Citation for published version:

Berenbrink, P, Khodamoradi, K, Sauerwald, T & Stauffer, A 2013, 'Balls-into-bins with nearly optimal load distribution', Paper presented at SPAA'13: 25th ACM symposium on Parallelism in Algorithms and Architectures, Montreal, Quebec, Canada, 23/07/13 - 25/07/13 pp. 326-335. <https://doi.org/10.1145/2486159.2486191>

DOI:

[10.1145/2486159.2486191](https://doi.org/10.1145/2486159.2486191)

Publication date:

2013

Document Version

Peer reviewed version

[Link to publication](#)

© ACM 2011. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in SPAA '13 Proceedings of the 25th ACM symposium on Parallelism in algorithms and architectures, <http://dx.doi.org/10.1145/2486159.2486191>

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Balls-into-Bins with Nearly Optimal Load Distribution

Petra Berenbrink¹, Kamyar Khodamoradi¹, Thomas Sauerwald², and Alexandre Stauffer³

¹ Simon Fraser University, School of Comp. Science, 8888 University Street, Burnaby, B.C., V5A 1S6, Canada
{petra, kka50}@cs.sfu.ca

² Max-Planck Institut for Informatics, Campus E1.4, 66123 Saarbrücken, Germany
sauerwal@mpi-inf.mpg.de

³ Computer Science Division, University of California, Berkeley, CA 94720-1776, USA,
stauffer@cs.berkeley.edu

Abstract We consider *balls-into-bins* processes that randomly allocate m balls into n bins. We analyze two allocation schemes that achieve a close to optimal maximum load of $\lceil m/n \rceil + 1$ and require only $\mathcal{O}(m)$ allocation time. These parameters should be compared with the classic d -choice-process which achieves a maximum load of $m/n + \log \log n/d + \mathcal{O}(1)$ and requires $m \cdot d$ allocation time.

1 Introduction

The goal of *balls-into-bins* processes is to allocate m balls into n bins. This is done by allowing each *ball* to choose its location among one or several randomly chosen bins. Two of the most important performance measures for these processes are the total number of random bin choices used for the allocation (which is usually a good measure for the allocation time) and the maximum load. In this paper we consider two allocation schemes that are almost optimal in *both* criteria. Both protocols achieve close to optimal maximum load of $\lceil m/n \rceil + 1$ and use only $\mathcal{O}(m)$ random bin choices. Balls-and-bins processes have a wide range of applications in the areas of hashing, load balancing, and resource allocation. In these applications, each ball represents a task or request, while each bin represents a server or processor.

The classical single-choice balls-into-bins process simply places each of the m balls into a bin chosen independently and uniformly at random. For $m = n$, it is well known that the maximum load is $\log n / \log \log n + \mathcal{O}(1)$ with high probability [15]. Karp et al. [11] observe that the maximum load can be vastly reduced if every ball is allowed to choose between two randomly chosen bins and is placed in the least loaded among them. This apparently small change reduces the maximum load to $\mathcal{O}(\log \log n)$ balls, with high probability, and has become widely known as “power of two choices”. Azar et al. [4] analyze a process called GREEDY[d] where every ball is allowed to choose the bin with minimum load among d randomly chosen bins. They prove that this process achieves a maximum load of $\ln \ln n / \ln d + \mathcal{O}(1)$, with high probability. In [5], the authors extend these results to the heavily loaded case where the number of balls m is much larger than the number of bins n . They show that the maximum load is upper bounded by $m/n + \ln \ln n / \ln d + \mathcal{O}(1)$. For the case $m = n$, Vöcking [16] proves a general lower bound on the maximum load of $\ln \ln n / (d \cdot \ln \Phi_d)$, where $1.61 \leq \Phi_d \leq 2$, if every ball is allowed to choose d random bins. He also presents the LEFT[d] protocol that uses an asymmetric tie-breaking rule. Surprisingly, the maximal load achieved by this process matches his lower bound up to an additive constant. Mitzenmacher, Prabhakar and Shah [14] consider a balls-into-bins process (for $m = n$) where every ball has some memory. They consider the (d, k) -memory model where each ball chooses $d + k$ bins, d of which are selected uniformly at random, whereas the other k are the least loaded bins among the ones picked for the previous ball. The ball is then allocated into the least loaded among the $d + k$ bins. For $d = k = 1$, they show that the maximum load is at most $\ln \ln n / (2 \cdot \log(\Phi_2)) + \mathcal{O}(1)$, which matches the aforementioned lower bound by Vöcking [16] up to an additive constant.

Note that in all above protocols, if $m = n$, the allocation time is $dm = dn$ while the maximum load is at least $\Omega(\log \log n/d)$. Therefore a natural question is whether there are allocation protocols that achieve a better tradeoff between allocation time and maximum load, e.g., protocols that achieve a constant maximum load and an allocation time of $\mathcal{O}(n)$.

Czumaj and Stemmann [7] study several balls-into-bins processes where the number of choices per ball depends on the load of the chosen bins of that ball. They show various tradeoffs between the average allocation time, maximum allocation time and the maximum load of a bin. For example, they consider the THRESHOLD protocol which is defined as follows. Every ball repeatedly samples bins until it finds a bin with load less than $m/n + 1$. They show that the (total) allocation time is $(1.146 + o(1))m$ with high probability, provided that $m = n$ [7, Theorem 4]. Note that this does not contradict the lower bound of Vöcking, since certain balls have to choose $\Omega(\log n)$ different bins before they are finally placed.

Czumaj, Riley and Scheideler [6] consider an algorithm that first calculates an initial allocation using GREEDY[d]. Then the algorithm performs iteratively so-called self-balancing steps where balls may switch between the two initial bin choices. Their algorithm achieves a maximum load of $\lceil m/n \rceil$ and uses $\mathcal{O}(m) + n^{\mathcal{O}(1)}$ reallocation steps. Note that this algorithm relies on reallocations which are typically expensive. Other allocation schemes that achieve a nearly perfect maximum load and rely on reallocations of the balls are [3, 13].

As a rather different line of research, several studies have considered the parallel version of balls-into-bins problem. This is the case where the process is divided into different rounds, at each of which a number of balls are allocated to bins simultaneously. This model was first introduced by Adler et al. [1]. Quite recently, Lenzen and Wattenhofer [12] studied parallel balls-into-bins processes for the special case where $m = n$. They present a symmetric and adaptive process of allocation of balls into bins that achieves a maximum bin load of 2 in $\log^*(n) + \mathcal{O}(1)$ rounds with a message complexity of $\mathcal{O}(n)$. Their algorithm works in rounds. All balls that are not allocated in round i access k_i bins in round $i + 1$ for increasing values of k_i . Bins that have at most 2 balls accept a randomly chosen ball, which is then allocated to the bin. The maximum number of bins accessed by any ball is $\mathcal{O}(\log n)$.

Similar reallocation schemes are also considered in the area of Cuckoo hashing. Here the goal is to allocate m data items (balls) into n buckets (bins) of size k , and to minimize kn . Every data item comes with d possible buckets. If a new data item arrives, it is allocated into one of its d bucket choices if one of them stores less than k items. Otherwise, one of the conflicting items, denote it by ℓ , is picked and reallocated. If, in turn, none of ℓ 's other bucket choices $b_1, \dots, b_k$ contains less than k items, another item in $b_1, \dots, b_k$ is reallocated. The process stops if all items are allocated to a bucket. There are many results dealing with the best choices for d and k . For an overview of results in the area of Cuckoo hashing see, for example, [8].

1.1 Our Results

In this paper we consider two balls-into-bins protocols. First, we study a new protocol called ADAPTIVE which works as follows. The i -th ball samples bins uniformly at random until it finds a bin with load less than $i/n + 1$ and is then placed into that bin. Unlike THRESHOLD from [7], the number of balls m does not need to be known in advance, as the threshold for each ball is *adaptive* to the current load distribution. From the definition of ADAPTIVE it follows directly that the maximum load is upper bounded by $\lceil m/n \rceil + 1$. We prove that ADAPTIVE requires only $\mathcal{O}(m)$ allocation time in expectation, which is asymptotically the same as THRESHOLD. Hence, our new protocol is nearly optimal in terms of maximum load and allocation time. It should be noted that during the execution of ADAPTIVE, each ball must know how many balls have been already placed. This assumption is comparable to the (d, k) -memory model of [14], where every ball communicates with the ball that comes right after it.

In the second part of the paper, we extend the analysis of THRESHOLD from [7] to the case where m is much larger than n . Note that using the proof technique of [7], we only get an upper bound of $\mathcal{O}(m)$ on the allocation time of THRESHOLD for the cases where $m = \mathcal{O}(n)$. Here in this paper, we prove that the allocation time is $m + \mathcal{O}(m^{3/4} \cdot n^{1/4})$, with high probability (see Theorem 4.1), which is an improvement over the upper bound of [7]. From the definition of THRESHOLD it follows directly that the maximum load is $\lceil m/n \rceil + 1$. Hence, our result shows that, even for $m > n$, THRESHOLD is nearly optimal both in terms of maximum load and allocation time.

The analysis of the ADAPTIVE protocol reveals an interesting result about the smoothness of the load distribution which might be of independent interest. We show that for the ADAPTIVE protocol the difference between the maximum and minimum load is at most $\mathcal{O}(\log n)$, with high probability (see Corollary 3.5). In

Algorithm	Allocation Time	Maximum Load	Conditions on m and n
GREEDY[d] [5]	$\Theta(md)$	$m/n + \frac{\ln \ln n}{\ln d} + \Theta(1)$	–
LEFT[d] [5]	$\Theta(md)$	$m/n + \frac{\ln \ln n}{d \cdot \ln(\Phi_d)} + \Theta(1)$	–
[14]	$\Theta(m)$	$\frac{\ln \ln n}{\ln(\Phi_2)} + \Theta(1)$	$m = n$
[6]	$\mathcal{O}(m) + n^{\mathcal{O}(1)}$	$\lceil m/n \rceil$	$m = \omega(n^6 \log n)$
[6]	$m^{\mathcal{O}(1)}$	$\lceil m/n \rceil + 1$	$m = \mathcal{O}(n \log n)$
THRESHOLD [7, Theorem 4]	$1.146194m + o(m)$	$\lceil m/n \rceil + 1$	$m = n$
THRESHOLD [7, Theorem 11]	$\mathcal{O}(m)$	$\lceil m/n \rceil + 1$	$m = \mathcal{O}(n)$
ADAPTIVE ★	$\mathcal{O}(m)$	$\lceil m/n \rceil + 1$	–
THRESHOLD ★	$m + \mathcal{O}(m^{3/4} \cdot n^{1/4})$	$\lceil m/n \rceil + 1$	–

Table 1: Comparison of the allocation time and maximum load of various allocation schemes for m balls and n bins. The rows with ★ are our new results. Note that $1.6 < \Phi_d < 2$, see [5] for the precise definition.

contrast to that, the difference is polynomial in n for the THRESHOLD protocol (see Lemma 4.2). Altogether, these results reveal that the load distribution of ADAPTIVE is much smoother than the one of THRESHOLD.

We close the paper by presenting some experimental results indicating that in practice ADAPTIVE requires only a slightly larger allocation time than THRESHOLD, but achieves a much smoother load distribution. Due to space limitations, some of the proofs are given in the appendix.

2 Algorithms and Notation

We first recall the following protocol from [7] which we call THRESHOLD. Clearly, the maximum load at the end of the THRESHOLD protocol is at most $\lceil m/n \rceil + 1$. We are interested in the *allocation time* of all balls, i.e., the total number of bins that have to be sampled in order to place all m balls.

```

for each ball  $i$  from 1 to  $m$  do
  repeat
    Choose a bin  $j$  independently and uniformly at random from  $\{1, \dots, n\}$ 
    if load of bin  $j$  is strictly less than  $i/n + 1$  then place ball  $i$  into bin  $j$ 
  until the ball is placed
end for

```

Figure 1: Our new ADAPTIVE protocol.

The ADAPTIVE protocol is very similar to the THRESHOLD protocol, except that the “threshold” is a function of the number of balls placed so far.

Similar to the THRESHOLD protocol, the ADAPTIVE protocol ensures that the maximum load of any bin is bounded by $\lceil m/n \rceil + 1$. Note that if we replace in the ADAPTIVE protocol the threshold $i/n + 1$ by i/n , then the allocation time of each batch of n consecutive balls is basically a coupon collector process which translates into an overall allocation time of $\Theta((m/n) \cdot n \log n) = \Theta(m \log n)$.

We denote by $L^t = (L_1^t, \dots, L_n^t)$ the random vector giving the load distribution at the end of step t , i.e., $L_i^t = x$ means that, among the first t balls, x have been allocated to bin i . We will use $\ell^t = (\ell_1^t, \dots, \ell_n^t)$ to denote a fixed load distribution; i.e., ℓ^t will be an instantiation of L_t .

```

for each ball  $i$  from 1 to  $m$  do
  repeat
    Choose a bin  $j$  independently and uniformly at random from  $\{1, \dots, n\}$ 
    if load of bin  $j$  is strictly less than  $m/n + 1$  then place ball  $i$  into bin  $j$ 
    until the ball is placed
  end for

```

Figure 2: The THRESHOLD algorithm from [7].

As mentioned in the introduction, the load distribution of balls-and-bins processes has been mostly analyzed in terms of the maximum load. However, the maximum load disregards by how much certain bins are underloaded. Therefore, we will also examine the smoothness of the final load distribution by means of the following two potential functions. One of the most widely used potential functions is the so-called *quadratic* potential [3], which is defined by

$$\Psi_t(\ell^t) := \sum_{i=1}^n (\ell_i^t - t/n)^2,$$

where we will in the following omit the subscript t for simplicity and just write $\Psi(\ell^t)$. We will also use the following *exponential* potential function [10].

$$\Phi_t(\ell^t) = \Phi(\ell^t) := \sum_{i=1}^n (1 + \varepsilon)^{t/n + 2 - \ell_i^t}.$$

In this paper we choose $\varepsilon := 1/200$. It is easy to verify that for any load vector ℓ^t such that $\max_i \ell_i^t \leq t/n + \mathcal{O}(1)$, $\Psi(\ell^t) = \mathcal{O}(\Phi(\ell^t))$. Another nice feature of the exponential potential function is that if L^t is a random load vector, then the probability that there exists a bin with load less than $t/n - 2 - \log_{1+\varepsilon}(\mathbf{E}\{\Phi(L^t)\} \cdot n^2)$ is at most n^{-1} ; this follows directly from

$$\begin{aligned} \mathbf{E}\{\Phi(L^t)\} &\geq \Pr\{L_i^t \leq t/n - 2 - \log_{1+\varepsilon}(\mathbf{E}\{\Phi(L^t)\} \cdot n^2)\} \cdot (1 + \varepsilon)^{t/n + 2 - (t/n - 2 - \log_{1+\varepsilon}(\mathbf{E}\{\Phi(L^t)\} \cdot n^2))} \\ &\geq \Pr\{L_i^t \leq t/n - 2 - \log_{1+\varepsilon}(\mathbf{E}\{\Phi(L^t)\} \cdot n^2)\} \cdot \mathbf{E}\{\Phi(L^t)\} \cdot n^2, \end{aligned}$$

and taking the union bound over all n bins.

Throughout this paper, $\text{Poi}(\lambda)$ denotes the Poisson distribution with parameter (and expected value) λ , and $\text{Bin}(n, p)$ denotes the binomial distribution with n trials and success probability p .

3 Analysis of ADAPTIVE

Theorem 3.1 *Assume $m \geq 8$. Then the expected allocation time of the ADAPTIVE protocol is $\mathcal{O}(m)$.*

3.1 Proof of Theorem 3.1

We assume in the following that $m = \varphi n$, where φ is lower bounded by a sufficiently large integer. Results for $\varphi \notin \mathbb{N}$ can be easily obtained by resorting to the case $m = \lceil \varphi \rceil n$. Moreover, results for the case φ being bounded from above by an integer also follow directly from the observation that there are always $n/\varphi = \Omega(n)$ bins in each step where a ball can be placed to.

Consider now the protocol ADAPTIVE. Since the load of every bin is an integer, the threshold of $i/n + 1$ in the line 4 of the protocol (Figure 1) only changes after n balls are allocated. Therefore, it is natural to divide

the analysis into φ stages of length n each. Stage τ ($1 \leq \tau \leq \varphi$) is responsible for the allocation of the balls $(\tau - 1) \cdot n + 1, \dots, \tau n$. Let $\ell^\tau = (\ell_1^\tau, \dots, \ell_n^\tau)$ be an arbitrary, but fixed load vector at the end of stage τ (note that ℓ^τ is different from L^τ ; in particular, ℓ^t is *not* a random variable). Then we define the potential for ℓ^τ as,

$$\Phi(\ell^\tau) = \sum_{i=1}^n (1 + \varepsilon)^{\tau+2-\ell_i^\tau},$$

where again $\varepsilon = 1/200$. We also define $\Phi_i(\ell^\tau) := (1 + \varepsilon)^{\tau+2-\ell_i^\tau}$.

Since $L_i^{\tau+1} \geq L_i^\tau$, it always holds that

$$\Phi(L^{\tau+1}) \leq (1 + \varepsilon) \cdot \Phi(L^\tau).$$

We also define $\Delta\Phi(L^\tau) := \Phi(L^\tau) - \Phi(L^{\tau+1})$ and $\Delta\Phi_i(L^\tau) := \Phi_i(L^\tau) - \Phi_i(L^{\tau+1})$.

Since the proof of Theorem 3.1 is rather lengthy and technical we break down the proof into 4 lemmas. The proof of Theorem 3.1 follows almost directly from these results. The proofs of the 4 lemmas can be found in Section 3.2.

In the following we call a bin *underloaded* at the end of stage τ if its load is less than $\tau + 2 - C_1$ at the end of the stage, where $C_1 > 0$ is a constant that will be fixed later. Let i be a bin that is underloaded at the end of stage τ . The next lemma shows that the expected number of balls allocated to bin i in stage $\tau + 1$ is slightly larger than 1. Hence, bin i has a good chance of “catching up” with its load.

Lemma 3.2 *Assume $L^\tau = \ell^\tau$ is fixed and let i be a bin with $\ell_i^\tau \leq \tau + 2 - C_1$, where $C_1 > 0$ is a sufficiently large constant. Let $Y_i^{\tau+1}$ be the number of balls allocated to bin i in stage $\tau + 1$. Then for any $0 \leq k \leq C_1$,*

$$\Pr\{Y_i^{\tau+1} \geq k \mid L^\tau = \ell^\tau\} \geq \Pr\{\text{Poi}(199/198) \geq k\} - 2 \cdot 10^{-10}.$$

With the help of Lemma 3.2 and some algebraic manipulation one can show that the potential contributed by underloaded bins expectedly decreases in stage $\tau + 1$ (recall that we choose $\varepsilon = 1/200$).

Lemma 3.3 *Assume $L^\tau = \ell^\tau$ is fixed and let i be a bin with $\ell_i^\tau \leq \tau + 2 - C_1$, where $C_1 > 0$ is the constant from Lemma 3.2. Then there is a constant $\kappa = \kappa(C_1) > 0$ with*

$$\mathbf{E}\{\Delta\Phi_i(\ell^\tau)\} \geq (1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot \kappa = \Phi_i(\ell^\tau) \cdot \kappa.$$

The next lemma shows the potential decrease due to underloaded bins is already sufficient to conclude that $\mathbf{E}\{\Phi(L^{\tau+1}) \mid L^\tau = \ell^\tau\} < \Phi(\ell^\tau)$, if $\Phi(\ell^\tau)$ is sufficiently large.

Lemma 3.4 *Let $\rho_n := (\varepsilon + \kappa)/(\kappa/2) \cdot (1 + \varepsilon)^{C_1} \cdot n$, where κ and C_1 are the constants from Lemma 3.3. For any load vector ℓ^τ with $\Phi(\ell^\tau) \geq \rho_n$,*

$$\mathbf{E}\{\Phi(L^{\tau+1}) \mid L^\tau = \ell^\tau\} \leq \left(1 - \frac{\kappa}{2}\right) \cdot \Phi(\ell^\tau),$$

The next result should be compared to Lemma 4.2 for the THRESHOLD protocol.

Corollary 3.5 *For any stage τ with $1 \leq \tau \leq \varphi$,*

$$\mathbf{E}\{\Phi(L^\tau)\} = (1 + \varepsilon)^2 \rho_n / (\kappa/2) = \mathcal{O}(n),$$

where ρ_n is the value defined in Lemma 3.4. Hence $\Pr\{\max_{i,j} (L_i^\tau - L_j^\tau) = \mathcal{O}(\log n)\} \geq 1 - n^{-1}$. In addition, $\mathbf{E}\{\Psi(L^\tau)\} = \mathcal{O}(n)$, where Ψ is the quadratic potential function (see Section 2).

Finally we relate the runtime of stage τ to $\Phi(L^{\tau-1})$ and make use of the corollary above.

Lemma 3.6 *Let T_τ be the runtime of stage τ . Then,*

$$\mathbf{E}\{T_\tau\} = \mathcal{O}(n) \cdot \mathbf{E}\{\log_{1+\varepsilon}(\Phi^{\tau-1}/n)\} = \mathcal{O}(n).$$

Now Theorem 3.1 follows directly from Lemma 3.6 since

$$\mathbf{E}\left\{\sum_{\tau=1}^{\varphi} T_\tau\right\} = \sum_{\tau=1}^{\varphi} \mathbf{E}\{T_\tau\} = \varphi \cdot \mathcal{O}(n) = \mathcal{O}(m).$$

□

3.2 Proof of the Lemmas

In this section we prove some of the lemmas from the last section, the remaining proofs are given in the appendix.

Proof (Proof of Lemma 3.2). Recall that $m = \varphi n$ with φ being an integer. We assume that $\tau n, \tau \in \mathbb{N}$ balls have been allocated and we analyze the process of allocating the next n balls. Using the Pigeonhole Principle and the fact that $\ell_i^\tau \leq \tau + 1$ for every i , it follows that there are at least $(1/2) \cdot n$ bins j with $\ell_j^\tau \geq \tau$. Let us call these bins *overloaded* in the following discussion. Recall that bins with load less than $\tau + 2 - C_1$ are called underloaded. Roughly speaking, the idea of the proof is now to show that the expected number of balls allocated to *underloaded* bins is larger than one since each overloaded bin can receive at most two balls.

Let us make this idea more formal. We divide the allocation of the next n balls into two phases, each phase allocates $n/2$ balls. Let $Y_i \in [0, n]$ be the number of balls that are placed into bin i during the execution of stage $\tau + 1$ (we drop the superscripts τ and $\tau + 1$ in the following for simplicity). Moreover, let $Y_i(1)$ and $Y_i(2)$ be the number of balls that are placed into bin i in the first and the second phase, so $Y_i = Y_i(1) + Y_i(2)$. Throughout we implicitly condition on L^τ being equal to ℓ^τ , but for simplicity we drop the conditioning.

Let us compute the probability that a particular bin j with $\ell_j = \tau$ reaches a load of $\tau + 2$ during the first phase. For each ball in the first phase, the probability of placing it into bin j is at least $1/n$, unless bin j has already reached a load of $\tau + 2$. Moreover, we need to sample at least $n/2$ bins randomly to place the first $n/2$ balls. Therefore,

$$\begin{aligned} \Pr\{Y_j(1) = 2\} &\geq \Pr\{\text{Bin}(n/2, 1/n) \geq 2\} \\ &\geq \binom{n}{2} \cdot \left(\frac{1}{n}\right)^2 \cdot \left(1 - \frac{1}{n}\right)^{n-2} \\ &= \frac{1}{2} \cdot \left(1 - \frac{1}{n}\right)^{n-1} \gg \frac{1}{20}. \end{aligned}$$

Obviously, the same bound also holds for bins j with $\ell_j = \tau + 1$. Hence the expected number of bins with load $\tau + 2$ after the first phase is at least $n/2 \cdot 1/20 = n/40$. Let $\mathcal{A}$ be the event that at least $n/100$ bins have load $\tau + 2$ after the first phase. To prove that $\mathcal{A}$ occurs with high probability, consider instead the event $\mathcal{B}$ that occurs if at least $n/100$ bins are chosen at least twice in the first $n/2$ samples of the incoming balls. Clearly, $\Pr\{\mathcal{A}\} \geq \Pr\{\mathcal{B}\}$. For bounding $\mathcal{B}$, we can use Azuma's inequality (Theorem A.3 with all c_i 's being 1) to obtain

$$\Pr\{\mathcal{A}\} \geq \Pr\{\mathcal{B}\} = 1 - e^{-\Omega(n)}.$$

Now we consider underloaded bins. Fix an arbitrary bin i with $\ell_i \leq \tau + 2 - C_1$, where $C_1 > 0$ is a sufficiently large constant to be chosen later. Then for $0 \leq k \leq C_1 + 2$,

$$\begin{aligned}
& \Pr\{Y_i \geq k\} \\
& \geq \Pr\{\mathcal{A}\} \cdot \Pr\{Y_i \geq k \mid \mathcal{A}\} \\
& \geq \Pr\{Y_i \geq k \mid \mathcal{A}\} - \Pr\{\neg\mathcal{A}\} \\
& \geq \sum_{k_1=0}^{C_1+2} \Pr\{Y_i(1) = k_1 \wedge Y_i(2) \geq (k - k_1) \mid \mathcal{A}\} - e^{-\Omega(n)} \\
& = \sum_{k_1=0}^{C_1+2} \Pr\{Y_i(1) = k_1 \mid \mathcal{A}\} \cdot \Pr\{Y_i(2) \geq (k - k_1) \mid \mathcal{A} \wedge Y_i(1) = k_1\} - e^{-\Omega(n)} \\
& \geq \sum_{k_1=0}^{C_1+2} \Pr\{Y_i(1) = k_1 \wedge \mathcal{A}\} \cdot \Pr\left\{\text{Bin}\left(\frac{n}{2}, \frac{1}{n - \frac{n}{100}}\right) \geq (k - k_1)\right\} - e^{-\Omega(n)} \\
& \geq \sum_{k_1=0}^{C_1+2} (\Pr\{Y_i(1) = k_1\} - \Pr\{\neg\mathcal{A}\}) \cdot \Pr\left\{\text{Bin}\left(\frac{n}{2}, \frac{1}{n - \frac{n}{100}}\right) \geq (k - k_1)\right\} - e^{-\Omega(n)} \\
& \geq \sum_{k_1=0}^{C_1+2} \Pr\{Y_i(1) = k_1\} \cdot \Pr\left\{\text{Bin}\left(\frac{n}{2}, \frac{1}{n - \frac{n}{100}}\right) \geq (k - k_1)\right\} - e^{-\Omega(n)} - n \cdot e^{-\Omega(n)},
\end{aligned}$$

where the last line uses $\Pr\{\neg\mathcal{A}\} = e^{-\Omega(n)}$. Note that for any $0 \leq r \leq C_1 + 2$, $\Pr\{Y_i(1) \geq r\} \geq \Pr\{\text{Bin}(n/2, 1/n) \geq r\}$. Since $\Pr\left\{\text{Bin}\left(\frac{n}{2}, \frac{1}{n - \frac{n}{100}}\right) \geq (k - k_1)\right\}$ is increasing in k_1 , we obtain that

$$\begin{aligned}
& \Pr\{Y_i \geq k\} \\
& \geq \sum_{k_1=0}^{C_1+2} \Pr\left\{\text{Bin}\left(\frac{n}{2}, \frac{1}{n}\right) = k_1\right\} \cdot \Pr\left\{\text{Bin}\left(\frac{n}{2}, \frac{1}{n - \frac{n}{100}}\right) \geq (k - k_1)\right\} - (n+1) \cdot e^{-\Omega(n)}.
\end{aligned}$$

Now recall that for any three values a, b ($a \rightarrow \infty$) and k such that $a \cdot b$ and k are fixed, it holds that $\Pr\{\text{Bin}(a, b) = k\} = \Pr\{\text{Poi}(ab) = k\} + o(1)$. In particular, for any constant C , $\Pr\{\text{Bin}(a, b) \geq C\} = \Pr\{\text{Poi}(a, b) \geq C\} + o(1)$. Hence since $k \leq C_1$

$$\begin{aligned}
& \sum_{k_1=0}^{C_1+2} \Pr\left\{\text{Bin}\left(\frac{n}{2}, \frac{1}{n}\right) = k_1\right\} \cdot \Pr\left\{\text{Bin}\left(\frac{n}{2}, \frac{1}{n - \frac{n}{100}}\right) \geq (k - k_1)\right\} - (n+1) \cdot e^{-\Omega(n)} \\
& = \sum_{k_1=0}^{C_1+2} \Pr\left\{\text{Poi}\left(\frac{1}{2}\right) = k_1\right\} \cdot \Pr\left\{\text{Poi}\left(\frac{100}{198}\right) \geq (k - k_1)\right\} - (C_1 + 3) \cdot o(1) - (n+1) \cdot e^{-\Omega(n)}.
\end{aligned}$$

Moreover, let the constant $C_1 > 0$ be chosen large enough such that $\sum_{k_1=C_1+3}^{\infty} \Pr\{\text{Poi}(\frac{1}{2}) = k_1\} \leq 10^{-10}$. Then

$$\begin{aligned}
& \sum_{k_1=0}^{C_1+2} \Pr\left\{\text{Poi}\left(\frac{1}{2}\right) = k_1\right\} \cdot \Pr\left\{\text{Poi}\left(\frac{100}{198}\right) \geq (k - k_1)\right\} - (C_1 + 3) \cdot o(1) - (n+1) \cdot e^{-\Omega(n)} \\
& = \sum_{k_1=0}^{\infty} \Pr\left\{\text{Poi}\left(\frac{1}{2}\right) = k_1\right\} \cdot \Pr\left\{\text{Poi}\left(\frac{100}{198}\right) \geq (k - k_1)\right\} \\
& \quad - \sum_{k_1=C_1+3}^{\infty} \Pr\left\{\text{Poi}\left(\frac{1}{2}\right) = k_1\right\} \cdot \Pr\left\{\text{Poi}\left(\frac{100}{198}\right) \geq (k - k_1)\right\} - (C_1 + 3) \cdot o(1) - (n+1) \cdot e^{-\Omega(n)}
\end{aligned}$$

$$\begin{aligned}
&\geq \sum_{k_1=0}^{\infty} \mathbf{Pr} \left\{ \text{Poi} \left(\frac{1}{2} \right) = k_1 \right\} \cdot \mathbf{Pr} \left\{ \text{Poi} \left(\frac{100}{198} \right) \geq (k - k_1) \right\} - 10^{-10} - (C_1 + 3) \cdot o(1) \\
&= \mathbf{Pr} \left\{ \text{Poi} \left(\frac{199}{198} \right) \geq k \right\} - 10^{-10} - (C_1 + 3) \cdot o(1),
\end{aligned}$$

since $\text{Poi}(\lambda_1) + \text{Poi}(\lambda_2)$ has the same distribution as $\text{Poi}(\lambda_1 + \lambda_2)$. Hence for any $0 \leq k \leq C_1$,

$$\mathbf{Pr} \{ Y_i^{\tau+1} \geq k \mid L^\tau = \ell^\tau \} \geq \mathbf{Pr} \left\{ \text{Poi} \left(\frac{199}{198} \right) \geq k \right\} - 2 \cdot 10^{-10}, \quad (3.1)$$

as needed.

Proof (Proof of Lemma 3.3). Consider now the expected change of the potential Φ w.r.t. to bin i , i.e.,

$$\begin{aligned}
\mathbf{E} \{ \Delta \Phi_i(\ell^\tau) \} &= \Phi_i(\ell^\tau) - \mathbf{E} \{ \Phi_i(L^{\tau+1}) \mid L^\tau = \ell^\tau \} \\
&= (1 + \varepsilon)^{\tau+2-\ell_i^\tau} - \mathbf{E} \left\{ (1 + \varepsilon)^{\tau+3-\ell_i^\tau - Y_i^{\tau+1}} \mid L^\tau = \ell^\tau \right\} \\
&= (1 + \varepsilon)^{\tau+2-\ell_i^\tau} - \sum_{k=0}^{\tau+3-\ell_i^\tau} \mathbf{Pr} \{ Y_i^{\tau+1} = k \mid L^\tau = \ell^\tau \} \cdot (1 + \varepsilon)^{\tau+3-\ell_i^\tau - k} \\
&= (1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot \left(1 - \sum_{k=0}^{\tau+3-\ell_i^\tau} \mathbf{Pr} \{ Y_i^{\tau+1} = k \mid L^\tau = \ell^\tau \} \cdot (1 + \varepsilon)^{1-k} \right).
\end{aligned}$$

Now we define $p_k := \mathbf{Pr} \{ Y_i^{\tau+1} = k \mid L^\tau = \ell^\tau \}$, $r_k := (1 + \varepsilon)^{1-k}$ (which is non-increasing in k), and $q_k := \mathbf{Pr} \{ \text{Poi}(199/198) = k \} + 2 \cdot 10^{-10}$. By inequality (3.1), for every $1 \leq k \leq C_1$, $\mathbf{Pr} \{ Y_i^{\tau+1} \leq k-1 \mid L^\tau = \ell^\tau \} \leq \mathbf{Pr} \{ \text{Poi}(199/198) \leq k-1 \} + 2 \cdot 10^{-10}$. For every $0 \leq r \leq C_1 - 1$, $\mathbf{Pr} \{ Y_i^{\tau+1} \leq r \mid L^\tau = \ell^\tau \} = \sum_{k=0}^r p_k \leq \sum_{k=0}^r q_k$. Now we choose $C_1 > 0$ large enough such that $\sum_{k=0}^{C_1-1} \mathbf{Pr} \{ \text{Poi}(199/198) = k \} \geq 1 - 2 \cdot 10^{-10}$ and therefore $\sum_{k=0}^{C_1-1} q_k \geq 1$. Since $\sum_{k=0}^{\tau+3-\ell_i^\tau} p_k \leq 1$, we obtain that, for any $1 \leq r \leq \tau + 3 - \ell_i^\tau$, $\sum_{k=0}^r p_k \leq \sum_{k=0}^{\max\{r, C_1-1\}} q_k$. Define $s_k := q_k$ for $0 \leq k \leq C_1 - 1$ and $s_k := 0$ otherwise. Hence, $\sum_{k=0}^r p_k \leq \sum_{k=0}^r s_k$. Applying now the second statement of Lemma A.1 to the sequences p_k, r_k and s_k we get

$$\begin{aligned}
&(1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot \left(1 - \sum_{k=0}^{\tau+3-\ell_i^\tau} \mathbf{Pr} \{ Y_i^{\tau+1} = k \mid L^\tau = \ell^\tau \} \cdot (1 + \varepsilon)^{1-k} \right) \\
&\geq (1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot \left(1 - \sum_{k=0}^{C_1-1} \left(\mathbf{Pr} \left\{ \text{Poi} \left(\frac{199}{198} \right) = k \right\} + 2 \cdot 10^{-10} \right) \cdot (1 + \varepsilon)^{1-k} \right) \\
&\geq (1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot \left(1 - \sum_{k=0}^{C_1-1} e^{-(199/198)} \frac{(199/198)^k}{k!} \cdot (1 + \varepsilon)^{1-k} - 2 \cdot 10^{-10} \frac{1 + \varepsilon}{1 - \frac{1}{1+\varepsilon}} \right) \\
&= (1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot \left(1 - e^{-(199/198)} \sum_{k=0}^{C_1-1} \frac{(199/198)^k}{k!} \cdot (1 + \varepsilon)^{1-k} - 2 \cdot 10^{-10} \cdot \frac{(1 + \varepsilon)^2}{\varepsilon} \right),
\end{aligned}$$

and since $\varepsilon = 1/200$,

$$\begin{aligned}
&= (1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot \left(1 - e^{-(199/198)} \cdot \sum_{k=0}^{C_1-1} \frac{(1 + \varepsilon) \cdot \left(\frac{(199/198)}{1+\varepsilon} \right)^k}{k!} - 2 \cdot 10^{-10} \cdot \frac{(1 + \varepsilon)^2}{\varepsilon} \right) \\
&\geq (1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot \left(1 - e^{-(199/198)} \cdot \frac{201}{200} \cdot \sum_{k=0}^{C_1-1} \frac{\left(\frac{200}{201} \cdot \frac{199}{198} \right)^k}{k!} - 2 \cdot 10^{-7} \right) \\
&\geq (1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot \left(1 - e^{-(199/198)} \cdot \frac{201}{200} \cdot e^{\frac{200}{201} \cdot \frac{199}{198}} - 2 \cdot 10^{-7} \right),
\end{aligned}$$

and an evaluation of these expressions numerically yields an upper bound of $(1 + \varepsilon)^{\tau+2-\ell_i^\tau} \cdot (\beta - 2 \cdot 10^{-7})$, where $\beta > 0.000012 \dots > 2 \cdot 10^{-7}$.

The proof of Lemma 3.4 and Corollary 3.5 can be found in the appendix.

Proof (Proof of Lemma 3.6).

Fix any load distribution $L^\tau = \ell^\tau$ at the end of stage τ . Let $\beta := \max\{\Phi(\ell^\tau)/n, 16/\varepsilon^2\}$. Our goal is to prove that, with probability $1 - n^{-2}$, ADAPTIVE needs $\mathcal{O}(n \cdot (\ln \beta))$ steps to place the n balls of stage $\tau + 1$. Let $A_k := \{i \in \{1, \dots, n\} : \ell_i^\tau = \tau + 2 - k\}$ be the set of bins with k holes. Then,

$$\beta \cdot n \geq \Phi(\ell^\tau) = \sum_{i \in A_k} (1 + \varepsilon)^k,$$

so that $|A_k| \leq \beta n \cdot (1 + \varepsilon)^{-k} \leq 2^{-k\varepsilon} \cdot \beta n$. Therefore the sum of the holes in bins having at least $4 \ln(\beta)/\varepsilon$ holes is

$$\sum_{k=4 \ln(\beta)/\varepsilon}^{\tau+2} |A_k| \cdot k \leq \sum_{k=4 \ln(\beta)/\varepsilon}^{\tau+2} 2^{-k\varepsilon} \cdot \beta n \cdot k = \frac{\beta n}{\varepsilon} \cdot \sum_{k=4 \ln(\beta)/\varepsilon}^{\tau+2} 2^{-k\varepsilon} \cdot k\varepsilon = \frac{\beta n}{\varepsilon} \cdot \sum_{k=4 \ln(\beta)/\varepsilon}^{\tau+2} e^{-k\varepsilon/2} \cdot \left(\frac{4}{e} \right)^{-k\varepsilon/2} \cdot k\varepsilon$$

Since $(4/e)^{-x/2} \cdot x \leq 2$ for any $x \geq 0$ we obtain

$$\begin{aligned}
\frac{\beta n}{\varepsilon} \cdot \sum_{k=4 \ln(\beta)/\varepsilon}^{\tau+2} e^{-k\varepsilon/2} \cdot \left(\frac{4}{e} \right)^{-k\varepsilon/2} \cdot (k\varepsilon) &\leq 2 \frac{\beta n}{\varepsilon} \cdot \sum_{k=4 \ln(\beta)/\varepsilon}^{\tau+2} e^{-k\varepsilon/2} \leq 2 \frac{\beta n}{\varepsilon} \cdot \frac{e^{-2 \ln(\beta)}}{1 - e^{-\varepsilon/2}} \\
&\leq 8 \frac{\beta n}{\varepsilon} \cdot \frac{\beta^{-2}}{\varepsilon} = 8 \cdot \frac{n}{\beta \varepsilon^2},
\end{aligned}$$

where in the second last inequality we have used that $e^{-x} \leq 1 - x/2$ for any $x \in (0, 1)$.

Now let $A' := \{i \in \{1, \dots, n\} : \tau + 2 - 4 \ln(\beta)/\varepsilon \leq \ell_i^\tau \leq \tau + 1\}$ be the set of bins with at least 1 and at most $4 \ln(\beta)/\varepsilon$ many holes. Since the total sum of holes is $2n$, it follows that the sum of holes of bins in A' is at least $2n - 8n \cdot \frac{1}{\beta \varepsilon^2} > (3/2) \cdot n$ by our choice of β . Hence, even after the allocation of $n - 1$ balls, the total number of holes of the bins in A' will still be at least $(3/2) \cdot n - (n - 1) \geq n/2$. Since every of these bins has at most $4 \ln(\beta)/\varepsilon$ holes it follows that after the allocation of $n - 1$ balls there are still $\Omega(n/(\ln(\beta)/\varepsilon)) = \Omega(n/(\ln \beta))$ bins with at least one hole. This means that the time to place n balls is stochastically smaller than the sum of n independent geometric random variables each with expectation $\mathcal{O}(\ln(\beta))$. Therefore, Theorem A.6 implies that with probability at least $1 - n^{-2}$, we need at most $\mathcal{O}(n \cdot (\ln \beta))$ steps to place n balls.

Corollary 3.5 establishes that $\mathbf{E}\{\Phi^\tau\} = \mathcal{O}(n)$. Using Jensen's inequality for concave functions, we obtain that

$$\mathbf{E}\{\log_{1+\varepsilon}(\Phi^\tau/n)\} \leq \log_{1+\varepsilon} \mathbf{E}\{\Phi^\tau/n\} \leq \log_{1+\varepsilon} \mathcal{O}(1) = \mathcal{O}(1).$$

Thus, for any stage $\tau + 1$, $\mathbf{E} \{ \log_{1+\varepsilon}(\Phi^\tau/n) \} = \mathcal{O}(1)$. Recall that Φ^τ can take finitely many values between n and $(1 + \varepsilon)^n$. Let $\mathcal{S}$ be the set of possible values that Φ^τ can attain. Therefore,

$$\begin{aligned}
\mathbf{E} \{ T_{\tau+1} \} &= \sum_{k \in \mathcal{S}} \mathbf{Pr} \{ \Phi^\tau = k \} \cdot \mathbf{E} \{ T_{\tau+1} \mid \Phi^\tau = k \} \\
&= \sum_{k \in \mathcal{S}} \mathbf{Pr} \{ \Phi^\tau = k \} \cdot \left((1 - n^{-2}) \cdot \mathcal{O}(n \cdot \ln(k/n)) + \frac{1}{n^2} \cdot n^2 \right) \\
&= \mathcal{O}(n) \cdot \sum_{k \in \mathcal{S}} \mathbf{Pr} \{ \Phi^\tau = k \} \cdot \ln(k/n) + \sum_{k \in \mathcal{S}} \mathbf{Pr} \{ \Phi^\tau = k \} \\
&= \mathcal{O}(n) \cdot \sum_{k \in \mathcal{S}} \mathbf{Pr} \{ \Phi^\tau = k \} \cdot \log_{1+\varepsilon}(k/n) / \log_{1+\varepsilon}(e) + 1 \\
&= \mathcal{O}(n) \cdot \mathbf{E} \{ \log_{1+\varepsilon}(\Phi^\tau/n) \} + 1 = \mathcal{O}(n) \cdot \mathcal{O}(1) = \mathcal{O}(n).
\end{aligned}$$

4 Analysis of THRESHOLD

In this section we analyze THRESHOLD, extending results from [7] for the case $m = \mathcal{O}(n)$ to the case where m is asymptotically greater than n . The proof is in the appendix.

Theorem 4.1 *The allocation time of the THRESHOLD protocol is $m + \mathcal{O}(m^{3/4} \cdot n^{1/4})$, with probability at least $1 - 8n^{-2}$. Moreover, the expected allocation time is also bounded by $m + \mathcal{O}(m^{3/4} \cdot n^{1/4})$.*

We can also prove a lower bound on the quadratic potential of the final load distribution. For simplicity we focus on the case $m = n^2$ here, but an inspection of the proof shows that we can also obtain lower bounds for other relations between m and n .

Lemma 4.2 *Consider the THRESHOLD protocol with $m = n^2$ and let $t = m$. Then L^t satisfies the following three relations with probability at least $1 - \exp(-\Omega(n^{1/2}))$,*

- (1) $\Psi(L^t) = \Omega(n^{9/8})$,
- (2) $\max_{i,j} (L_i^t - L_j^t) = \Omega(n^{1/8})$, and
- (3) $\Phi(L^t) = 2^{\Omega(n^{1/8})}$.

This result should be compared to Corollary 3.5 for the ADAPTIVE protocol.

5 Experiments

We also compare the ADAPTIVE protocol to the THRESHOLD protocol experimentally (Figure 3). Inline with our theoretical results, we observe that the runtime of THRESHOLD quickly converges to m . Moreover, the runtime of ADAPTIVE seems to converge to a small constant times m . To measure the smoothness of the final load distribution, we also compute the average values of the quadratic potential function. Here we observe that this value is significantly smaller in case of the ADAPTIVE protocol, as it converges quickly to a value that is independent of m (note that this principal convergence is also guaranteed by Lemma 3.4, since the quadratic potential function is asymptotically at most the exponential potential function).

Bibliography

- [1] M. Adler, S. Chakrabarti, M. Mitzenmacher, and L. E. Rasmussen. Parallel randomized load balancing. *Random Struct. Algorithms*, 13(2):159–188, 1998.

Figure 3: Simulation results for the THRESHOLD and ADAPTIVE protocol. Every point in the chart represents the average allocation time (or average value of the final load distribution) taken over 100 simulations.

- [2] N. Alon and J. Spencer. *The Probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, 2nd edition, 2000.
- [3] B. Awerbuch, Y. Azar, S. A. Plotkin, and O. Waarts. Competitive routing of virtual circuits with unknown duration. In *SODA*, pages 321–327, 1994.
- [4] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal. Balanced allocations. *SIAM J. Comput. (conference version in STOC'94)*, 29(1):180–200, 1999.
- [5] P. Berenbrink, A. Czumaj, A. Steger, and B. Vöcking. Balanced allocations: The heavily loaded case. *SIAM J. Comput. (conference version in STOC'00)*, 35(6):1350–1385, 2006.
- [6] A. Czumaj, C. Riley, and C. Scheideler. Perfectly balanced allocation. In *RANDOM-APPROX*, pages 240–251, 2003.
- [7] A. Czumaj and V. Stemann. Randomized allocation processes. *Random Struct. Algorithms*, 18(4):297–331, 2001.
- [8] M. Dietzfelbinger, A. Goerdts, M. Mitzenmacher, A. Montanari, R. Pagh, and M. Rink. Tight thresholds for cuckoo hashing via XORSAT. In *ICALP (1)*, pages 213–225, 2010.
- [9] D. Dubhashi and A. Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- [10] B. Ghosh, F. T. Leighton, B. M. Maggs, S. Muthukrishnan, C. G. Plaxton, R. Rajaraman, A. W. Richa, R. E. Tarjan, and D. Zuckerman. Tight analyses of two local load balancing algorithms. *SIAM J. Comput.*, 29(1):29–64, 1999.
- [11] R. M. Karp, M. Luby, and F. Meyer auf der Heide. Efficient PRAM simulation on a distributed memory machine. *Algorithmica (conference version in STOC'92)*, 16(4/5):517–542, 1996.
- [12] C. Lenzen and R. Wattenhofer. Tight bounds for parallel randomized load balancing. *to appear in the proceedings of 43rd ACM Symposium on Theory of Computing (STOC)*, 2011.
- [13] R. Lüling and B. Monien. A dynamic distributed load balancing algorithm with provable good performance. In *SPAA*, pages 164–172, 1993.
- [14] M. Mitzenmacher, B. Prabhakar, and D. Shah. Load balancing with memory. In *FOCS*, pages 799–808, 2002.
- [15] M. Raab and A. Steger. “balls into bins” - a simple and tight analysis. In M. Luby, J. D. P. Rolim, and M. J. Serna, editors, *RANDOM*, volume 1518 of *Lecture Notes in Computer Science*, pages 159–170. Springer, 1998.
- [16] B. Vöcking. How asymmetry helps load balancing. *J. ACM (conference version in FOCS'99)*, 50(4):568–589, 2003.

A Probabilistic Tools

We also observe the following simple fact about convolution of sequences. Recall the definition of convolution: If $p = (p_k)_{k=0}^n, q = (q_k)_{k=0}^n$ are two sequences, then the sequence $p \star q = ((p \star q)_k)_{k=0}^{2n}$ defined by $(p \star q)_k := \sum_{i=\max\{0, k-n\}}^{\min\{k, n\}} p_i \cdot q_{k-i}$ is the convolution of p and q . We say that a sequence p majorizes a sequence q if for every $0 \leq j \leq n$, $\sum_{k=j}^n p_k \geq \sum_{k=j}^n q_k$.

Lemma A.1 *Let p, q and r three sequences of n numbers. If p majorizes q and r is non-increasing, then $\sum_{k=0}^n p_k \cdot r_k \leq \sum_{k=0}^n q_k \cdot r_k$.*

We also list several concentration inequalities that are used throughout the paper.

Theorem A.2 (Hoeffdings inequality) *Let X_k , $1 \leq k \leq n$, be n independent binary random variables. Let $X := \sum_{k=1}^n X_k$. Then for any $\lambda > 0$,*

$$\Pr \{|X - \mathbf{E}\{X\}| \geq \lambda\} \leq 2e^{-\lambda^2/n}.$$

Theorem A.3 (Azuma's inequality) *Let X_k , $0 \leq k \leq n$, be a martingale such that $|X_k - X_{k-1}| \leq c_k$ for every $1 \leq k \leq n$. Then for any $\varepsilon > 0$,*

$$\Pr \{|X_n - X_0| \geq \varepsilon\} \leq 2e^{-\frac{\varepsilon^2}{2 \sum_{i=0}^n c_i^2}}.$$

Theorem A.4 ([2]) *Let $\text{Poi}(\mu)$ be the Poisson distribution with mean μ . Then for any $\varepsilon > 0$,*

$$\begin{aligned} \Pr \{\text{Poi}(\mu) \leq (1 - \varepsilon)\mu\} &\leq e^{-\varepsilon^2 \mu / 2} \\ \Pr \{\text{Poi}(\mu) \geq (1 + \varepsilon)\mu\} &\leq \left[e^\varepsilon (1 + \varepsilon)^{-(1+\varepsilon)} \right]^\mu \end{aligned}$$

We note the following standard Chernoff bound for sum of geometric random variables which can be easily derived by using a Chernoff bound for a sum of Bernoulli random variables.

Theorem A.5 ([9]) *Suppose that $X_1, \dots, X_n$ are independent geometric random variables on $\mathbb{N}$ with parameter δ , so $\mathbf{E}\{X_i\} = 1/\delta$ for each i . Let $X := \sum_{i=1}^n X_i$, $\mu = \mathbf{E}\{X\} = n/\delta$. Then it holds for any $\varepsilon > 0$ that*

$$\Pr \{X \geq (1 + \varepsilon)\mu\} \leq e^{-\varepsilon^2 n / 2(1+\varepsilon)}.$$

From Theorem A.5 and a simple majorization argument, one can obtain the following extension:

Theorem A.6 *Suppose that $X_1, \dots, X_n$ are independent random variables on $\mathbb{N}$, such that there is a value $0 < \delta < 1$ with $\Pr \{X_i = k + 1\} \leq (1 - \delta) \cdot \Pr \{X_i = k\}$ for all $k \geq 1$. Let $X := \sum_{i=1}^n X_i$, $\mu = \mathbf{E}\{X\}$. Then it holds for any $\varepsilon > 0$ that*

$$\Pr \{X \geq (1 + \varepsilon)\mu\} \leq e^{-\varepsilon^2 n / 2(1+\varepsilon)}.$$

Moreover, $\mu \leq \frac{1}{\delta}$.

The following lemma due to Adler et al. is an extremely powerful tool for analyzing balls-into-bins processes.

Lemma A.7 ([1, Corollary 13]) *Let process P_1 be the placement of m balls into n bins where each bin is chosen independently and uniformly at random. Let P_2 the process where the load of every bin is an independent Poisson random variable with expectation m/n . Let $\mathcal{A}$ be any event concerning the final load distribution.*

- (1) *Then any event that occurs in P_2 with probability at most p , holds in P_1 with probability at most $p \cdot \sqrt{n}$.*
- (2) *Moreover, any event that is increasing w.r.t. adding balls and holds in P_2 with probability at most p holds in P_1 with probability at most $4p$.*

B Omitted Proofs from Section 4

Proof of Theorem 4.1

Proof. Let $\varphi := \frac{m}{n}$. We first observe that we can focus on the case where φ is an integer. If φ is not an integer, we can resort to the statement of theorem for $m = \lceil \varphi \rceil n$. We can also assume that φ is bounded from below by a sufficiently large constant for the following reason. If φ is upper bounded by some constant, then one can easily prove that in each step t there are at least $\Omega(n/\varphi)$ bins with load less than $t/n + 1$. Hence the expected time to place m balls is stochastically smaller than the sum of m geometric random variables each with constant mean.

For the proof we assume that all the bin choices of the balls are fixed in advance. Let C be a vector of infinite length. Every entry in C is a number in $[n]$ chosen uniformly and independently at random. The first ball uses the first i entries of C as random bin choices, until it is allocated to bin $C[i]$ (using the rules of THRESHOLD). The next ball then uses $C[i+1], \dots, C[i+j]$ until it is allocated to bin $C[i+j]$, and so on. Our goal is now to upper bound the number of entries of C that are used until THRESHOLD allocates all m balls into the n bins. The main idea of our proof is to upper bound the total number of *holes* W_t in the bins at time t . The number of holes of a bin is defined as follows. A bin with ℓ balls has $\varphi + 1 - \ell$ holes. The number of holes of the bin is the maximal load of $\varphi + 1$ minus the number of balls in that bin. If the number of remaining holes is less than or equal to n , then all m balls are allocated, since THRESHOLD never allocates a ball into a bin that has already $\varphi + 1$ balls. Hence, at time t there are $(\varphi + 1) \cdot n - W_t$ many balls allocated. We will show that for $T = \alpha n$, $\alpha := \varphi + \varphi^{3/4} + 1$, we have $W_t \leq n$ with high probability. We note that our proof is inspired by the proof of the corresponding result from [7].

We define $X_1^t, \dots, X_n^t$ as the *access* distribution at the end of step t , i.e., $X_i^t = x$ means that bin i occurs x times in $C[1], \dots, C[t]$. Moreover, let $L_1^t, \dots, L_n^t$ be the *load* distribution when the first t entries of C are used, i.e., $L_i^t = x$ means that exactly x balls are allocated to bin i when we have gone through $C[1], \dots, C[t]$. By definition of THRESHOLD

$$L_i^t = \min\{\varphi + 1, X_i^t\}.$$

Our aim is to prove for our choice of T that w.h.p.

$$W_T = \sum_{i=1}^n (\varphi + 1 - L_i^T) = \sum_{i=1}^n \max\{\varphi + 1 - X_i^T, 0\} \leq n.$$

Similar to the analysis of Czumaj and Stemmann in [7] we use Poisson distributed random variables as an approximation for $X_1^t, \dots, X_n^t$. We also fix $t = T$ and omit the superscript T in the following, as we only consider the access and load distribution at step T . Let $Y_1, \dots, Y_n$ be n independent Poisson random variables with expectation α modeling the access distribution in step T . Using the Chernoff bound from Theorem A.4 with $\varepsilon = \frac{\varphi^{3/4}}{\alpha}$, we obtain for any i that

$$\begin{aligned} \Pr\{Y_i \leq \varphi + 1\} &= \Pr\left\{Y_i \leq \alpha - \varphi^{3/4}\right\} = \Pr\left\{Y_i \leq \left(1 - \frac{\varphi^{3/4}}{\alpha}\right) \alpha\right\} \\ &\leq e^{-\left(\frac{\varphi^{3/4}}{\alpha}\right)^2 \cdot \alpha/2} \leq e^{-\alpha^{1/2}/4}, \end{aligned}$$

since φ is sufficiently large. Note that if $\alpha^{1/2} \geq 8 \log n$, then the above probability is smaller than n^{-2} , so that with probability at least $1 - n^{-1}$, every Y_i is larger than $\varphi + 1$. Therefore, we assume in the remainder of the proof that $\alpha^{1/2} \leq 8 \log n$. Let $\mathcal{I} := \{i \in \{1, \dots, n\} : Y_i \leq \varphi + 1\}$. Then $|\mathcal{I}|$ is a random variable with $\mathbf{E}\{|\mathcal{I}|\} \leq n \cdot e^{-\alpha^{1/2}/4}$. Using the inequality of Hoeffding, we obtain that

$$\Pr\left\{|\mathcal{I}| \geq \mathbf{E}\{|\mathcal{I}|\} + \sqrt{n \log n}\right\} \leq e^{-\frac{2(\sqrt{n \log n})^2}{n}} = n^{-2}.$$

Let $Y := \sum_{i=1}^n \max\{(\varphi + 1) - Y_i, 0\}$. Our goal is now to prove that Y is not too large with high probability. To this end, let Z_i be a random variable whose distribution is that of $(\varphi + 1) - Y_i$ conditioned on the event $Y_i \leq (\varphi + 1)$ ($i \in \mathcal{I}$). For any $0 \leq k \leq \varphi + 1$,

$$\Pr\{Z_i = k\} = \frac{\Pr\{Y_i = \varphi + 1 - k\}}{\Pr\{Y_i \leq \varphi + 1\}}. \quad (\text{B.1})$$

We now expose Y in two stages. In the first stage, we consider, for each $1 \leq i \leq n$, the event $\{Y_i \leq \varphi + 1\}$. Let I_i be the indicator variable with $I_i = 1$ if $Y_i \leq \varphi + 1$ and $I_i = 0$ otherwise. Recall that $\mathcal{I} := \{i \in \{1, \dots, n\} : Y_i \leq \varphi + 1\}$. With these notations we have

$$\Pr\{Y \geq \lambda\} = \sum_{I \subseteq \{1, \dots, n\}} \Pr\{\mathcal{I} = I\} \cdot \Pr\left\{\sum_{i \in I} Z_i \geq \lambda \mid \mathcal{I} = I\right\}$$

For any $1 \leq k \leq \varphi + 1$, define

$$p_k := \Pr\{Z_i = k\} = \frac{\Pr\{Y_i = \varphi + 1 - k\}}{\Pr\{Y_i \leq \varphi + 1\}} = e^{-\alpha} \cdot \frac{\alpha^{\varphi+1-k}}{(\varphi + 1 - k)!} \Big/ \Pr\{Y_i \leq \varphi + 1\}.$$

Hence,

$$\frac{p_k}{p_{k-1}} = \frac{\varphi + 2 - k}{\alpha},$$

that is, $p_k = \frac{\varphi+2-k}{\alpha} \cdot p_{k-1} \leq \frac{\varphi+2}{\varphi+\varphi^{3/4}+2} \cdot p_{k-1} \leq (1 - \varphi^{-1/4}/2) \cdot p_{k-1}$, where the last inequality holds if φ is sufficiently large.

Therefore, for any fixed set $I \subseteq \{1, \dots, n\}$, the random variables Z_i with $i \in I$ satisfy the preconditions of Theorem A.6 with $\delta = \varphi^{-1/4}/2$, so that $\mathbf{E}\{Z_i\} \leq 2\varphi^{1/4}$ and

$$\Pr\left\{\sum_{i \in I} Z_i \geq (1 + \varepsilon) \cdot |I| \cdot 2\varphi^{1/4}\right\} \leq e^{-\varepsilon^2 |I| / 2(1 + \varepsilon)}.$$

Then

$$\begin{aligned} & \Pr\{Y \geq n\} \\ & \leq \Pr\left\{Y \geq 2 \cdot e^{-\alpha^{1/2}/4} \cdot n \cdot 2\varphi^{1/4} + 2 \cdot \sqrt{n \log n} \cdot 2\varphi^{1/4}\right\} \\ & \leq \Pr\left\{Y \geq 2 \cdot \mathbf{E}\{|\mathcal{I}|\} \cdot 2\varphi^{1/4} + 2 \cdot \sqrt{n \log n} \cdot 2\varphi^{1/4}\right\} \\ & \leq \Pr\left\{|\mathcal{I}| \geq \mathbf{E}\{|\mathcal{I}|\} + \sqrt{n \log n}\right\} \\ & \quad + \Pr\left\{Y \geq 2 \cdot \left(\mathbf{E}\{|\mathcal{I}|\} + \sqrt{n \log n}\right) \cdot 2\varphi^{1/4} \mid |\mathcal{I}| \leq \mathbf{E}\{|\mathcal{I}|\} + \sqrt{n \log n}\right\} \\ & \leq n^{-2} + \Pr\left\{\sum_{i \in \mathcal{I}} Z_i \geq 2 \cdot |\mathcal{I}| \cdot 2\varphi^{1/4} \mid |\mathcal{I}| = \mathbf{E}\{|\mathcal{I}|\} + \sqrt{n \log n}\right\} \\ & \leq n^{-2} + e^{-(\mathbf{E}\{|\mathcal{I}|\} + \sqrt{n \log n})/4} \leq 2 \cdot n^{-2}, \end{aligned}$$

where the first inequality uses the assumptions $\alpha^{1/2} \leq 8 \log n$, $\alpha \geq \varphi$ and φ being bounded below by a sufficiently large constant.

By the relation between the Poisson model and the original balls-and-bins model (second statement in Lemma A.7),

$$\Pr\left\{\sum_{i=1}^n \max\{(\varphi + 1) - L_i, 0\} \geq n\right\} \leq 4 \cdot \Pr\left\{\sum_{i=1}^n \max\{(\varphi + 1) - Y_i, 0\} \geq n\right\} \leq 8n^{-2}.$$

But from $\sum_{i=1}^n \max\{(\varphi+1) - L_i, 0\} \leq n$, we can conclude that $\sum_{i=1}^n L_i \geq (\varphi+1)n - n = \varphi n = m$. Hence the THRESHOLD protocol finishes the placement of m balls before step t with probability at least $1 - 8n^{-2}$. This finishes the proof of the first statement. For the second statement, we simply divide the infinite vector C into consecutive sections of length αn . Then the probability that the choices in one section suffice to terminate the process is at least $1 - 8n^{-2}$, regardless of the outcomes in the previous sections.

Proof of Lemma 4.2

Proof. The proof of Lemma 4.2 is similar to the proof of Theorem 4.1. In order to cope with the dependencies with the balls-into-bins process, we analyze a corresponding setting with independent Poisson random variables (cf. Lemma A.7). To prove that the final load distribution is not too balanced, we consider the access distribution $X_1^t, \dots, X_n^t$ for all time-steps between 1 and $2\varphi n$ (we know from Theorem 4.1 that larger time need not to be considered). Every access distribution is analyzed by examining n independent Poisson variables $Y_1^t, \dots, Y_n^t$ with expectation t/n . Translating the analysis for the Poisson variables back to the access distribution X^t and then to the load distribution L^t , we shall establish that for all these time steps t , the potential $\Psi(L^t)$ is large provided that not more than m balls have been placed.

Let us define the following event $\mathcal{A} := \mathcal{A}(t)$, where $1 \leq t \leq 2\varphi n$,

$$\mathcal{A}(t) := \left\{ \left(\sum_{i=1}^n L_i^t \leq m \right) \wedge \left(\tilde{\Psi}(L^t) \leq n^{9/8}/2 \right) \right\},$$

where as in Theorem 4.1, $L_i^t = \min\{X_i^t, \varphi + 1\}$ denotes the load of bin i after t random choices of the THRESHOLD protocol and $\tilde{\Psi}(L^t) := \sum_{i=1}^n (L_i^t - \varphi)^2$. The event $\mathcal{A}(t)$ occurs if not more than m balls have been placed and the potential of the load vector L^t is small. The corresponding event $\mathcal{B} := \mathcal{B}(t)$ is defined as

$$\mathcal{B}(t) := \left\{ \left(\sum_{i=1}^n \min\{Y_i^t, \varphi + 1\} \leq m \right) \wedge \left(\tilde{\Psi}(\min\{Y^t, \varphi + 1\}) \leq n^{9/8}/2 \right) \right\},$$

where $\min\{Y^t, \varphi + 1\}_i := \min\{Y_i^t, \varphi + 1\}$. For the lemma to follow it suffices to prove

$$\Pr \left\{ \bigvee_{t=1}^{2\varphi n} \mathcal{A}(t) \right\} \leq e^{-\Omega(n^{1/2})}, \quad (\text{B.2})$$

since by Theorem 4.1, we will place all m balls within at most $2\varphi n$ steps with probability at least $1 - 4n^{-2}$. But then $\bigwedge_{t=1}^{2\varphi n} (\neg \mathcal{A}(t))$ implies that when the process stops at some step $T \in [1, 2\varphi n]$, $\tilde{\Psi}(L^T) \geq n^{9/8}/2$, as desired. Now once we have established that for every $1 \leq t \leq 2\varphi n$,

$$\Pr \{\mathcal{B}(t)\} \leq e^{-\Omega(n^{1/2})} \quad (\text{B.3})$$

we can infer B.2 as follows. By the first statement of Lemma A.7, for any $t \in \mathbb{N}$,

$$\Pr \{\mathcal{A}(t)\} \leq 2\sqrt{n} \cdot \Pr \{\mathcal{B}(t)\},$$

and by the union bound and the above two inequalities,

$$\begin{aligned} \Pr \left\{ \bigvee_{t=1}^{2\varphi n} \mathcal{A}(t) \right\} &\leq \sum_{t=1}^{2\varphi n} \Pr \{\mathcal{A}(t)\} \\ &\leq 2\varphi n \cdot 2\sqrt{n} \cdot e^{-\Omega(n^{1/2})} = e^{-\Omega(n^{1/2})}. \end{aligned}$$

Hence it remains to show B.3, which is done by a case analysis.

Case 1: $1 \leq t \leq \varphi n/2$. Note that $\sum_{i=1}^n Y_i^t$ is itself a Poisson random variable with expectation $t \leq \varphi n/2 = m/2$. Hence it follows by Theorem A.4 that

$$\Pr \left\{ \sum_{i=1}^n Y_i^t \geq (3/4)\varphi n \right\} \leq \left[e^{1/2} \cdot (3/2)^{-(3/2)} \right]^{\varphi n/2} = e^{-\Omega(\varphi n)}.$$

Note that the event $\sum_{i=1}^n Y_i^t \leq (3/4)\varphi n$ implies the existence of at least $n/9$ bins j with $Y_j^t \leq (7/8) \cdot (\varphi n/n) = (7/8)\varphi$. Hence,

$$\tilde{\Psi}(\min\{Y^t, \varphi + 1\}) \geq (n/9) \cdot (\varphi + 1 - (7/8)\varphi)^2 \gg n^{9/8}/2,$$

since $\varphi = n$. This proves B.3 for every t with $1 \leq t \leq \varphi n/2$.

Case 2: $\varphi n/2 \leq t \leq 2\varphi n$. Let us now analyze the distribution of $Y_1^t, \dots, Y_n^t$. Let $\alpha := t/n$ and $p_k := \Pr\{Y_i = k\} = e^{-\alpha} \cdot \alpha^k/k!$. Note that one can easily verify that p_k is maximized for $k = \lceil \alpha \rceil - 1$, in which case Stirling's approximation for $k!$ implies that $p_k \leq C/\sqrt{\alpha}$ for some constant $C > 0$.

For any $0 \leq k \leq \varphi + 1$, let $Z_k^t := |\{i \in \{1, \dots, n\} : Y_i^t = k\}|$. Note that $\mathbf{E}\{Z_k^t\} \leq Cn/\sqrt{\alpha}$. Moreover, we have $Z_k^t = \sum_{i=1}^n I_i$, where $I_i = 1$ if $Y_i = k$ and 0 otherwise. Therefore we can apply the following Chernoff bound for binomial random variables,

$$\Pr\{Z_k^t \geq (1 + \varepsilon) \mathbf{E}\{Z_k^t\}\} \leq e^{-\min\{\varepsilon^2, \varepsilon\} \mathbf{E}\{Z_k^t\}/3},$$

to obtain that

$$\Pr\{Z_k^t \geq 2 \cdot (Cn/\sqrt{\alpha})\} \leq \Pr\{Z_k^t \geq 2 \cdot \mathbf{E}\{Z_k^t\} \cdot (Cn/(\sqrt{\alpha} \mathbf{E}\{Z_k^t\}))\} \leq e^{-(C/3)n/\sqrt{\alpha}} = e^{-\Omega(n^{1/2})},$$

so by the union bound,

$$\Pr \left\{ \bigvee_{\varphi - \alpha^{1/8} \leq i \leq \varphi} (Z_i^t \geq 2 \cdot (Cn/\sqrt{\alpha})) \right\} = \alpha^{1/8} \cdot e^{-\Omega(n^{1/2})} = e^{-\Omega(n^{1/2})}.$$

Hence,

$$\Pr \left\{ \sum_{i=\varphi - \alpha^{1/8}}^{\varphi} (\varphi + 1 - i) \cdot Z_i^t \leq (\alpha^{1/8} + 1)^2 \cdot 2 \cdot (Cn/\sqrt{\alpha}) \right\} = 1 - e^{-\Omega(n^{1/2})}$$

Suppose now that the event in the above probability occurs. By definition of the event $\mathcal{B}(t)$, we may assume that $\sum_{i=1}^n \min\{Y_i^t, \varphi + 1\} \leq m$ which is equivalent to $\sum_{i=0}^{\varphi} (\varphi + 1 - i) Z_i^t \geq n$, which in turn implies that

$$\sum_{i=0}^{\varphi - \alpha^{1/8} - 1} (\varphi + 1 - i) \cdot Z_i^t \geq n - \sum_{i=\varphi - \alpha^{1/8}}^{\varphi} (\varphi + 1 - i) \cdot Z_i^t \geq n - \mathcal{O}(n/\alpha^{1/4}) \geq n/2,$$

since $\alpha \geq \varphi/2 = n/2$. Clearly, the quadratic potential function can be lower bounded as follows,

$$\tilde{\Psi}(\min\{Y^t, \varphi + 1\}) = \sum_{k=1}^n (\min\{Y_k, \varphi + 1\} - \varphi)^2 \geq \sum_{i=0}^{\varphi - \alpha^{1/8} - 1} Z_i^t \cdot (\varphi - i)^2.$$

It is easy to check that the last term is minimized by the configuration where $Z_{\varphi - \alpha^{1/8} - 1}^t$ is as large as possible, which translates into the following lower bound on $\tilde{\Psi}$:

$$\tilde{\Psi}(\min\{Y^t, \varphi + 1\}) \geq \frac{n/2}{\alpha^{1/8} + 2} \cdot (\alpha^{1/8} + 1)^2 \geq n^{9/8}/4.$$

Altogether this implies equation (B.3), i.e., that for any $\varphi n/2 \leq t \leq 2\varphi n$,

$$\Pr\{\mathcal{B}(t)\} \leq e^{-\Omega(n^{1/2})}.$$

This finishes the proof of the first statement. The second part follows by identical arguments and the third statement is an immediate consequence of the second statement. This completes the proof of Lemma 4.2.

C Omitted Proofs from Section 3

Proof of Lemma 3.4

Proof. Consider now the beginning of stage τ where $\Phi^\tau \geq \rho_n$, where $\rho_n := (\varepsilon + \kappa)/(\kappa/2) \cdot (1 + \varepsilon)^{C_1} \cdot n$. Let $A \subseteq \{1, \dots, n\}$ be the set of bins with a hole of size at most C_1 , i.e. with load larger than $\tau + 2 - C_1$. Note that for every $i \in \{1, \dots, n\}$, $\mathbf{E}\{\Phi_i(L^{\tau+1}) \mid L^\tau = \ell^\tau\} \leq (1 + \varepsilon) \cdot \Phi_i(\ell^\tau)$, and therefore,

$$\sum_{i \in A} \mathbf{E}\{\Phi_i(L^{\tau+1}) \mid L^\tau = \ell^\tau\} \leq (1 + \varepsilon) \cdot \sum_{i \in A} \Phi_i(\ell^\tau).$$

On the other hand, for every bin $i \notin A$, Lemma 3.3 implies that

$$\mathbf{E}\{\Phi_i(L^{\tau+1}) \mid L^\tau = \ell^\tau\} \leq (1 - \kappa) \cdot \Phi_i(\ell^\tau),$$

for some constant $\kappa < 1$. Hence,

$$\begin{aligned} \sum_{i \notin A} \mathbf{E}\{\Phi_i(L^{\tau+1}) \mid L^\tau = \ell^\tau\} &\leq \sum_{i \notin A} (1 - \kappa) \cdot \Phi_i(\ell^\tau) = (1 - \kappa) \cdot \sum_{i \notin A} \Phi_i(\ell^\tau) \\ &= (1 - \kappa) \cdot \left(\sum_{i=1}^n \Phi_i(\ell^\tau) - \sum_{i \in A} \Phi_i(\ell^\tau) \right). \end{aligned}$$

Putting everything together we get

$$\begin{aligned} \mathbf{E}\{\Phi(L^{\tau+1}) \mid L^\tau = \ell^\tau\} &= \sum_{i \in A} \mathbf{E}\{\Phi_i(L^{\tau+1}) \mid L^\tau = \ell^\tau\} + \sum_{i \notin A} \mathbf{E}\{\Phi_i(L^{\tau+1}) \mid L^\tau = \ell^\tau\} \\ &\leq (1 + \varepsilon) \cdot \sum_{i \in A} \Phi_i(\ell^\tau) + (1 - \kappa) \cdot \left(\sum_{i=1}^n \Phi_i(\ell^\tau) - \sum_{i \in A} \Phi_i(\ell^\tau) \right) \\ &= (\varepsilon + \kappa) \cdot \sum_{i \in A} \Phi_i(\ell^\tau) + (1 - \kappa) \cdot \sum_{i=1}^n \Phi_i(\ell^\tau) \\ &\leq (\varepsilon + \kappa) \cdot n \cdot (1 + \varepsilon)^{C_1} + (1 - \kappa) \cdot \Phi(\ell^\tau) \\ &\leq \frac{\kappa}{2} \cdot \Phi(\ell^\tau) + (1 - \kappa) \cdot \Phi(\ell^\tau) \\ &\leq \left(1 - \frac{\kappa}{2}\right) \cdot \Phi(\ell^\tau), \end{aligned}$$

where in the second to last line we used the lower bound on $\Phi(\ell^\tau)$.

Proof of Corollary 3.5

Proof. We show this result by induction on τ . The basic idea is that if the expected potential value is already small for step τ , then it can only increase by a factor of at most $(1 + \varepsilon)$. On the other hand, if the expected potential value is large, then Lemma 3.4 implies that the expected potential will decrease in step $\tau + 1$. Combining these insights yields the first statement of Corollary 3.5.

Let us now turn to the formal proof. Clearly, $\mathbf{E}\{\Phi(\ell^0)\} = \mathcal{O}(n)$. We break the proof into two cases:

- (1) As the first case, assume now that $\mathbf{E}\{\Phi(L^{\tau-1})\} \geq (1 + \varepsilon) \cdot \rho_n/(\kappa/2)$ and also, similar to the proof of Lemma 3.6, let $\mathcal{S}$ be the set of possible values that $\Phi(L^{\tau-1})$ can attain. Observe that $\sum_{k \in \mathcal{S}} k \cdot$

$\Pr\{\Phi(L^{\tau-1}) = k\} \leq \rho_n$. Consequently

$$\begin{aligned}
\mathbf{E}\{\Phi(L^\tau)\} &= \sum_{k \in \mathcal{S}} \Pr\{\Phi(L^{\tau-1}) = k\} \cdot \mathbf{E}\{\Phi(L^\tau) \mid \Phi(L^{\tau-1}) = k\} \\
&= \sum_{\substack{k \in \mathcal{S}: k \leq \rho_n \\ k \leq \rho_n}}^{\rho_n} \Pr\{\Phi(L^{\tau-1}) = k\} \cdot \mathbf{E}\{\Phi(L^\tau) \mid \Phi(L^{\tau-1}) = k\} \\
&\quad + \sum_{k \in \mathcal{S}: k > \rho_n} \Pr\{\Phi(L^{\tau-1}) = k\} \cdot \mathbf{E}\{\Phi(L^\tau) \mid \Phi(L^{\tau-1}) = k\} \\
&\leq \sum_{k \in \mathcal{S}: k \leq \rho_n} \Pr\{\Phi(L^{\tau-1}) = k\} \cdot (1 + \varepsilon) \cdot k + \sum_{k \in \mathcal{S}: k > \rho_n} \Pr\{\Phi(L^{\tau-1}) = k\} \cdot (1 - \kappa) \cdot k \\
&\leq (1 + \varepsilon) \cdot \rho_n + \sum_{k \in \mathcal{S}: k > \rho_n} \Pr\{\Phi(L^{\tau-1}) = k\} \cdot (1 - \kappa) \cdot k \\
&\leq (\kappa/2) \cdot \mathbf{E}\{\Phi(L^{\tau-1})\} + (1 - \kappa) \cdot \sum_{k \in \mathcal{S}} \Pr\{\Phi(L^{\tau-1}) = k\} \cdot k \\
&\leq (1 - \kappa/2) \cdot \mathbf{E}\{\Phi(L^{\tau-1})\},
\end{aligned}$$

where the second last line uses the assumption that $\mathbf{E}\{\Phi(L^{\tau-1})\} \geq (1 + \varepsilon)\rho_n/(\kappa/2)$.

(2) On the other hand, if $\mathbf{E}\{\Phi(L^{\tau-1})\} \leq (1 + \varepsilon) \cdot \rho_n/(k/2)$, then

$$\mathbf{E}\{\Phi(L^\tau)\} \leq (1 + \varepsilon) \cdot \mathbf{E}\{\Phi(L^{\tau-1})\} \leq (1 + \varepsilon)^2 \rho_n/(\kappa/2).$$

Hence for all stages τ , the load vector L^τ at the end of stage τ fulfills $\mathbf{E}\{\Phi(L^\tau)\} = \mathcal{O}(n)$.