

The $(1 + \beta)$ -Choice Process and Weighted Balls-into-Bins

Yuval Peres*, Kunal Talwar†, Udi Wieder‡

Abstract

Suppose m balls are sequentially thrown into n bins where each ball goes into a random bin. It is well-known that the gap between the load of the most loaded bin and the average is $\Theta(\sqrt{\frac{m \log n}{n}})$, for large m . If each ball goes to the lesser loaded of two random bins, this gap dramatically reduces to $\Theta(\log \log n)$ independent of m . Consider now the following “ $(1 + \beta)$ -choice” process for some parameter $\beta \in (0, 1)$: each ball goes to a random bin with probability $(1 - \beta)$ and the lesser loaded of two random bins with probability β . *How does the gap for such a process behave? Suppose that the weight of each ball was drawn from a geometric distribution. How is the gap (now defined in terms of weight) affected?*

In this work, we develop general techniques for analyzing such balls-into-bins processes. Specifically, we show that for the $(1 + \beta)$ -choice process above, the gap is $\Theta(\log n / \beta)$, irrespective of m . Moreover the gap stays at $\Theta(\log n / \beta)$ in the weighted case for a large class of weight distributions. No non-trivial explicit bounds were previously known in the weighted case, even for the 2-choice paradigm.

1 Introduction

We consider balls-into-bins processes in which m balls are sequentially placed into n bins by some randomized placement procedure. The goal of the placement procedure is to create an allocation which is as balanced as possible while being simple and efficient. Balls-into-bins processes are a powerful and simple way to mathematically model many load balancing and resource allocation tasks common in computer science. The balls typically represent a set of independent items or tasks that need to be allocated to resources which are modeled by the bins.

One well known family of placement algorithms is the d -choice paradigm where $d > 0$ is an integer. In these processes each ball is placed in the less loaded of d uniformly sampled bins. (We remark that when these schemes are implemented in practice, the identifier of the items (balls) is hashed to produce the identifiers of

the sampled bins and the hash functions are assumed to produce a uniform and independent sample. In this paper we ignore the issue of providing an explicit hash function and always assume all samples are independent and uniform.) The ‘classical’ solution of setting $d = 1$, also known as the *single choice scheme*, is well understood and relatively easy to analyze, as the location of each ball is an independent random variable. It is well known (e.g. [6]) that if $m = n$ the maximum load is $\Theta(\frac{\log n}{\log \log n})$ with high probability, and if $m \gg n \log n$ the maximum load is with high probability $\frac{m}{n} + \Theta(\sqrt{\frac{m \log n}{n}})$. In a seminal paper Azar *et al.* [1] show that when $d \geq 2$ and $m = n$ the maximum load is $\frac{\log \log n}{\log d} + O(1)$ w.h.p. (the case $d = 2$ was implicitly shown by Karp *et al.* in [3]). There are many variants and applications of this result, see for instance the survey [5] and the references therein. Berenbrink *et al.* [2] generalize the bound for arbitrarily large m showing the load to be $\frac{m}{n} + \frac{\log \log n}{\log d} + O(1)$ with probability $1 - 1/\text{poly}(n)$. Thus, surprisingly, the additive gap between the heaviest bin and the average is independent of the number of balls.

1.1 Characterization by a Probability Vector

The d -choice scheme can be characterized by a probability vector $\mathbf{p} = (p_1, \dots, p_n)$, where p_i denotes the probability a ball falls in the i 'th most loaded bin; i.e. we order the bins from the most loaded to the least loaded (ties are broken arbitrarily) and we set p_1 to be the probability the most loaded bin receives the ball, p_2 is the probability the second bin receives it, and so on. Note that in the d -choice scheme $p_i = \binom{i}{n}^d - \binom{i-1}{n}^d$. So for $d = 1$, $p_i = \frac{1}{n}$ while if $d > 1$ it holds that $p_i > p_j$ for $i > j$. In other words, when $d > 1$ the process has a *bias towards the light bins*. At a high level, this bias explains the stark differences between the cases $d = 1$ and $d > 1$.

The $(1 + \beta)$ -choice process. In light of this characterization it is natural to investigate processes with different probability placement vectors. As a leading example of such a process consider the case where each time a ball is placed with probability $\beta \geq 0$ (for some fixed β) a two-choice scheme is used and with probability $1 - \beta$ a single choice scheme is used. We call this

*Microsoft Research Redmond, peres@microsoft.com

†Microsoft Research Silicon Valley, kunal@microsoft.com

‡Microsoft Research Silicon Valley, uwieder@microsoft.com

process *the* $(1 + \beta)$ -choice scheme. This process is motivated by the case where the cost assigned to querying the load of the bins is relatively high. Consider for instance a distributed storage system in which a front-end server places data items in back-end servers. Using the 2-choice scheme requires querying two servers for their load, possibly locking them until the data item is placed. Moreover, a lookup for the item results in two queries. In the $(1 + \beta)$ -choice scheme with probability $1 - \beta$ the assignment is done without performing this query, and the expected lookup cost is only $(1 + \beta)$. The question of course is how much imbalance is introduced due to this modified procedure. It is easy to verify that in this case $p_i = \frac{\beta(2i-1)}{n^2} + \frac{1-\beta}{n}$. Mitzenmacher in his thesis [4] suggested studying the $(1 + \beta)$ -process in the context of queueing theory. He observes that this process may also serve to model the case where all balls perform the best of 2 strategy, but some fraction of them are ‘misinformed’ by, say, an unreliable load reporting mechanism, and make the wrong decision.

Weighted balls. In practical applications it is often the case that items have heterogeneous sizes. Consider for instance the case of a storage system which places files (balls) into servers (bins). In this case each file has a different size, and the load of the server is the total size of files assigned to it. We model this case by letting each ball receive a weight by independently sampling from a distribution \mathcal{D} . Talwar and Wieder show [7] that in the two choice scheme, if the weight distribution has a finite variance and is ‘smooth’ in some mild sense, then the gap is independent of the number of balls thrown. However, no non-trivial upper bounds on the gap for specific distributions were known prior to this work.

1.2 Summary of Results and Techniques The intuition of our approach is drawn from viewing the process as an n -dimensional version of a biased random walk on a line: Consider a random walk on the line which starts at 0 and each step moves either right or left with equal probability. It is well known that after m steps, the expected distance from 0 is roughly \sqrt{m} . If however the random walk has a *bias towards* 0 then the expected distance from 0 is bounded by some constant independent of m . When $n = 2$, the difference between the loads of the two bins is indeed a random walk on the line: an unbiased one in the one choice case and a biased one in the $(1 + \beta)$ -choice case. We reduce the process at hand to a one dimensional process. Our approach is to find a potential function $\Gamma : \mathbb{R}^n \rightarrow \mathbb{R}$ from the set of allocations to the real numbers such that (i) if Γ is small then the allocation is balanced, and (ii) the expectation of Γ is bounded by a constant independent of m . Using

this technique we obtain the following results:

Generalization over the placement procedure: We can upper bound the expected gap of a large family of placement processes (defined by the vector \mathbf{p}). In particular for the $(1 + \beta)$ -choice scheme we show the expected gap to be $O(\frac{\log n}{\beta})$. We show that this bound is tight up to constants for β bounded away from 1. We then employ the concept of majorizations to generalize the bound further. As a case in point we analyze processes that are biased against the maximal bin (or biased in favor of the minimal bin) and are uniform everywhere else. We show that even this small bias in the allocation rule suffices to have the gap in the allocation independent of m . Thus, interestingly, the gap has three regimes as a function of β . When $\beta = 0$ it is $\Theta(\sqrt{\frac{m \log n}{n}})$ which goes to infinity with m . When $0 < \beta < 1$ the gap is $\Theta((\log n)/\beta)$ and when $\beta = 1$ it is $\Theta(\log \log n)$.

Weighted Balls: We generalize to the case where each ball samples a *weight* from a weight distribution \mathcal{D} , and a bin’s load is measured by the sum of weights of the balls it contains. We provide bounds for large families of weight distribution. We show that if \mathcal{D} ’s exponential moment generating function is finite for some positive value, then the expected gap remains $O((\log n)/\beta)$.

1.3 The Markov Chain Recall that we characterize a balls-and-bins process by a distribution vector (p_1, \dots, p_n) , where p_i is the probability a ball is placed in the i ’th most loaded bin. Further, let the size of the ball be an independent random variable drawn from a distribution \mathcal{D} on $[0, \infty)$. Let $x(t)$ be the vector denoting the load of each bin minus the average load, so $\sum x_i(t) = 0$. We assume that x is sorted so that $x_i \geq x_{i+1}$ for $i \in [n - 1]$. The process defines a Markov chain over the vectors $x(t)$ as follows:

- sample $j \in_p [n]$
- sample $W \in \mathcal{D}$
- set $y_i = x_i(t) + W - \frac{W}{n}$ for $i = j$ and $y_i = x_i(t) - \frac{W}{n}$ for $i \neq j$
- obtain $x(t + 1)$ by sorting y

We will prove our bounds for a family of distribution vectors \mathbf{p} . We make the following assumptions:

$$(1.1) \quad p_i \leq p_{i+1} \text{ for } i \in [n - 1]$$

For some $\epsilon > 0$ it holds that

$$(1.2) \quad p_{\frac{n}{3}} \leq \frac{1 - 4\epsilon}{n} \text{ and } p_{\frac{2n}{3}} \geq \frac{1 + 4\epsilon}{n}$$

The first assumption says that the allocation rule is no worse than the 1-choice scheme. The second assumption states that the allocation rule strictly prefers the least loaded third of bins over the most loaded third. Note that these assumptions imply that $\sum_{i \geq \frac{3n}{4}} p_i \geq \frac{1}{4} + \epsilon$ and $\sum_{i \leq \frac{n}{4}} p_i \leq \frac{1}{4} - \epsilon$, a fact that turns out to be very useful.

For the distribution \mathcal{D} , we assume that there is a $\lambda > 0$ such that the moment generating function $M[\lambda] = \mathbb{E}[e^{\lambda W}] < \infty$. Further, without loss of generality, $\mathbb{E}[W] = 1$. Note that

$$M''(z) = \mathbb{E}[W^2 e^{zW}] \leq \sqrt{\mathbb{E}[W^4] \mathbb{E}[e^{2zW}]}.$$

The above assumption implies that there is a $S \geq 1$ such that for every $|z| < \lambda/2$ it holds that $M''(z) < 2S$.

The $(1+\beta)$ -choice process: Note that the $(1+\beta)$ -choice process satisfies the assumptions for $\epsilon = \beta/12$ since $p_{\frac{n}{3}} \leq \frac{(1-\beta)}{n} + \frac{2\beta}{3n} = \frac{1-\beta}{n}$, and similarly $p_{\frac{2n}{3}} \geq \frac{1+\beta}{n}$.

2 The Upper Bound

Let $\alpha = \min(\frac{\epsilon}{6S}, \lambda/2)$. We can assume that $\epsilon \leq 1/4$ and thus that $\alpha \leq 1/6$. Define the following potential functions

$$\begin{aligned} \Phi(t) &= \Phi(x(t)) := \sum_{i=1}^n \exp(\alpha x_i) \\ \Psi(t) &= \Psi(x(t)) := \sum_{i=1}^n \exp(-\alpha x_i) \\ \Gamma(t) &= \Gamma(x(t)) := \Phi(t) + \Psi(t) \end{aligned}$$

Note that $\Gamma(0) = 2n$. We show that if $\Gamma(x(t)) \geq an$ for some $a > 0$ then $\mathbb{E}[\Gamma(x(t+1)) | x(t)] \leq \Gamma(x(t))(1 - \frac{\alpha\epsilon}{4n})$. We will use this to show that for every given t , $\mathbb{E}[\Gamma(t)] \in O(n)$, which implies that the maximum gap is $O(\log n)$ w.h.p.

REMARK 1. *The function $\Phi(x(t))$ by itself is a more natural potential function. The problem with using Φ is that it may not decrease on expectation, even if it takes an arbitrarily high value. To see this let a be arbitrary and consider a load vector where $x_i = a$ for $i \leq n-1$ and $x_n = -(n-1)a$. We have $\Phi(x) = (n-1)e^{\alpha a} + e^{-(n-1)\alpha a}$, so Φ is arbitrarily large. Assume $p_n = 2/n$. Now the expected value of Φ after we throw one ball is at least*

$$\left(1 - \frac{2}{n}\right) e^{-\frac{\alpha}{n}} \left(e^{\alpha a} + (n-2)e^{\alpha a}\right) + \frac{2}{n} e^{-\frac{\alpha}{n}} (n-1)e^{\alpha a}$$

Since $e^\alpha \geq 1 + \alpha + \frac{\alpha^2}{2}$ this is at least

$$\begin{aligned} &\left(1 - \frac{\alpha}{n}\right) e^{\alpha a} \left(\left(1 - \frac{2}{n}\right) \left(1 + \alpha + \frac{\alpha^2}{2} + n - 2\right) + \frac{2}{n}(n-1)\right) \\ &\geq \left(1 - \frac{\alpha}{n}\right) e^{\alpha a} \left(n - 1 + \alpha + \frac{\alpha^2}{2}\right) \end{aligned}$$

which is larger than $\Phi(x)$ for large a . One may observe that the reason Φ fails in this case is that it is insensitive to the load of the light bins. In this example the fact that the lightest bin had a load of $-(n-1)a$ barely affected the value of Φ . The function Ψ is designed to mitigate this problem by being sensitive to the light bins, while hardly being affected by the heavy bins. The sum of these two functions turns out to be suitable for the analysis.

We start by calculating the expected change of Φ and Ψ individually. We write Φ instead of $\Phi(t)$ when t is clear from context.

LEMMA 2.1. *For Φ defined as above,*

$$\begin{aligned} (2.3) \quad &\mathbb{E}[\Phi(t+1) - \Phi(t) | x(t)] \\ &\leq \sum_{i=1}^n \left(p_i(\alpha + S\alpha^2) - \left(\frac{\alpha}{n} - S\frac{\alpha^2}{n^2}\right)\right) e^{\alpha x_i}. \end{aligned}$$

Proof. Let Δ_i be the expected change in Φ if the ball is put in bin i , i.e. if $y_i(t+1) = x_i(t) + W - \frac{W}{n}$ and for $j \neq i$, $y_j(t+1) = x_j(t) - \frac{W}{n}$, (recall that $x(t+1)$ is obtained by sorting $y(t+1)$ and that $\Phi(x) = \Phi(y)$). Now the expected contribution of i is

$$\begin{aligned} &\mathbb{E}[e^{\alpha(x_i + W - \frac{W}{n})}] - e^{\alpha x_i} \\ &= e^{\alpha x_i} \left(M(\alpha(1 - \frac{1}{n})) - 1\right) \\ &= e^{\alpha x_i} \left(M(0) + M'(0)\alpha(1 - \frac{1}{n}) + M''(\zeta)(\alpha(1 - \frac{1}{n}))^2/2 - 1\right) \end{aligned}$$

for some $\zeta \in [0, \alpha(1 - \frac{1}{n})]$. By the assumption on \mathcal{D} and α , $M''(\zeta) \leq 2S$. Moreover, $M(0) = 1$ and $M'(0) = E[W] = 1$. Thus the above expression can be bounded from above by

$$e^{\alpha x_i} \left(\alpha(1 - \frac{1}{n}) + S\alpha^2\right)$$

The contribution of j is similarly computed to be bounded by $(-\frac{\alpha}{n} + S\frac{\alpha^2}{n^2})e^{\alpha x_j}$. Thus

$$\Delta_i \leq (\alpha + S\alpha^2)e^{\alpha x_i} - \left(\frac{\alpha}{n} - S\frac{\alpha^2}{n^2}\right)\Phi$$

so we have

$$\begin{aligned} &\mathbb{E}[\Phi(t+1) - \Phi(t) | x(t)] \\ &\leq \sum_{i=1}^n p_i e^{\alpha x_i} \left(\alpha(1 - \frac{1}{n}) + S\alpha^2 + (n-1)\left(-\frac{\alpha}{n} + S\frac{\alpha^2}{n^2}\right)\right) \\ &\leq \sum_{i=1}^n p_i e^{\alpha x_i} \left(\alpha + S\alpha^2 + n\left(-\frac{\alpha}{n} + S\frac{\alpha^2}{n^2}\right)\right) \end{aligned}$$

and the claim follows

COROLLARY 2.1.

$$(2.4) \quad \mathbb{E}[\Phi(t+1) - \Phi(t) \mid x(t)] \leq \frac{2\alpha}{n}\Phi(t)$$

Proof. Note that $S\alpha \leq \frac{1}{6} < 1$ so that

$$(2.5) \quad \mathbb{E}[\Phi(t+1) - \Phi(t) \mid x(t)] \leq \sum_{i=1}^n 2\alpha p_i e^{\alpha x_i}.$$

The claim follows by observing that p_i 's are increasing and x_i 's are decreasing, so that the expression is maximized when p_i 's are all equal.

Similar arguments show that

LEMMA 2.2. *Let Ψ be defined as above. Then*

$$(2.6) \quad \mathbb{E}[\Psi(t+1) - \Psi(t) \mid x(t)] \leq \sum_{i=1}^n \left(p_i(-\alpha + S\alpha^2) + \left(\frac{\alpha}{n} + S\frac{\alpha^2}{n^2} \right) \right) e^{-\alpha x_i}.$$

COROLLARY 2.2.

$$(2.7) \quad \mathbb{E}[\Psi(t+1) - \Psi(t) \mid x(t)] \leq \frac{2\alpha}{n}\Psi(t)$$

Proof. This follows immediately as $p_i > 0$ and $S\alpha < \frac{1}{6}$.

We start by showing that for reasonably balanced configurations, both Φ and Ψ have the right decrease in expectation. More precisely, if $x_{\frac{3n}{4}} \leq 0$, then Φ decreases in expectation, and if $x_{\frac{n}{4}} \geq 0$, then Ψ decreases in expectation.

LEMMA 2.3. *Let Φ be defined as above. If $x_{\frac{3n}{4}}(t) \leq 0$, then $\mathbb{E}[\Phi(t+1) \mid x(t)] \leq (1 - \frac{\alpha\epsilon}{n})\Phi(t) + 1$.*

Proof. We upper bound $\sum_{i=1}^n p_i(\alpha + S\alpha^2)e^{\alpha x_i}$ for a fixed $\Phi(x)$, for x which is non increasing with $\sum_i x_i = 0$. We first write

$$(2.8) \quad \begin{aligned} & \sum_{i=1}^n p_i(\alpha + S\alpha^2)e^{\alpha x_i} \\ & \leq \sum_{i < \frac{3n}{4}} p_i(\alpha + S\alpha^2)e^{\alpha x_i} + \sum_{i \geq \frac{3n}{4}} p_i(\alpha + S\alpha^2)e^0 \\ & \leq \sum_{i < \frac{3n}{4}} p_i(\alpha + S\alpha^2)e^{\alpha x_i} + 1 \end{aligned}$$

since $\alpha + S\alpha^2 \leq \frac{6\epsilon + \epsilon^2}{36S} \leq 1$ by our assumptions that $\epsilon \leq 1$ and $S \geq 1$.

Now set $y_i := e^{\alpha x_i}$. The first term above is no larger than the maximum value of

$$(\alpha + S\alpha^2) \sum_{i < \frac{3n}{4}} p_i y_i$$

subject to

$$\sum_{i < \frac{3n}{4}} y_i \leq \Phi$$

$$y_{i-1} \geq y_i \quad \forall 1 < i < \frac{3n}{4}.$$

Since \mathbf{p} is non-decreasing and \mathbf{y} is non-increasing, the maximum is achieved when $y_i = \frac{4\Phi}{3n}$ for each i , and is at most $(\alpha + S\alpha^2)(\frac{3}{4} - \epsilon)\frac{4\Phi}{3n}$.

We can now plug this bound in (2.8), and substituting in (2.3) we upper-bound the expected change in Φ .

$$\begin{aligned} & \mathbb{E}[\Phi(t+1) - \Phi(t) \mid x(t)] \\ & \leq (\alpha + S\alpha^2)\left(\frac{3}{4} - \epsilon\right)\frac{4\Phi}{3n} - \left(\frac{\alpha}{n} - S\frac{\alpha^2}{n^2}\right)\Phi + 1 \\ & \leq \frac{\alpha\Phi}{n}\left((1 + S\alpha)\left(1 - \frac{4\epsilon}{3}\right) - 1 + S\frac{\alpha}{n}\right) + 1 \end{aligned}$$

Assuming $S\alpha \leq \epsilon/6$ we have

$$\begin{aligned} & \leq \frac{\alpha}{n}\Phi\left(\frac{\epsilon}{6} - \frac{4\epsilon}{3} + \frac{\epsilon}{6n}\right) + 1 \\ & \leq -\frac{\alpha\epsilon}{n}\Phi + 1 \end{aligned}$$

The claim follows.

LEMMA 2.4. *Let Ψ be defined as above. If $x_{\frac{n}{4}}(t) \geq 0$, then $\mathbb{E}[\Psi(t+1) \mid x(t)] \leq (1 - \frac{\alpha\epsilon}{n})\Psi(t) + 1$.*

Proof. The proof is similar to the previous lemma. We first upper bound $\sum_{i=1}^n p_i(-\alpha + S\alpha^2)e^{-\alpha x_i}$ for a fixed $\Psi(x)$, for x which is non increasing with $\sum_i x_i = 0$. Since $(-\alpha + S\alpha^2)$ is negative, we have

$$\sum_{i=1}^n p_i(-\alpha + S\alpha^2)e^{-\alpha x_i} \leq (-\alpha + S\alpha^2) \sum_{i \geq \frac{n}{4}} p_i e^{-\alpha x_i}$$

Now set $z_i := e^{-\alpha x_i}$. Under the assumption on $x_{\frac{n}{4}}$, the sum $\sum_{i \geq \frac{n}{4}} z_i$ is at least $\Psi - \frac{n}{4}$. Since $(-\alpha + S\alpha^2)$ is negative, to upper bound the second term, we need to find the minimum value of

$$\sum_{i \geq \frac{n}{4}} p_i z_i$$

subject to

$$\sum_{i \geq \frac{n}{4}} z_i \geq \Psi - \frac{n}{4}$$

$$z_{i-1} \geq z_i \quad \forall i > \frac{n}{4}.$$

Since both \mathbf{p} and \mathbf{z} are (weakly) increasing, the minimum is achieved when $z_i = \frac{4(\Psi - \frac{n}{4})}{3\frac{3}{4}}$ for each i . Using the assumption that $\sum_{i \geq n/4} p_i \geq \frac{3}{4} + \epsilon$ we can bound the expression above by $(-\alpha + S\alpha^2)(\frac{3}{4} + \epsilon)\frac{4(\Psi - \frac{n}{4})}{3n}$. We can now upper-bound the expected change in Ψ by plugging this bound in (2.6).

$$\begin{aligned} & \mathbb{E}[\Psi(t+1) - \Psi(t) \mid x(t)] \\ & \leq (-\alpha + S\alpha^2)(\frac{3}{4} + \epsilon)\frac{4(\Psi - \frac{n}{4})}{3n} + \frac{\alpha}{n}(1 + S\frac{\alpha}{n})\Psi \\ & = \frac{\alpha}{n}((1 + S\frac{\alpha}{n})\Psi + (-1 + S\alpha)(\frac{3}{4} + \epsilon)\frac{4\Psi - n}{3}) \\ & = \frac{\alpha}{n}((1 + S\frac{\alpha}{n})\Psi + S\alpha(\frac{3}{4} + \epsilon)\frac{4\Psi - n}{3} - (\frac{3}{4} + \epsilon)\frac{4\Psi - n}{3}) \\ & \leq \frac{\alpha\Psi}{n}(1 + S\frac{\alpha}{n} + S\alpha(\frac{3}{4} + \epsilon)\frac{4}{3} - (\frac{3}{4} + \epsilon)\frac{4}{3}) + \frac{\alpha}{3}(\frac{3}{4} + \epsilon) \\ & \leq -\frac{\alpha\epsilon}{n}\Psi + 1 \end{aligned}$$

where the last inequality follows since $\epsilon \leq \frac{1}{4}$ and $S\alpha \leq \frac{\epsilon}{6}$.

The next lemma will be useful in the case that $x_{\frac{3n}{4}} > 0$.

LEMMA 2.5. *Suppose that $x_{\frac{3n}{4}} > 0$ and $\mathbb{E}[\Delta\Phi|x(t)] \geq -\frac{\alpha\epsilon}{4n}\Phi$. Then either $\Phi < \frac{\epsilon}{4}\Psi$ or $\Gamma < cn$ for some $c = \text{poly}(\frac{1}{\epsilon})$.*

Proof. First note that the expected increase in Φ is at most

$$\begin{aligned} & \sum_i (p_i(\alpha + S\alpha^2) - \frac{\alpha}{n} + S\frac{\alpha^2}{n^2})e^{\alpha x_i} \\ & \leq \sum_{i \leq n/3} (p_i(\alpha + S\alpha^2) - \frac{\alpha}{n} + S\frac{\alpha^2}{n^2})e^{\alpha x_i} \\ & \quad + (\alpha + S\alpha^2) \sum_{i > n/3} p_i e^{\alpha x_i} \\ & \leq -\frac{\alpha\epsilon}{2n}\Phi_{\leq n/3} + \frac{2\alpha}{n}\Phi_{> n/3} \\ (2.9) \quad & \leq -\frac{\alpha\epsilon}{2n}\Phi + \frac{3\alpha}{n}\Phi_{> n/3} \end{aligned}$$

where in the next to last inequality we used that for $i \leq n/3$, $p_i \leq \frac{1-4\epsilon}{n}$ and that for given Φ , $\sum p_i e^{\alpha x_i}$ is maximized when \mathbf{p} is uniform.

Thus $\mathbb{E}[\Delta\Phi|x(t)] \geq -\frac{\alpha\epsilon}{4n}\Phi$ implies that

$$\frac{3\alpha}{n}\Phi_{> \frac{n}{3}} \geq \frac{\alpha\epsilon}{4n}\Phi.$$

Let $B = \sum_i \max(0, x_i) = \frac{1}{2}\|x\|_1$. Note that $\Phi_{\geq \frac{n}{3}}$ is upper bounded by $\frac{2n}{3}e^{\frac{3\alpha B}{n}}$. Thus

$$(2.10) \quad \Phi \leq \frac{12}{\epsilon}\Phi_{> \frac{n}{3}} \leq \frac{8n}{\epsilon}e^{\frac{3\alpha B}{n}}.$$

On the other hand, $x_{\frac{3n}{4}} > 0$ implies that $\Psi \geq \frac{n}{4}e^{\frac{4\alpha B}{n}}$. If $\Phi < \frac{\epsilon}{4}\Psi$, we are already done. Otherwise,

$$\frac{8n}{\epsilon}e^{\frac{3\alpha B}{n}} \geq \Phi \geq \frac{\epsilon}{4}\Psi \geq \frac{\epsilon n}{16}e^{\frac{4\alpha B}{n}}$$

so that $e^{\frac{\alpha B}{n}} \leq \frac{128}{\epsilon^2}$. It follows that

$$\Gamma \leq \frac{5}{\epsilon}\Phi \leq \frac{40n}{\epsilon^2}(\frac{128}{\epsilon})^3 \leq cn.$$

Similarly,

LEMMA 2.6. *Suppose that $x_{\frac{n}{4}} < 0$ and $\mathbb{E}[\Delta\Psi|x(t)] \geq -\frac{\alpha\epsilon}{4n}\Psi$. Then either $\Psi < \frac{\epsilon}{4}\Phi$ or $\Gamma < cn$ for some $c = \text{poly}(\frac{1}{\epsilon})$.*

Proof. First observe that for any $i > \frac{2n}{3}$, $p_i > \frac{1+\epsilon}{n}$ so that $p_i(-\alpha + S\alpha^2) + (\frac{\alpha}{n} + S\frac{\alpha^2}{n^2}) \leq -\frac{\alpha\epsilon}{2n}$. Since $p_i \geq 0$ it holds that $p_i(-\alpha + S\alpha^2) + (\frac{\alpha}{n} + S\frac{\alpha^2}{n^2}) \leq \frac{2\alpha}{n}$ for every i . Using the upper bound from (2.6) we get

$$\begin{aligned} \mathbb{E}[\Delta\Psi \mid x(t)] & \leq -\frac{\alpha\epsilon}{2n}\Psi_{> \frac{2n}{3}} + \frac{2\alpha}{n}\Psi_{\leq \frac{2n}{3}} \\ & = -\frac{\alpha\epsilon}{2n}\Psi + \frac{4\alpha + \alpha\epsilon}{2n}\Psi_{\leq \frac{2n}{3}} \\ & \leq -\frac{\alpha\epsilon}{2n}\Psi + \frac{3\alpha}{n}\Psi_{\leq \frac{2n}{3}}. \end{aligned}$$

Thus $\mathbb{E}[\Delta\Psi \mid x(t)] \geq -\frac{\alpha\epsilon}{4n}\Psi$ implies that

$$\frac{3\alpha}{n}\Psi_{\leq \frac{2n}{3}} \geq \frac{\alpha\epsilon}{4n}\Psi.$$

Let $B = \sum_i \max(0, x_i) = \frac{1}{2}\|x\|_1$. Note that $\Psi_{\leq \frac{2n}{3}}$ is upper bounded by $\frac{2n}{3}e^{\frac{3\alpha B}{n}}$. Thus

$$(2.11) \quad \Psi \leq \frac{12}{\epsilon}\Psi_{\leq \frac{2n}{3}} \leq \frac{8n}{\epsilon}e^{\frac{3\alpha B}{n}}.$$

On the other hand, $x_{\frac{n}{4}} < 0$ implies that $\Phi \geq \frac{n}{4}e^{\frac{4\alpha B}{n}}$. If $\Psi < \frac{\epsilon}{4}\Phi$, we are already done. Otherwise,

$$\frac{8n}{\epsilon}e^{\frac{3\alpha B}{n}} \geq \Psi \geq \frac{\epsilon}{4}\Phi \geq \frac{n\epsilon}{16}e^{\frac{4\alpha B}{n}}$$

so that $e^{\frac{\alpha B}{n}} \leq \frac{128}{\epsilon^2}$. It follows that

$$\Gamma \leq \frac{5}{\epsilon}\Psi \leq \frac{40n}{\epsilon^2}(\frac{128}{\epsilon})^3 \leq cn.$$

We are now ready to prove the supermartingale-type property of Γ .

THEOREM 2.1. *Let Γ be as above. Then $E[\Gamma(t+1) \mid x(t)] \leq (1 - \frac{\alpha\epsilon}{4n})\Gamma(t) + c$, for a constant $c = c(\epsilon) = \text{poly}(\frac{1}{\epsilon})$.*

Proof. The proof proceeds via a case analysis. In case the conditions, $x_{\frac{n}{4}} \geq 0$ and $x_{\frac{3n}{4}} \leq 0$ hold, we show both Φ and Ψ decrease in expectation. If one of these is violated Lemmas 2.5 and 2.6 come to the rescue.

Case 1: $x_{\frac{n}{4}} \geq 0$ and $x_{\frac{3n}{4}} \leq 0$. In this case the theorem follows from Lemmas 2.3 and 2.4.

Case 2: $x_{\frac{n}{4}} \geq x_{\frac{3n}{4}} > 0$. Intuitively, this means that the allocation is very non symmetric with big holes in the less loaded bins. While Φ may sometimes grow in expectation, we will show that if that happens, then the asymmetry implies that Γ is dominated by Ψ which decreases. Thus the decrease in Ψ offsets the increase in Φ and the expected change in Γ is negative.

Formally, if $\mathbb{E}[\Delta\Phi|x] \leq -\frac{\alpha\epsilon}{4n}\Phi$, Lemma 2.4 implies the result. Otherwise, by Lemma 2.5 there are two subcases:

Case 2.1: $\Phi < \frac{\epsilon}{4}\Psi$. In this case, using Lemma 2.4 and Corollary 2.1

$$\begin{aligned} \mathbb{E}[\Delta\Gamma|x] &= \mathbb{E}[\Delta\Phi|x] + \mathbb{E}[\Delta\Psi|x] \\ &\leq \frac{2\alpha}{n}\Phi - \frac{\alpha\epsilon}{n}\Psi + 1 \leq -\frac{\alpha\epsilon}{2n}\Psi + 1 \leq -\frac{\alpha\epsilon}{4n}\Gamma + 1 \end{aligned}$$

Case 2.2: $\Gamma < cn$. In this case, Corollaries 2.1 and 2.2 imply that

$$\mathbb{E}[\Delta\Gamma|x] \leq \frac{2\alpha}{n}\Gamma \leq 2c\alpha.$$

On the other hand, $c - \frac{\alpha\epsilon}{4n}\Gamma \geq c(1 - \frac{\alpha\epsilon}{4}) > 2c\alpha$.

Case 3: $x_{\frac{3n}{4}} \leq x_{\frac{n}{4}} < 0$. This case is similar to case 2. If $\mathbb{E}[\Delta\Psi|x] \leq -\frac{\alpha\epsilon}{4n}\Psi$, Lemma 2.3 implies the result. Otherwise, by Lemma 2.6 there are two subcases:

Case 3.1: $\Psi < \frac{\epsilon}{4}\Phi$. In this case, using Lemma 2.3 and Corollary 2.2, the claim follows.

Case 3.2: $\Gamma < cn$. This case is the same as case 2.2.

Once we have shown that Γ decreases in expectation when large, we can use that to bound the expected value of Γ .

THEOREM 2.2. For any $t \geq 0$, $\mathbb{E}[\Gamma(t)] \leq \frac{4c}{\alpha\epsilon}n$.

Proof. We show the claim by induction. For $t = 0$, it is trivially true. By Theorem 2.1, we have

$$\begin{aligned} E[\Gamma(t+1)] &= E[E[\Gamma(t+1) | \Gamma(t)]] \\ &\leq E[(1 - \frac{\alpha\epsilon}{4n})\Gamma(t) + c] \\ &\leq \frac{4c}{\alpha\epsilon}n(1 - \frac{\alpha\epsilon}{4n}) + c \\ &\leq \frac{4c}{\alpha\epsilon}n - c + c \end{aligned}$$

The claim follows.

Finally let $Gap(t)$ denote the gap between the maximum loaded bin and the average, i.e. $x_1(t)$. Then we have

COROLLARY 2.3. $\mathbb{E}[Gap(t)] \leq \frac{\log n}{\alpha} + O(\frac{\log(1/\alpha\epsilon)}{\alpha})$. Moreover, $\Pr[Gap(t) > \frac{2\log n}{\alpha} + O(\frac{\log(1/\alpha\epsilon)}{\alpha})] \leq \frac{1}{n}$.

Proof. Note that $\Gamma(t) \geq e^{\alpha Gap(t)}$. Since $e^{\alpha x}$ is convex:

$$\begin{aligned} \mathbb{E}[Gap(t)] &\leq \frac{1}{\alpha} \log \mathbb{E}[\Gamma(t)] \\ &\leq \log n/\alpha + \log(4c/\alpha\epsilon)/\alpha \\ &= \log n/\alpha + O(\log(1/\alpha\epsilon)/\alpha). \end{aligned}$$

Similarly, $\Pr[Gap(t) > 2\log n/\alpha + O(\log(1/\alpha\epsilon)/\alpha)] \leq \Pr[\Gamma(t) \geq n\mathbb{E}[\Gamma(t)]] \leq 1/n$.

Since the $(1 + \beta)$ -choice process satisfies the conditions on \mathbf{p} with $\epsilon = \Theta(\beta)$, specializing the above result to the unweighted case, we get:

COROLLARY 2.4. For the $(1 + \beta)$ -choice process with unweighted balls, $Gap(t)$ is in $O(\log(n/\beta)/\beta)$ with high probability, for any t .

Let us check the bounds the above results give for some simple distributions.

EXAMPLE 2. Let \mathcal{D} be the distribution that is zero with probability $(1 - \frac{1}{K})$ and K with probability $\frac{1}{K}$. For this distribution, we can check that $M''(\frac{1}{K})$ is in $O(K)$, so that we can set $\alpha = \Theta(\frac{\epsilon}{K})$. This leads to a gap of $O(K \log n/\epsilon)$, which is tight up to constants.

EXAMPLE 3. Let \mathcal{D} be the exponential distribution with parameter 1. Then $M''(1/2)$ is finite. Thus the gap is $O(\log n/\epsilon)$. This is tight up to constants.

Moreover, note that the distribution may be different in each time step, as long as the independence is preserved and the bound on the moment generating function holds.

2.1 A Note on Majorization In some cases it is possible to say that one process dominates another. Establishing a partial order of stochastic dominance extends the analysis for a larger set of processes. Let p and q be two vectors. We say that p is majorized by q denoted by $p \preceq q$ if for every i ,

$$\sum_{j \leq i} p_{\pi(j)} \leq \sum_{j \leq i} q_{\sigma(j)}$$

where π is a permutation sorting p in decreasing order and σ is a permutation sorting q in decreasing order.

A function $\gamma : \mathbb{R}^n \rightarrow \mathbb{R}$ is said to be Schur-convex if $p \preceq q$ implies $\gamma(p) \leq \gamma(q)$. For two distributions over

vectors x and y , we say that x is *majorized* by y , written as $x \preceq y$ if for every Schur-convex function γ it holds that $\mathbb{E}[\gamma(x)] \leq \mathbb{E}[\gamma(y)]$.

The following theorem is implicit in [1] and [8], it establishes a way of comparing different processes in the case of uniform weights.

THEOREM 2.3. *Let \mathbf{p} and \mathbf{q} be probability vectors associated with processes $x(\cdot)$ and $y(\cdot)$ respectively, and assume all balls are of weight 1. If $\mathbf{q} \preceq \mathbf{p}$ then for every t , $x(t) \preceq y(t)$.*

This enables us to extend our results to a larger family of processes. Consider for example the bias-away-from-max and bias-towards-min processes, defined as follows:

$$p_i^{\text{away-from-max}} = \begin{cases} \frac{1-\eta}{n} & i = 1 \\ \frac{1}{n} + \frac{\eta}{n(n-1)} & \text{otherwise} \end{cases}$$

and

$$p_i^{\text{towards-min}} = \begin{cases} \frac{1+\eta}{n} & i = n \\ \frac{1}{n} - \frac{\eta}{n(n-1)} & \text{otherwise} \end{cases}$$

It is easy to check that both these processes are majorized by the $(1 + \frac{\eta}{4n})$ -choice process. Thus we conclude that the expected gap for these processes is bounded by $O(n \log(n/\eta)/\eta)$. Thus even a small bias suffices for the gap to stay bounded independent of m . We note that it is not obvious how to apply the coupling techniques in [2, 7] to these processes.

3 The Lower Bound

We show that for ϵ bounded away from 1, our analysis is essentially tight up to constants. Consider the $(1 + \beta)$ -choice process, and recall that $\epsilon \in \Theta(\beta)$.

Throw $\frac{an \log n}{\beta^2}$ balls into n bins using the $(1 + \beta)$ choice process, where a is to be defined later. The expected load is thus $\frac{a \log n}{\beta^2}$. The expected number of balls thrown using one choice is $\frac{an(1-\beta) \log n}{\beta^2}$. Raab and Steger [6] show that the load of the most loaded bin when throwing $cn \log n$ balls into n bins using 1-choice is at least $(c + (\sqrt{c}/10)) \log n$. Plugging in $c = \frac{a(1-\beta)}{\beta^2}$, we get that the number of balls in most loaded bin is at least

$$\begin{aligned} & \left(\frac{a(1-\beta)}{\beta^2} + \sqrt{\frac{a(1-\beta)}{100\beta^2}} \right) \log n \\ & = \left(\frac{a}{\beta^2} + \frac{\sqrt{a(1-\beta)} - a}{10\beta} \right) \log n. \end{aligned}$$

It is easy to see that for $a < \frac{(1-\beta)}{2}$, this is $\Omega((1 - \beta) \log n / \beta)$ more than the average.

Acknowledgment We thank Ben Birnbaum and Tonci Antunovic for carefully reading an earlier version of the paper.

References

- [1] Yossi Azar, Andrei Broder, Anna Karlin, and Eli Upfal. Balanced allocations. *SIAM J. Computing*, 29(1):180–200, 1999.
- [2] Petra Berenbrink, Artur Czumaj, Angelika Steger, and Berthold Vöcking. Balanced allocations: The heavily loaded case. *SIAM J. Computing*, 35(6):1350–1385, 2006.
- [3] Richard M. Karp, Michael Luby, and Friedhelm Meyer auf der Heide. Efficient pram simulation on a distributed memory machine. In *STOC '92: Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 318–326, New York, NY, USA, 1992. ACM.
- [4] Michael Mitzenmacher. *The Power of Two Choices in Randomized Load Balancing*. PhD thesis, Harvard University, 1991.
- [5] Michael Mitzenmacher, Andréa W. Richa, and Ramesh Sitaraman. The power of two random choices: A survey of techniques and results. In *in Handbook of Randomized Computing*, pages 255–312. Kluwer, 2000.
- [6] Martin Raab and Angelika Steger. "balls into bins" - a simple and tight analysis. In *RANDOM '98: Proceedings of the Second International Workshop on Randomization and Approximation Techniques in Computer Science*, pages 159–170, London, UK, 1998. Springer-Verlag.
- [7] Kunal Talwar and Udi Wieder. Balanced allocations: the weighted case. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 256–265, New York, NY, USA, 2007. ACM.
- [8] Udi Wieder. Balanced allocations with heterogenous bins. In *SPAA '07: Proceedings of the nineteenth annual ACM symposium on Parallel algorithms and architectures*, pages 188–193, New York, NY, USA, 2007. ACM.