, for $m \geq n$, the number of bins with load at least $\frac{m}{n} + i \pm \Theta(1)$ is also bounded below by $n \cdot \exp(-d^i)$, w.h.p. In particular, our result yields an almost exact estimation for the number of balls in the fullest bin, that is, the max height above average is at most $\frac{\ln\ln n}{\ln d} \pm \Theta(1)$, w.h.p.

The result for the always-go-left scheme is even slightly better. The allocation is described in terms of Fibonacci numbers defined in the last section.

THEOREM 3. *Suppose we allocate m balls to n bins using* Left[d] *with $d \geq 2$. Then the number of bins with load at least $\frac{m}{n} + i + \gamma$ is bounded above by $n \cdot \exp(-\phi_d^{d \cdot i})$, w.h.p., where γ denotes a suitable constant.*

Also this bound is tight up to additive constants because the number of bins with load at least $\frac{m}{n} + i \pm \Theta(1)$ is lower bounded by $n \cdot \exp(-\phi_d^{d \cdot i})$, w.h.p., too. In particular, the max height above average produced by Left[d] is only $\frac{\ln \ln n}{d \cdot \ln \phi_d} + \Theta(1)$, w.h.p.
In addition to these quantitative results, we investigate the relationship between the greedy and the always-go-left scheme directly.

THEOREM 4. Left[d] *is majorized by* Greedy[d].

In other words, we show that the always-go-left scheme produces a better load balancing than the greedy scheme for any choices of n, m, and d.

1.3 Outline

First, we will present the analysis for the greedy process. In Section 2, we will show that Greedy[d] has short memory. Based on this property, we will show in Section 3 that one only needs to consider a polynomial number of balls in order to analyze the allocation for an arbitrary number of balls. In Section 4, we will analyze the allocation generated by Greedy[d] assuming a polynomial number of balls.
Second, we will present the analysis for the always-go-left process. Here we do not prove the short memory property explicitly. Instead our main tool is majorization of Left[d] by Greedy[d]. In Section 5, we will show this majorization. In Section 6, we will analyze the allocation obtained by Left[d] based on the knowledge about the allocation of Greedy[d].

2. GREEDY HAS SHORT MEMORY

In this section we prove Theorem 1. For every $k \geq 0$, let Ω_k denote the set of normalized load vectors with k balls. Let X and Y denote two vectors from Ω_M. We add further balls on top of the allocations described by these vectors and asked how many balls do we have to add until the two allocations are almost indistinguishable.
In our analysis, we shall investigate the following Markov chain $\mathfrak{M}[d] = (\mathcal{M}_t)_{t \in \mathbb{N}}$, which models the behavior of protocol Greedy[d]:

<u>**Input:**</u> *$\mathcal{M}_0$ is any load vector in Ω_M*

<u>**Transitions $\mathcal{M}_t \Rightarrow \mathcal{M}_{t+1}$:**</u>

Pick $q \in [n]$ at random such that

$$\mathbf{Pr}[q = k] = \frac{(n+1-k)^d - (n-k)^d}{n^d}$$

$\mathcal{M}_{t+1}$ is obtained from $\mathcal{M}_t$ by adding a new ball to the qth fullest bin

Let us remark that the choice of q is equivalent to the choice obtained by the following simple randomized process: Pick $q_1, q_2, \ldots, q_d \in [n]$ i.u.r. and then set $q = \max\{q_i : 1 \leq i \leq d\}$. Therefore, using the notation from Theorem 1, conditioned on $\mathcal{M}_0 = X$, it holds $X_t = \mathcal{M}_t$, and similarly, conditioned on $\mathcal{M}_0 = Y$ it holds $Y_t = \mathcal{M}_t$.
Our main tool in the analysis of this Markov chain is a variant of the path coupling argument of Bubley and Dyer [4]. This variant was introduced in [5] and is described in the following lemma.

LEMMA 1. **(Neighboring-Coupling Lemma)** *Let $\mathfrak{M} = (\mathbf{Y}_t)_{t \in \mathbb{N}}$ be a discrete-time Markov chain with a state space Ω. Let $\Omega^* \subseteq \Omega$. Let Γ be any subset of $\Omega^* \times \Omega^*$ (elements $(X,Y) \in \Gamma$ are called neighbors). Suppose that there is an integer D such that for every $(X,Y) \in \Omega^* \times \Omega^*$ there exists a sequence $X = \Lambda_0, \Lambda_1, \ldots, \Lambda_r = Y$, where $(\Lambda_i, \Lambda_{i+1}) \in \Gamma$ for $0 \leq i < r$, and $r \leq D$.*
If there exists a coupling $(X_t, Y_t)_{t \in \mathbb{N}}$ for $\mathfrak{M}$ such that for some $\mathbb{T} \in \mathbb{N}$ it holds $\mathbf{Pr}[X_\mathbb{T} \neq Y_\mathbb{T} \mid (X_0, Y_0) = (X,Y)] \leq \frac{\varepsilon}{D}$ for all $(X,Y) \in \Gamma$, then

$$\|\mathcal{L}(X_t|X_0 = X) - \mathcal{L}(Y_t|Y_0 = Y) \leq \varepsilon\|,$$

for every $(X,Y) \in \Omega^* \times \Omega^*$.

PROOF. For any pair of neighbors $(\Lambda, \Lambda') \in \Gamma$ we have

$$\|\mathcal{L}(Z_t|Z_0 = \Lambda) - \mathcal{L}(Z_t|Z_0 = \Lambda')\| \leq \frac{\varepsilon}{D}$$

by the well known coupling lemma (see, e.g., [3, Lemma 3.6]). As a consequence,

$$\begin{aligned} &\|\mathcal{L}(Z_t|Z_0 = X) - \mathcal{L}(Z_t|Z_0 = Y)\| \\ &\leq \sum_{i=1}^{r} \|\mathcal{L}(Z_t|Z_0 = \Lambda_i) - \mathcal{L}(Z_t|Z_0 = \Lambda_{i-1})\| \\ &\leq r \cdot \frac{\varepsilon}{D} \leq \varepsilon . \end{aligned}$$

□

Thus, if we can find a neighboring-coupling, we obtain immediately a bound on the total variation distance in terms of the tail probabilities of the *coupling time*, i.e., a random time $\mathbb{T}$ for which $X_t = Y_t$ for all $t \geq \mathbb{T}$.
In order to apply Lemma 1, we must first define the notion of neighbors. Let us fix M and n. Let us define $\Omega^* = \Omega_M$ and let Γ to be the set of pairs of those load vectors from Ω_M which correspond to the balls' allocations that differ in exactly one ball. In that case, if X can be obtained from Y by moving a ball from the ith fullest bin into the jth fullest bin, then we shall write $X = Y - \mathbf{e}_i + \mathbf{e}_j$. Thus,

$$\Gamma = \{(X,Y) \in \Omega_M \times \Omega_M \mid X = Y - \mathbf{e}_i + \mathbf{e}_j \text{ for certain } i,j \in [n], i \neq j\} .$$

Clearly, for each $X, Y \in \Omega_M$ there exists a sequence $X = Z^{\langle 0 \rangle}, Z^{\langle 1 \rangle}, \ldots, Z^{\langle l-1 \rangle}, Z^{\langle l \rangle} = Y$, where l is the number of balls on which X and Y differ, and $(Z^{\langle i \rangle}, Z^{\langle i+1 \rangle}) \in \Gamma$ for every i, $0 \leq i \leq l-1$. Notice further that $l \leq M$. Thus, we can apply the Neighboring-Coupling Lemma with $D = M$. In this way, it only remains to show the following lemma in order to complete the proof of Theorem 1.

LEMMA 2. *Let $\varepsilon > 0$. Let $d \geq 2$ be integer. Then, there exists $\mathbb{T} = \Theta(M\,n^2 + n^4 \cdot \ln(M/\varepsilon))$ such that for any $\tau \geq \mathbb{T}$ and any $X, Y \in \Gamma$ it holds that*

$$\mathbf{Pr}[X_\tau \neq Y_\tau] \leq \frac{\varepsilon}{M}.$$

In the rest of this section, we deal with the proof of Lemma 2. For simplicity, we shall assume $d = 2$. (In fact, the case $d > 2$ requires some more arguments.) For any load vectors $X = (x_1, \ldots, x_n)$ and $Y = (y_1, \ldots y_n)$ with $X = Y - \mathbf{e}_i + \mathbf{e}_j$, $i, j \in [n]$, let us define the *distance* function $\Delta(X,Y)$ to be the maximum of $|x_i - x_j|$ and $|y_i - y_j|$.
Observe that $\Delta(X,Y)$ is always a non-negative integer, it is zero only if $X = Y$, and that it never takes the value of 1. The following lemma describes main properties of the desired coupling.

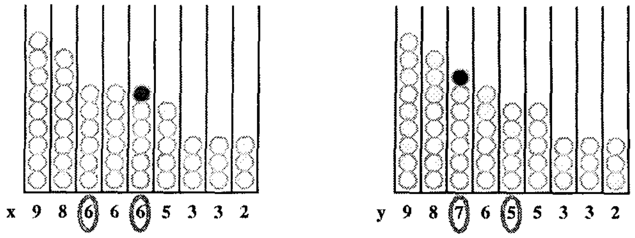

Figure 1: An example of vectors X and Y that differ only in one ball. In this case $X = Y - \mathbf{e}_3 + \mathbf{e}_5$ so that $\Delta(X,Y) = \max\{|x_3 - x_5|, |y_3 - y_5|\} = 2$.

LEMMA 3. *If* $(X,Y) \in \Gamma$ *then there exists a coupling* $(X_t, Y_t)_{t\in\mathbb{N}}$ *for* $\mathfrak{M}[2]$ *that, conditioned on* $(X_0, Y_0) = (X, Y)$, *possesses the following properties:*

- *for every* $t \in \mathbb{N}$, *if* $X_t = Y_t$ *then* $X_{t+1} = Y_{t+1}$,
- *for every* $t \in \mathbb{N}$, *if* $X_t \neq Y_t$ *then* X_t *and* Y_t *differ in at most one ball,*
- $\Delta(X_{t+1}, Y_{t+1}) - \Delta(X_t, Y_t) \in \{-2, -1, 0, 1\}$ *for every* $t \in \mathbb{N}$, *and*
- *for every* $t \in \mathbb{N}$, *if* $X_t \neq Y_t$ *then we have*
$$\mathbf{E}[\Delta(X_{t+1}, Y_{t+1}) \mid X_t, Y_t] \leq \Delta(X_t, Y_t) - \xi\ ,$$
where $\xi = \min\{\mathbf{Pr}[\text{Greedy}[2]$ *picks the jth fullest bin*$] - \mathbf{Pr}[\text{Greedy}[2]$ *picks the ith fullest bin*$]\ :\ i, j \in [n],\ i < j\}$ *and* $\xi \geq 1/n^2$.

PROOF. We use the following natural coupling: each time we increase the vectors X and Y by one ball, we use the same random choice. That is, in each step the obtained load vectors will be obtained from X and Y, respectively, by allocating a new ball to the qth fullest bin for certain $q \in [n]$.
The lemma follows directly from the following properties of the coupling. Consider X, Y from Ω_M with $X = Y - \mathbf{e}_i + \mathbf{e}_j$ for certain $i < j$. Let X_q and Y_q be obtained from X and Y, respectively, by allocating a new ball to the qth fullest bin. Then,

① either $X_q = Y_q$ and $\Delta(X_q, Y_q) = \Delta(X, Y) - 2$, or

② (X_q, Y_q) differ in one ball and

$$\Delta(X_q, Y_q) = \begin{cases} \Delta(X,Y) - 1 & \text{if and only if } q = j \\ \Delta(X,Y) + 1 & \text{if and only if } q = i \\ \Delta(X,Y) & \text{otherwise.} \end{cases}$$

The proof for these properties is by case analysis which is tedious but otherwise straightforward, and therefore we omit it here. □

Finally, we define $\Delta_t = \Delta(X_t, Y_t)$, for $t \geq 0$. From Lemma 3, we obtain that Δ_t behaves like a random walk with drift towards 0. Analyzing this random walk yields that $\Delta_t = 0$ with probability p, for $t = \Theta(M\,n^2 + n^4 \cdot \ln p)$. This implies Lemma 2 and, hence, Theorem 1.

3. A REDUCTION TO A POLYNOMIAL NUMBER OF BALLS

Now we show how to use the short memory property for the analysis of Greedy$[d]$ in the heavily loaded case. We use the following corollary which follows directly from Theorem 1 assuming that n is sufficiently large.

COROLLARY 1. *Suppose* $X_0 = (x_1, \ldots, x_n)$ *is any normalized load vector describing an allocation of some number of balls to* n *bins. Define* $\Delta = x_n - x_1$ *to be the maximum load difference in* X_0. *Let* Y_0 *be the load vector describing the optimal allocation of the same number of balls to* n *bins. Let* X_k *and* Y_k, *respectively, denote the vectors obtained after inserting* $k \geq 1$ *further balls to both systems using* Greedy$[d]$. *Then*
$$||\mathcal{L}(X_k) - \mathcal{L}(Y_k)|| \leq k^{-\alpha},$$
for $k \geq n^5\Delta$, *where* α *denotes an arbitrary constant.*

Using this corollary, we present a general transformation which shows that the allocation obtained by an allocation process with short memory is more or less independent of the number of balls. The following theorem shows that the allocation is basically determined after inserting a polynomial number of balls.

THEOREM 5. *Suppose* $\mathcal{P}$ *is an allocation protocol that has short memory and is majorized by the single-choice process. Let* $X_m = (x_1^{(m)}, \ldots, x_n^{(m)})$ *be a load vector obtained after allocating* m *balls with* $\mathcal{P}$. *Define* $\tilde{X}_m = (x_1^{(m)} - \frac{m}{n}, \ldots, x_n^{(m)} - \frac{m}{n})$. *Then for every* $M \geq N = n^{25}$ *and being a multiple of* n,
$$||\mathcal{L}(\tilde{X}_M) - \mathcal{L}(\tilde{X}_N)|| \leq N^{-\alpha},$$
where α *denotes an arbitrary constant.*

The following lemma shows that the variation distance between two systems with M and $M^{0.8}$ balls, respectively, is very small.

LEMMA 4. *Suppose* M *and* m *being multiples of* n *with* $M \geq n^{25}$ *and* $M \geq m \geq M^{0.8}$. *Then*
$$||\mathcal{L}(\tilde{X}_M) - \mathcal{L}(\tilde{X}_m)|| \leq M^{-\alpha}.$$

PROOF. Set $m' = M - m$. We use the majorization from the single-choice process to describe the situation after inserting m' balls. With probability p, each bin contains $m' \pm O(\sqrt{pm' \ln n/n})$ balls. Let $\Delta = M^{0.6}$. Applying $M \geq m' \geq n$ and doing some calculations yields that every bin contains between $\frac{m'}{n} - \Delta/2$ and $\frac{m'}{n} + \Delta/2$ balls, with probability $1 - p$, for $p = M^{-\alpha}/2$.
For the time being, let us assume that the entries in $\tilde{X}_{m'}$ are in the Δ-range specified above. Let Y describe another system in which the first m' balls are inserted in an optimal way, that is, $Y_{m'} = (\frac{m'}{n}, \ldots, \frac{m'}{n})$. Now we add m balls using protocol $\mathcal{P}$ on top of $X_{m'}$ and $Y_{m'}$, respectively. Observe that $m \geq M^{0.8} \geq n^5 M^{0.6} \geq n^5\Delta$. Thus, applying Corollary 1, we obtain $||\mathcal{L}(X_M) - \mathcal{L}(Y_M)|| \leq m^{2\alpha} + p \leq M^{-\alpha}/2$. □

For $M \leq N^{1/0.8}$, Lemma 4 directly implies Theorem 5. Otherwise, we have to apply the lemma repeatedly as follows. Let $m_0, m_1, \ldots, m_k$ denote a sequence of integers such that $m_0 = N$, $m_k = M$, $m_i^{0.8} \leq m_{i-1}$, and $m_i^{\alpha} \geq 2m_{i-1}^{\alpha}$. Then

$$\begin{aligned} ||\mathcal{L}(\tilde{X}_M) - \mathcal{L}(\tilde{X}_N)|| &\leq \sum_{i=1}^{k} ||\mathcal{L}(\tilde{X}_{m_i}) - \mathcal{L}(\tilde{X}_{m_{i-1}})|| \\ &\leq \sum_{i=1}^{k} m_i^{-\alpha} \leq N^{-\alpha}\ , \end{aligned}$$

where the last equation follows because $m_i^{-\alpha} \leq 2^{-i} m_0^{-\alpha} = 2^{-i} N^{-\alpha}$. This completes the proof of Theorem 5.

4. THE ALLOCATION GENERATED BY GREEDY

In this section, we investigate the allocation obtained by Greedy[d] in the heavily loaded case. In particular, we prove the bounds given in Theorem 2. Our arguments in the previous sections, 2 and 3, show that we can restrict ourselves to a polynomial number of balls in order to analyze the allocation for an arbitrary number of balls. In particular, we assume $m \leq n^{25}$. Furthermore, we assume w.l.o.g. that m is a multiple of n. We will show that the number of bins with load $\frac{m}{n} + i + \gamma$ is bounded above by $n \cdot \exp(-d^{\,i})$, w.h.p., where γ is a suitable constant.

For the analysis, we divide the set of balls into *batches* of size n each. The allocation *at time* t describes the number of balls in the bins after we have inserted the balls of the first t batches, i.e., after placing tn balls, starting with a set of empty bins at time 0.

We prove the theorem by an induction on t. Our induction must hold only for a polynomial number of steps. Nevertheless, we are not allowed to weaken the constraints on the allocation even by only one ball per step, as this would result in a too large deviation after a polynomial number of steps. Our trick that solves this problem is considering not only the balls lying above the average load but also the "holes" below the average load.

Obviously, the average number of balls per bin at time t is t. The bins with less than t balls are called *light bins* and the bins with more than t balls are called *heavy bins*. The *number of holes* at time t is defined as the number of balls one has to add to the bins so that each bin has load at least t.

We investigate the number of holes in the light bins and the number of balls in the heavy bins batch by batch in an interleaved induction. The analyses for the light and the heavy bins are almost independent from each other. Each of them uses only one simple but crucial induction assumption provided by the other. These assumptions are given by the following two invariants.

- $L(t)$: at time t, there are at most $2\,n$ holes below height t.
- $H(t)$: at time t, there are at most $n/32000$ balls with height $t + 12$ or larger.

Clearly, since t is the average number of balls per bin at time t, the number of holes below height t corresponds to the number of balls above height t. Thus, Invariant $L(t)$ implies that there are at most $2n$ balls with height $t + 1$ or larger at time t. Obviously, this is a very helpful assumption for bounding the height of the heavy bins.

4.1 Analysis for the light bins (Invariant L)

In order to show the simple bound on the total number of holes given in invariant L, we have to give almost exact bounds on the distribution of the holes among the light bins. A simple coupling argument (cf. also [2, Theorem 3.5]) shows that the number of holes generated by Greedy$[d]$ with $d \geq 2$ is majorized by the number of holes generated by Greedy[2]. The same argument shows that the tie breaking mechanism is irrelevant in case of the greedy algorithm. Therefore, we only need to consider Greedy[2] using a randomized tie breaking mechanism.

Let $\nu^t_{\leq i}$ denote the number of bins with load at most $t - i$ at time t. Define $\alpha_1 = 0.80$, $\alpha_2 = 0.55$, and $\alpha_i = 1.9^{-i+1}$, for $i \geq 3$. We will show the following invariants by induction:

- $L_1(t)$: $\nu^t_{\leq i} \leq \alpha_i n$, for $1 \leq i \leq c_1 \ln n$,
- $L_2(t)$: $\nu^t_{\leq i} = 0$, for $i \geq c_2 \ln n$,

where c_1, c_2 denote suitable constants, $c_1 \leq c_2$. Conditioning on the fact that L_1, L_2, and H hold up to time $t - 1$, we show that $L(t)$, $L_1(t)$, and $L_2(t)$ follow, w.h.p. Notice that invariant $L(t)$ is implied by $L_1(t)$ and $L_2(t)$ because these invariants yield that the number of holes at time t is at most

$$\sum_{i \geq 1} \alpha_i n + n \cdot (1.9)^{-c_1 \ln(n)+1} \cdot c_2 \ln n,$$

which is bounded above by $2n$.(Throughout the analysis, we assume w.l.o.g. that n is sufficiently large.) Hence, it remains to show only $L_1(t)$ and $L_2(t)$. In the rest of this subsection we shall condition on $L_1(T)$, $L_2(T)$, and $H(T)$ for all $T < t$.

We start the analysis with a simple observation. The major reason why the number of holes is very limited is, that bins with fewer balls are more likely to get a ball than bins with more balls. This can be formalized as follows.

OBSERVATION 1. *Let ℓ be an arbitrary integer and assume that at some point of time there exist at most $a_0 n$ bins with at most ℓ balls and at most $a_1 n$ bins with less than ℓ balls. Suppose that b is a bin with load exactly ℓ. Then the probability that the next ball allocated by* Greedy[2] *will be placed into bin b is at least $(2 - a_0 - a_1)/n$.*

Combining the bound in Observation 1 with invariant $L_1(T-1)$, for $1 \leq T \leq t$, we conclude that the probability that a ball from batch T falls into a fixed bin b holding $T - i$ or less balls is at least $(2 - 1.9^{-i+2} - 1.9^{-i+3})/n$. For example, if b contains $T - 11$ or less balls then this probability is larger than $1.99/n$ which is almost twice the average probability over all bins. This gives a clear intuition why none of the bins falls far behind, which is formalized in the following lemma.

LEMMA 5. *Let c_0, c_1, c_2 denote suitable constants. For any i with $c_0 < i \leq c_1 \ln n$, at most $n \cdot 1.9^{-i+1}$ bins contain $t - i$ or less balls at time t, w.h.p. Furthermore, every bin contains at least $t - c_2 \ln n$ balls, w.h.p.*

PROOF. Let $c_0 < i \leq c_1 \ln n$, where c_0 and c_1 will be specified later, and consider a bin b. Let q_T denote the number of holes in bin b below height T after round $T - 11$. Assume $t' < t$ is such that

$$q_{t'} = 0, \quad q_t \geq i, \quad \text{and} \quad q_T \geq 0, \quad \text{for all } t' < T \leq t.$$

Then, as outlined above, Observation 1 together with invariant $L_1(T)$ implies that the probability that a ball from a batch $t' < T \leq t$ falls into a bin b is at least at least $1.99/n$. That is, the number of balls which are placed into bin b during rounds $t' + 1$ to t is stochastically dominated by a binomially distributed random variable $\mathrm{BIN}((t - t') \cdot n, 1.99/n)$. Hence,

$$\begin{aligned} \mathbf{Pr}[q_t \geq i] &\leq \sum_{t'=0}^{t-1} \mathbf{Pr}[\mathrm{BIN}((t-t') \cdot n, 1.99/n) \leq (t-i) - t'] \\ &\leq \sum_{\tau \geq 1} \mathbf{Pr}[\mathrm{BIN}(\tau \cdot n, 1.99/n) \leq \tau - i] \ . \end{aligned}$$

We claim that, for n sufficiently large, $\mathbf{Pr}[\mathrm{BIN}(\tau \cdot n, 1.99/n) \leq \tau - i] \leq 1.05^{-\tau} \cdot 4^{-i}$. This is easily shown by induction on τ (condition on the outcome of the first n events). Hence,

$$\mathbf{Pr}[q_t \geq i] \leq \sum_{\tau \geq 1} 1.05^{-\tau} \cdot 4^{-i} \leq 20 \cdot 4^{-i} \ .$$

Consequently, $q_t = O(\log n)$, w.h.p., so that we can conclude that, for some suitable constant c_2, every bin includes at least $T - c_2 \ln n$ balls, w.h.p.

Next we show that for any i with $c_0 < i \leq c_1 \ln n$ at most $n \cdot 1.9^{-i+1}$ bins include $T - i$ or less balls at time T, w.h.p. Let X_b denote the event that bin b includes $T - i$ or less balls. From the above analysis we can conclude that for any $i > c_0 = 24$

$$\begin{aligned} \mathbf{E}[\sum X_b] &\leq n \cdot \mathbf{Pr}[q_t \geq i - 11] \\ &\leq n \cdot 20 \cdot 4^{-i+11} \\ &\leq \frac{1}{2} \cdot n \cdot 1.9^{-i+1} \ . \end{aligned}$$

for any $i > c_0 = 24$. Applying the zero-one lemma for balls and bins [6] we conclude that the random variables X_b are "negatively associated" so that we can apply a Chernoff bound, which yields

$$\sum X_b \leq \max\{n \cdot 1.9^{-i+1}, \sqrt{n}\} \ ,$$

w.h.p. We set $c_1 = 1/(2\ln(1.9)) = 0.77\ldots$. Then $\sum X_b \leq n \cdot 1.9^{-i+1}$, w.h.p, for $i \leq c_1 \ln n$, which yields the lemma. □

Lemma 5 yields invariant $L_2(t)$ and invariant $L_1(t)$ but only for $i > c_0$. Thus, it remains to show invariant $L_1(t)$ for $1 \leq i \leq c_0$. The following lemma estimates how the allocation of balls changes when placing n balls with Greedy[2] on the top of some previously placed balls.

LEMMA 6. *Let $\epsilon > 0$ and $a_0, \ldots, a_4$ with $0 < a_4 < \cdots a_0 < 1$ be constant reals. Let k and ℓ denote any integers. Suppose for $i = 0, \ldots, 4$ there are at most $a_i n$ bins with load at most $\ell - i$ balls at time $t - 1$. Then, at time t, the number of bins with load at most ℓ is less than or equal to $g(0,k) \cdot n$, w.h.p., where the function g is defined by $g(i,j) = a_i$, if $j = 0$ or $i = 4$, and otherwise*

$$\begin{aligned} &g(i,j) = \\ &\quad (1+\epsilon) \cdot (g(i+1,j-1) + (g(i,j-1) - g(i+1,j-1)) \cdot E) \ , \end{aligned}$$

where

$$E = \exp\left(-\frac{2 - g(i+1,j-1) - g(i,j-1)}{k}\right) \ .$$

The function g is monotonically increasing in each of the (implicit) parameters $a_0, \ldots, a_k$.

PROOF. We divide the allocation of the n balls into k phases in each of which we insert n/k balls using Greedy[2]. (For simplicity we assume that n is a multiple of k.) For $0 \leq i \leq 4$ and $0 \leq j \leq k$, we show that $n \cdot g(i,j)$ is an upper bound on the number bins with load at most $\ell - i$ after phase j.
For $j = 0$ or $i = 4$ the statement above holds trivially. Now suppose the statement is true for $g(0,j-1), \ldots, g(4,j-1)$. Consider the allocation of the n/k balls in phase j. Suppose b is a bin having load $\ell - i$ $(0 \leq i \leq 3)$ at the beginning of that phase. Observation 1 yields that the probability that b receives none of the next n/k balls is at most

$$\left(1 - \frac{2 - g(i,j-1) - g(i+1,j-1)}{n}\right)^{n/k} \leq E \ .$$

Thus, the expected number of bins that include at most $\ell - i$ balls at the end of phase k is upper bounded by

$$n \cdot g(i+1,j-1) + n \cdot (g(i,j-1) - g(i+1,j-1)) \cdot E$$

for $0 \leq i \leq 3$, which, by our definition, is equivalent to $n \cdot g(i,j)/(1+\epsilon)$. Applying Azuma's inequality, we can observe that the deviation from this expectation is only $o(n)$, w.h.p. Furthermore, as $n \cdot g(i,j) \geq n \cdot a_4 = \Theta(n)$, we conclude that we deviate only by a factor of $(1+\epsilon)$ from the expected value, so that the number of bins that include at most $\ell - i$ balls at the end of phase k is at most $n \cdot g(i,j)$, w.h.p.
Finally, we show the monotonicity properties of g. We observe that $g(i,j)$ is monotonically increasing in $g(i,j-1)$ and $g(i+1,j-1)$, for any $0 \leq i \leq 3$ and $1 \leq j \leq k$. Thus, $g(0,k)$ is monotonically increasing in $a_0, \ldots, a_4$, which completes the proof of the lemma. □

For $2 \leq i \leq c_0$, setting $a_0 = \alpha_{i-1}, a_1 = \alpha_i, \ldots, a_4 = \alpha_{i+3}$, the recurrence in Lemma 6 yields invariant $L_1(t)$. Notice that $a_0, \ldots, a_4$ fulfill the assumptions made in the lemma because we condition on invariant $L_1(t-1)$. The respective calculations are done numerically with Maple using Using $k = 40$ and $\epsilon = 10^{-4}$.
To show $L_1(t)$ for $i = 1$, we use the recurrence in Lemma 6, too. Here the challenging task, however, is to find an appropriate value for a_0. On the one hand, we require $a_0 n \geq \nu_{\leq 0}^{t-1}$. On the other hand, we want to show that $g(0,k) \leq \alpha_1$. For example, we may set $a_0 = 1.0$ as this is a trivial upper bound on $\nu_{\leq 0}^{t-1}/n$. It turns out, however, that this value is too large so that $g(0,k) > \alpha_1$. Thus, we have to use a more clever way to upper bound $\nu_{\leq 0}^{t-1}$.
The number of holes below height $t-1$ at time $t-1$ is $A := \sum_{j \geq 1} \nu_{\leq j}^{t-1}$. As the number of balls above the average height is equal to the number of holes below the average height, we can conclude that the number of balls above height $t-1$ is at least A, too. Furthermore, we can conclude from invariant $H(t-1)$ that, at the same time, there are at most $B := n/32000$ balls of height $(t-1) + 12 = t + 11$ or larger. Combining these two bounds, the number of balls which have height from t to $t + 10$ is at least $A - B$. This, however, requires that at least $(A-B)/11$ bins are filled with at least t balls, which gives us the desired upper bound on $\nu_{\leq 0}^{t-1}$, i.e.,

$$\begin{aligned} \nu_{\leq 0}^{t-1} &\leq n - \frac{A - B}{11} \\ &= n - \frac{\left(\sum_{j \geq 1} \nu_{\leq i}^{t-1}\right) - n/32000}{11} \\ &\leq n - \frac{\left(\sum_{j=1}^{4} \nu_{\leq i}^{t-1}\right) - n/32000}{11} \ . \end{aligned}$$

Now, we check all possible choices for $a_0, \ldots, a_4$ such that $a_j \leq \alpha_{j-1}$, for $1 \leq j \leq 4$, and $a_0 \leq 1 - ((\sum_{i=1}^{4} a_j) - 1/32000)/11$. In order to do so we make use of the monotonicity properties of $g(0,k)$. On one hand, the function $g(0,k)$ is monotonically increasing in a_0 and a_0 is monotonically decreasing in each of the parameters $a_1, \ldots, a_4$. On the other hand, for fixed a_0, the function $g(0,k)$ is monotonically increasing in each of the parameters $a_1, \ldots, a_4$. Therefore, it is sufficient to check the parameters $a_1, \ldots, a_4$ in steps of 0.02 while assuming

$$\begin{aligned} a_0 &\leq 1 - \frac{(\sum_{i=1}^{4}(a_j - 0.02)) - 1/32000}{11} \\ &\leq 1 - \frac{(\sum_{i=1}^{4} a_j) - 0.09}{11} \ . \end{aligned}$$

We do all these computations assuming $k = 40$ and $\epsilon = 10^{-4}$. For all possible choices, we obtain the desired result, i.e., $\nu_{\leq 1}^{t} \leq g(0,k) \leq \alpha_1$.

4.2 Analysis for the heavy bins (Invariant H)

In order to prove the bounds on the allocation of balls in the heavy bins, we use the upper bound on the number of holes given by invariant L. Let $\nu_{\geq i}^{t}$ denote the number of bins with load at least $t + 8 + i$ at time t. We will estimate these numbers using a function

f, which is defined as follows. Let $\kappa = \lceil \log_d(0.9\log_2 n - 1) \rceil$, and let $b \geq 1$ denote a suitable constant, which will be specified later. Define

$$\begin{aligned} f(i) &= \exp_2\left(-d^{\,i} - 1\right) \text{ , for } 0 \leq i < \kappa \\ f(\kappa) &= \max\left\{\tfrac{1}{4}n^{-0.9}, \exp_2\left(-d^{\,\kappa} - 1\right)\right(\ , \\ f(\kappa+1) &= b/n \ , \end{aligned}$$

For every point of time t, we will show the following invariants.

- $H_1(t)$: $\nu_{\geq i}^t \leq f(i) \cdot n$, for $0 \leq i \leq \kappa$;
- $H_2(t)$: $\sum_{i>\kappa} \nu_{\geq i}^t \leq b$;

Roughly speaking, invariant H_1 states that the sequence $\nu_{\geq 0}, \nu_{\geq 1}, \dots, \nu_{\geq \kappa}$ decreases doubly exponentially from n down to $\approx n^{0.1}$, and invariant H_2 states that there is only a constant number of balls above the last layer considered in this sequence. Clearly, these invariants yield the bounds given in Theorem 2.

We show the invariants H_1 and H_2 by induction. Our induction assumptions are $H_1(0), \dots, H_1(t-1)$, $H_2(t-1)$, and $L(t)$. We show that these assumptions imply $H_1(t)$, $H_2(t)$, and $H(t)$, w.h.p. Observe that $H_1(t)$ immediately implies $H(t)$ because it states that the number of balls of height 12 or larger is at most $\sum_{i\geq 12} \nu_{\geq i-8} \leq \sum_{i\geq 4} f(i) \cdot n \leq n/32000$. Thus, it remains only to show $H_1(t)$ and $H_2(t)$. We start our analysis by summarizing some properties of the function f.

OBSERVATION 2.

A1) $f(0) = 0.25$;

A2) $f(i) \geq 4f(i+1)$, *for* $1 \leq i \leq \kappa$;

A3) $f(i) \geq 2f(i-1)^d$, *for* $1 \leq i \leq \kappa+1$

A4) $f(i) \geq 0.25n^{-0.9}$, *for* $0 \leq i \leq \kappa$.

Property A2 requires that n is sufficiently large so that $f(\kappa) \geq 4b/n = 4f(\kappa+1)$. Property A3 follows from $2\exp_2\left(-d^{\,i-1} - 1\right)^d \leq \exp_2\left(-d^{\,i} - 1\right)$. Property A4 holds because κ is defined to be the smallest integer such that $\exp_2\left(-d^{\,\kappa} - 1\right) \leq n^{-0.9}$, so that for all all $i < \kappa$, $f(i) > n^{-0.9}$.

First, we show $H_1(t)$ using a "layered induction" on i, similar to the analysis presented in [2]. For the base case (i.e., i=0) we apply invariant $L(t)$. This invariant yields that, at time t, there are at most $2n$ balls of height larger than t. Consequently, the number of bins with $t+8$ or more balls is at most $2n/8 = n/4$. Applying property A1 yields $\nu_{\geq 0} \leq n/4 = f(0) \cdot n$. Thus, invariant $H_1(t)$ is shown for the case $i = 0$.

Now we show $H_1(t)$ for $i \geq 1$. We assume that $H_1(t)$ holds for $i-1$. Let $h(i)$ denote the number of bins that hold already $t+8+i$ or more balls at the beginning of round t, and let $h'(i)$ denote the number of balls from batch t that are placed into a bin containing at least $t+8+i-1$ balls. Observe that $\nu_{\geq i}^t \leq h(i) + h'(i)$. Thus, we only have to show that $h(i) + h'(i) \leq f(i) \cdot n$.

Applying induction assumption $H_1(t-1)$, we immediately obtain

$$\begin{aligned} h(i) &\leq \nu_{\geq i+1}^{t-1} \\ &\leq f(i+1) \cdot n \\ &\overset{(A2)}{\leq} f(i) \cdot n/4 \ , \end{aligned}$$

for $1 \leq i \leq \kappa$.

Bounding above $h'(i)$ requires some further arguments. For $1 \leq i \leq \kappa$, the probability that a fixed ball of batch t is allocated to height $t+8+i$ is at most $f(i-1)^d$. This is because each of its locations has to point to one of the bins with $t+8+i-1$ or more balls. By our induction on i, the number of these bins is bounded above by $f(i-1)$. Taking into account all n balls of round t, we obtain

$$\begin{aligned} \mathbf{E}[h(i)] &\leq f(i-1)^d \cdot n \\ &\overset{(A3)}{\leq} f(i) \cdot n/2 \ . \end{aligned}$$

Applying a Chernoff bound yields

$$\begin{aligned} \mathbf{Pr}[h(i) \geq 1.5(f(i) \cdot n/2)] &\leq \exp\left(0.5^2(f(i) \cdot n/2)/2\right) \\ &\overset{(A4)}{\leq} \exp\left(n^{0.1}/32\right) \ , \end{aligned}$$

for $1 \leq i \leq \kappa$. Consequently, $h'(i) \leq 0.75 f(i) \cdot n$, w.h.p., so that $\nu_{\geq i}^t \leq h(i) + h'(i) \leq f(i) \cdot n$. Hence, invariant $H_1(t)$ is shown.

Finally, we prove invariant $H_2(t)$. For $0 \leq T \leq t$, let x_T denote a random variable which is one if at least one ball of round T is allocated into a bin with load larger than $T+8+\kappa$, and zero, otherwise. Furthermore, let h_T denote the number of balls that are allocated into a bin with load larger than $T+8+\kappa$ in round T. Because of the invariants $H_1(1), \dots, H_1(t)$, the probability for a fixed ball from batch T to fall into a bin with more than $T+8+\kappa$ balls is at most $f(\kappa)^d \leq (n^{-0.9})^d \leq n^{-1.8}$. Therefore, the probability for $x_T = 1$ is bounded above by $n \cdot n^{-1.8} = n^{-0.8}$. Furthermore, the probability that $h_T \geq r$, for any integer $r \geq 3$, is at most

$$\binom{n}{r}\left(\frac{1}{n^{0.8}}\right)^r \leq \left(\frac{\mathrm{e}}{rn^{0.8}}\right)^r \leq n^{-0.8r} \ .$$

Thus, $h_T \leq r$, w.h.p., for some suitable constant r. Therefore, we may assume $h_T \leq r$, for $1 \leq T \leq t$. A violation of $H_2(t)$ implies that the bins with load at least $t+9+\kappa$ contain more than b balls of height at least $t+9+\kappa$. Observe that these balls must have been placed during the last b rounds, as otherwise one of the invariants $H_2(1), \dots, H_2(t-1)$ would be violated. That is, a violation of $H_2(t)$ implies implies that $r \cdot \sum_{T=t-b}^{t} x_T \geq b$. Consequently,

$$\begin{aligned} \mathbf{Pr}[\neg H_2(t)] &\leq \Pr\left\{\sum_{T=t-b}^{t} x_T \geq b/r\right\} \\ &\leq \binom{b}{b/r} \cdot \left(\frac{1}{n^{0.8}}\right)^{b/r} \\ &\leq \left(\frac{\mathrm{e}r}{n^{0.8}}\right)^{b/r} \ . \end{aligned}$$

For n sufficiently large, we obtain

$$\mathbf{Pr}[\neg H_2(t)] \leq n^{-0.5b/r} \ .$$

In other words, choosing b sufficiently large ensures that invariant H_2 holds, w.h.p., over all rounds. This completes the proof of Theorem 2.

5. GREEDY MAJORIZES ALWAYS-GO-LEFT

In this section, we prove Theorem 4, that is, we show that Left$[d]$ is majorized by Greedy$[d]$.

Let u denote the load vector obtained after inserting some number of balls with Left$[d]$, and let v denote the load vector obtained after inserting the same number of balls with Greedy$[d]$. W.l.o.g., we assume that u and v are normalized, i.e., $u_1 \geq u_2 \geq \cdots \geq u_n$ and $v_1 \geq v_2 \geq \cdots \geq v_n$. (Notice that the normalization of u jumbles the bins in the different groups used by Left$[d]$ in some way which,

however, we do not need to specify here. Further, observe that the normalized vector does not specify the always-go-left system completely.) By induction we assume that $u \geq v$.
Furthermore, let u' and v' denote the load vectors obtained by adding another ball b with Left[d] and Greedy[d], respectively. For $1 \leq i \leq n$, let e_i denote the ith unit vector, and define $a_i = \mathbf{Pr}[u' = u + e_i]$ and $b_i = \mathbf{Pr}[v' = v + e_i]$. For $0 \leq i \leq n$, set $A_i = \sum_{j=1}^{i} a_i$ and $B_i = \sum_{j=1}^{i} b_i$.

We use the following coupling of Left[d] and Greedy[d]. The random selection of the d locations of b and b's allocation to one of these locations is simulated by the following experiment. We choose uniformly at random some real number x from $[0 \ldots 1]$. Left[d] allocates ball b into the ith bin (with respect to the normalization) if $A_{i-1} < x \leq A_i$. Greedy[d] allocates b into the ith bin if $B_{i-1} < x \leq B_i$. By definition, the probabilities for these assignments correspond to the probabilities for the same assignments of the original schemes. Thus, the coupling is well defined. We have to show that $u' \leq v'$.

Suppose $u' = u + e_i$ and $v' = v + e_j$ for some i and j, that is, i and j specify the bins in which Left[d] and Greedy[d], respectively, placed the ball b. First, we assume that the initial vectors u and v are equal. In this case, we have to show that $u + e_i \leq u + e_j$.
Consider the *plateaus* of u, i.e., index sets of bins with same height. The first plateau U_1 includes all bins with load u_1, and the kth plateau U_k, for $k \geq 2$, includes all bins with load $u_{\max\{U_{k-1}\}+1}$. Let I and J denote the index of the plateau that contain i and j, respectively. Then $I \geq J$ implies $u + e_i \leq u + e_j$ because adding a ball to different positions of the same plateau results in the same normalized vector. Thus, it remains to show $I \geq J$.
Let $\ell = \max\{U_I\}$. Then the randomly selected value x satisfies $x \leq A_\ell \leq (\ell/n)^d$, which can be seen as follows. From $i \leq \ell$ we can conclude $x \leq A_\ell$. Further, A_ℓ corresponds to the probability P that Left[d] places b in a location with index smaller than or equal to ℓ. P depends on the distribution of the balls among the different groups. Let ℓ_k, for $1 \leq k \leq d$, denote the number of bins in group k with load u_ℓ or larger. Then $P = \prod_{k=1}^{d} (\ell_k \cdot d)/n$ because the kth location of b has to point to one of those ℓ_k bins among the d/n bins in group k that have load at least u_ℓ. Notice that P is maximized if we set $\ell_1 = \cdots = \ell_d = \ell/d$, because of the constraint $\sum_{k=1}^{d} \ell_k = \ell$. Hence, $x \leq A_\ell = P \leq (\ell/n)^d$.
Next we investigate the value of j in dependence from this bound on x. The probability that Greedy[d] places b in a bin with index smaller than or equal to ℓ is $(\ell/n)^d$ because all d locations of b must have an index in $1, \ldots, \ell$. Consequently, $B_\ell = (\ell/n)^d$, so that $x \leq (\ell/n)^d$ implies $j \leq \ell$. Now, because $\ell = \max\{U_I\}$, we obtain $J \leq I$.
Until now we have shown only $u + e_i \leq u + e_j$. This, however, yields the lemma almost immediately because for any two normalized load vectors w and w', $w \leq w'$ implies $w + e_i \leq w' + e_i$ (see [2, Lemma 3.4]). Consequently, we can conclude from $u + e_i \leq u + e_j$ that $u' = u + e_i \leq u + e_j \leq v + e_j = v'$, which yields Theorem 4.

6. ANALYSIS OF ALWAYS-GO-LEFT

In this section, we investigate the allocation generated by Left[d]. In particular, we prove Theorem 3, that is, we show that the number of bins with load at least $\frac{m}{n} + i + \gamma$ is $n \cdot \exp(-\phi_d^{d \cdot i})$, w.h.p., where γ is a suitable constant.
Similar to the proof for Greedy[d], we divide the set of balls into batches of size n, and we apply an induction on the number of batches. On the one hand, the proof for Left[d] is slightly more complicated as we have to take into account that the set of bins is partitioned into d groups. On the other hand, we can avoid the detour through analyzing the holes below average height as we can use the majorization of Left[d] by Greedy[d].
In the following, we assume that $m \geq n \log_2 n$. (It is easy to check, however, that a simplified variant of the following analysis works for the case $m < n \log_2 n$, too.) Basically, our analysis starts after the insertion of the first $m' = m - n \log_2 n$ balls, that is, we consider only the insertion of the last $n \log_2 n$ balls. We divide the set of these balls into $\log_2 n$ batches of size n each. The t-th batch is inserted in round t, for $1 \leq t \leq \log_2 n$. Let time 0 denote the point of time at the beginning of round 1, and let time t, for $1 \leq t \leq \log_2 n$, denote the point of time after inserting batch t. Furthermore, set $\Gamma = m'/n + 12$ and let $\nu_{\geq i}^{(j,t)}$ denote the number of bins with load at least $\Gamma + t + i$ in group j at time t, for $i, t \geq 0$ and $1 \leq j \leq d$.
We use majorization from Greedy[d] do estimate the allocation at time 0. (Notice that already m' balls are inserted at time 0.) For $i \geq 0$, define

$$f_0(i) \quad = \quad \frac{1}{4^i \cdot 16d} \ .$$

The following lemma gives a bound on the allocation of Left[d] obtained by the majorization from Greedy[d] at time 0. Based on this relatively weak bound, however, we will be able to prove the strong bounds on the allocation at the end of the process described in Theorem 3. (Later we will use the same lemma to estimate parts of the allocation also for other time steps $t \geq 1$.)

LEMMA 7.
$\nu_{\geq i}^{(j,t)} \leq f_0(i) \cdot n/d$, *w.h.p., for any* $i, j, t \geq 0, j < d$.

PROOF. Fix a time step t. The analysis of Greedy[d] ensures that, for $0 \leq i \leq \kappa$, the number of bins with $\Gamma + t + i$ or more balls at time t is bounded above by $\exp_2(-d^{\,4+i} - 1) \cdot n$, w.h.p., where $\kappa = \log_d \ln n + O(1)$. Furthermore, the number of balls above height κ is bounded above by a constant b. Thus, for $i > \kappa$, the number of bins with height $\Gamma + t + i$ or larger is bounded above by $b/(i - \kappa)$. Consequently, using Greedy[d], the number of bins with $\Gamma + t + i$ or more balls is at most

$$\begin{aligned} \max\left\{\exp_2\left(-d^{\,4+i} - 1\right) \cdot n, \frac{b}{i - \kappa}\right\} &\leq \frac{n}{4^i \cdot 32d^2} \\ &= \frac{f_0(i) \cdot n}{2d} \ , \end{aligned}$$

assuming that n is sufficiently large. Unfortunately, we need a bound on the number of balls above some given height rather than a bound on the number of bins above the height in order to apply the majorization. However, as the bound given above decreases geometrically in i, we obtain that the number of balls of height at least $\Gamma + t + i$ when using Greedy[d] is bounded above by $2 \cdot f_0(i) \cdot n/(2d) = f_0(i) \cdot n/d$. Now, because of the majorization, this result holds for Left[d], too. As the number of balls above height $\Gamma + t + i - 1$ upper bounds the number of bins with $\Gamma + t + i$ or more balls, we obtain that the number of bins with $\Gamma + t + i$ or more balls is bounded above by $f_0(i) \cdot n/d$. $\square$

Based on the knowledge about the allocation at time 0 obtained by the majorization, we analyze the allocation generated by Left[d] at any point of time t with $1 \leq t \leq \log_2 n$. For any $k, t \geq 0$, define

$$f_1(k) \quad = \quad \frac{\exp_2(-F_d(k - d + 1))}{16d} \ ,$$

and set

$$f'(k,t) = \max\{f_0(\lfloor k/d \rfloor)\cdot 2^{-t}, f_1(k)\} \ .$$

($F_d(k)$ denotes the kth d-ary Fibonacci number as defined in the Introduction.) Observe that $f'(k,0) = f_0(\lfloor k/d \rfloor)$ and $f'(k,\log_2 n) \leq f_1(k) + 1/(2n)$, for any $k \geq 0$, that is, f_0 determines f' at time 0 and f_1 determines f' at time $\log_2 n$, which is the time after inserting all balls.

Let κ_t denote the smallest integer such that $f'(\kappa_t, t) \leq n^{-0.9}$. For $0 \leq k < \kappa_t$, define $f(k,t) = f'(k,t)$. For $\kappa_t \leq k < \kappa_t + d$, define $f(k,t) = \max\{n^{-0.9}/4, f'(\kappa,t)\}$. Finally, for $k \geq \kappa_t + d$, set $f(k,t) = c/n$, where c denotes a suitable constant that will be specified later.

For every point of time t, we will show that the following invariants hold w.h.p. For $i,j \geq 0$, $j < d$, let $k(i,j) = i \cdot d + j$.

- $H_1(t)$: $\nu_{\geq i}^{(j,t)} \leq f(k(i,j),t)\cdot n/d$, for all $i,j \geq 0$, $j < d$ with $0 \leq k(i,j) \leq \kappa_t$.
- $H_2(t)$: $\sum_{i,j \text{ with } k(i,j) > \kappa_t} \nu_{\geq i}^{(j,t)} \leq c$.

Consider the invariants for the time $t = \log_2 n$. At this point of time the term $f_1(k)$ determines $f(k,t)$. This function decreases "Fibonacci exponentially", that is, the invariants state that the number of bins with $\Gamma + t + i = m/n + 13 + i$ or more balls is at most $\lceil \exp_2(-F_d(i\cdot d)\cdot n \rceil$. Combining the two invariants yields the bounds given in Theorem 3.

We show the invariants by induction on the number of rounds t. Lemma 7 yields that the invariants hold at time 0, that is, $H_1(0)$ and $H_2(0)$ are fulfilled. In the following, we assume that $H_1(t')$ and $H_2(t')$ are shown for any $t' < t$, and we show that these induction assumptions imply $H_1(t)$ and $H_2(t)$. We use the following properties of the function f.

LEMMA 8.

B1) $f(k,t) = f_0(0)$, *for* $0 \leq k < d$, $t \geq 0$.

B2) $f(k,t) \geq 2\cdot f(k+d, t-1)$, *for* $d-1 \leq k < \kappa_t + d$, $t \geq 1$*;*

B3) $f(k,t) \geq (4d)\cdot \prod_{\ell=1}^{d} f(k-\ell,t)$, *for* $d \leq k < \kappa_t + 2d$, $t \geq 0$*;*

B4) $f(k,t) \geq n^{-0.9}/2$, *for* $0 \leq k < \kappa_t + d$, $t \geq 0$.

PROOF. We start with the proof of property B1. For $0 \leq k < d$, $f_1(k) = \exp_2(-3)/d = 1/(8d)$. Thus, for $t \geq 0$,

$$\begin{aligned} f(k,t) &= f'(k,t) = \max\{f_0(\lfloor k/d \rfloor)\cdot 2^{-t}, f_1(k)\} \\ &= 1/(8d) \ = \ f_0(0) \ . \end{aligned}$$

Next we show property B2. For $k \geq d-1$,

$$\begin{aligned} f'(k,t) &= \max\{f_0(\lfloor k/d \rfloor)\cdot 2^{-t}, f_1(k)\} \\ &\geq \max\{2f_0(\lfloor (k+d)/d \rfloor)\cdot 2^{-(t-1)}, 2f_1(k+d)\} \\ &= 2f'(k+d,t-1) \ . \end{aligned}$$

For $d-1 \leq k < \kappa_t - d$, this implies

$$f(k,t) = f'(k,t) \geq 2f'(k+d,t-1) = 4f(k+d,t-1).$$

For $\kappa_t - d \leq k < \kappa_t$, the last equation may not hold, that is, $f'(k+d,t-1) < f(k+d,t-1)$. In this case, however, the definition of f ensures that $f(k+d,t-1) = n^{-0.9}/4$. Now we obtain directly from the definition of κ_t that

$$f(k,t) \geq n^{-0.9} = 2f(k+d,t-1).$$

For $\kappa_t \leq k < \kappa_t + d$, $f(k,t) \geq n^{-0.9}/4 \geq 2c/n = 2f(k+d,t-1)$, for n sufficiently large.

Property B3 can be shown as follows. Fix $d \leq k < \kappa_t \cdot d + 2d$. Depending on the outcome of $f(k-1,t)$, we distinguish three cases.

- Suppose $f(k-1,t) = f_0(\lfloor (k-1)/d \rfloor)\cdot 2^{-t}$. In this case,

$$\begin{aligned} \prod_{\ell=1}^{d} f(k-\ell,t) &= f_0\left(\left\lfloor \frac{k-1}{d} \right\rfloor\right)\cdot 2^{-t}\cdot \prod_{\ell=2}^{d} f(k-\ell,t) \\ &\leq 4f_0\left(\left\lfloor \frac{k}{d} \right\rfloor\right)\cdot 2^{-t}\cdot \frac{1}{16d} \\ &= \frac{f_0(\lfloor k/d \rfloor)\cdot 2^{-t}}{4d} \\ &\leq \frac{f(k,t)}{4d} \ . \end{aligned}$$

- Suppose $f(k-1,t) = f_1(k-1)$. Then, for $k' < k$, $f(k') = f_1(k')$, too. Thus,

$$\begin{aligned} \prod_{\ell=1}^{d} f(k-\ell,t) &= \prod_{\ell=1}^{d} \frac{\exp_2(-F_d(k-\ell-d+1))}{16d} \\ &= \frac{\exp_2(-\sum_{\ell=1}^{d} F_d(k-\ell-d+1))}{(16d)^d} \\ &\leq \frac{\exp_2(-F_d(k-d+1))}{(4d)(16d)} \\ &= \frac{f_1(k)}{4d} \\ &\leq \frac{f(k,t)}{4d} \ . \end{aligned}$$

- Suppose $f(k-1,t) = n^{-0.9}/4$. Then $f(k,t) \leq n^{-0.9}/4$. In this case either $f(k,t) = n^{-0.9}/4$ or $f(k,t) = c/n$. If $f(k,t) = n^{-0.9}/4$ then

$$\begin{aligned} \prod_{\ell=1}^{d} f(k-\ell,t) &\leq \frac{n^{-0.9}}{4}\cdot \prod_{\ell=2}^{d} f(k-\ell,t) \\ &\leq \frac{n^{-0.9}}{4}\cdot \frac{1}{4d} \\ &= \frac{f(k,t)}{4d} \ . \end{aligned}$$

If $f(k,t) = c/n$ then $k \geq \kappa_t + d$ so that $f'(k-\ell,t) \leq f(\kappa_t,t) \leq n^{-0.9}$, for $1 \leq \ell \leq d$. Thus, $f(k,t) = c/n \geq (n^{-0.9})^d/(4d) \geq (\prod_{\ell=1}^{d} f(k-\ell,t))/(4d)$.

Finally, we show property B4. For $0 \leq k < \kappa_t$, this property follows immediately from the definition of κ_t, and, for $\kappa \leq k < \kappa + d$, the property is ensured by the definition of f.

□

Now, exploiting the properties B1 to B4, we show that $H_1(t-1)$ implies $H_1(t)$. We use an induction on $k = i\cdot d + j$. First, we show that $H_1(t)$ holds for $k < d$, corresponding to $i = 0$. In this case, Lemma 7 yields $\nu_{\geq i}^{(j,t)} \leq f_0(0)\cdot n/d$. Furthermore, Property $B1$ yields $f_0(0) = f(k,t)$. Thus, $\nu_{\geq i}^{(j,t)} \leq f(k,t)\cdot n/d$, for $k < d$, $i = 0$.

Now assume that $H_1(t)$ is shown for all $k' < k$. Let $h(k) = h(i,j)$ denote the number of bins of group j that include $\Gamma + t + i$ balls already at the beginning of round t, and let $h'(k) = h'(i,j)$ denote the number of balls from batch t that are placed into a bin of group j containing at least $\Gamma + t + i - 1$ balls. Clearly,

$$\nu_{\geq i}^{(j,t)} \leq h(k) + h'(k).$$

Therefore, we calculate upper bounds for $h(k)$ and $h'(k)$.
By definition, $h(k) = h(i,j)$ is equal to $\nu_{\geq i+1}^{(j,t-1)}$. Thus, applying induction assumption $H_1(t-1)$, we can conclude that

$$\begin{aligned} h(k) &\leq \nu_{\geq i+1}^{(j,t-1)} \\ &\leq f((i+1)\cdot d + j, t-1)\cdot n/d \\ &= f(k+d, t-1)\cdot n/d \\ &\overset{(B2)}{\leq} 0.5 f(k)\cdot n/d \ . \end{aligned}$$

The term $h'(k) = h'(i,j)$ can be estimated as follows. If a ball is placed into a bin of group j with $\Gamma + t + i - 1$ balls, the d possible locations for that ball fulfill the following constraints. The location in group ℓ, for $0 \leq \ell < j$, points to a bin with load at least $\Gamma + t + i$. (Otherwise, the always-go-left scheme would assign the ball to that location instead of location j!) The number of these bins is $\nu_{\geq i}^{(\ell,t)}$. By the induction on k, $\nu_{\geq i}^{(\ell,t)} \leq f(i\cdot d + \ell, t)\cdot n/d$. Thus, the probability that the location points to a suitable bin is at most $f(i\cdot d+\ell, t)$. Besides, the location in group ℓ, for $j \leq \ell < d$, points to a bin with load at least $\Gamma + t + i - 1$. The number of these bins is $\nu_{\geq i-1}^{(\ell,t)}$. Thus, the probability for this event is at most $f((i-1)\cdot d + \ell, t)$. Now multiplying the probabilities for all d locations yields that the probability that a fixed ball is allocated to group j with height $\Gamma + t + i$ or larger is at most

$$\begin{aligned} &\prod_{\ell=0}^{j-1} f(i\cdot d+\ell, t)\cdot \prod_{\ell=j}^{d-1} f((i-1)\cdot d + \ell, t) \\ &\leq \prod_{\ell=1}^{d} f(k-\ell, t) \\ &\overset{(B3)}{\leq} \frac{f(k,t)}{4d} \ . \end{aligned}$$

Taking into account all n balls of round t, we obtain $\mathbf{E}[h'(k)] \leq f(k,t)\cdot n/(4d)$. Applying a Chernoff bound yields

$$\begin{aligned} &\mathbf{Pr}[h'(k) \geq 2\cdot f(k,t)\cdot n/(4d)] \\ &\quad\leq \exp(-f(k,t)\cdot n/(8d)) \\ &\quad\overset{(B4)}{\leq} \exp(-n^{0.1}/(16d)) \ . \end{aligned}$$

As a consequence, $h'(k) \leq 0.5 f(k,t)\cdot n/d$, w.h.p.
Combining both bounds, we obtain

$$\nu_{\geq i}^{(j,t)} \leq h(k) + h'(k) \leq f(k,t)\cdot n/d,$$

for $i \geq 0$, $0 \leq j < d$, and $k = i\cdot d + j$. Thus, invariant $H_1(t)$ is shown.
Invariant $H_2(t)$ states that there are at most c balls above layer $\Gamma + t + \kappa$. This layer is reached by at most $n^{0.1}/4$ bins. Because of the invariants $H_1(1), \ldots, H_1(t)$, the probability for a fixed ball from batch t to fall above this layer (i.e., to be placed into a bin with more than $\Gamma + t + \kappa$ balls) is at most $(n^{-0.9})^d \leq n^{-1.8}$. Thus, the probability that there are more than c balls above height $\Gamma + t + \kappa$ at the end of the process is at most

$$\binom{n\log_2 n}{c}\left(\frac{1}{n^{1.8}}\right)^c \leq \left(\frac{\mathrm{e}\log n}{cn^{0.8}}\right)^c \leq n^{-0.5c} \ .$$

Hence, $H_2(t)$ is shown as well. This completes the proof of Theorem 3.

7. REFERENCES

[1] Y. Azar, A. Z. Broder, and A. R. Karlin. On-line load balancing. *Theoretical Computer Science*, 130:73–84, 1994.

[2] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, September 1999. A preliminary version appeared in *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, pages 593–602, Montreal, Quebec, Canada, May 23–25, 1994. ACM Press, New York, NY.

[3] D. Aldous. Random walks of finite groups and rapidly mixing Markov chains. In J. Azéma and M. Yor, editors, *Séminaire de Probabilités XVII, 1981/82*, volume 986 of *Lecture Notes in Mathematics*, pages 243–297. Springer-Verlag, Berlin, 1983.

[4] R. Bubley and M. Dyer. Path coupling: A technique for proving rapid mixing in Markov chains. In *Proceedings of the 38th IEEE Symposium on Foundations of Computer Science*, pages 223–231, Miami Beach, FL, October 19–22, 1997. IEEE Computer Society Press, Los Alamitos, CA.

[5] A. Czumaj. Non-Markovian couplings and generating permutations via random transpositions. Manuscript, February 2000.

[6] D. Dubhashi and D. Ranjan. Balls and bins: A study in negative dependence. *Random Structures & Algorithms*, 13(2):99–124 1998.

[7] M. Mitzenmacher. Load balancing and density dependent jump Markov processes. In *Proceedings of the 37th IEEE Symposium on Foundations of Computer Science*, pages 213–222, Burlington, Vermont, October 14–16, 1996. IEEE Computer Society Press, Los Alamitos, CA.

[8] M. D. Mitzenmacher. *The Power of Two Choices in Randomized Load Balancing*. PhD thesis, Computer Science Department, University of California at Berkeley, September 1996.

[9] M. Raab and A. Steger. "Balls into bins" — a simple and tight analysis. In M. Luby, J. Rolim, and M. Serna, editors, *Proceedings of the 2nd International Workshop on Randomization and Approximation Techniques in Computer Science*, volume 1518 of *Lecture Notes in Computer Science*, pages 159–170, Barcelona, Spain, October 8–10, 1998. Springer-Verlag, Berlin.

[10] N. D. Vvedenskaya, R. L. Dobrushin, and F. I. Karpelevich. Queueing system with selection of the shortest of two queues: An assymptotic approach. *Problems of Information Transmission*, 32(1):15–27, January-March 1996.

[11] B. Vöcking. How asymetry helps load balancing. In *Proceedings of the 40th IEEE Symposium on Foundations of Computer Science*, pages 131–141, New York City, NY, October 17–19, 1999. IEEE Computer Society Press, Los Alamitos, CA.

[12] N. D. Vvedenskaya and Y. M. Suhov. Dobrushin's mean-field approximation for queue with dynamicTechnical Report N° 3328, INRIA, France, December 1997.