

Chains-into-Bins Processes

Tuğkan Batu*

Petra Berenbrink†

Colin Cooper‡

Abstract

The study of *balls-into-bins processes* or *occupancy problems* has a long history. These processes can be used to translate realistic problems into mathematical ones in a natural way. In general, the goal of a balls-into-bins process is to allocate a set of independent objects (tasks, jobs, balls) to a set of resources (servers, bins, urns) and, thereby, to minimize the maximum load. In this paper, we analyze the maximum load for the *chains-into-bins* problem, which is defined as follows. There are n bins, and m objects to be allocated. Each object consists of balls connected into a chain of length ℓ , so that there are $m\ell$ balls in total. We assume the chains cannot be broken, and that the balls in one chain have to be allocated to ℓ consecutive bins. We allow each chain d independent and uniformly random bin choices for its starting position. The chain is allocated using the rule that the maximum load of any bin receiving a ball of that chain is minimized. We show that, for $d \geq 2$ and $m \cdot \ell = O(n)$, the maximum load is $((\ln \ln m) / \ln d) + O(1)$ with probability $1 - \tilde{O}(1/m^{d-1})$.

*Department of Mathematics, London School of Economics, London WC2A 2AE, UK. Email: t.batu@lse.ac.uk.

†School of Computing Science, Simon Fraser University, Burnaby, BC V5A 1S6, Canada. Email: petra@cs.sfu.ca.

‡Department of Computer Science, King's College London, London WC2R 2LS, UK. Email: colin.cooper@kcl.ac.uk.

1 Introduction

The study of *balls-into-bins processes* or *occupancy problems* has a long history. These models are commonly used to derive results in probability theory. Furthermore, balls-into-bins processes can be used as a means of translating realistic load-balancing problems into mathematical ones in a natural way. In general, the goal of a balls-into-bins process is to allocate a set of independent objects (tasks, jobs, balls) to a set of resources (servers, bins, urns). It is assumed that the balls are independent and do not know anything about the other balls. Each ball is allowed to choose a subset of the bins independently and uniformly at random (i.u.r.) in order to be allocated into one of these bins. The performance of these processes is usually analyzed in terms of the maximum load of any bin.

One extreme solution is to allow each ball to communicate with every bin. Thus, it is possible to query the load of every bin and to place the ball into the bin that is least loaded. This allocation process always yields an optimum allocation of the balls. However, the time and the number of communications needed to allocate the balls are extremely large. The opposite approach is to allow every ball to communicate with only one bin. The usual model is for every ball to be thrown into one bin chosen independently and uniformly at random. For the case of m balls allocated to n bins i.u.r., it is well known that a bin that receives $m/n + \Theta\left(\sqrt{(m \log n)/n}\right)$ balls exists with high probability (w.h.p.).¹ An alternative approach which lies between these two extremes, is to allow every ball to select one of $d \geq 2$ i.u.r. chosen bins. The GREEDY[d] process, studied by Azar et al. [1], chooses d i.u.r. bins per ball, and the ball is allocated into the least loaded among these bins. For this process, and $m \geq n$, the maximum number of balls found in any bin, i.e., the *maximum load*, is $m/n + \ln \ln n / \ln d + O(1)$ w.h.p. (see, for example, [1], [2]). Thus, even a small amount of additional random choice can decrease the maximum load drastically compared to a single choice. This phenomenon is often referred to as the “power of two random choices” (see [7]).

In this paper we consider the chains-into-bins problem which can be regarded as a generalization of the balls-into-bins problem. We are given m chains consisting of ℓ balls each. The balls of any chain have to be allocated to ℓ consecutive bins. We allow each chain d i.u.r. bin choices, and allocate the chain using the rule that the maximum load of any bin receiving a ball of that chain is minimized. In this paper, we show that, for $d \geq 2$ and $m \cdot \ell = o(n \cdot (\ln \ln m)^{1/2})$, the maximum load achieved by this algorithm is at most $(\ln \ln m / \ln d) + O((m\ell/n)^2) + O(1)$, with probability $1 - O((\ln m)^d / m^{d-1})$. This result shows that for a fixed number of balls, the maximum load decreases with increasing chain length. The maximum load depends on the number of chains only in the following sense. Allocating $m = n/\ell$ chains of length ℓ (with a total number of n balls) into n bins will, w.h.p., result in a maximum load of at most $\ln \ln(n/\ell) / \ln d + O(1)$. It follows that if $\ell = O((\ln n)^a)$ for any $a > 0$, our result is asymptotically the same as that for allocating n/ℓ balls into n/ℓ bins using GREEDY[d] protocol of [1].

We also prove that the naive heuristic that ‘allocates the chain headers using GREEDY[d] and hopes for the best’ performs badly for some values of m, ℓ as one might expect. Indeed, if $m \geq \ln^2 n$ and $\ell \geq (\ln m) / \ln \ln m$, then the maximum occupancy of this heuristic is at least $(\ln m) / (2 \ln \ln m)$, w.h.p.

¹A sequence of events A_n occurs with high probability if $\lim_{n \rightarrow \infty} P(A_n) = 1$.

Clustering. It can be seen that, provided we can make some extra assumptions, the results of [1] can be applied to the chains-into-bins problem for m chains of length ℓ . Suppose we are *allowed to cluster the bins* into $N = n/\ell$ clusters of ℓ successive bins, and each chain *can be allocated directly to one cluster*. This is now equivalent to allocating m balls into N bins. Thus, we would get a maximum load of $\Theta((\ln \ln N)/\ln(d) + 1)$ with GREEDY[d] and $(\ln \ln N)/d + O(1)$ using the ALWAYS-GO-LEFT[d] protocol. However, this solution essentially ignores the model under consideration, and is equivalent in a hashing context to saying that we do not need to hash the data item at the given location, but rather somewhere in the next ℓ cells at our convenience. If we have this freedom to ignore the locations we are given and pack the balls into $N = n/\ell$ clusters of length ℓ , then why not $N = n/(2\ell)$ clusters of length 2ℓ , placing the chains one after the other? Indeed why not arbitrary N , or even $N = 1$ and pack the chains cyclically in a round-robin fashion? That would be even more efficient. Obviously if such clustering were available it would be easier to organize the behavior of the balls. We assume henceforth that we have to put the chains where we are instructed, rather than where we would like to. Finally, we remark that, provided $m = N = n/\ell$, we can get the same results without restructuring the problem, and thus, the extra provisions are unnecessary.

Applications. Our model can be viewed as a form of hashing in which the first data item of the chain is placed in the selected position of the hash table, and the remaining items overflow into the neighbouring positions of the table.

The chains-into-bins problem has several important applications. One example is data storage on disk arrays, such as RAID systems (see [9]). Here, each data item is stored on several neighbouring disks in order to increase the data transfer rate. In this case, the bins model the disks from the storage array, and the chains model data requests which are directed to several neighbouring “bins.” A second application is the scheduling of reconfigurable embedded platforms (see [4, 11]). Here, the tasks and the reconfigurable chip are modelled as rectangles with integral dimensions. All tasks have the same height but different length. The chip is modelled by a much larger rectangle that can hold several tasks in both dimensions. The goal is to allocate the tasks to a chip with a fixed length such that the required height is minimized. In this case, the tasks are modelled by the chains and the chip is modelled by the bins. The problem also models a scheduling problem where m allocated items persist in the system for ℓ time steps. For example, imagine a train travelling in a circle with n station stops. The bins represent stations and the length of a chain represent the number of stops travelled by a passenger.

1.1 Related Work

Azar et al. [1] introduced GREEDY[d] to allocate n balls into n bins. GREEDY[d] chooses d bins i.u.r. for each ball and allocates the ball into a bin with the minimum load. They show that after placing n balls, the maximum load is $\Theta((\ln \ln n)/\ln d)$, w.h.p. Compared to single-choice processes, this is an *exponential* decrease in the maximum load. For the case where $m < n$, their results can be extended to show a maximum load of at most $(\ln \ln n - \ln \ln(n/m))/\ln d + O(1)$, w.h.p. Vöcking [13] introduced the ALWAYS-GO-LEFT[d] protocol, which clusters the bins into d clusters of n/d consecutive bins each. Every ball now chooses i.u.r. one bin from every cluster

and is allocated into a bin with the minimum load. If several of the chosen bins have the same minimum load, the ball is allocated into the “leftmost” bin. The protocol yields a maximum load of $(\ln \ln n)/d + O(1)$, w.h.p. In [5], Kenthapadi and Panigrahy suggest an alternative protocol yielding the same maximum load. They cluster the bins into $2n/d$ clusters of $d/2$ consecutive bins each. Every ball now randomly chooses 2 of these clusters and it is allocated into the cluster with the smallest total load. In the chosen cluster, the ball is then allocated into the bin with minimum load again. The authors also argue in that paper that clustering is essential to reduce the load to $(\ln \ln n)/d + O(1)$. In [2], the authors analyse GREEDY[d] for $m \gg n$. It is shown that the maximum load is $m/n + \ln \ln(n)$, w.h.p. Mitzenmacher et al. [8] show that a similar performance gain occurs if the process is allowed to memorize a constant number of bins with small load.

In [10], Sanders and Vöcking consider the *random arc allocation problem*, which is closely related to the chains-into-bins problem. In their model, they allocate arcs of an arbitrary length to a cycle. Every arc is assigned a position i.u.r. on the cycle. The chains-into-bins problem with $d = 1$ can be regarded as a special discrete case of their problem, where the cycle represents the n bins and the arcs represent the chains (in [10], different arc lengths are allowed). Translated into the chains-into-bins setting, the authors show the following result. If $m = n/\ell$ chains of length ℓ are allocated to n bins ($m \rightarrow \infty$), then the maximum load is at most $(\ln(n/\ell))/(\ln \ln(n/\ell))$, w.h.p. Note that their result is asymptotically the same as that for allocating n/ℓ balls into n/ℓ bins, provided that $n/\ell \rightarrow \infty$. In [3], the author shows that the expected maximum load is smaller if we allocate $n/2$ chains of length 2 with one random choice per chain, compared to n balls into n bins with $d = 2$.

2 Model and Results

Assume m chains of length ℓ are allocated i.u.r. to bins wrapped cyclically round $1, \dots, n$. A chain contains ℓ balls linked together sequentially. The first ball of a chain is called *header*; the remaining balls comprise the *tail* of the chain. If chain i (meaning the header of chain i) is allocated to bin j , then the balls of the chain occupy bins $j, j + 1, \dots, j + \ell - 1$, where counting is modulo n . We define the *h-load* of a bin as the number of headers allocated to the bin. This is to be distinguished from the *load* of bin j . The load is the total number of balls allocated to bin j ; that is, the number of chain headers allocated to bins $j - \ell + 1, \dots, j - 1, j$.

We consider the case where each chain header randomly chooses d bins j_1, \dots, j_d . For random choice j_k it computes the maximum load of bins $j_k, j_k + 1, \dots, j_k + \ell - 1$. The chain header is allocated to the bin $j_k \in j_1, \dots, j_d$ such that the maximum load is minimized. This allocation process is called GREEDY_CHAINS[d, ℓ].

We show the following result, which is proved in Section 3.

Theorem 1 *Let $m \leq n$, $\ell \geq 1$, $d \geq 2$, and assume $m \cdot \ell \geq n/(2e)$ and $m \cdot \ell = o(n(\ln \ln m)^{1/2})$. Let m chains of length ℓ be allocated to n bins with d i.u.r. bin choices per chain header. The maximum load of any bin obtained by GREEDY_CHAINS[d, ℓ] is at most*

$$\frac{\ln \ln m}{\ln d} + O\left(\left(\frac{m \cdot \ell}{n}\right)^2\right),$$

with probability $1 - O((\ln m)^d/m^{d-1})$.

Note that when $m \cdot \ell = O(n)$, $\text{GREEDY_CHAINS}[d, \ell]$ achieves a maximum load of $(\ln \ln m)/(\ln d) + O(1)$, with high probability. In order to make a direct comparison, we extend of the results of [1] on the algorithm $\text{GREEDY}[d]$ to the case where $m < n$.

Theorem 2 *Assume that $d \geq 2$, $m \leq n$, and that c is an arbitrary constant. Then, the maximum load achieved by $\text{GREEDY}[d]$ after the allocation of m balls is at most*

$$\frac{\ln \ln n - \ln \ln(n/m)}{\ln d} + c$$

with a probability of $1 - O(n^{-s})$, where s is a constant depending on c .

Theorem 2 gives the bin load arising from chain headers (ignoring the rest of the chain). Since other collisions can occur, for example, between chain headers and internal links of the chain, this will always be a lower bound on the maximum load. Then, provided $m\ell/n = O(1)$,

$$\frac{\ln \ln n - \ln \ln(n/m)}{\ln d} + O(1) \leq \max \text{load} \leq \frac{\ln \ln m}{\ln d} + O(1).$$

We see that, provided that $\ell = e^{(\ln n)^{o(1)}}$ (in particular, ℓ is poly-logarithmic in n), the ratio of the upper and lower bounds on the maximum load is $(1 + o(1))$.

Finally, suppose we allocate chain headers using $\text{GREEDY}[d]$ but ignore the effect of this allocation on the rest of the chain. The following theorem, proved in Section 4, shows that this approach leads to a large maximum occupancy.

Theorem 3 *Assume that $m \cdot \ell \geq n/(2e)$, that $m < n/(2e)$, that $m \geq \ln^2 n$, and that $\ell \geq (\ln m)/(\ln \ln m)$. Then, the maximum occupancy of any bin based on $\text{GREEDY}[d]$ allocation of chain headers is at least $(\ln m)/(2 \ln \ln m)$, w.h.p.*

The proofs of Theorem 2 and Theorem 3 can be found in Section 4.

3 Analysis of $\text{GREEDY_CHAINS}[d, \ell]$

In this section, we prove Theorem 1. The proof uses layered induction. In the case of $\text{GREEDY}[d]$, Azar et al. [1] use variables γ_i as a high-probability upper bound on the number of bins with i or more balls, where $\gamma_6 = n/2e$ and, for $i > 6$, $\gamma_i = e \cdot n \cdot (\gamma_{i-1}/n)^d$.

Since we allocate chains into bins, we cannot consider only the number of bins with i or more chain headers, we have to consider both the chain headers and tails. Hence, to calculate the load of a bin, we have to consider the chain headers allocated to neighbouring bins. To do so, we define the

set S_i which can be thought of as the set of bins which will result in a maximum load of at least $i + 1$ if one of the bins in S_i is chosen for a chain header. The set S_i contains all bins j with load (at least) i and the bins at distance at most $\ell - 1$ in front of bin j . We emphasize that not all bins in S_i have load of i themselves. We use variables β_i as high-probability upper bounds and show that, for i large enough, $|S_i| \leq \beta_i = 2e \cdot m \cdot \ell \cdot (\beta_{i-1}/n)^d$, w.h.p., in our induction. In the following, we define some sets and random variables that are used in our analysis.

- Let $\lambda_j(t)$ be a random variable counting the h -load of bin j . That is, $\lambda_j(t)$ is the number of chain headers allocated to bin j at (the end of) step t for $t = 0, 1, \dots, m$.
- For given $A \subseteq [n]$, define $\lambda_A(t) = \sum_{j \in A} \lambda_j(t)$. Thus, $\lambda_A(t)$ is a random variable counting the total h -load of the bins in A at the end of step t .
- Let $R_j = \{j - \ell + 1, \dots, j - 1, j\}$, the set of bins that will increase the load of bin j if a chain header is allocated to them.
- Let $L_j(t)$ be a random variable counting the load of bin j at (the end of) step t . Thus, $L_j(t) = \lambda_{R(j)}(t)$, the load arising from the chain headers allocated to the ℓ bins of $R(j)$.
- Let $Q_i(t) = \{j : L_j(t) \geq i\}$ be the set of labels of bins whose load is at least i at (the end of) step t .
- Let $S_i(t) = \cup_{j \in Q_i(t)} R_j$. Thus, $S_i(t)$ contains the labels of bins such that an allocation of a chain header to one of these bins will increase the load of a bin with a load of at least i by 1.
- Let $\theta_{\geq i}(t) = |S_i(t)|$.
- Let h_t be a random variable counting the height of chain t . The algorithm GREEDY_CHAINS $[d, \ell]$ allocates the header to the bin which minimizes the maximum total load h_t , where

$$h_t = 1 + \min_{i=1, \dots, d} \max \{L_{j_i+k}(t-1), k = 0, \dots, \ell - 1\} \quad (1)$$

and j_1, \dots, j_d are the bins chosen i.u.r. at step t .

Our method of proving Theorem 1 uses an approach developed in [1], but incorporates the added complexity of considering the maximum load over the chain length. For consistency, we have preserved notation as far as possible.

Let $\alpha = m\ell/n$ with $\alpha \geq 1/2e$ and $k = \lceil 8\alpha^2 e \rceil$. First, we show that i chains can contribute a block of bins of length at most $2\ell - 1$ to $S_i(m)$.

Lemma 1 For $i \geq k$, (1) $\theta_{\geq i}(m) \leq 2m\ell/i$, (2) $\theta_{\geq i}(m) \leq n/2$.

Proof To prove part (1) we first consider a the following worst case scenario. Suppose at step t bin j contains i chain headers, bins $j - \ell + 1, \dots, j - 1$ are empty, and bins $j + 1, \dots, j + \ell - 1$ do not contain any chain headers. Then, $\{j, j + 1, \dots, j + \ell - 1\} \subseteq Q_i(t)$, $\{j - \ell + 1, \dots, j, j + 1, \dots, j + \ell - 1\} \subseteq S_i(t)$ and $|S_i(t)| \geq 2\ell - 1$.

Now suppose that $|S_i(t)| = r$ and ask the question how many chain headers do we need for that. This number is minimized when the chain headers are aligned, as demonstrated in our worst case example. Then, every set of i chain headers covers $2\ell - 1$ bins. This means that we need

$$r \geq \frac{i \cdot |S_i(t)|}{2\ell - 1}$$

chain headers. In general, for $t \leq m$, we get

$$\lambda_{S_i(t)}(t) \geq \frac{i \cdot |S_i(t)|}{2\ell - 1} > \frac{i \cdot \theta_{\geq i}(t)}{2\ell}$$

since $\theta_{\geq i}(t) = |S_i(t)|$. For $t = m$, we have

$$\frac{i \cdot \theta_{\geq i}(m)}{2\ell} \leq \lambda_{S_i(m)}(m) \leq m. \quad (2)$$

Part (2) follows from part (1). Since $i \geq k$, $\alpha = m\ell/n$, and $\alpha \geq 1/(2e)$, we get

$$\theta_{\geq i}(m) \leq \theta_{\geq k}(m) \leq \frac{2m\ell}{k} \leq \frac{2m\ell}{8\alpha^2 e} = \frac{2m\ell}{8 \cdot (m\ell/n) \cdot \alpha e} = \frac{n}{4\alpha e} \leq \frac{n}{2}.$$

□

To prove Theorem 1, we define

$$\beta_i = \begin{cases} n & i = 1, \dots, k-1; \\ \frac{n}{4e\alpha} & i = k; \\ 2em\ell \cdot \left(\frac{\beta_{i-1}}{n}\right)^d & i > k. \end{cases}$$

For $i \geq 0$ and $j = k + i$, it follows from the definition of β_i that

$$\beta_j = \beta_{k+i} = n \cdot \frac{(2e)^{\frac{d^i-1}{d-1}}}{(4e\alpha)^{d^i}} = n \cdot 2^{-d^i} \cdot (2e\alpha)^{-(d^i(d-2)+1)/(d-1)}, \quad (3)$$

and, thus, provided $2e\alpha \geq 1$ (i.e., $m\ell/n \geq 1/2e$), we have $\beta_{k+i} \leq n \cdot 2^{-d^i}$.

Define $\mathcal{E}_i(t) = \{\theta_{\geq i}(t) \leq \beta_i\}$ and let

$$\mathcal{E}_i = \mathcal{E}_i(m) = \{\theta_{\geq i}(m) \leq \beta_i\} \quad (4)$$

be the event that $|S_i(t)|$ is bounded by β_i throughout the process. From the discussion following (2), we have that \mathcal{E}_k holds with certainty. Our goal is to obtain a value for i such that $\Pr(\mathcal{E}_i)$ is close to 1 and, given \mathcal{E}_i , no bin receives more than i balls, with high probability.

We next state a standard lemma (see Lemma 3.1 in [1]), proof of which is omitted.

Lemma 2 *Let X_1, X_2, \dots, X_m be a sequence of random variables with values in an arbitrary domain, and let Y_1, Y_2, \dots, Y_m be a sequence of binary random variables with the property that $Y_t = Y_t(X_1, \dots, X_t)$. If*

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

then

$$\Pr\left(\sum_{t=1}^m Y_t \geq k\right) \leq \Pr(B(m, p) \geq k),$$

where $B(m, p)$ denotes a binomially distributed random variable with parameters m and p .

As the d choices of bins for a chain header are independent, we have that

$$\Pr(h_t \geq i + 1 \mid \theta_{\geq i}(t-1) = r) \leq \left(\frac{r}{n}\right)^d,$$

where h_t is given by Eq. (1). For chain t and integer i , let $Y_t^{(i)}$ be an indicator variable given by

$$Y_t^{(i)} = 1 \iff \{h_t \geq i + 1, \theta_{\geq i}(t-1) \leq \beta_i\}.$$

Let $X_i = (x_i^1, \dots, x_i^d)$ denote the set bin choices of i th chain header, and let $X_{1,t} = (X_1, \dots, X_t)$ be the choices of the first t chains. We define $\mathbf{X}_{1,t}$ as the event $\{\mathbf{X}_{1,t} = (X_1, \dots, X_t)\}$.

Assume $\mathbf{X}_{1,t-1} \in \mathcal{E}_i(t-1)$, meaning after the allocation of the first $t-1$ chains, we have at most β_i bins that would result in a load of at least $i+1$ when hit by chain t . Then,

$$\Pr(Y_t^{(i)} = 1 \mid \mathbf{X}_{1,t-1}) \leq \left(\frac{\beta_i}{n}\right)^d,$$

and, if $\mathbf{X}_{1,t-1} \notin \mathcal{E}_i(t-1)$, then $\Pr(Y_t^{(i)} = 1 \mid \mathbf{X}_{1,t-1}) = 0$. Either way,

$$\Pr(Y_t^{(i)} = 1 \mid \mathbf{X}_{1,t-1}) \leq \left(\frac{\beta_i}{n}\right)^d \triangleq p_i.$$

We can apply Lemma 2 to conclude that

$$\Pr\left(\sum_{t=1}^m Y_t^{(i)} \geq r\right) \leq \Pr(B(m, p_i) \geq r). \quad (5)$$

Considering the extreme case discussed above in Lemma 1, we see that each event $\{Y_t^{(i)} = 1\}$ adds at most an extra $2\ell - 1$ bins to $S_{i+1}(t)$. Thus for $\mathbf{X}_{1,m} \in \mathcal{E}_i$,

$$\theta_{\geq i+1}(m) \leq 2\ell \cdot \sum_{t=1}^m Y_t^{(i)}. \quad (6)$$

Let $r_i = e \cdot m \cdot p_i$. Then, provided that $\sum_{t=1}^m Y_t^{(i)} \leq r_i$, we have

$$\theta_{\geq i+1}(m) \leq 2\ell r_i = 2\ell e m \cdot p_i = 2em\ell \cdot \left(\frac{\beta_i}{n}\right)^d = \beta_{i+1}. \quad (7)$$

From (5) and (6), we have

$$\Pr\left(\theta_{\geq i+1}(m) > 2\ell \cdot r_i \mid \mathcal{E}_i\right) \leq \Pr\left(\sum_{t=1}^m Y_t^{(i)} > r_i \mid \mathcal{E}_i\right) \leq \frac{\Pr(B(m, p_i) \geq r_i)}{\Pr(\mathcal{E}_i)}. \quad (8)$$

Provided that $m \cdot p_i \geq 2\ln \omega$ (where the precise value of ω is established below in (13)), using the Chernoff bounds, we get

$$\Pr(B(m, p_i) \geq em \cdot p_i) \leq e^{-m \cdot p_i} \leq \frac{1}{\omega^2}. \quad (9)$$

Recall that $\Pr(\neg \mathcal{E}_k) = 0$. Assume inductively that $\Pr(\neg \mathcal{E}_i) \leq i/\omega^2$, for $i \geq k$. Since

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

we have, from (4), (7), (8), and (9), that

$$\Pr(\neg \mathcal{E}_{i+1}) \leq \frac{i+1}{\omega^2}.$$

Choose i^* as the smallest i such that $p_i = \left(\frac{\beta_i}{n}\right)^d \leq \frac{2\ln \omega}{m}$. From (3),

$$i^* - k \leq \frac{\ln \ln(m/\ln \omega)}{\ln d} + O(1). \quad (10)$$

Also, as $\alpha = m\ell/n = o((\ln \ln m)^{1/2})$, we have that $k = o(i^*)$ so that the induction is not empty. Since $p_{i^*} \leq (2\ln \omega)/m$, we have that

$$\begin{aligned} \Pr(\theta_{\geq i^*+1}(m) \geq (2\ell) \cdot 6\ln \omega \mid \mathcal{E}_{i^*}) &\leq \frac{\Pr(B(m, p_{i^*}) \geq 6\ln \omega)}{\Pr(\mathcal{E}_{i^*})} \\ &\leq \frac{\Pr(B(m, (2\ln \omega)/m) \geq 6\ln \omega)}{\Pr(\mathcal{E}_{i^*})} \\ &\leq \frac{1}{\omega^2 \cdot \Pr(\mathcal{E}_{i^*})}, \end{aligned}$$

and, thus,

$$\Pr(\theta_{\geq i^*+1}(m) \geq (2\ell) \cdot 6\ln \omega) \leq \frac{1}{\omega^2 \cdot \Pr(\mathcal{E}_{i^*})} \cdot \Pr(\mathcal{E}_{i^*}) + \Pr(\neg \mathcal{E}_{i^*}) \leq (i^* + 1)/\omega^2. \quad (11)$$

Conditioned on $\theta_{\geq i^*+1}(m) \leq 12\ell \ln \omega$, the probability that a chain is placed at height at least $i^* + 2$ is at most $(12\ell \ln \omega/n)^d$. Given that $Y \sim B(m, (12\ell \ln \omega/n)^d)$, $\Pr(Y \geq 1) \leq m(12\ell \ln \omega/n)^d$ by Markov's Inequality. Thus applying Lemma 2, we get

$$\Pr\left(\sum_{t=1}^m Y_t^{(i^*+1)} \geq 1 \mid \theta_{\geq i^*+1}(m) \leq (2\ell) \cdot 6\ln \omega\right) \leq \frac{m \cdot \left(\frac{12\ell \cdot \ln \omega}{n}\right)^d}{\Pr(\theta_{\geq i^*+1} \leq 12\ell \cdot \ln \omega)}. \quad (12)$$

Let ω satisfy

$$\omega = \left(\frac{m^{d-1} \cdot \ln \ln m}{(\ln m)^d} \right)^{1/2}. \quad (13)$$

Using $m \cdot \ell = o(n(\ln \ln m)^{1/2})$, (13), (12) and (11), the probability that there is a bin with load at least $i^* + 2$ is bounded by a term of order

$$\frac{i^* + 1}{\omega^2} + O\left(\frac{(\ln \omega)^d}{m^{d-1}}\right) = O\left(\frac{(\ln m)^d}{m^{d-1}}\right).$$

By plugging in ω and k into (10), we get

$$i^* \leq \frac{\ln \ln m}{\ln d} + O(\alpha^2) + O(1).$$

Thus, w.h.p, no bin receives more than $i^* + 1$ balls. \square

4 Allocating Chain Headers

In this section we present the proofs of Theorem 2 and Theorem 3.

4.1 Proof of Theorem 2

This theorem can be shown similarly to the proof of Theorem 4 in [1]. We define $\gamma_1 = \gamma_2 = \dots = \gamma_5 = n$, $\gamma_6 = m/(2e)$, and

$$\gamma_i = em \cdot \left(\frac{\gamma_{i-1}}{n} \right)^d \text{ for } i > 6.$$

Thus $\gamma_i = C n(m/(n2e))^{d^i}$ for some $C > 1$ constant. Integer i^* is defined as the smallest i such that $em(\gamma_i/n)^d \leq 6 \ln n$, which holds for

$$i^* \leq \frac{\ln \ln n - \ln \ln(n/m)}{\ln d} + O(1).$$

It can be shown that the maximum load is bounded by $i^* + 2 = (\ln \ln n - \ln \ln(n/m))/\ln d + O(1)$, w.h.p.

4.2 Proof of Theorem 3

The idea of the proof is as follows. Let U_m be the number of bins with load at least one after the allocation of m chains by GREEDY[d] applied to the chain headers. We show that, with a good probability, there exists a strip of ℓ consecutive bins which i) is used by one chain, and ii) at least t of its bins are in U_m . This gives us a bin with load at least t .

First, we find a lower bound on U_m . Azar et al. [1] show that the protocol GREEDY[d] for $d \geq 1$ is majorized by GREEDY[1] in the following sense. Let x_i be the load of the bin with the i th largest load after allocation of m balls with GREEDY[d], and let x'_i be the load of the bin with the i th largest load after allocation of m balls with GREEDY[1]. It is shown in [1] that there exists a one-to-one mapping between the random choices of GREEDY[1] and GREEDY[d] such that for all $1 \leq j \leq n$,

$$\sum_{i=1}^j x_i \leq \sum_{i=1}^j x'_i.$$

From this it follows that the number of empty bins in an allocation with GREEDY[d] is smaller than the one in GREEDY[1]. Let $f(m)$ be the number of occupied bins in an allocation generated by GREEDY[1]. When $m = n/(2e)$ we get

$$E[f(m)] = n \left(1 - \left(1 - \frac{1}{n} \right)^m \right) \geq 0.9m.$$

As m decreases, $m - E[f(m)]$ decreases, too. Thus, provided $m \leq n/(2e)$, the expected number of occupied cells with $d \geq 1$ choices per ball is always at least $0.9m$. By concentration, we can assume $U_m \geq m/2 + 1$, provided $m \geq \ln^2 n$.

Given U_m , we can assume that the locations of the bins occupied by chain headers are sampled uniformly without replacement from $[1, \dots, n]$. We fix one of the m chains and consider the strip of ℓ consecutive bins occupied by the chain. Assuming $\ell \geq 2t$, the probability of at least t bins occupied by additional chain headers in that strip is at least

$$\binom{U_m - 1}{t} \cdot (\ell)_t \cdot \left(\frac{1}{n} \right)^t \geq \binom{m/2}{t} \cdot (\ell)_t \cdot \left(\frac{1}{n} \right)^t \geq \left(\frac{m\ell}{2etn} \right)^t \geq \left(\frac{1}{4e^2 t} \right)^t,$$

as $m\ell \geq n/(2e)$ and $(\ell)_t = \ell(\ell-1)\cdots(\ell-t+1) \geq (\ell/e)^t$.

Let $c = 1/(4e^2)$. Then, the expected number of chains allocated into a strip with load at least $t = \ln m/(2 \ln \ln m)$ is at least

$$\begin{aligned} m \cdot \left(\frac{c}{t} \right)^t &= \exp(\ln m - t \ln t/c) \\ &\geq m^{1/3}, \end{aligned}$$

for large enough m . The probability that such an event does not occur tends to zero by the Chebychev's inequality. \square

5 Conclusions and Open Problems

In this paper we analyse the maximum load for the *chains-into-bins* problem where $m \cdot \ell$ balls are connected in m chains of length ℓ . We show that, provided $m\ell \geq n/2e$ and $m\ell = o(n(\ln \ln m)^{1/2})$, the maximum load is at most $\frac{\ln \ln m}{\ln d} + O((m\ell/n)^2)$, with probability $1 - O((\ln m)^d/m^{d-1})$. This shows that the maximum load is going down with increasing chain length.

Surprisingly, there are many open questions in the area of balls-into-bins processes. Only very few results are known for weighted balls-into-bins processes, where the balls come with weights and the load of a bin is the sum of the weights of the balls allocated to it. Here, it is even not known if two or more random choices improve the maximum load, compared to the simple process where every ball is allocated to a randomly chosen bin (see [12]). Also, it would be interesting to get tight results for the maximum load and results specifying “worst-case” weight distributions for the balls. Something in the flavor “given that the total weight of the balls is fixed, it is better to allocate lots of small balls, compared to fewer bigger ones.” Another interesting problem is to show results relating the maximum load to the order in which the balls are allocated. For example, is it always better to allocate balls in the order of decreasing ball weight, compared to the order of increasing ball weight?

For chains-into-bins problem, an open question is to prove Knuth’s [6] conjecture stating that breaking chains into two parts only increases the maximum load. This question still open for a single choice and also for several random choices per ball. See [3] for a first progress in this direction. Another question is if similar results to the one we showed in this paper for GREEDY[d] applied to chains also holds for the ALWAYS-GO-LEFT protocol from [13] applied on chains.

Finally, we note that it would be interesting to generalize the problem to two dimensional packing, and consider online allocation of m objects of length ℓ and width w to the cells of a toroidal grid of length n and width h .

References

- [1] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, Eli Upfal: Balanced Allocations. *SIAM J. Computing* **29** (1999), pp. 180–200.
- [2] Petra Berenbrink, Artur Czumaj, Angelika Steger, Berthold Vöcking: Balanced Allocations: The Heavily Loaded Case. *SIAM J. Computing* **35** (2006), pp. 1350–1385.
- [3] Matthias Englert: Chains of Length Two into Bins. *Manuscript, University of Aachen*.
- [4] S. Fekete, E. Köhler, J. Teich. Optimal FPGA module placement with temporal precedence constraints. Proc. of the Conference on Design, Automation and Test in Europe (DATE01), pp. 658-667, 2001.
- [5] Krishnaram Kenthapadi, Rina Panigrahy: Balanced allocation on graphs. *Proc. of 17th Annual Symposium on Discrete Algorithms (SODA 06)*, pp. 434-443.
- [6] D. E. Knuth: Sorting and Searching, vol. 3 of *The Art of Computer Programming*, Addison-Wesley, 2nd edition, 1998.
- [7] Michael Mitzenmacher, Andrea W. Richa, Ramesh Sitaraman: The Power of Two Random Choices: A Survey of Techniques and Results. *Handbook of Randomized Computing*, 2000.
- [8] Michael Mitzenmacher, Balaji Prabhakar, Devarat Shah: Load Balancing with Memory. *Proc. of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2002)*, pp. 799–808.

- [9] David A. Patterson, Garth A. Gibson, Randy H. Katz: A Case for Redundant Arrays of Inexpensive Disks (RAID). In Proc. of SIGMOD International Conference on Management of Data, pp. 109-116, 1988.
- [10] Peter Sanders, Berthold Vöcking: Tail Bounds And Expectations For Random Arc Allocation And Applications. In *Combinatorics, Probability & Computing* 12(3), 2003.
- [11] Christoph Steiger, Herbert Walder, Marco Platzner: Operating Systems for Reconfigurable Embedded Platforms: Online Scheduling of Real-Time Tasks. In *IEEE Trans. Computers* 53(11): 1393-1407, 2004.
- [12] Kunal Talwar, Udi Wieder: Balanced allocations: the weighted case. *Proc. of the 39th Symposium on Theory of Computing (STOC)*, pp 256-265, 2007.
- [13] Berthold Vöcking: How Asymmetry Helps Load Balancing. *J. ACM* **50** (4) (2003), pp. 568–589.