

Balanced Allocations: The Weighted Case

Kunal Talwar
Microsoft Research Silicon Valley
kunal@microsoft.com

Udi Wieder
Microsoft Research Silicon Valley
uwieder@microsoft.com

June 15, 2008

Abstract

We investigate balls-and-bins processes where m weighted balls are placed into n bins using the “power of two choices” paradigm, whereby a ball is inserted into the less loaded of two randomly chosen bins. The case where each of the m balls has unit weight had been studied extensively. In a seminal paper Azar *et al.* [2] showed that when $m = n$ the most loaded bin has $\Theta(\log \log n)$ balls with high probability. Surprisingly, the gap in load between the heaviest bin and the average bin does not increase with m and was shown by Berenbrink *et al.* [4] to be $\Theta(\log \log n)$ with high probability for arbitrarily large m . We generalize this result to the weighted case where balls have weights drawn from an arbitrary weight distribution. We show that as long as the weight distribution has finite second moment and satisfies a mild technical condition, the gap between the weight of the heaviest bin and the weight of the average bin is independent of the number balls thrown. This is especially striking when considering heavy tailed distributions such as Power-Law and Log-Normal distributions. In these cases, as more balls are thrown, heavier and heavier weights are encountered. Nevertheless with high probability, the imbalance in the load distribution does not increase. Furthermore, if the fourth moment of the weight distribution is finite, the expected value of the gap is shown to be independent of the number of balls.

1 Introduction

Suppose m balls are to be put one by one into n bins such that the final allocation is as balanced as possible. The well-known ‘power of two choices’ algorithm, a.k.a GREEDY[2], inserts each ball into the less loaded among two randomly chosen bins. The case where all balls are of uniform weight had been extensively studied. In a seminal paper Azar *et al.* [2] showed that when $m = n$ the heaviest bin has $\ln \ln n / \ln 2 + O(1)$ balls, compared with $(1 + o(1)) \ln n / \ln \ln n$ when the 1-choice algorithm is used. Berenbrink *et al.* [4] have shown that when $m \gg n$, with probability $1 - o(1)$ the heaviest bin has at most $m/n + O(\ln \ln n)$ balls, compared with $m/n + \Omega(\sqrt{m \log n/n})$ for the one choice algorithm. Note that the additive gap between the maximum load and the average load does not depend on the number of balls thrown! The two choice paradigm had been investigated in a variety of models and scenarios, see [12] for a survey.

Our goal in this paper is to prove similar bounds for the case where balls are weighted. In our model there is a weight distribution \mathcal{W} . In each round i a weight w_i is sampled from \mathcal{W} . A ball of weight w_i is then inserted into the less loaded of two uniformly sampled bins. The main contribution of this paper is to show that under reasonable assumption on the weight distribution, the additive gap between the loads of the maximum and average bins is not a function of m , but depends only on n and the weight distribution.

The idea of allowing two (or more) choices to improve load balancing is known to be useful in many contexts and spawned a large body of literature (c.f. [12] and the references therein). The two most common applications are probably hashing and online load balancing. While the assumption of uniform weights is justified when hashing is considered, it is often the case in load balancing scenarios that the elements to be balanced are many and vary in weight. Consider for instance a distributed storage system in which there are n servers and whenever a data item is to be inserted into the system it is assigned to the less loaded among two random servers¹ (c.f. [7] [8]). It is known that many types of data items such as files in a PC file system and multimedia files have sizes distributed by a heavy tailed distribution [10]. While splitting the files into fixed-sized chunks is a natural way to reduce the problem to the uniform-weights case, it introduces failure dependencies and increases the lookup cost. If data items are not split, the online load balancing algorithm must accommodate variable weights. The same line of argument holds for other types of resources such as computational load, bandwidth etc.

The $m = n$ Case

Assume for simplicity the weight distribution is identical in all rounds. Denote by $M := \max\{x : \Pr[w \geq x] \geq \frac{1}{n}\}$. M is a natural lower bound on the weight of the heaviest bin, since when throwing n balls, with constant probability a ball of weight at least M is encountered. It turns out that for most interesting distributions this lower bound is tight up to constants. Consider for instance the case where weights come from the Geometric distribution. The largest ball is of weight $\Omega(\log n)$ with high probability, while the expected load on each bin is $O(1)$. Thus any allocation algorithm, including GREEDY[2] will have a gap of $\Omega(\log n)$. Yet, the sum of $\log n / \log \log n$ independent Geometric variables is $O(\log n)$ with high probability. Thus, even if the weight-oblivious one-choice paradigm is used, the maximum bin would still have a weight of $O(\log n)$. We conclude that in this case GREEDY[2] does not perform significantly better than GREEDY[1]. Clearly, the same argument holds for distributions which are more heavy tailed. It may well be that constants and low order terms are improved when the two choice algorithm is used. Nevertheless, the domain where the two choice algorithm is fundamentally different is where the number of balls is much larger than the number of bins.

¹In practice the two servers are often chosen by hashing the data item’s identifier.

The Heavily Loaded Case - A Toy Example

Consider the case where $m \gg n$ and the weight distribution assigns a weight of 1 with probability $\frac{1}{2}$ and a weight of 2 with probability $\frac{1}{2}$. This seems like a simple case which should somehow be reduced to the unweighted case. Yet in order for the allocation to be balanced, an allocation algorithm must take the weights of balls into consideration. Indeed, in a weight-oblivious allocation scheme, even if each bin receives m/n balls, the weight distribution may cause the allocation to be unbalanced. The weight distribution alone allows the weight of m/n random balls to be $\frac{3m}{2n} \pm \sqrt{\frac{m}{n}}$. Thus, if GREEDY[2] considers the number of balls in each bin and not their weights then the obtained allocation is not much better than the one obtained by the one-choice algorithm! The behavior of a weight-oblivious algorithm deteriorates as the weight distribution becomes more heavy tailed. It is therefore essential to analyze the weighted case separately.

1.1 Related Work

As mentioned previously, the unweighted case had been the object of extensive study in many contexts (c.f. [14],[1], [11] [12]) and to a large extent is well understood. Vöcking [13] proved the surprising result that a similar process with an asymmetric tie breaking rule called GLEFT obtains a better bound in the $m = n$ case, and is majorized by GREEDY[2] in the heavily loaded case ([4]). In the weighted case it may be extremely unlikely that a tie is ever encountered², thus we do not investigate the GLEFT algorithm.

Little is known about the weighted case. Berenbrink *et al.* show in [5] that if the weights are arbitrary then several natural conjectures turn out to be false. In particular they show that replacing two balls of different weights by two balls with weight equal to their average does not necessarily improve the balance of the allocation. They also show that the majorization order is not preserved under weighted balls. Majorization is a partial order which captures the degree to which an allocation is well balanced. It is commonly used for showing stochastic dominance of one balls-and-bins process over another, see [5] for more details. The break of the majorization order implies that the known techniques are unable to reduce the weighted case to an unweighted instance even for simple weight distributions. Indeed, we do not even know how to show directly that GREEDY[2] is more balanced than GREEDY[1], a statement which is trivially true in the unweighted case. The weighted case had been investigated in a somewhat different model, where the balls arrive in parallel and are allowed to communicate with one another prior to making the allocation decision [1],[3]. In this model, the weights are arbitrary, but the additive gap may be large.

1.2 Our Contributions

Our main result shows that under mild assumptions on the weight distribution, the differences from average in the allocation after m steps are essentially distributed the same as the differences in the allocation after $poly(n)$ steps. Thus the difference in load between the heaviest and the lightest bin is *independent* of the number of balls thrown.

Define $Gap(t)$ to be the excess weight the heaviest bin has over the lightest at time t using GREEDY[2] Assume the expectation and variance of the weight distribution are finite. We also assume the distribution \mathcal{W} is ‘smooth’³. The following two theorems are the main contribution of this paper.

Theorem 1.1. *For any t , the probability that $Gap(t) > k$ is at most $Pr[gap(n^c) > k] + \frac{1}{n^c}$ where c is some constant depending on \mathcal{W} alone.*

²Our result holds also when the weight distribution is continuous, in which case the probability of a tie may be 0.

³The exact definition of ‘smooth’ is deferred to Section 3.1. At this point it suffices to say that the definition covers most natural distributions. Section 6.1 discusses some distributions excluded by this definition.

Theorem 1.2. *If the fourth moment of the weight distribution \mathcal{W} is finite, then for any t , $\mathbb{E}[\text{Gap}(t)] \leq n^c$ where c is some constant depending on \mathcal{W} alone.*

Theorem 1.1 reduces the case where m is arbitrary large to the case where $m \leq n^c$. Theorem 1.2 shows that the expected value of the gap does not increase with m . It is conceivable that the case where $m \leq n^c$ can be analyzed, possibly for specific weight distributions, using known techniques such as layered induction. In [4], it is shown that in the unweighted case the gap in the polynomial case is $\Theta(\log \log n)$. It will be interesting to find tight bounds on the gap for other weight distributions.

We stress that the weight distribution is not required to be over the integers and may take its values over the reals.

1.3 The Outline of the Proof

The general outline of the proof draws from the work of Berenbrink *et al.* [4]. There are many obstacles when applying the technique to the weighted case, which requires many new ideas, some of which may be of interest in their own right. It is shown in [4] that two theorems are needed in order to prove Theorem 1.1.

The first is a **weak gap theorem** which proves that w.h.p $\text{Gap}(t) \leq t^{2/3}$. In the unweighted case the weak gap theorem is trivial and follows immediately from the fact that GREEDY[2] is majorized by the one choice algorithm. As mentioned previously, in the weighted case it is not the case that GREEDY[2] is dominated by the one choice algorithm. We therefore prove the weak gap Theorem in Section 2 via a potential function argument.

The second theorem needed is a **short memory theorem**. In this theorem we show that given some initial configuration with gap Δ , after adding $\Delta \text{poly}(n)$ more balls the initial configuration is ‘forgotten’. In the unweighted case the short memory theorem is proven via coupling. Our proof uses similar coupling arguments but is considerably more involved technically. In particular, we need to define a somewhat different distance function and use a sophisticated argument to show that the coupling converges. The short memory theorem is proved in Section 3

Theorem 4.2 in [4] proves that a weak gap theorem and a short memory theorem implies a stronger theorem such as Theorem 1.1. For the sake of completeness, we present a somewhat simplified version of the proof in Section 4. The proof in Section 3 assumes that the weight distribution is over the integers. In Section 5 we show a reduction from the real-weighted case to the integer-weighted case. The reduction turns out to be non-trivial and requires the introduction of dependencies between the weights of balls.

1.4 Basic Notations and Definitions

We model the state of the system by *load vectors*. A load vector $x = (x_1, x_2, \dots, x_n)$ specifies the load in each bin where x_i specifies the total weight of all balls assigned to bin i . We assume that vectors are *normalized*, i.e. that $x_1 \geq x_2 \geq \dots \geq x_n$. Note that after an insertion of a ball the order may change and we may need to rename the bins. Denote by β_i the probability bin i is chosen, so that $\beta_i = \frac{i^2 - (i-1)^2}{n^2} = \frac{2i-1}{n^2}$. In each step of the process a weight W is sampled from the distribution \mathcal{W} and a ball of weight W is put in bin i with probability β_i .

Denote by $x(t)$ the allocation after throwing t balls. The random process $(x(t))_{t \in \mathbb{N}}$ is therefore a Markov chain with transition probabilities defined by the allocation rule and the weight distribution. For two random variables $x(t), y(t)$ we abuse notation and write $\|x(t) - y(t)\|$ for the *variation distance* between their respective distributions.

2 The Weak Gap Theorem

In this section we prove that the gap after throwing t balls is at most $t^{2/3}$ w.h.p. Such a weak gap is then iteratively sharpened using the short memory theorem of Section 3 to obtain the main result. Recall that $Gap(t)$ denotes the excess weight the heaviest bin has over the lightest one at time t . Denote by M_2 the second moment of the weight distribution \mathcal{W} .

Theorem 2.1. *For all t and every $k > 0$, it holds that $\Pr[Gap(t) \leq 2\sqrt{ktM_2}] \geq 1 - \frac{1}{k}$. In particular, $\Pr[Gap(t) \leq 2t^{\frac{2}{3}}\sqrt{M_2}] \geq 1 - \frac{1}{t^{\frac{1}{3}}}$*

A routine use of Chebyshev's inequality would prove Theorem 2.1 for the case where each ball is thrown independently to a random location. It is tempting to claim that GREEDY[2] must do better and indeed that would be correct in the unweighted case. However, when weights are introduced it is no longer true that GREEDY[2] dominates the one-choice algorithm, hence a different argument should be used:

Proof. Let $x(t)$ be the normalized vector at time t and $\bar{x}(t) = \frac{1}{n} \sum_i x(t)_i$ be the average load at time t . Define $V(t)$ to be the variance of the allocation at time t , i.e., $V(t) := \sum_i (x(t)_i - \bar{x}(t))^2$.

Lemma 2.2. $\mathbb{E}[V(t+1) - V(t) \mid V(t)] \leq M_2$ where the expectation is taken over the random choices of the algorithm and the weight distribution.

Proof. First we calculate the expectation given that the weight of the ball at time t is w . Recall that β_i denotes the probability the ball is put in the i 'th bin. Denote by $\delta_{i,j}$ the function which is 1 iff $i = j$ and 0 otherwise.

$$\begin{aligned}
& \mathbb{E}[V(t+1) - V(t) \mid V(t), w] \\
&= \sum_i \beta_i \sum_j (x_j + \delta_{i,j}w - \bar{x}(t+1))^2 - \sum_j (x_j - \bar{x}(t))^2 \\
&= \sum_i \beta_i \left[\sum_j (x_j + \delta_{i,j}w - \frac{w}{n} - \bar{x}(t))^2 - \sum_j (x_j - \bar{x}(t))^2 \right] \\
&= \sum_i \beta_i \left[\sum_j (\delta_{i,j}w - \frac{w}{n})(\delta_{i,j}w - \frac{w}{n} + 2x_j - 2\bar{x}(t)) \right] \\
&= \sum_i \beta_i \left[\sum_j (\delta_{i,j}w)^2 - \frac{2\delta_{i,j}w^2}{n} + (\frac{w}{n})^2 + 2\delta_{i,j}wx_j - 2\delta_{i,j}w\bar{x}(t) - \frac{2x_jw}{n} + \frac{2w\bar{x}(t)}{n} \right] \\
&= \frac{w^2}{n} + 2w\bar{x}(t) + \sum_i \beta_i \left[\sum_j (\delta_{i,j}w)^2 - \frac{2\delta_{i,j}w^2}{n} + 2\delta_{i,j}wx_j - 2\delta_{i,j}w\bar{x}(t) - \frac{2x_jw}{n} \right] \\
&= \frac{w^2}{n} + 2w\bar{x}(t) - \sum_i \beta_i \sum_j \frac{2x_jw}{n} + \sum_i \beta_i (w^2 + 2wx_i - \frac{2w^2}{n} - 2w\bar{x}(t))
\end{aligned}$$

Now, $\sum_i \beta_i \sum_j \frac{2x_jw}{n} = 2w\bar{x}(t)$ therefore we have

$$\begin{aligned}
&= \frac{w^2}{n} + \sum_i \beta_i (w^2 + 2wx_i - \frac{2w^2}{n} - 2w\bar{x}(t)) \\
&= w^2 - \frac{w^2}{n} - 2w\bar{x}(t) + \sum_i \beta_i 2wx_i
\end{aligned}$$

Since $\sum \beta_i x_i$ is a weighted average of the x_i 's which is *biased* towards the smaller elements we have that $\sum \beta_i x_i \leq \bar{x}(t)$. We conclude that:

$$\mathbb{E}[V(t+1) - V(t) \mid V(t), w] \leq w^2 - \frac{w^2}{n}$$

We have that for every ball weight w it holds that

$$\mathbb{E}[V(t+1) - V(t) \mid V(t), w] \leq w^2$$

The Lemma is proven by taking the expectation over w on both sides. □

It holds that $\mathbb{E}[V(t)] \leq tM_2$ so by Markov's inequality $\Pr[V(t) \geq ktM_2] \leq \frac{1}{k}$. It also holds that

$$\begin{aligned}
\text{Gap}(t) &\leq \max_i x(t)_i - \min_i x(t)_i \\
&\leq 2 \max_i |x(t)_i - \bar{x}(t)| \\
&\leq 2\sqrt{V(t)}
\end{aligned}$$

therefore we have that with probability $1 - \frac{1}{k}$, $\text{Gap}(t) \leq 2\sqrt{ktM_2}$. The second part follows by substituting $k = t^{\frac{1}{3}}$. □

Define M_s to be the s 'th moment of \mathcal{W} ; i.e. $M_s := \sum_w w \cdot p_w^s$ where p_w is the probability w is sampled. Note that since the weight distribution is non-negative M_s is well defined for any $s > 0$. Now, if M_4 is finite, Markov's inequality can be applied on a higher moment deriving a stronger bound. A stronger bound would be needed to prove Theorem 1.2.

Lemma 2.3. *Suppose that M_4 is finite. Then there is a constant $c = c(M_2, M_4)$ such that for every t it holds that $\Pr[\text{Gap}(t) \leq ct^{\frac{4}{5}}] \geq 1 - \frac{1}{t^{\frac{6}{5}}}$.*

Proof. We shall first upper bound $E[(V(t))^2]$. Towards this end, we bound the increments $E[V(t+1)^2 - V(t)^2]$. First observe that $\sum_i (x(t+1)_i - z)^2$ is minimized when $z = \bar{x}(t+1)$, therefore $V(t+1) = \sum_i (x(t+1)_i - \bar{x}(t+1))^2 \leq \sum_i (x(t+1)_i - \bar{x}(t))^2$. Denote this latter expression by $\hat{V}(t+1)$. Then we have:

$$\begin{aligned}
&\mathbb{E}[V^2(t+1) - V^2(t) \mid V(t), w] \\
&\leq \mathbb{E}[\hat{V}^2(t+1) - V^2(t) \mid V(t), w] \\
&= \mathbb{E}[(\hat{V}(t+1) - V(t))^2 \mid V(t), w] + 2\mathbb{E}[V(t)(\hat{V}(t+1) - V(t)) \mid V(t), w] \tag{1}
\end{aligned}$$

We now proceed to bound the two terms of Equation (1) separately. First observe that

$$\begin{aligned}
& \mathbb{E} \left[(\hat{V}(t+1) - V(t)) \mid V(t), w \right] \\
&= \sum_i \beta_i \left(\sum_j (x_j + \delta_{i,j} w - \bar{x}(t))^2 - \sum_j (x_j - \bar{x}(t))^2 \right) \\
&= \sum_i \beta_i ((x_i + w - \bar{x}(t))^2 - (x_i - \bar{x}(t))^2) \\
&= \sum_i \beta_i w^2 + \sum_i \beta_i 2w(x_i - \bar{x}(t)) \\
&\leq w^2
\end{aligned}$$

The first term is bound as follows,

$$\begin{aligned}
& \mathbb{E} \left[(\hat{V}(t+1) - V(t))^2 \mid V(t), w \right] \\
&= \sum_i \beta_i \left(\sum_j (x_j + \delta_{i,j} w - \bar{x}(t))^2 - \sum_j (x_j - \bar{x}(t))^2 \right)^2 \\
&= \sum_i \beta_i ((x_i + w - \bar{x}(t))^2 - (x_i - \bar{x}(t))^2)^2 \\
&= \sum_i \beta_i (w^2 + 2w(x_i - \bar{x}(t)))^2 \\
&= \sum_i \beta_i w^4 + \sum_i \beta_i 4w^2(x_i - \bar{x}(t))^2 + \sum_i \beta_i 4w^3(x_i - \bar{x}(t)) \\
&\leq w^4 + 4w^2 V(t)
\end{aligned}$$

since $\beta_i < 1$ and $\sum_i \beta_i (x_i - \bar{x}(t)) \leq 0$.

Plugging these bounds into Equation (1) we have that

$$\mathbb{E} [V^2(t+1) - V^2(t) \mid V(t), w] \leq w^4 + 6w^2 V(t)$$

Taking expectation over w we conclude that

$$\mathbb{E} [V^2(t+1) - V^2(t) \mid V(t)] \leq M_4 + 6M_2 V(t)$$

and thus combined with Lemma 2.2

$$\mathbb{E} [V^2(t+1) - V^2(t)] \leq M_4 + 6t(M_2)^2.$$

Thus $\mathbb{E} [V^2(t)] \leq tM_4 + 3t(t+1)(M_2)^2$.

Finally

$$Pr[Gap(t) > 2k] \leq Pr[(V^2(t) \geq k^4] \leq \frac{\mathbb{E}[V^2(t)]}{k^4} \leq \frac{ct^2}{k^4}$$

where c depends only on M_2 and M_4 . Plugging in $k = c^{\frac{1}{4}} t^{\frac{4}{5}}$ completes the proof of Lemma 2.3. \square

3 The Short Memory Theorem

In this section we generalize Lemma 1.2 in [4] to the case of weighted balls. In Section 3.1 we assume the weight distribution \mathcal{W} is over the integers. Restricting \mathcal{W} to be over the integers allows us to use the Neighbor Coupling approach of [4] which simplifies the proof. In Section 5 we prove the more general case where \mathcal{W} is over the reals.

Definition 3.1. Let p_i denote the probability i is sampled by \mathcal{W} . For every $\ell > 0$ define δ_ℓ to be \max_i s.t. $p_i \geq \frac{1}{\ell}$. For $\alpha > 0$ let $\mathcal{W}(\alpha)$ denote the sum of α independent samples from \mathcal{W} and $\mathcal{W}(0)$ is concentrated at 0. A distribution \mathcal{W} is said to be (α, β) -smooth if for every ℓ and for every $i < \delta_\ell$ there exists $\alpha_0, \alpha_1 \geq 0$ such that $\alpha_0 + \alpha_1 \leq \alpha$ and $\Pr[\mathcal{W}(\alpha_0) - \mathcal{W}(\alpha_1) = i] \geq \left(\frac{1}{\delta_\ell}\right)^\beta$.

Lemma 3.2. If there exists k such that p_i is non-increasing in $[k, \infty]$ then \mathcal{W} is smooth.

Proof. Let S denote all i for which $p_i > 0$. If $|S|$ is finite then w.l.o.g the greatest common divider of all elements in S is 1, otherwise normalize by the gcd . There is therefore a choice of $w_1 \dots w_\alpha \in S$ such that $\sum a_j w_j = 1$ and repeating this sequence i times and noticing that δ_ℓ is no bigger than the largest element in S implies smoothness.

Now assume that $|S| = \infty$. By assumption it holds that $p_k > 0$ and $p_{k+1} > 0$, thus for $i < k$ we use the previous argument and the fact that $gcd(k, k+1) = 1$. If $k \leq i < \delta_\ell$ then by assumption $p_i > \frac{1}{\ell}$. \square

The lemma implies that most natural distributions are smooth, including heavy tailed distributions such as Power-Law and Log-Normal. For many distribution it holds that α and β are small. For instance, it is easily observed that the Geometric distribution and Power-Law distributions are $(1, 1)$ -smooth. In fact, it is rather difficult to come up with a distribution that is not smooth. See Section 6.1 for details.

3.1 The Integer Case

For two vectors x, y define $\gamma_{x,y} := \max_{i,j} \{|x_i - x_j|, |y_i - y_j|\}$. We write γ when the context is clear. The following is the main Theorem of this section.

Theorem 3.3. Let \mathcal{W} be an (α, β) -smooth distribution over the integers, with a finite second moment. Let x, y be any two load vectors such that $\sum_i x_i = \sum_i y_i$. Let $x(t), y(t)$ be the random variables describing the load vector after allocating t more balls. Then, there exists t , with $t = O(\gamma^{7/5} \cdot 2^\alpha n^{6\beta+10} \log^2(n\gamma))$, such that $\|x(t) - y(t)\| \leq (\gamma_{x,y})^{-1/5}$.

The proof will show a coupling between $(x(t))_{t \in \mathbb{N}}$ and $(y(t))_{t \in \mathbb{N}}$ such that $\Pr[x(t) \neq y(t)] \leq (\gamma_{x,y})^{-1/5}$ for t as above. The coupling Lemma would then imply the theorem. As in [4] we use neighbor-coupling, which is a variant of the well known path coupling technique [6]. In the following we define the graph we work with, the distance function and the coupling itself.

3.1.1 The Graph

Recall that a vector $x \in \mathbb{R}^n$ is *normalized* if $x_1 \geq x_2 \geq \dots \geq x_n$. Clearly each configuration of loads in bins corresponds to a normalized vector by setting x_i to be the load in the i 'th most loaded bin. Let Ω be the set of all normalized vectors. The neighbor set $\Gamma \subset \Omega \times \Omega$ is defined as follows: $(x, y) \in \Gamma$ iff there exists $i, j \in [n]$ such that $x = y + e_i - e_j$ where e_i is the vector with 1 at the i 'th location and 0 everywhere else. The graph we use is therefore identical to the one used in [4]. Note that $(x, y) \in \Gamma$ implies that $(y, x) \in \Gamma$.

We therefore think of the graph as undirected. Often when the path coupling technique is used, the neighbor graph spans the entire state space. Denote by Ω_W all the normalized integer vectors with total weight exactly W . We show below that for every W , the sub-graph (Ω_W, Γ) is connected, which suffices in our case.

3.1.2 The Distance Function

Let $(x, y) \in \Gamma$ and assume that $x = y + e_i - e_j$ with $i < j$ (otherwise switch the roles of x and y). We define the distance $\Delta(x, y) := x_i - y_j$. We assign to each edge in $(x, y) \in \Gamma$ the length $\Delta(x, y)$.

Remark In [4] the distance function used is $\gamma_{x,y} := \max\{|x_i - x_j|, |y_i - y_j|\}$. Our distance function has nicer properties (as would be seen below). In particular it is drawn from the following physical intuition: Imagine that per-unit cost of moving an infinitesimal amount of mass from i to j is equal to the height difference between i and j . Thus it costs $x_i - x_j$ at the beginning and the cost decreases as more mass moves from i to j . Our distance function $\Delta(x, y) = x_i - y_j$ then captures exactly the cost of moving one unit of weight from i to j .

Lemma 3.4. *If $x, y \in \Omega_W$, i.e. the total weight of both x and y is W then x and y are connected via a path with at most $n \cdot \gamma_{x,y}$ edges where all nodes along the path belong to Ω_W . Furthermore, if Δ is the length of the longest edge in the path then $\Delta \leq \gamma_{x,y} \leq n\Delta$.*

Proof. A path from x to the completely balanced allocation can be found by repeatedly moving one unit of weight from the smallest bin with above average load to the largest bin with below average load. The fully balanced allocation is obtained after at most $\frac{n\gamma_{x,y}}{2}$ moves. The same holds for y so the path between x and y has at most $n\gamma_{x,y}$ edges, each of which of length at most $\gamma_{x,y}$. □

3.1.3 The Coupling

Recall that β_i denotes the probability a ball falls in the i 'th largest bin. Let $(x, y) \in \Gamma$ and denote by (x', y') the configuration obtained after performing one step. The coupling we use is essentially similar to the one used in [4]. First sample a weight w from the weight distribution. Both x and y would receive a ball of weight w . Then sample a bin to put the ball in x , say the k 'th largest bin. Add the ball to bin k both in x and in y . Note that after the insertion of the ball it may be the case that x or y should be sorted in order for them to maintain the invariant that bins are ordered by decreasing weight.

It is straightforward to verify that this is indeed a valid coupling; i.e. that each vector receives a ball distributed according to \mathcal{W} and that for every k a ball falls in bin k with probability β_k . The following lemma summarizes the properties of the coupling.

Lemma 3.5. *Let $(x, y) \in \Gamma$ such that $x = y + e_i - e_j$ where w.l.o.g $i < j$ and let x', y' be the two vectors obtained after one step of the coupling. The coupling has the following two properties:*

1. *It holds that either $x' = y'$ or $(x', y') \in \Gamma$. In other words, the coupling preserves the neighbor relation.*
2. *If the ball falls in bin i then $\Delta(x', y') = \Delta(x, y) + w$.*
3. *If the ball falls in bin j then $\Delta(x', y') = |\Delta(x, y) - w|$.*
4. *If the ball falls in bin $k \neq i, j$ then $\Delta(x', y') = \Delta(x, y)$.*

Proof. The proof is a straightforward case analysis on the bin k into which the ball is put.

case (1) $1 \leq k < i$. For every ball weight w the insertion of the ball would cause bins $1 \dots k$ to be reordered in the same way both in x and in y . Thus it holds that $x' = y' + e_i - e_j$ and $\Delta' = \Delta$.

case (2) $k = i$. We set as a convention that $x_0 = y_0 = \infty$. Let ℓ be such that $x_{\ell-1} \geq x_i + w > x_\ell$. Since $x_i = y_i + 1$, the weight w is an integer and $x_{\ell-1} = y_{\ell-1}$, it holds that $y_{\ell-1} > y_i + w \geq y_\ell$. The resorting of the vector moves bin i to the ℓ location. It follows that $x' = y' + e_\ell - e_j$ and $\Delta' = x'_\ell - y'_j = x_i + w - y_j = \Delta + w$

case (3) $i < k < j$. Let ℓ be such that $x_{\ell-1} > x_k + w \geq x_\ell$. Since $x_k = y_k$ it holds that $y_{\ell-1} \geq y_k + w \geq y_\ell$. The normalization of the vector moves bin k to the ℓ location. If $\ell \geq i$ then $x' = y' + e_{i+1} - e_j$ and if $\ell < i$ then $x' = y' + e_i - e_j$. In both cases $\Delta' = \Delta$.

case (4) $k = j$. As in the previous cases let ℓ be such that $x_{\ell-1} \geq x_k + w > x_\ell$ and the normalization of x is done by moving the j 'th bin into location ℓ . Now, if $w \leq \Delta$ then $x_j + w \leq x_j + \Delta = x_i - 1 = y_i$. Therefore, $\ell < i$ and $y_j + w \leq y_\ell$. It follows that $x' = y' + e_i - e_\ell$ and $\Delta' = x'_i - y'_\ell = x_i - (y_j + w) = \Delta - w$. If $w \geq \Delta + 1$ then $x_j + w \geq x_i$ so $\ell \geq i$. Further, $y_j + w > y_\ell$. It follows that $y' = x' + e_\ell - e_{i+1}$ and $\Delta' = y'_\ell - x'_{i+1} = y_j + w - x_i = w - \Delta$.

case (5) $k > j$. Let ℓ be such that $x_{\ell-1} \geq x_k + w > x_\ell$. If $x_k + w \leq x_j$ then $\ell > j$ and also $y_{\ell-1} > y_w + w \geq y_\ell$. It follows that $x' = y' + e_i - e_j$. If $i < \ell \leq j$ then $x' = y' + e_i - e_{j+1}$ and if $\ell \leq i$ then $x' = y' + e_{i+1} - e_{j+1}$. In all cases $\Delta' = \Delta$. \square

Note that $\beta_j \geq \beta_i + \frac{1}{n^2}$ so case (3) above is more likely than case (2). This bias is the reason the chain mixes fast. While couplings often do not preserve the neighbor relation, our coupling does; this is an artifact of the weight distribution being over the integers. This property would later allow us to use the neighbor coupling lemma. We define an *active step* of the coupling to be a step of type (2) or (3) above.

Remark In the unweighted case it is possible to use a different coupling which altogether avoids case (2) above and instead has $x' = x + e_i$ and $y' = y + e_j$. In this coupling it holds that $\Delta' = \Delta - 1$ with probability at least $1/n^2$ and remains Δ otherwise, i.e. the distance never increases. Thus it is possible to prove the unweighted case using a straightforward path coupling argument. This somewhat simplifies the proof in [4].

3.1.4 The Neighbor Coupling Approach

We are now ready to resume with the proof of Theorem 3.3. Recall that x, y are two initial configurations and that Δ denotes the longest edge in the path between x and y .

Lemma 3.6. *If $(x, y) \in \Gamma$ are two initial configurations and $\Delta_0 := \Delta(x, y)$, it holds after t active steps $\Pr[x(t) \neq y(t) | x(0) = x, y(0) = y] \leq \frac{1}{n^{11/5} \Delta_0^{6/5}}$ for $t \in O(2^\alpha n^{6\beta+8} \Delta_0^{7/5})$.*

Denote by D the number of edges in this path and recall that by Lemma 3.4 it holds that $D \leq n\gamma$ where $\gamma = \max\{|x_i - x_j|, |y_i - y_j|\}$ and that $\Delta_0 \leq \gamma \leq n\Delta_0$. It follows that $\frac{1}{n^{11/5} \Delta_0^{6/5}} \leq \frac{1}{D\gamma^{1/5}}$, thus Lemma 3.6 directly implies Theorem 3.3 by union bounding the failure probability of Lemma 3.6 along the D edges on the path from x to y in Γ and by noticing that there are t active steps within $O(tn^2 \log(n\Delta_0))$ steps with probability $1 - \frac{1}{(\Delta_0 n)^3}$.

The remainder of the section is dedicated to the proof of Lemma 3.6. We need to show that after enough steps the value of $\Delta(x(t), y(t))$ decreases to 0. Denote by Δ the current distance and by Δ' the distance after one *active* step of the coupling; note that steps that are not active do not change Δ . If it had been the case that for all Δ it holds that $\mathbb{E}[\Delta'] < \Delta$, then the standard path coupling lemma would have sufficed.

Indeed this is the case when balls have uniform weight. In the weighted case we have some bound Δ^* such that for $\Delta \geq \Delta^*$, $\mathbb{E}[\Delta'] < \Delta$. However, for $\Delta < \Delta^*$, it may be the case that $\mathbb{E}[\Delta'] \geq \Delta$. We overcome this difficulty by showing that when Δ is small we have a $1/\text{poly}(n)$ probability of hitting 0 in the next few steps. In the case that the distance does not hit zero, it does not increase by much and we can repeat the argument.

We start by identifying the threshold above which Δ decreases on expectation. Let p_j denote $\Pr_{\mathcal{W}}[w = j]$. Define

$$\Delta^* := \max_j \text{s.t. } p_j \geq \frac{1}{n^6} \quad (2)$$

Chebyshev's inequality implies that $\Delta^* \leq cn^3$ for a constant c depending only on the distribution.

We prove that $E[\Delta']$ is smaller than Δ for any $\Delta \geq \Delta^*$.

Lemma 3.7. *Denote by Δ the current distance and by Δ' the distance after one active step of the coupling. If $\Delta \geq \Delta^*$, then $\mathbb{E}[\Delta'] \leq \Delta - \frac{\mu}{8n}$.*

Proof. We first show that

$$\sum_{i \geq \Delta^*} ip_i \leq \frac{\mu}{4n} \quad (3)$$

we will then show that (3) implies the lemma. For every integer $A > \Delta^*$ we have

$$\begin{aligned} \sum_{i \geq \Delta^*} ip_i &= \sum_{i=\Delta^*}^{A-1} ip_i + \sum_{i \geq A} ip_i \\ &= \sum_{i=\Delta^*}^{A-1} ip_i + (A-1) \sum_{i \geq A} p_i + \sum_{i \geq A} \sum_{j \geq i} p_j \\ &\leq \frac{A(A-\Delta^*)}{n^6} + (A-1) \frac{c}{A^2} + \sum_{i \geq A} \frac{c}{i^2} \end{aligned}$$

we take $A = \Delta^* + n^{5/4}$

$$\leq \frac{c}{n^{3/2}} + \frac{c}{n^{5/4}} + \frac{c}{n^{5/4}} \leq \frac{\mu}{4n}$$

The first inequality follows from Chebyshev's inequality, and the last step assumes that n is large enough. We now have:

$$\begin{aligned} \mathbb{E}[\Delta'] &= \frac{\beta_i}{\beta_i + \beta_j} \mathbb{E}[\Delta + w] + \frac{\beta_j}{\beta_i + \beta_j} \mathbb{E}[|\Delta - w|] \\ &= \Delta + \frac{\beta_i}{\beta_i + \beta_j} \mu - \frac{\beta_j}{\beta_i + \beta_j} (\Delta - \mathbb{E}[|\Delta - w|]) \end{aligned} \quad (4)$$

Now

$$\begin{aligned}
\mathbb{E}[|\Delta - w|] &= \sum_{i=1}^{\Delta} (\Delta - i)p_i + \sum_{i=\Delta+1}^{\infty} (i - \Delta)p_i \\
&\leq \Delta + \sum_{i=\Delta+1}^{\infty} ip_i - \sum_{i=1}^{\Delta} ip_i \\
&\leq \Delta - \left(1 - \frac{1}{2n}\right)\mu
\end{aligned}$$

where the last inequality is due to (3).

Note that $\frac{\beta_i}{\beta_i + \beta_j} \leq \frac{1}{2} - \frac{1}{4n}$. We plug both bounds in Equation (4) and have:

$$\mathbb{E}[\Delta'] \leq \Delta + \left(\frac{1}{2} - \frac{1}{4n}\right)\mu - \left(\frac{1}{2} + \frac{1}{4n}\right)\left(1 - \frac{1}{2n}\right)\mu < \Delta - \frac{\mu}{8n}$$

□

For notational simplicity, from this point on we count only active steps. Define t^* to be the first time for which $\Delta(x(t^*), y(t^*)) \leq \Delta^*$.

Lemma 3.8. *For every $c > 0$, with probability $\geq 1 - \frac{1}{\Delta_0^{13/10} n^c}$ it holds that $t^* \leq O(c\Delta_0^{7/5} n^{c+3})$.*

Proof. Consider an edge $(x(0), y(0)) \in \Gamma$ starting out at distance Δ_0 . Let $(x(s), y(s))$ be the state of the pair after s active steps in our coupling. We argue that with high probability, there is some $t \in O(\Delta_0^{7/5} n^{c+3})$ such that within t active steps the distance is reduced to Δ^* .

For brevity, let us denote $\Delta(x(s), y(s))$ by Δ_s . Recall that if $(x(s), y(s))$ are such that $x(s) = y(s) + e_i - e_j$, then

$$\Delta_{s+1} = \begin{cases} \Delta_s + w & \text{with probability } \frac{\beta_i}{\beta_i + \beta_j} \\ |\Delta_s - w| & \text{with probability } \frac{\beta_j}{\beta_i + \beta_j} \end{cases}$$

We would like to show that Δ_s decreases fast enough, as long as it has not hit Δ^* . However, our quest is complicated by the fact that the random variables $\Delta_{s+1} - \Delta_s$ depend on i, j , and as a result on all the previous steps. To handle this dependence, we shall argue that even conditioned on the worst history (and hence the worst i, j), the decrement is expected to be large enough as long as we have not hit $[0, \Delta^*]$ already.

We first define a distribution Z' as follows: we first sample a w from the weight distribution. With probability $\frac{\beta_{n-1}}{\beta_{n-1} + \beta_n}$, we set $Z' = w$. With probability $\frac{\beta_n}{\beta_{n-1} + \beta_n}$, we set $Z' = |\Delta^* - w| - \Delta^*$.

Let Z_s denote the random variable defined as follows

$$Z_s = \begin{cases} \Delta_{s+1} - \Delta_s & \text{if } \Delta_k > \Delta^* \text{ for } k = 0, 1, \dots, s \\ \text{an independent sample from } Z' & \text{otherwise} \end{cases}$$

Thus until Δ_s hits $[0, \Delta^*]$, Z_s is the decrement in Δ_s . After Δ_s hits $[0, \Delta^*]$ for the first time, Z_s is distributed like an independent copy of Z' .

In the first case, we have

$$Z_s = \begin{cases} w & \text{with probability } \frac{\beta_i}{\beta_i + \beta_j} \\ |\Delta_s - w| - \Delta_s & \text{with probability } \frac{\beta_j}{\beta_i + \beta_j} \end{cases}$$

where i and j are the indices $x(s)$ and $y(s)$ differ in. Note that i and j are dependent on the Z_0, \dots, Z_{s-1} and that $i < j$. Note also that if $\sum_{i=1}^t Z_s < -\Delta_0$, then there must be an $s \leq t$ such that $\Delta_s \in [0, \Delta^*]$. Thus it suffices to show that with high probability, $\sum_{i=1}^t Z_s < -\Delta_0$.

For any s , let \mathbf{Z}_s denote the vector Z_1, \dots, Z_s . Now consider the random variable $Z_s | (\mathbf{Z}_{s-1} = \mathbf{z}_{s-1})$ for some vector $\mathbf{z}_{s-1} \in \mathbb{R}^s$.

Lemma 3.9. *For any $\mathbf{z}_{s-1} \in \mathbb{R}^s$, the random variable $Z_s | (\mathbf{Z}_{s-1} = \mathbf{z}_{s-1})$ is stochastically dominated by the random variable Z' .*

Proof. If $\Delta_k \leq \Delta^*$ for some $k \leq s$, Z_s is distributed as Z' and there is nothing to prove. So we assume otherwise. Then the natural coupling works: we couple w with w and match up as much of the mass in case 1 for Z_s (probability $\frac{\beta_i}{\beta_i + \beta_j}$) with the corresponding case in Z' (probability $\frac{\beta_{n-1}}{\beta_{n-1} + \beta_n}$). The main observation is that the ratio $\frac{\beta_i}{\beta_i + \beta_j}$ is minimized for $(i, j) = (n-1, n)$. It is easy to see that for any $\Delta \geq \Delta^*$ and for any w , the number $|\Delta - w| - \Delta$ is no larger than $|\Delta^* - w| - \Delta^*$. Moreover, for any Δ and any w , $|\Delta - w| - \Delta$ is no larger than w . The claim follows. \square

Note that stochastic dominance implies the following: for any real valued random variables A_1, A_2 , if $A_2 | (A_1 = a)$ is stochastically dominated by B for every $a \in \mathbb{R}$, then

$$\begin{aligned} \Pr[A_1 + A_2 > a'] &= \int \Pr[A_2 > a' - a | A_1 = a] \mu_{A_1}(a) da \\ &\leq \int \Pr[B > a' - a] \mu_{A_1}(a) da = \Pr[A + B > a'] \end{aligned}$$

Thus $\sum_{s=0}^t Z_s | \mathbf{Z}_{s-1}$ is stochastically dominated by $\sum_{s=0}^t Z'_s$, where each Z'_s is an **independent** copy of Z' . We conclude that the probability that we do not hit Δ^* in t steps is at most the probability that the sum of t independent copies of Z' is larger than $-\Delta_0$.

Now note that from Lemma 3.7, $E[Z'] \leq -\frac{\mu}{8n}$. Moreover, $\text{Var}[Z'] \leq E[Z'^2] \leq M_2(\mathcal{W})$. Let us take $t = \frac{8\Delta_0^{7/5} n^{c+3}}{\mu}$ so that $E[\sum_{i=1}^t Z'_i] \leq -n^{c+2} \Delta_0^{7/5}$.

Let A be the event that all t balls have weight below $\Delta_0^{27/20} n^{(2c+3)/2}$. By Chebyshev's inequality we have that for each sample $\Pr[w \geq \Delta_0^{27/20} n^{(2c+3)/2}] \leq \frac{2M_2}{\Delta_0^{27/10} n^{2c+3}}$. Union bounding over the t balls we have $\Pr[A] \leq \frac{2M_2}{\Delta_0^{27/10} n^{2c+3}} \cdot \frac{8\Delta_0^{7/5} n^{c+3}}{\mu} = \frac{16M_2}{\Delta_0^{13/10} n^c} \in O\left(\frac{1}{\Delta_0^{13/10} n^c}\right)$.

We next condition on A and use Bernstein's inequality (c.f. Theorem 2.7 in [9]).

Theorem 3.10 (Bernstein's inequality). *Let the random variables X_1, \dots, X_n be independent with $X_i - E[X_i] \leq b$ for each $i \in [n]$. Let $X := \sum_i X_i$ and let $\sigma^2 := \sum_i \sigma_i^2$ be the variance of X . Then, for any $\delta > 0$,*

$$\Pr[X > E[X] + \delta] \leq \exp\left(-\frac{\delta^2}{2\sigma^2(1 + b\delta/3\sigma^2)}\right).$$

In our case we have $\sigma^2 = tM_2$, $b = \Delta_0^{27/20} n^{2c+3/2}$ and $\delta = (n^{c+2} - 1)\Delta_0^{7/5}$. Plugging in Theorem 3.10 we have

$$\Pr\left[\sum_{i=1}^t Z'_i > E\left[\sum_{i=1}^t Z'_i\right] + (n^{c+2} - 1)\Delta_0^{7/5}\right] \leq \exp\left(-\frac{n^{2c+4}\Delta_0^{14/5}}{2tM_2 + \Delta_0^{11/4} n^{2c+5/2}}\right)$$

where $t = 8\Delta_0^{7/5}n^{c+3}/\mu$. Since μ and $M_2[\mathcal{W}]$ are constants, this failure probability is exponentially small in $n^{1/2}\Delta_0^{1/20}$. The conditioning on A adds $O\left(\frac{1}{\Delta_0^{13/10}n^c}\right)$ to the error probability. This completes the proof of Lemma 3.8. \square

We can now complete the proof of Lemma 3.6. By the definition of smoothness, if $\Delta < \Delta^*$ then there is a sequence of at most α active steps that would bring the distance to 0. Thus, once $\Delta < \Delta^*$, the probability the distance would hit 0 within the next α steps is at least $\frac{1}{2^\alpha} \frac{1}{(\Delta^*)^\beta} \geq \frac{1}{2^\alpha n^{3\beta}}$. If the distance didn't hit 0 within α steps then the process reiterates and after a while the distance is below Δ^* again. We need $m := 3 \log(n\Delta_0)2^\alpha n^{3\beta}$ iterations to succeed with probability at least $1 - \frac{1}{(n\Delta_0)^3}$.

If the distance did not hit 0 in the i 'th attempt, it reaches some distance k_i which is dominated by a sum of α independent samples from \mathcal{W} . Let c' be such that $k_i^{13/10}n^{c'} = 3\Delta_0^{13/10}n^{3\beta+3}2^\alpha$ and let $c_i = \max\{c', 1\}$.

By Lemma 3.8 it takes $c_i k_i^{7/5}n^{c+3} \log nk_i$ steps to reach Δ^* with probability $1 - \frac{1}{k_i^{13/10}n^{c_i}}$. Thus, the total number of active steps needed is

$$\begin{aligned} \sum_{i=1}^m c_i k_i^{7/5} n^{c_i+3} \log(nk_i) &= n^3 \sum_{i=1}^m (k_i^{3/10} n^{c_i}) k_i^{1/10} \log(nk_i) \\ &\leq 2^{\alpha+2} \Delta_0^{13/10} n^{3\beta+7} \sum_{i=1}^m k_i^{1/10} \end{aligned}$$

Now all that remains is to show that

$$\Pr \left[2^{\alpha+2} \Delta_0^{13/10} n^{3\beta+7} \sum_{i=1}^m k_i^{1/10} \leq 2^{\alpha+2} \Delta_0^{7/5} n^{6\beta+8} \right] = \Pr \left[\sum_{i=1}^m k_i^{1/10} \leq \Delta_0^{1/10} n^{3\beta+1} \right]$$

is at most $\frac{1}{\Delta_0^{6/5}n^{11/5}}$. Indeed recall that k_i is dominated by a sum of α independent samples from \mathcal{W} . We have

$$\Pr \left[\sum_{i=1}^m k_i^{1/10} \leq \Delta_0^{1/10} n^{3\beta+1} \right] \leq \Pr \left[\left(\sum_{i=1}^m k_i^{1/10} \right)^{20} \leq \Delta_0^2 n^{60\beta+20} \right] \leq \frac{\alpha^2 \mu^2 m^{20}}{\Delta_0^2 n^{60\beta+20}} \leq \frac{1}{\Delta_0^{6/5} n^{11/5}}$$

this completes the proof of Lemma 3.6.

4 Putting it together

In this section, we show how the weak gap Theorem 2.1 and the short memory theorem 3.3 together imply a strong gap theorem. More precisely, we show that with probability $(1 - \frac{1}{\text{poly}(n)})$, the gap at the end of t steps is independent of t . This part of the proof is similar to Berenbrink *et al.* [4]. We assume for simplicity that the bound on the mixing time in the short memory theorem is at most $\gamma^{7/5}n^c$ for concreteness, where $c \geq 5$ is some constant.

The following is a restatement of Theorem 1.1.

Theorem 1.1. *For any t , the probability that $\text{gap}(t) > k$ is at most $\Pr[\text{gap}(n^{30c}) > k] + \frac{1}{n^{10c}}$*

Proof. We show a slightly stronger result by induction on t : we show that for any $t \geq n^{30c}$, and any $k > 0$, we have that

$$Pr[\text{gap}(t) > k] \leq Pr[\text{gap}(n^{30c}) > k] + \frac{6}{n^{\frac{29c}{3}}} - \sum_{j=0}^{\infty} \frac{2}{t^{\frac{1}{3}} \left(\frac{30}{29}\right)^{j+1}}$$

For any t such that $n^{\frac{29}{30}30c} < t \leq n^{30c}$, this is trivially true. Suppose that for some integer $s \geq 0$, it is true for all t such that $n^{\left(\frac{30}{29}\right)^{s-1}30c} < t \leq n^{\left(\frac{30}{29}\right)^s30c}$. We now argue that the claim also holds for all t such that $n^{\left(\frac{30}{29}\right)^s30c} < t \leq n^{\left(\frac{30}{29}\right)^{s+1}30c}$. Indeed, let t be in such a range. Consider the process at the end of $(t - t^{\frac{14}{15}}n^c)$ steps. By the weak gap theorem, at the end of these many steps, the gap in X is at most $t^{\frac{2}{3}}$ with probability at least $(1 - \frac{2}{t^{\frac{1}{3}}})$; let us assume that the gap is indeed at most $t^{\frac{2}{3}}$. Consider a process Y that at time $t - t^{\frac{14}{15}}n^c$ is balanced, and continues like X from this point on. By the coupling lemma, the difference of $t^{\frac{2}{3}}$ between X and Y is forgotten in time $t^{\frac{14}{15}}n^c$; i.e. the probability that X and Y are different is bounded by $\frac{1}{t^{\frac{2}{3}}}$. Thus $\text{gap}_X(t)$ and $\text{gap}_Y(t)$ differ with probability at most $\frac{1}{t^{\frac{2}{3}}}$. However, $\text{gap}_Y(t)$ is distributed exactly as $\text{gap}(t^{\frac{14}{15}}n^c)$, since at time $t - t^{\frac{14}{15}}n^c$, the process Y was fully balanced. Moreover, note that since $t \geq n^{30c}$, $t^{\frac{14}{15}}n^c \leq t^{\frac{29}{30}}$. Thus by the induction hypothesis, $Pr[\text{gap}_Y(t^{\frac{14}{15}}n^c) > k]$ is bounded by

$$Pr[\text{gap}(n^{30c}) > k] + \frac{6}{n^{\frac{29c}{3}}} - \sum_{j=0}^{\infty} \frac{3}{t^{\left(\frac{29}{30}\frac{1}{3}\right)\left(\frac{30}{29}\right)^{j+1}}$$

Thus, we have

$$\begin{aligned} & Pr[\text{gap}_X(t) > k] \\ & \leq Pr[\text{gap}_X(t - t^{\frac{14}{15}}n^c) > t^{\frac{2}{3}}] \\ & \quad + Pr[\text{gap}_Y(t^{\frac{14}{15}}n^c) > k] + Pr[X(t^{\frac{14}{15}}n^c) \neq Y(t^{\frac{14}{15}}n^c)] \\ & \leq \frac{2}{t^{\frac{1}{3}}} + Pr[\text{gap}(n^{30c}) > k] + \frac{6}{n^{\frac{29c}{3}}} - \sum_{j=0}^{\infty} \frac{3}{\left(t^{\frac{29}{30}\frac{1}{3}}\right)\left(\frac{30}{29}\right)^{j+1}} + \frac{1}{t^{\frac{4}{5}}} \\ & = Pr[\text{gap}(n^{30c}) > k] + \frac{6}{n^{\frac{29c}{3}}} - \sum_{j=0}^{\infty} \frac{3}{\left(t^{\frac{1}{3}}\right)\left(\frac{30}{29}\right)^{j+1}} \end{aligned}$$

Hence the induction holds. \square

We now sketch the proof of Theorem 1.2.

Theorem 1.2. *If \mathcal{B} has a finite fourth moment, then for any t , $\mathbb{E}[\text{Gap}(t)] \leq n^c$ where c is some number depending on \mathcal{B} alone.*

We wish to argue that for some $\epsilon > 0$, $Pr[\text{gap}(t) \geq y] \leq \frac{\text{poly}(n)}{y^{1+\epsilon}}$ for all y ; the bound on the expectation would follow immediately. For $y < n^c$ for some constant c , there is nothing to prove, so we assume $y \geq n^c$ for a large enough c .

First note that under the finite fourth moment assumption, the bound in the weak gap theorem is improved to $\frac{1}{t^{\frac{6}{5}}}$ instead of $\frac{1}{t^{\frac{1}{3}}}$ above.

Thus for $y > t^{\frac{4}{5}}$, the required bound follows directly from the weak gap theorem. For smaller y , we use the above induction approach, except that the base case is at some $t_0 \in [y^{\frac{15}{16}}, y)$, where $t_0 = t^{\left(\frac{15}{16}\right)^j}$ for some

integer j . This would then imply that $\Pr[\text{gap}(t) \geq y] \leq \Pr[\text{gap}(t_0) \geq y] + \frac{4}{t_0^{6/5}}$. It is then easy to see that the desired probability is at most $\frac{4}{y^8}$.

5 The Real Valued Case

We now consider the case when the weight distribution \mathcal{B} is a distribution over non-negative reals. We make the following assumptions on the distribution \mathcal{B} :

- Finite variance: $M_2(\mathcal{B})$ is finite.
- For any $C > 0$, there is a ϵ_C such that $f_{\mathcal{B}}(x) \geq \epsilon_C$ for all $x \in [0, C]$.⁴ where $f_{\mathcal{B}}(x)$ denotes the probability density function of \mathcal{B} .

Our proof basically works via a reduction to a *dependent* version of the integer case. More precisely, let process X_s denote the evolution of the load vector. We define an auxiliary process X'_s in which each ball has an integer weight and the load of each bin in X_s is close to that in X'_s . In the process, we lose the independence of the weight samples. We then show how the argument in the integer case extends to this setting.

5.1 The reduction

The most natural way to randomly round a real weight value W to an integer is to set W' to be $\lceil W \rceil$ with probability $(W - \lfloor W \rfloor)$, and set it to $\lfloor W \rfloor$ otherwise. This has the nice property that $E[W'] = W$ and that $|W' - W| < 1$.

Doing this independently for each weight however will lead to a large discrepancy between the sum of W 's and the sum of W' 's. Indeed, this difference is expected to be about \sqrt{m} if we have thrown m balls. Thus the implied bounds on the gap will be not much stronger than those obtained in the weak gap theorem.

The situation changes dramatically however once we allow the rounding of various weight values to be dependent! We shall round the size of the ball based on which bin it is placed, while ensuring that the total (true) weight of the balls in it is within one of the total weight of the rounded values of the balls in it.

More precisely, let $T_{s-1}(i) = \sum_{t \in B_i} W_t$ be the total weight of the balls in bin i at time step $(s-1)$ and suppose that the process X_s places a ball of weight W_s in bin i . Let $T_s(i) = T_{s-1}(i) + W_s$ be the new total weight in bin i . The process X' mimics the process X , except that when X places a ball of weight W_s in bin i , the process X' places a ball of weight $W'_s = \lceil KW_s \rceil - \lceil KT_{s-1}(i) \rceil$ in bin i , where $K := \frac{8n}{\mu}$ is a scaling factor. The weights are scaled up for a technical reason to be clarified later. We observe the following:

- $W'_s \in \{\lfloor KW_s \rfloor, \lceil KW_s \rceil\}$
- $T'_s(i) = \sum_{t \in B_i} W'_t$ is (inductively) equal to $\lceil T_s(i) \rceil$
- If $T_s(i) \geq T_s(j)$ then $T'_s(i) \geq T'_s(j)$

Thus it suffices to show that Theorem 1.1 holds for process X' . Note that in process X' the weights of balls are not independent of one another. In particular, the decision whether KW_s should be rounded up or down depends upon the history of the process.

⁴This assumption can be significantly relaxed to a similar assumption on the distribution of $\sum_{i=1}^k a_i W_i$, for some k and some a_i 's in $\{-1, 1\}$.

5.2 Strengthening the Integer case

We wish to argue that under the assumptions above, the induced weight distribution W' over integers has all the right properties to prove the strong gap property. We only sketch the modifications needed in the proof above, and omit the details from this extended abstract.

Lemma 5.1. *Let X' be as above. Then for any integer w , the following holds:*

$$\Pr[W'_s = w | (\mathbf{W}_{s-1} = \mathbf{w}_{s-1}, \mathbf{A}_s = \mathbf{a}_s)] \geq \min_{x \in (\frac{w-1}{K}, \frac{w+1}{K})} f_{\mathcal{B}}(x)/K$$

where \mathbf{W}_{s-1} and \mathbf{A}_s denote the vectors of random variables corresponding to the weights of the balls and their allocations respectively in the first $(s-1)$ (respectively s) steps and \mathbf{w}_{s-1} and \mathbf{a}_s are arbitrary values for these variables.

Proof. Note that for any setting of i and any value of $T_{s-1}(i)$, and for any integer w , the rounded weight value $W'_s = w$ whenever $\lceil K(T_{s-1}(i) + W) \rceil = w + \lceil KT_{s-1}(i) \rceil$. This happens whenever $KW \in (w + \lceil KT_{s-1}(i) \rceil - ST_{s-1}(i) - 1, w + \lceil KT_{s-1}(i) \rceil - KT_{s-1}(i))$. Since the corresponding interval for W is a subrange on $(\frac{w-1}{K}, \frac{w+1}{K})$ and has length $\frac{1}{K}$, the claim follows. \square

First note that $EW'^2 \leq K^2EW^2$. Since K is polynomial in n , this only affects the bounds by a *poly*(n) factor. The weak gap theorem continues to hold, and the definition of the graph and the coupling remains unchanged. Moreover, under the assumptions on the distribution, the induced distribution over integers has the properties we need.

The dependency of the weight value on the past and the choice of the bin in this step creates some subtle problems. In particular, when the sampled real weight value is W , the value $T_s(i)$ may increase by $\lceil KW \rceil$ when the ball falls in bin i , but the increase in $T_s(j)$ may be only $\lfloor KW \rfloor$ in case it falls in bin j . Thus the increase in Δ , even though it is less likely than a decrease, may be *larger* than the decrease that would happen if the same ball were to fall in bin j . However, this can decrease $E[\Delta']$ in Lemma 3.7 by at most one, whereas our choice of the scaling factor $K = 8n/\mu$ ensures that the decrease in expectation is at least two if there were no rounding (and hence at least one). This is the only place where we need the scaling.

Moreover, the dependency on the past is easily conditioned out by the dominance in lemma 5.1. Lemmas 3.8 continues to hold, and thus the claim follows, albeit with worse constants.

6 Open Problems

6.1 Distributions that Fall through the Cracks

Recall that we assumed that if $|\text{Support}(\mathcal{B})|$ is infinite, the distribution is (α, β) -smooth. In fact it is a rather challenging exercise to find a distribution for which this property is not self-evident. One such distribution is the following. Say X is distributed Geometrically with parameter $\frac{1}{2}$. Define $Y := \lceil 2^{X/3} \rceil$. Y has finite expectation and variance. Now, given that $\Delta = \sum_{i \leq 3 \log n} \lceil 2 \rceil^{i/3}$ (which is $O(n)$), the probability the next $3 \log n$ active steps bring Δ to 0 is $1/n^{\Omega(\log n)}$. Note however that the probability that we end up with a gap of $\Delta = \sum_{i \leq 3 \log n} \lceil 2 \rceil^{i/3}$ is very small to begin with. In other words, in this case the ‘‘complexity’’ of a configuration is not accurately captured by an upper bound on the magnitude of the gap, and thus not captured by the distance function either. It may well be the case that a different distance function would take care of these types of distributions. The fate of such distributions therefore remains open.

6.2 Lower Bounds

We assume that the weight distribution has finite variance and finite expectation. If a distribution has infinite expectation then it may be the case the gap between maximum and average would be increase with m even if the insertion algorithm has perfect information and complete freedom:

Lemma 6.1. *If the weight distribution is $2^{G(\frac{1}{2})}$ where $G(\frac{1}{2})$ is the geometric distribution with parameter $\frac{1}{2}$, then with probability $\geq \frac{1}{2}$, after throwing m balls the gap between maximum and average is $\Omega(m)$ as long as $m \in O(2^n)$.*

Proof. The probability a ball is of weight $\geq m$ is the probability a Geometric variable is of size $\geq \log m$ which is at least $1/e$. Next we show that with high probability it holds that the sum of the m balls is $O(m \log m)$ w.h.p. This implies that the average bin is $O(\frac{m \log m}{n})$, thus the gap between maximum and average is $\Omega(m)$ as long as $m \in O(2^n)$. Consider the set of m Geometric variables X_1, \dots, X_m such that $W_i = 2^{X_i}$. Define S_i to be the number of X variables that were sampled to be i . We have that S_i is distributed Binomially with parameters $(m, 2^{-i})$, so $\mu(S_i) = \frac{m}{2^i}$. Chernoff bound implies that $\Pr[S_i \geq \frac{2m}{2^i}] \leq \frac{1}{10m}$ as long as $i \leq \log m - \log \log m$. Chernoff bound also implies that $\Pr[\sum_{i \geq \log m - \log \log m} S_i \geq \log m] \leq \frac{1}{10m}$. We conclude that the contribution of each S_i to the total sum of the weights is at most $S_i 2^i \leq 2m$ and therefore the total sum of weights is at most $O(m \log m)$. \square

Lemma 6.1 holds for any allocation algorithm and uses the fact that the expectation is infinite. It would be interesting to demonstrate that a finite second moment is also necessary. Such a lower bound may require the use of specific properties of the algorithm.

6.3 Better Bounds for Specific Distributions

We have reduced the case where m is arbitrarily large to the case where $m \leq poly(n)$. It is therefore interesting to derive tight bounds for interesting distributions such as the Geometric, the Log-Normal distribution and so on. In particular it may be possible to prove a general bound which is tighter than ours and that would explicitly use the moments of the distribution, perhaps using the layered induction technique.

Acknowledgments

The authors wish to express their gratitude to Ittai Abraham.

References

- [1] Micah Adler, Soumen Chakrabarti, Michael Mitzenmacher, and Lars Rasmussen. Parallel randomized load balancing. *Random Structures and Algorithms*, 13(2):159 – 188, 1998.
- [2] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. *SIAM Journal Computing*, 29(1):180–200, 1999.
- [3] Petra Berenbrink, Friedhelm Meyer auf der Heide, and Klaus Schröder. Allocating weighted jobs in parallel. *Theory of Computing Systems*, 32(3):281–300, 1999.
- [4] Petra Berenbrink, Artur Czumaj, Angelika Steger, and Berthold Vöcking. Balanced allocations: the heavily loaded case. *SIAM Journal Computing*, 35(6):1350–1385, 2006.

- [5] Petra Berenbrink, Tom Friedetzky, Zengjian Hu, and Russel Martin. On weighted balls-into-bins games. In *Proc. of Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 231–243, 2005.
- [6] Russ Bubley and Martin E. Dyer. Path coupling: A technique for proving rapid mixing in markov chains. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 223–231, 1997.
- [7] John Byers, Jeffrey Considine, and Michael Mitzenmacher. Simple load balancing for distributed hash tables. In *Proc. of Intl. Workshop on Peer-to-Peer System (IPTPS)*, pages 80–87, 2003.
- [8] John W. Byers, Jeffrey Considine, and Michael Mitzenmacher. Geometric generalizations of the power of two choices. In *Proc. of Symp. on Parallelism in Algorithms and Architectures (SPAA)*, pages 54–63, 2004.
- [9] Colin McDiarmid. *Probabilistic Methods for Algorithmic Discrete Mathematics*, chapter Concentration, pages 195–248. Springer, 1998.
- [10] Michael Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, 1(2):225–251.
- [11] Michael Mitzenmacher. Load balancing and density dependent jump markov processes. In *Proc. of Symp. on Foundations of Computer Science (FOCS)*, pages 213–222, 1996.
- [12] Michael Mitzenmacher, Andrèa Richa, and Ramesh Sitaraman. *Handbook of Randomized Computing*, chapter The power of two random choices: A survey of the techniques and results. Kluwer, 2000.
- [13] Berhold Vöcking. How asymmetry helps load balancing. *Journal of ACM*, 50(4):568–589, 2003.
- [14] Udi Wieder. Balanced allocations with heterogenous bins. In *Proc. of Symp. on Parallelism in Algorithms and Architectures (SPAA)*, 2007.