On Balls and Bins with Deletions

Richard Cole* Alan Frieze† Bruce M. Maggs‡ Michael Mitzenmacher§
Andréa W. Richa¶ Ramesh K. Sitaraman|| Eli Upfal**

Abstract

We consider the problem of extending the analysis of balls and bins processes where a ball is placed in the least loaded of d randomly chosen bins to cover deletions. In particular, we are interested in the case where the system maintains a fixed load, and deletions are determined by an adversary before the process begins. We show that with high probability the load in any bin is $O(\log\log n)$. In fact, this result follows from recent work by Cole et al. concerning a more difficult problem of routing in a butterfly network.

The main contribution of this paper is to give a different proof of this bound, which follows the lines of the analysis of Azar, Broder, Karlin, and Upfal for the corresponding static load balancing problem. We also give a specialized (and hence simpler) version of the argument from the paper by Cole et al. for the balls and bins scenario. Finally, we provide an alternative analysis also based on the approach of Azar, Broder, Karlin, and Upfal for the special case where items are deleted according to their age. Although this analysis does not yield better bounds than our argument for the general case, it is interesting because it utilizes a two-dimensional family of random variables in order to account for the age of the items. This technique may be of more general use.

1 Introduction

A standard question in load balancing is to consider what the distribution of balls in bins looks like when m balls are thrown into n bins. In particular, when n balls are thrown into n bins, it is well known that the maximum load is approximately $\log n/\log\log n$ with high probability. The seminal paper of Azar, Broder, Karlin, and Upfal asked a related question: suppose the balls are placed sequentially, and each ball is placed in the least loaded of d bins chosen independently and uniformly at random [4]. In this case, they find that the maximum load is $\log\log n/\log d + O(1)$ with high probability; more detailed analysis of the distribution in this case is undertaken in [10]. This work has led to a number papers analyzing related load balancing schemes, including for example [1, 2, 5, 6, 7, 8, 9, 11, 12].

Note that the above result applies to a *static* problem, where a fixed number of balls are distributed. An interesting related question is to consider the *dynamic* situation where balls can be deleted as well as inserted into the system over time. Indeed, the original paper by Azar, Broder, Karlin, and Upfal examines the dynamic situation where at each step a random ball is deleted and re-inserted in the system [4]. Related work by, for example, Mitzenmacher [8, 9, 11] and Adler et al. [1] examines deletions via connections with queueing theoretical models.

Here we focus on a model where an adversary may specify a deletion sequence in advance. Our first and main result is to extend the proof of [4] to handle a polynomial length sequence of insertions and deletions, where the maximum load in the system is always at most n balls. We then note that an even more general result, in which re-insertions can occur, is already essentially contained in the results of [6] (a re-insertion causes a ball to choose among the same bins as on its first insertion). This work considered a similar problem related to routing on a butterfly network. We restate this proof for the balls and bins setting, where it becomes significantly simpler. Finally, we consider a special case in which deletions are always of the item that has been longest in the system. We again use a variant of the two-choice argument from [4], this time making use of a two-dimensional family of random variables, similar in spirit to the work of [11].

We emphasize that the interest of this work lies in the techniques used rather than the result, which is already implicit in the work of [6].

2 Adversarial deletions: polynomially many steps

In this section, we demonstrate that the original proof of Azar, Broder, Karlin, and Upfal in [4] can be extended to handle deletions under an appropriate adversarial model for polynomially many steps. We first define the underlying process.

For a vector $\vec{v} = (v_1, v_2, \ldots)$, let $P_d(\vec{v})$ be the following process: at time steps 1 through n, n balls are placed into n bins sequentially, with each ball going into the least loaded of d bins chosen independently and uniformly at random. After these balls are placed, deletions and insertions alternate, so that at each subsequent time step $n + j$, first the ball inserted at time v_j is removed, and then a new ball is placed into the least loaded of d bins chosen independently and uniformly at random. (Actually we do not require this alternation; the main point is that we have a bound, n, on the number of balls in the system at any point. The alternation merely makes notation more convenient.)

We assume the vector $\vec{v}$ is suitably defined so that at each step an actual deletion occurs; that is, the v_j are unique and $v_j \leq n+j-1$. Otherwise $\vec{v}$ is arbitrary, although we emphasize that it is chosen before the process begins and does not depend on the random choices made during the process.

We adopt some of the notation of [4]. Each ball is assigned a fixed *height* upon entry, where the height is the number of balls in the bin, including itself. The height of the ball placed at time t is denoted by $h(t)$. The load of a bin at time t refers to the number of balls in the bin at that time. We let $\mu_{\geq k}(t)$ denote the number of balls that have height at least k at time t, and $\nu_{\geq k}(t)$ be the number of bins that have load at least k at time t. Note that if a bin has load k, it must contain some ball of height at least k. Hence $\mu_{\geq k}(t) \geq \nu_{\geq k}(t)$ for all times t. Finally, $B(n, p)$ refers to a binomially distributed random variable based on n trials each with probability p of success.

Before giving the proof, we sketch the main ideas. The flavor of the proof is to show that with high probability the number of bins containing at least i balls is doubly exponentially decreasing for sufficiently large i. The bound on the number of bins containing at least $i + 1$ balls is obtained from the bound on the number of bins containing at least i balls. Establishing the proper conditioning between the number of bins with i and $i + 1$ balls makes the proof challenging.

A key idea is to avoid seeking a direct bound on the number of bins containing at least i balls. Rather, following [4], we use the fact that the number of balls of height at least i bounds the number of bins containing at least i balls. This leads us to obtain bounds on the distribution of ball heights which, with high probability, hold for polynomially many steps. A concern is that here the adversarial choice of deletions might lead this bound to be too weak. On the other hand, the adversary is constrained, for the full sequence of deletions must be chosen up front, and this allows the result.

The key difference between our result on that of [4] is that they find a dominating distribution of heights on one set of n balls, whereas we use a distribution that applies to every set of n balls present in the system as it evolves. As it happens, the bounds are essentially the same; the most significant changes lie in the end game, where we must bound the number of bins containing more than $\log\log n / \log d$ balls.

Theorem 1 *For any fixed constants c_1 and c_2, with probability at least $1 - o(1/n^{c_1})$ the maximum load of a bin achieved by process $P_d(\vec{v})$ over $T = n^{c_2}$ steps is* $\log\log n / \log d + O((c_1 + c_2)/d)$

Proof: The argument extends the original Theorem 4 of [4], by determining a distribution on the heights of the balls that holds for polynomially many steps, regardless of which n balls are in the system at any point in time.

Let $\mathcal{E}_i$ be the event that $\nu_{\geq i}(t) \leq \beta_i$ for time steps $t = 1, \ldots, T$, where the β_i will be revealed shortly. We want to show that at time t, $1 \leq t \leq T$,

$$\Pr(\mu_{\geq i+1} > \beta_{i+1} \mid \mathcal{E}_i)$$

is sufficiently small. That is, given $\mathcal{E}_i$, we want $\mathcal{E}_{i+1}$ to hold as well. This probability is hard to estimate directly. However, we know that since the d choices for a ball are independent, we have

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

We would like to bound for each time t the distribution of the number of time steps j such that $h(j) \geq i+1$ and the ball inserted at time step j has not been deleted by time t. In particular, we would like to bound this distribution by a binomial distribution over n events with success probability $(\beta_i/n)^d$. But this is difficult to do directly as the events are not independent.

Instead, we fix i and define the binary random variables Y_t for $t = 1, \ldots, T$, where

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

(The value Y_t is 1 if and only if the height of the ball t is at least $i+1$ despite the fact that the number of boxes that have load at least i is currently below β_i.)

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

$$\Pr(Y_t = 1 \mid \omega_1, \ldots, \omega_{t-1}, v_1, \ldots, v_{t-n}) \leq \frac{\beta_i^d}{n^d} \stackrel{\text{def}}{=} p_i.$$

Consider the situation immediately after a time step t' where a new ball has entered the system. Then there are n balls in the system, that entered at times $u_1, u_2, \ldots, u_n$. Let $I(t')$ be the set of times $\{u_1, u_2, \ldots, u_n\}$. Then

$$\sum_{t \in I(t')} Y_t = \sum_{i=1}^{n} Y_{u_i};$$

that is, the summation over $I(t')$ is implicitly over the values of Y_t for the balls in the system at time t'. (This statement differs from the result of [4]; the important point here is that we can bound $\sum_{t \in I(t')} Y_t$ regardless of what n balls are in the system. Note these balls are fixed for any time t by the deletion sequence $\vec{v}$.)

We may conclude that at any time $t' \leq T$

$$\Pr\left(\sum_{t \in I(t')} Y_t \geq k\right) \leq \Pr(B(n, p_i) \geq k). \tag{1}$$

Observe that conditioned on $\mathcal{E}_i$, we have $\mu_{\geq i+1}(t') = \sum_{t \in I(t')} Y_t$. Therefore

$$\Pr(\mu_{\geq i+1}(t') \geq k \mid \mathcal{E}_i) = \Pr\left(\sum_{t \in I(t')} Y_t \geq k \mid \mathcal{E}_i\right) \leq \frac{\Pr(B(n, p_i) \geq k)}{\Pr(\mathcal{E}_i)}. \tag{2}$$

Thus:

$$\Pr(\neg \mathcal{E}_{i+1} \mid \mathcal{E}_i) \leq \frac{T \Pr(B(n, p_i) \geq k)}{\Pr(\mathcal{E}_i)}$$

Since

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

we have

$$\Pr(\neg\mathcal{E}_{i+1}) \leq T\Pr(B(n,p_i) \geq k) + \Pr(\neg\mathcal{E}_i). \tag{3}$$

We can bound large deviations in the binomial distribution with the formula (see for instance [3], Appendix A.)

$$\Pr(B(n,p_i) \geq ep_i n) \leq e^{-p_i n}. \tag{4}$$

We may then set $\beta_6 = \frac{n}{2e}$, and subsequently

$$\beta_i = \frac{e\beta_{i-1}^d}{n^{d-1}} \text{ for } i \geq 7.$$

Note that the β_i are chosen so that $\Pr(B(n,p_i) \geq \beta_{i+1}) \leq e^{-p_i n}$.

With these choices $\mathcal{E}_6$ holds with certainty, as there cannot be more than $n/2e$ bins with 6 balls. For $i \geq 6$,

$$\Pr(\neg\mathcal{E}_{i+1}) \leq \frac{T}{n^{c_1+c_2+1}} + \Pr(\neg\mathcal{E}_i) \leq \frac{1}{n^{c_1+1}} + \Pr(\neg\mathcal{E}_i),$$

provided that $p_i n \geq (c_1+c_2+1)\ln n$.

Let i^* be the smallest value for which $p_{i^*-1}n \leq (c_1+c_2+1)\ln n$. Note that $i^* \leq \ln\ln n/\ln d + O(1)$. Note that the preceding argument can not be used to bound the number of bins with height at least i^*, as the Chernoff bounds are no longer powerful enough; hence we must tackle the tail more directly. This requires some careful attention to detail. In fact we will show that, given $\mathcal{E}_{i^*-1}$, then with probability $1 - o(1/n^{c_1})$, there are no balls with height $i^* + \lceil (c_1+c_2+2)/(d-1)\rceil + 1$ over the course of the entire process.

Let $\mathcal{F}_1$ be the event that $\nu_{\geq i^*}(t) \leq (ec_1 + ec_2 + e)\ln n$ for all times $t \leq T$. In other words, at every time there are not too many bins with height at least i^*. Then

$$\Pr(\neg\mathcal{F}_1) \leq \Pr(\neg\mathcal{F}_1 \mid \mathcal{E}_{i^*-1})\Pr(\mathcal{E}_{i^*-1}) + \Pr(\neg\mathcal{E}_{i^*-1}),$$

and again by (1), (2), and (4)

$$\Pr(\neg\mathcal{F}_1 \mid \mathcal{E}_{i^*-1})\Pr(\mathcal{E}_{i^*-1}) \leq T\Pr(B(n,(c_1+c_2+1)\ln n/n) \geq (ec_1+ec_2+e)\ln n) \leq \frac{1}{n^{c_1+1}}.$$

Let $\mathcal{F}_2$ be the event that $\nu_{\geq i^*+1}(t) \leq c_1 + c_2 + 2$ for all times $t \leq T$. In other words, at every time there are no more than a constant number of bins with load at least i^*+1. Again,

$$\Pr(\neg\mathcal{F}_2) \leq \Pr(\neg\mathcal{F}_2 \mid \mathcal{F}_1)\Pr(\mathcal{F}_1) + \Pr(\neg\mathcal{F}_1).$$

Here

$$\Pr(\neg\mathcal{F}_2 \mid \mathcal{F}_1)\Pr(\mathcal{F}_1) \leq T\Pr(B(n,((ec_1+ec_2+e)\ln n/n)^d) \geq c_1+c_2+2).$$

The binomial expression can be checked to be $o(1/n^{c_1+c_2+1})$, and hence this last term is also $O(1/n^{c_1+1})$.

Conditioned on $\mathcal{F}_2$, we must now show that throughout the process there are no bins with load $i^* + \lceil (c_1+c_2+2)/(d-1)\rceil + 1$ with sufficiently high probability. Let this event be given by $\mathcal{G}$. Consider any specific time step z. For there to be be any bins with load $i^* + \lceil (c_1+c_2+2)/(d-1)\rceil + 1$ at time z, at least $\lceil (c_1+c_2+2)/(d-1)\rceil$ of the balls in the system must have landed in bins with height at least i^*+1. Hence

$$\Pr(\neg\mathcal{G}) \leq \Pr(\neg\mathcal{G} \mid \mathcal{F}_2)\Pr(\mathcal{F}_2) + \Pr(\neg\mathcal{F}_2).$$

Here

$$\Pr(\neg\mathcal{G} \mid \mathcal{F}_2)\Pr(\mathcal{F}_2) \leq T\Pr\left(B(n, ((c_1+c_2+2)/n)^d) \geq \lceil (c_1+c_2+2)/(d-1)\rceil\right).$$

The binomial expression can be checked to be $o(1/n^{c_1+c_2+1})$, and hence this last term is also $O(1/n^{c_1+1})$.

To conclude (abusing notation), we have

$$\begin{aligned}
\Pr(\neg\mathcal{G}) &\leq O\left(\frac{1}{n^{c_1+1}}\right) + \Pr(\neg\mathcal{F}_2) \\
&\leq O\left(\frac{1}{n^{c_1+1}}\right) + \Pr(\neg\mathcal{F}_1) \\
&\leq O\left(\frac{1}{n^{c_1+1}}\right) + \Pr(\neg\mathcal{E}_{i^*-1}) \\
&\leq O\left(\frac{1}{n^{c_1+1}}\right) + \sum_{i=1}^{i^*-1}\frac{1}{n^{c_1+1}}.
\end{aligned}$$

Here the last bound comes from the recurrence (3). That is, $\Pr(\neg\mathcal{G})$ is dominated by the sum of $O(\log\log n)$ events each with with probability $O(1/n^{c_1+1})$; the theorem follows. ■

Note that the probability bounds of Theorem 1 can be improved by choosing a smaller value of i^*, so that the strong Chernoff bounds hold, and then increasing the maximum load allowed appropriately. Similarly, we can improve the theorem so that a superpolynomial number of steps can be handled, although this requires increasing the bound on the maximum load above $\log\log n/\log d + O(1)$ to for example $(1+o(1))\log\log n/\log d$.

3 Adversarial deletions: a witness tree argument

In this section, we provide a simple witness tree argument for the balls-and-bins problem with deletions and re-insertions. A similar argument appears in [6] for the more difficult problem of routing circuits in a butterfly network. Therefore, our result is not new, in that it follows naturally from the argument in [6]. Rather, our goal is to present a self-contained and simplified version of the proof for the simpler balls-and-bins situation. For convenience, we focus on the case $d = 2$.

We consider a variation $Q_d(\vec{v}, \vec{w})$ of the process $P_d(\vec{v})$. Again, the process begins with n insertions, followed by alternating insertions and deletions, with $\vec{v}$ specifying the balls to be deleted. We now use $\vec{w}$ to represent insertions, however. We assign each ball an *identification number*, and without loss of generality we assume the first n balls have ID numbers 1 through n. At time $n + j$, the ball with ID number w_j is inserted. If this ball has never been inserted before, then it is placed in the least loaded of d bins chosen independently and uniformly at random. If the ball has been inserted before, it is placed in the least loaded of the d bins chosen for its first insertion – that is the bin choices of a ball are fixed after it is first inserted in the system. We assume that $\vec{v}$ and $\vec{w}$ are consistent, so there is only one ball with a given ID number in the system at a time. Note also that $\vec{v}$ and $\vec{w}$ must be chosen by the adversary before the process begins, without reference to the random choices made during the course of the process.

This scenario appears when, for example, we use a (random) hash function for the two bin choices of every ball. As before, when a ball is (re-)inserted, the algorithm places the ball in the bin with the smaller load.

The main theorem of this section is stated below. The constants of the theorem have been chosen for convenience and have not been optimized. Note that the techniques used to prove this theorem can be generalized to show that if each ball makes d bin choices for some constant d, then the maximum load of any bin is $O(\log\log n)$ with high probability. The result can also be extended for non-constant d as well.

Theorem 2 *At any time t, with probability at least $1 - 1/n^{\Omega(\log\log n)}$, the maximum load of a bin achieved by process $Q_2(\vec{v}, \vec{w})$ is $4 \log\log n$.*

Proof: We prove the theorem in two parts. First, we show that if there is a bin r at time t with 4ℓ balls, where $\ell = \log\log n$, there exists a degree ℓ pruned witness tree. Next, we show that, with high probability, no degree ℓ pruned witness tree exists.

Constructing a witness tree. In a witness tree, each node represents a bin and each edge (r_i, r_j, t_e) represents a ball that was inserted at time t_e whose two bin choices are r_i and r_j. Suppose that some bin r has load 4ℓ at time t. We construct the witness tree as follows. The root of the tree corresponds to bin r. Let $b_1, \ldots, b_{4\ell}$ be the balls in r at time t. Let r_i be the other bin choice associated with ball b_i (one of the choices is bin r). The root r has 4ℓ children, one corresponding to each bin r_i. Let $t_i < t$ be last time b_i was (re-)inserted into system. Without loss of generality, assume that $t_1 < t_2 < \ldots < t_{4\ell}$. Note that the height of ball b_i when it was inserted at time t_i is at least i. Therefore, the load of bin r_i, the other choice of b_i, is at least $i - 1$ at time t_i. We use this fact to recursively grow a tree rooted at each r_i.

The witness tree we have described is irregular. However, it contains as a subgraph an ℓ-ary tree of height ℓ such that

- The root in level 0 has ℓ children that are internal nodes.
- Each internal node on levels 1 to $\ell - 2$ has two children that are internal nodes and $\ell - 2$ children that are leaves.
- Each internal node on level $\ell - 1$ has ℓ children that are leaves.

For convenience we refer to this subtree as the actual witness tree henceforth.

Constructing a pruned witness tree. If the nodes of the witness tree are guaranteed to represent distinct bins, proving our probabilistic bound is a relatively easy matter. However, this is not the case; a bin may reappear several times in a witness tree, leading to dependencies that are difficult to resolve. This makes it necessary to *prune* the tree so that each node in the tree represents a distinct bin. Consequently, the balls represented by the edges of the pruned witness tree are also distinct. In this regard, note that a ball appears at most once in a pruned witness tree, even if it was (re-)inserted multiple times in the sequence.

We visit the nodes of the witness tree iteratively in *breadth-first* search order starting at the root. As we proceed, we remove (i.e., prune) some nodes of the tree and the subtrees rooted at these nodes – what remains is the pruned witness tree. We start by visiting the root. In each iteration, we visit the next node v in breadth-first order that has not been pruned. Let $B(v)$ denote the nodes visited *before* v.

- If v represents a bin that is different from the bins represented by nodes in $B(v)$, we do nothing.
- Otherwise, prune all nodes in the subtree rooted at v. Then, we mark the edge from v to its parent as a *cutoff edge*.

The procedure continues until either no more nodes remain to be visited or there are ℓ cutoff edges. In the latter case, we apply a final pruning by removing all nodes that are yet to be visited. The tree that results from this pruning process is the pruned witness tree. After the pruning is complete, we make a second pass through the tree and construct a set C of *cutoff balls*. Initially, C is set to $\emptyset$. We visit the cutoff edges in BFS order and for each cutoff edge (u, v) we add the ball corresponding to (u, v) to C, if this ball is distinct from all balls currently in C and if $|C| \le \lceil p/2 \rceil$, where p is the total number of cutoff edges.

Lemma 3 *The pruned witness tree constructed above has the following properties.*

1. *All nodes in the pruned witness represent distinct bins.*
2. *All edges in the pruned witness tree represent distinct balls. (Note that cutoff edges are not included in the pruned witness tree.)*
3. *The cutoff balls in C are distinct from each other, and from the balls represented in the pruned witness tree.*
4. *There are $\lceil p/2 \rceil$ cutoff balls in C, where p is the number of cutoff edges.*

Proof: The first three properties follow from the construction. We prove the fourth property as follows. Let b be a ball represented by some cutoff edge, and let v and w be its bin choices. Since v and w can appear at most once as nodes in the pruned witness tree, ball b can be represented by at most two cutoff edges. Thus, there are $\lceil p/2 \rceil$ distinct cutoff balls in C. ■

Enumerating pruned witness trees.

We bound the probability that a pruned witness tree exists by bounding both the number of possible pruned witness trees and the probability that each such tree could arise. First, we choose the shape of the pruned witness tree. Then, we traverse the tree in breadth-first order and bound the number of choices for the bins for each tree node and the balls for each tree edge; we also bound the associated probability that these choices came to pass. Finally, we consider the number of choices for cutoff balls in C and the corresponding probability that they arose. Multiplying these quantities together yields the final bound – it is important to note here that we can multiply term together only because all the balls and the bins in the pruned witness tree and the cutoff balls in C are distinct.

Choosing the shape of the pruned witness tree. Assume that there are p cutoff edges in the pruned tree. The number of ways of selecting the p cutoff edges is at most

$$\binom{\ell^2 2^\ell}{p} \leq \ell^{2p} 2^{\ell p},$$

since there are at most $\ell^2 2^\ell$ nodes in the pruned witness tree.

Ways of choosing balls and bins for the nodes and edges of the pruned witness tree. The enumeration proceeds by considering the nodes in BFS order. The number of ways of choosing the bin associated with the root is n. Assume that you are considering the i^{th} internal node v_i of the pruned witness tree whose bin has already been chosen to be r_i. Let v_i have δ_i children. We evaluate the number of ways of choosing a distinct bin for each of the δ_i children of v_i and choosing a distinct ball for each of the δ_i edges incident on v_i and weight it by multiplying by the appropriate probability. We call this product E_i. There are at most $\binom{n}{\delta_i}$ ways of choosing distinct bins for each of the δ_i children of v_i. The number of ways in which you can choose distinct balls for the δ_i edges incident on v_i is also at most $\binom{n}{\delta_i}$. Furthermore, since there are at most n balls in the system at any time, the probability that a chosen ball chooses bin r_i and one of δ_i bins chosen above is at most $2\delta_i/n^2$. Thus,

$$E_i \leq \binom{n}{\delta_i}\binom{n}{\delta_i}\left(\frac{2\delta_i}{n^2}\right)^{\delta_i} \leq (2e)^{\delta_i}/\delta_i!. \tag{5}$$

Let m be number of internal nodes v_i in the pruned witness tree such that $\delta_i = \ell$. Using the bound in Equation 5 for only these m nodes, the number of ways of choosing the bins and balls for the nodes and edges respectively of the pruned witness tree weighted by the probability these choices occurred is at most $n \cdot ((2e)^\ell/\ell!)^m$.

Ways of choosing the cutoff balls in C. Using Lemma 3, we know that there are $\lceil p/2 \rceil$ distinct cutoff balls in C. The number of ways of choosing the balls in C is at most $n^{\lceil p/2 \rceil}$, since at any time step there are at most n balls in the system to choose from. Note that a cutoff ball has both its bin choices in the pruned witness tree. Therefore, the probability that a given ball is a cutoff ball is at most

$$\binom{\ell^2 2^\ell}{2} \frac{2}{n^2} \leq \ell^4 2^{2\ell}/n^2.$$

Thus the number of choices for the $\lceil p/2 \rceil$ cutoff balls in C weighted by the probability these cutoff balls occurred is at most

$$n^{\lceil p/2 \rceil} (\ell^4 2^{2\ell}/n^2)^{\lceil p/2 \rceil} \leq (\ell^4 2^{2\ell}/n)^{\lceil p/2 \rceil}.$$

Putting it all together. The probability at time t of there existing a pruned witness tree with p cutoff edges, and m internal nodes with $\ell = \log\log n$ children, is at most

$$\begin{aligned} \ell^{2p} 2^{\ell p} \cdot n \cdot ((2e)^\ell/\ell!)^m \cdot (\ell^4 2^{2\ell}/n)^{\lceil p/2 \rceil} &\leq n \cdot ((2e)^\ell/\ell!)^m \cdot (\ell^8 2^{4\ell}/n)^{\lceil p/2 \rceil} \\ &\leq n \cdot (2e^2/\log\log n)^{m \log\log n} \cdot (\log\log^8 n \log^4 n/n)^{\lceil p/2 \rceil}. \end{aligned} \tag{6}$$

Observe that either the number the cutoff edges, p, equals ℓ or the number of internal nodes with ℓ children, m, is at least $2^{\ell-2} = \log n/4$. Thus, in either case, the bound in Equation 6 is $1/n^{\Omega(\log\log n)}$. Further, since there are at most $\ell^2 2^\ell$ values for p, the total probability of a pruned witness tree is at most $\ell^2 2^\ell \cdot 1/n^{\Omega(\log\log n)}$ which is $1/n^{\Omega(\log\log n)}$. This completes the proof of the theorem. ∎

4 Deletions based on item age

We now consider an alternative scenario, where items are deleted in order of their insertion time. We define a process P'_d based on phases: in the first phase, there are n insertions, where again all insertions are made by putting a ball into the least loaded of d bins chosen independently and uniformly at random. In subsequent phases, there are first n insertions, and then the items inserted in the previous phase are deleted.

One way to view this process is as a modified version of the process $P_d(\vec{v})$ in the case where $\vec{v} = (1, 2, 3, \ldots)$. The difference here is that deletions and insertions do not alternate; instead, deletion requests are batched and acted on at the end of a phase. (Unfortunately, we do not yet have an argument showing bounds no this modified version of $P_d((1, 2, \ldots))$ necessarily also hold on the original process, even though it seems natural to conclude that batching deletions until a later time can only worsen performance.)

This phase-based system allows us to regard the state of each bin as a two-dimensional variable. A bin is said to be in state (i, j) if it has i balls that will be deleted in the next deletion phase and at least j balls that have been inserted in the current insertion phase. Such two-dimensional models have previously proven useful for dynamic variations of load balancing problems [11]. We prove bounds for this phased-based system; for convenience we consider only the case $d = 2$.

Theorem 4 *For any fixed constant c, the maximum load of a bin achieved by process $P'_2(1, 2, \ldots)$ over $T = n^c$ steps is $\log\log n/\log 2 + O(1)$ with high probability.*

Proof: [Sketch] We vary Theorem 1 so that it works on phases. That is, consider a time interval of n insertions preceding a deletion phase. Let $X_{i,j}(t)$ be the number of bins with i balls that will be deleted and at least j newly inserted balls after the tth insertion. Note that $X_{i,j}(0) = 0$ unless $j = 0$.

Our goal is to show there is a simple "stable" bounding distribution with the following property: if we begin with $X_{i,0}(0) \leq \beta_i$ for some appropriate sequence β_i, then after n insertions and n deletions, we again have $X_{i,0}(0) \leq \beta_i$ with suitably high probability. This will imply that the process continues for polynomially many steps before the load becomes too high, assuming we begin properly.

Suppose

$$X_{i,0}(0) \leq \frac{\alpha n}{i}\gamma^{2^i}$$

for sufficiently large $i \geq L$, where L and γ are suitable constants. (For example, we may take $L = 20$, $\alpha = 1/20$, and $\gamma = (1 - \frac{1}{2^L})$; note $\gamma^{2^L} \approx \frac{1}{e}$.) It can be checked that this condition holds after the first n insertions in a straightforward manner. We will show that

$$X_{i,j}(n) \leq \frac{\alpha n}{i+j}\gamma^{2^{i+j}}$$

for $i + j \geq L$ with high probability.

Further, let $\hat{X}_{i,0} = \sum_{k=0}^{\infty} X_{k,i}(n)$; $\hat{X}_{i,0}$ is just the number of bins with i balls after the insertions and deletions complete. We will also show that

$$\hat{X}_{i,0}(0) \leq \frac{\alpha n}{i}\gamma^{2^i}$$

for $i \geq L$ with high probability; this will allow us to continue the process for polynomially many steps. (Actually, technically we only require these conditions hold for up to the point where $i + j$ is $\log\log n/\log 2 + O(1)$, so that there are still $\Omega(\log n)$ bins of this height. Once the number of bins at a state becomes sufficiently small so that Chernoff bounds no longer apply, we must use a more explicit tail argument, as in Theorem 1. This affects the $O(1)$ term of the theorem, which depends on c. We skip these details here.)

We prove this inductively on $i + j$ in a similar fashion to the induction in Theorem 1. Define

$$y_{i,j} \stackrel{\text{def}}{=} \frac{\alpha}{i+j}\gamma^{2^{i+j}}.$$

Let $\mathcal{E}_g$ be the event that $X_{i,j}(n) \leq y_{i,j}$ for all $i + j = g$. Now for a fixed pair $(i, j + 1)$ with $i + j = g$ consider a series of binary random variables Y_t for $t = 1, \ldots, n$, where $Y_t = 1$ iff the tth ball lands in a bin in state $(i, j + 1)$ (after its entry) and $\mathcal{E}_g$. (The value Y_t is 1 if the height of the ball t is at least $i + j + 1$ and i balls in its bin are to deleted, despite the fact that the number of such bins has not grown too large.)

Let ω_i represent the choices available to the ith ball. Then

$$\Pr(Y_t = 1 \mid \omega_1, \ldots, \omega_{t-1}) \leq y_{i,j}^2 + 2 \sum_{\substack{k+l=g \\ (k,l)\neq(i,j)}} y_{i,j}y_{k,l} + 2\sum_{k>g} y_{i,j}y_{k,0} = P_{i,j}.$$

We simplify the above expression:

$$\begin{aligned} P_{i,j} &\leq \left(\frac{\alpha^2\gamma^{2^{i+j+1}}}{(i+j)^2}\right)\left(1 + 2(i+j) + 2\sum_{k>g}\gamma^{2^k - 2^{i+j}}\right) \\ &\leq 3\alpha\left(\frac{\alpha\gamma^{2^{i+j+1}}}{i+j}\right). \end{aligned}$$

In the last inequality, we bounded the last summation using the specified value for γ and the fact that $g \geq L$.

By the Chernoff bound $\sum_{t=1}^{n} Y_t \leq \frac{1}{2}ny_{i,j}$ with probability at least $1 - 1/n^{2c+2}$ as long as $P_{i,j} \geq (2c + 2)\ln n$. But here $\sum_{t=1}^{n} Y_t = X_{i,j}(n)$ conditioned on $\mathcal{E}_g$. Hence, as long as $i + j = o(\log n)$,

$$\Pr(\neg\mathcal{E}_{g+1} \mid \mathcal{E}_g) = o\left(\frac{1}{n^{c+1}}\right).$$

Similarly, conditioned on all $\mathcal{E}_g$ for $g \geq L$, for $i \geq L+1$,

$$\begin{aligned}
\hat{X}_{i,0} &= \sum_{k\geq 0} X_{k,i}(n) \\
&\leq \sum_{k\geq 0} \frac{\alpha\gamma^{2^{i+k}}}{2(i+k)} \\
&\leq \frac{\alpha}{i}\gamma^{2^i}.
\end{aligned}$$

Hence, with high probability, the inductive argument goes through, handling all levels of height at least $L+1$. As the probability of failure is $o(\frac{1}{n^{c+1}})$ in any phase, we can go through n^c phases before a failure with high probability.

Note, however, than in using the assumption that the previous phase was well bounded to obtain the bound for the next one, we have "lost" the appropriate bound for $\hat{X}_{L,0}$. This is because to bound $\hat{X}_{L,0}(0)$ we would need to have an initial bound on $X_{L-1,0}(0)$, which we lack. This problem can be easily handled by noting that the number of bins with at least L balls at any point in the process is stochastically dominated by the number of bins with at least L balls when $4n$ bins are thrown into the n bins; this is because we could consider the distribution when each of the at most $2n$ balls in the system has a twin, and a ball or its twin goes into each of the two bins a ball chooses from. This distribution clearly dominates the distribution present in the system. For $L = 20$ and the parameters chosen we can conclude that at the beginning of every phase $\mathcal{E}_L$ holds with probability exponentially small in n, so the inductive proof goes through. ■

As for Theorem 1, the probability of failure can be made lower by increasing the bound on the maximum load appropriately.

References

[1] M. Adler, P. Berenbrink, and K. Schróder. Analyzing an infinite parallel job allocation process. To appear in *ESA 98*.

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

[3] N. Alon and J. H. Spencer. **The Probabilistic Method**. John Wiley and Sons, 1992.

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

[5] P. Berenbrink, F. Meyer auf der Heide, and K. Schróder. Allocating weighted jobs in parallel. In *Proceedings of the 9th Annual ACM Symposium on Parallel Algorithms and Architectures*, 1997, pp. 302–310.

[6] R. Cole, B. M. Maggs, F. Meyer auf der Heide, M. Mitzenmacher, A. W. Richa, K. Schröder, R. K. Sitaraman, and B. Vöcking. Randomized Protocols for Low-Congestion Circuit Routing in Multistage Interconnection Networks. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, 1998, pp. 378–388.

[7] A. Czumaj and V. Stemann. Randomized allocation processes. In *Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science*, 1997, pp. 194–203.

[8] M. Mitzenmacher. Density dependent jump markov processes and applications to load balancing. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, 1996, pp. 213–223.

[9] M. Mitzenmacher. On the analysis of randomized load balancing schemes. In *Proceedings of the 9th Annual ACM Symposium on Parallel Algorithms and Architectures*, 1997, pp. 292–301.

[10] M. Mitzenmacher. Studying balanced allocations with differential equations. Technical Note 1997-024, Digital Equipment Corporation Systems Research Center, Palo Alto, CA, October 1997.

[11] M. Mitzenmacher, How useful is old information? In *Proceedings of the 16th Annual ACM Symposium on Principles of Distributed Computing*, 1997, pp. 83-91. Extended version available as Digital Systems Research Center Technical Note 1998-003.

[12] V. Stemann. Parallel balanced allocations. In *Proceedings of the 8th Annual ACM Symposium on Parallel Algorithms and Architectures*, 1996, pp. 261–269.