

Balls into Bins via Local Search

Paul Bogdan* Thomas Sauerwald† Alexandre Stauffer‡ He Sun§

Abstract

We propose a natural process for allocating n balls into n bins that are organized as the vertices of an undirected graph G . Each ball first chooses a vertex u in G uniformly at random. Then the ball performs a local search in G starting from u until it reaches a vertex with local minimum load, where the ball is finally placed on. In our main result, we prove that this process yields a maximum load of only $\Theta(\log \log n)$ on expander graphs. In addition, we show that for d -dimensional grids the maximum load is $\Theta\left(\left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}\right)$. Finally, for almost regular graphs with minimum degree $\Omega(\log n)$, we prove that the maximum load is constant and also reveal a fundamental difference between random and arbitrary tie-breaking rules.

1 Introduction

It is well known that if each of n balls is placed sequentially into one of n bins chosen independently and uniformly at random, then the highest loaded bin is likely to contain $\Theta\left(\frac{\log n}{\log \log n}\right)$ balls. We call this process the 1-choice process. Alternatively, in the d -choice process, each ball is allowed to choose d bins independently and uniformly at random and is placed in the least loaded among the d bins. It was shown by Azar et al. [3] and Karp et al. [11] that the maximum load reduces drastically to $\Theta\left(\frac{\log \log n}{\log d}\right)$ in the d -choice process. The constants omitted in the Θ are known [3] and as shown by Vöcking [18], a slight modification of the d -choice process can decrease the maximum load further. Berenbrink et al. [4] extended

these results to the case where the number of balls is larger than the number of bins.

In some applications, it is important to allow each ball to choose bins in a *correlated* way. For example, such correlations occur naturally in distributed systems, where the bins represent processors that are interconnected as a graph and the balls represent tasks that need to be assigned to processors. From a practical point of view, letting each task choose d *independent* random bins may be undesirable, since the cost of accessing two bins which are far away in the graph may be higher than accessing two bins which are nearby. Furthermore, in some contexts, tasks are actually created by the processors, which are then able to forward tasks to other processors to achieve a more balanced load distribution. In such settings, allocating balls close to the processor that created it is certainly very desirable as it reduces the costs of probing the load of a processor and allocating the task.

We propose a very natural and simple process for allocating balls into bins that are interconnected as a graph. We refer to this process as *local search allocation*. At each step, a ball is “born” in a bin chosen independently and uniformly at random, which we call the *birthplace* of the ball. Then, starting from its birthplace, the ball performs a *local search* in the graph, where in each step the ball moves to the adjacent bin with the smallest load, provided that the load is strictly smaller than the load of the bin the ball is currently in. Unless otherwise stated, we assume that ties are broken uniformly at random. The local search ends when the ball visits the first vertex that is a *local minimum*, which is a vertex for which no neighbor has a smaller load. After that, the next ball is born and is allocated according to the procedure described above, and this process is repeated. See Figure 1 for an illustration.

In this paper we analyze the asymptotic behavior of the maximum load obtained when n balls are allocated to n bins as $n \rightarrow \infty$. The main question is whether the local search allocation ensures a small load amount like the d -choice process with $d \geq 2$. Our main result gives a positive answer for the important case where the n balls are allocated to n bins which are organized as the vertices of an *expander* graph.

^{*}Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA.

[†]Max Planck Institute for Informatics, Saarbrücken, Germany and Cluster of Excellence “Multimodal Computing and Interaction”, Universität des Saarlandes, Saarbrücken, Germany. Part of this work was done while visiting Microsoft Research Redmond.

[‡]Microsoft Research, Redmond WA, USA. Part of this work was done while visiting Max Planck Institute for Informatics, Saarbrücken, Germany.

[§]Max Planck Institute for Informatics, Saarbrücken, Germany, and Institute of Modern Mathematics and Physics, Fudan University, Shanghai, China.

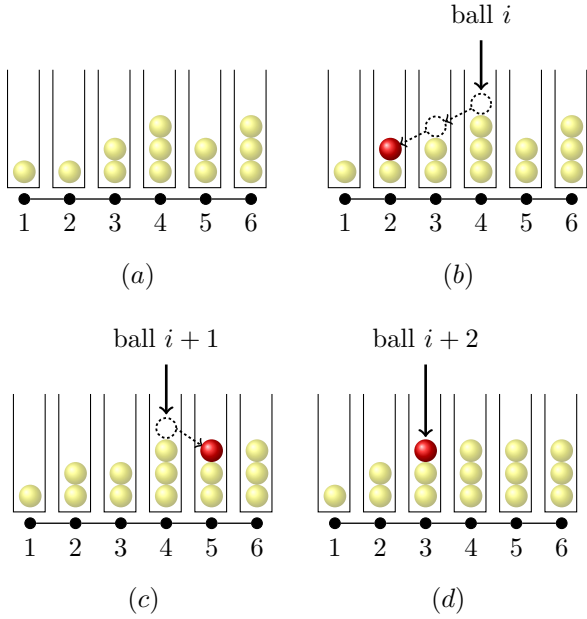


Figure 1: Illustration of the local search allocation. Black circles represent the vertices 1–6 arranged as a path, and the yellow circles represent the balls of the process (the most recently allocated ball is marked red). Figure (a) shows the configuration after placing $i - 1$ balls. As shown in Figure (b), ball i born at vertex 4 has two choices in the first step of the local search (vertices 3 or 5) and is finally allocated to vertex 2. Figure (c) and (d) shows the placement of ball $i + 1$ and $i + 2$.

We show, in this case, that the maximum load is $\Theta(\log \log n)$, which has the same order of magnitude as of the d -choice process (refer to Theorem 3.1 for a more general statement of Theorem 1.1 below).

THEOREM 1.1. (EXPANDER GRAPHS) *Let G be any expander graph with constant maximum degree. Then, as $n \rightarrow \infty$, the maximum load after n balls are allocated is $\Theta(\log \log n)$ with probability $1 - o(1)$.*

REMARK 1.2. Theorem 1.1 holds also when ties are not broken uniformly at random but by means of a fixed permutation. In this tie-breaking procedure, for each vertex $v \in V$, we associate an arbitrary permutation ξ_v of the neighbors of v . Then, whenever a ball is currently at v during its local search, the ball breaks ties by using the order in the permutation ξ_v . With this tie-breaking procedure, the local search allocation is a *deterministic* function of the birthplaces of the balls and the permutations $\{\xi_v\}_{v \in V}$.

Before discussing and comparing our main result with existing results, we proceed to state our other results. Recall that, unless otherwise stated, we assume that in the local search allocation ties are broken uniformly at random. An important instance for applications is when the bins are organized as a

ring or a grid. The theorem below establishes the maximum load in this case up to constant factors.

THEOREM 1.3. (GRID GRAPHS) *Let G be any d -dimensional grid graph, where $d \geq 1$ is any integer constant. Then, as $n \rightarrow \infty$, the maximum load after n balls are allocated is $\Theta\left(\left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}\right)$ with probability $1 - o(1)$.*

In comparison to Theorem 1.1, the above theorem shows that the maximum load can be quite high on graphs with small expansion. Besides the expansion, it is also conceivable that a large degree ensures a small maximum load. The next theorem confirms this intuition by showing that a constant maximum load is obtained for any almost regular graph with minimum degree $\Omega(\log n)$, where a graph is *almost regular* if the ratio between the minimum and maximum degrees is a constant. This class of graphs includes hypercubes and Erdős-Rényi random graphs with average degree $(1 + \varepsilon) \log n$, for any $\varepsilon > 0$.

THEOREM 1.4. (DENSE GRAPHS) *Let G be any almost regular graph with minimum degree $\Omega(\log n)$. Then, as $n \rightarrow \infty$, the maximum load after n balls are allocated is $\mathcal{O}(1)$ with probability $1 - o(1)$.*

In the next two theorems, we do not restrict the local search allocation to break ties uniformly at random. In particular, we show that deviating from the usual random tie breaking rule can dramatically increase the maximum load. Here, as in Remark 1.2, we define a tie-breaking rule as a collection of variables $\{\xi_v\}_{v \in V}$, where each ξ_v is a permutation of the neighbors of v .

THEOREM 1.5. (IMPACT OF TIE-BREAKING RULES) *For any $d = \omega(1)$ as $n \rightarrow \infty$, there is a d -regular graph and a choice of $\{\xi_v\}_{v \in V}$ for which the maximum load after n balls are allocated is at least $\Omega(\min\{d^{1/4}, \frac{\log n}{\log d}\})$ with probability $1 - o(1)$.*

To highlight the difference between random tie breaking and arbitrary tie breaking, Theorem 1.5 establishes that there exists a d -regular graph G with $d = \Theta(\log n)$ and a tie-breaking rule for which the maximum load in G is $\Omega((\log n)^{1/4})$. On the other hand, by Theorem 1.4 we have that breaking ties uniformly at random leads to a constant maximum load in G .

Our final result establishes some lower bounds on the maximum load. These lower bounds hold for arbitrary tie-breaking rules.

THEOREM 1.6. (LOWER BOUNDS) *For any graph with maximum degree Δ , the maximum load after n*

balls are allocated is at least $\Omega\left(\frac{\log \log n}{\log \Delta}\right)$ with probability $1 - o(1)$ as $n \rightarrow \infty$. Furthermore, for any integer $2 \leq d \leq \frac{\log n}{e}$, there is a d -regular graph for which the maximum load after n balls are allocated is at least $\Omega\left(\sqrt{\frac{\log n}{d \cdot \log\left(\frac{\log n}{d}\right)}}\right)$ with probability $1 - o(1)$ as $n \rightarrow \infty$.

The combination of Theorem 1.5 and the second statement of Theorem 1.6 shows that the two conditions concerning the degree and the random tie breaking in Theorem 1.4 are not only sufficient but also necessary for obtaining a constant maximum load.

While in the d -choice process a ball only probes the load of a constant number of bins (provided $d = \mathcal{O}(1)$), the local search allocation may probe the load of $\omega(1)$ bins for some balls. However, the number of bins which a given ball can probe is bounded above by Δ times the load of the birthplace of the ball, where Δ is the maximum degree of G . Therefore, the *expected* number of probed bins per ball is at most Δ and, for the case of expander graphs with constant maximum degree, Theorem 1.1 implies that the maximum number of probed bins per ball is $\mathcal{O}(\log \log n)$.

An important feature of our local search allocation is that it maintains a smooth load distribution (cf. Lemma 2.2); i.e., the load difference for each edge of G is at most one and balls are only placed in the bins whose load is a local minimum. Hence, if each ball is controlled by an agent who strives for a minimization of its load, then each agent has a natural incentive to follow the local search allocation.

It is also important to remark that our process uses only a small amount of randomness. For instance, when ties are broken by means of a *deterministic* collection of $\{\xi_v\}_{v \in V}$, the only randomness comes from the birthplaces of the balls, which requires only $n \log_2 n$ random bits. From this perspective, it is comparable to the process by Mitzenmacher et al. [15], which uses the same number of random bits and achieves a maximum load of $\mathcal{O}(\log \log n)$ as well.

Further Related Work. The work that is most related to ours is that of Kenthapadi and Panigrahy [12]. They analyzed a balls-into-bins model where each ball chooses a pair of adjacent bins uniformly at random and is placed in the bin with the smaller load. They proved that, for any d -regular graph, the maximum load after n balls are allocated is $\log \log n + \mathcal{O}\left(1 + \frac{\log n}{\log(d/\log^4 n)}\right)$, and also showed an almost matching lower bound. Hence, to retain the maximum load of $\mathcal{O}(\log \log n)$ from the 2-choice process, the degree d must be as large as $\Omega(n^{\Omega(1/(\log \log n))})$. In contrast, our local search

allocation achieves a maximum load of $\mathcal{O}(\log \log n)$ even for a large class of constant-degree graphs. The model from Kenthapadi and Panigrahy [12] was also studied implicitly in Peres et al. [16], where the authors analyzed the *gap* defined as the difference between the maximum and minimum load. They proved that, even if the number of balls m is much larger than the number of bins n , the gap is $\Theta(\log n)$ for expander graphs and, for the cycle, the gap is between $\Omega(\log n)$ and $\mathcal{O}(n \log n)$, i.e., independent of the number of balls.

In a different context, Adler et al. [1] studied a related graph-based coupon collector process. In this process, each ball performs a local search but distinguishes only between empty and nonempty bins. Therefore, each ball is allocated either to its birthplace or to one of its neighbors. For different graphs, the authors analyzed how many balls have to be allocated until all bins are non-empty.

There are several other variations of balls-into-bins models for which the power of two choices has been analyzed (we refer to [14] for a survey). For instance, Broder and Mitzenmacher [6] considered a multidimensional version where balls correspond to 0/1 vectors. Godfrey [9] analyzed balanced allocations on hypergraphs, extending the model of [12] for allocations on graphs. Among other results, the author obtained a constant maximum load for an almost regular hypergraph if each hyperedge consists of $d = \Theta(\log n)$ bins, which is comparable to our result for dense graphs (Theorem 1.4).

Local search is a generic method to solve optimization problems and several recent studies in algorithmic game theory involve local search-based processes. In contrast to our model, these processes usually start from a state where all tasks are assigned and allow, either sequentially or in parallel, tasks to be reallocated by using so-called improvement steps (a.k.a. selfish steps). For various settings, lower and upper bounds on the number of improvement steps until a Nash equilibrium is found have been shown [5, 8, 10]. We note that the total number of improvement steps in our process is always bounded by $\mathcal{O}(n)$ with probability $1 - o(1)$, regardless of the underlying graph.

2 Basic properties

We start with some notation. Let $G = (V, E)$ be an undirected, connected graph, where the n vertices represent n bins to which n balls should be allocated. For each node $v \in V$, denote by $X_v^{(i)}$ the load of v right after the i th ball is allocated. Thus, we initially have $X_v^{(0)} = 0$ for all $v \in V$. Let $X_{\max}^{(n)}$ be the

maximum load after n balls have been allocated; i.e.,

$$X_{\max}^{(n)} = \max_{v \in V} X_v^{(n)}.$$

Let $U_i \in V$ be the birthplace of ball i , so U_i is a uniformly random sample from V . Recall that in the 1-choice process, for all $i \geq 1$, ball i is allocated to U_i . For any $v \in V$, let $\bar{X}_v^{(n)}$ be the load of v after n balls are allocated according to the 1-choice process. In symbols, we have

$$\bar{X}_v^{(n)} = |\{i \in [1, n] : U_i = v\}|.$$

With this, define the maximum load for the 1-choice process as

$$\bar{X}_{\max}^{(n)} := \max_{v \in V} \bar{X}_v^{(n)}.$$

For vectors $A = (a_1, a_2, \dots, a_n)$ and $A' = (a'_1, a'_2, \dots, a'_n)$ such that $\sum_{i=1}^n a_i = \sum_{i=1}^n a'_i$, we say that A *majorizes* A' if, for each $\kappa = 1, 2, \dots, n$, the sum of the κ largest entries of A is at least the sum of the κ largest entries of A' . More formally, if j_1, j_2, \dots, j_n are distinct numbers such that $a_{j_1} \geq a_{j_2} \geq \dots \geq a_{j_n}$ and j'_1, j'_2, \dots, j'_n are distinct numbers such that $a'_{j'_1} \geq a'_{j'_2} \geq \dots \geq a'_{j'_n}$, then A majorizes A' if

$$\sum_{i=1}^{\kappa} a_{j_i} \geq \sum_{i=1}^{\kappa} a'_{j'_i}$$

for all $\kappa = 1, 2, \dots, n$. The lemma below establishes that the load vector obtained by the 1-choice process majorizes the load vector obtained by the local search allocation. As a consequence, the maximum load of our local search allocation process is $\mathcal{O}(\log n / (\log \log n))$ for any graph.

LEMMA 2.1. (COMPARISON WITH 1-CHOICE) *For any fixed $k \geq 0$, we can couple $X^{(k)}$ and $\bar{X}^{(k)}$ so that, with probability 1, $\bar{X}^{(k)}$ majorizes $X^{(k)}$. Consequently, we have that, for all $k \geq 0$, $\bar{X}_{\max}^{(k)}$ stochastically dominates $X_{\max}^{(k)}$.*

Proof. The proof is by induction on k . Clearly, for $k = 0$, we have $X_v^{(0)} = \bar{X}_v^{(0)} = 0$ for all $v \in V$. Now, assume that we can couple $X^{(k-1)}$ with $\bar{X}^{(k-1)}$ so that $\bar{X}^{(k-1)}$ majorizes $X^{(k-1)}$. Now let j_1, j_2, \dots, j_n be distinct elements of V so that $X_{j_1}^{(k-1)} \geq X_{j_2}^{(k-1)} \geq \dots \geq X_{j_n}^{(k-1)}$. Similarly, let j'_1, j'_2, \dots, j'_n be distinct elements of V so that $\bar{X}_{j'_1}^{(k-1)} \geq \bar{X}_{j'_2}^{(k-1)} \geq \dots \geq \bar{X}_{j'_n}^{(k-1)}$. Now let ℓ be a uniformly random integer from 1 to n . Then, for the process $(X_v^{(k)})_{v \in V}$, we

let the k th ball be born at vertex j_ℓ and define ι such that j_ℓ is the vertex to which the k th ball is allocated. Note that, $\iota \geq \ell$. For the process $(\bar{X}_v^{(k)})_{v \in V}$, we set the birthplace of the k th ball to be j'_ℓ . Therefore, for any $\kappa = 1, 2, \dots, n$, we have

$$\begin{aligned} & \sum_{i=1}^{\kappa} \bar{X}_{j'_i}^{(k)} \\ &= \sum_{i=1}^{\kappa} \bar{X}_{j'_i}^{(k-1)} + \mathbf{1}(\kappa \geq \ell) \geq \sum_{i=1}^{\kappa} X_{j_i}^{(k-1)} + \mathbf{1}(\kappa \geq \ell) \\ &\geq \sum_{i=1}^{\kappa} X_{j_i}^{(k-1)} + \mathbf{1}(\kappa \geq \iota) = \sum_{i=1}^{\kappa} X_{j_i}^{(k)}, \end{aligned}$$

where the first inequality follows by the induction hypothesis and the second inequality holds as $\iota \geq \ell$. \blacksquare

For any $v \in V$, let N_v be the set of neighbors of v in G . The next lemma establishes that the local search allocation always maintains a *smoothed* load vector in the sense that the load of any two adjacent vertices differs by at most 1.

LEMMA 2.2. (SMOOTHNESS) *For any $k \geq 0$, any $v \in V$ and any $u \in N_v$, we have that $|X_v^{(k)} - X_u^{(k)}| \leq 1$.*

Proof. In order to obtain a contradiction, suppose that $X_v^{(k)} \geq X_u^{(k)} + 2$, and let j be the last ball allocated to v . Then, we have that

$$X_v^{(j-1)} = X_v^{(k)} - 1 \geq X_u^{(k)} + 1 \geq X_u^{(j-1)} + 1.$$

Therefore, the moment the j th ball is born, vertex v has at least one neighbor with load strictly smaller than v . Therefore, ball j is not allocated to v , establishing a contradiction. \blacksquare

For any vertex $v \in V$ and integer $r \geq 0$, let N_v^r be the set of vertices of G whose distance to v is exactly r (in particular, $N_v^0 = \{v\}$), and let B_v^r be the set of vertices of G whose distance to v is at most r ; then $B_v^r = \bigcup_{i=0}^r N_v^i$. Below we show that, if for a given $v \in V$ we have an upper bound for the number of balls allocated to B_v^r , then we obtain an upper bound for the load of v .

LEMMA 2.3. (UPPER BOUND) *Let v be an arbitrary vertex of G . Suppose that there exists an integer $r \geq 1$ and a positive Ψ , that may depend on n , such that the total number of balls allocated to the vertices of B_v^r is at most $\Psi |B_v^r|$; i.e., $\sum_{u \in B_v^r} X_u^{(n)} \leq \Psi |B_v^r|$. Then, we have that*

$$X_v^{(n)} \leq \Psi + \sum_{i=0}^r i \frac{|N_v^i|}{|B_v^r|}.$$

Proof. Note that, by Lemma 2.2, for any $u \in N_v^i$, we have $X_u^{(n)} \geq X_v^{(n)} - i$. Using this and the condition of the lemma, we obtain

$$\begin{aligned} \Psi|B_v^r| &\geq \sum_{u \in B_v^r} X_u^{(n)} \geq \sum_{i=0}^r (X_v^{(n)} - i)|N_v^i| \\ &= X_v^{(n)}|B_v^r| - \sum_{i=0}^r i|N_v^i|. \end{aligned}$$

■

Complementing the previous lemma, we now prove a lower bound on the maximum load that depends only on the number of balls born at a subset of vertices and the cardinality of a small ball around that subset.

LEMMA 2.4. (LOWER BOUND) *For any subset $S \subseteq V$, let $\Phi_S := \sum_{i=1}^n \mathbf{1}(U_i \in S)$ be the number of balls born in S . Then, the maximum load $\beta := X_{\max}^{(n)}$ satisfies $\beta \cdot |B_S^\beta| \geq \Phi_S$, where $B_S^\beta := \bigcup_{s \in S} B_s^\beta$ is the set of vertices with distance at most β from S .*

Proof. If the maximum load is β , then every ball born at some vertex u is allocated in B_u^β , and clearly the load of any vertex in B_S^β is at most β . Combining these two insights yields

$$\Phi_S \leq \sum_{u \in B_S^\beta} X_u^{(n)} \leq \beta \cdot |B_S^\beta|.$$

Hence the maximum load β satisfies $\beta \cdot |B_S^\beta| \geq \Phi_S$. ■

For the next two lemmas, we need ties to be broken either uniformly at random or by means of a fixed permutation ξ_v of the neighbors of v for each $v \in V$. The next lemma establishes that the load vector $X^{(n)}$ satisfies a Lipschitz condition, which will turn out to be crucial in our proofs.

LEMMA 2.5. (LIPSCHITZ PROPERTY) *Let $k \geq 1$ be fixed and $u_1, u_2, \dots, u_k \in V$ be arbitrary. Let $(X_v^{(k)})_{v \in V}$ be the load of the vertices of G after the local search allocation places k balls with birthplaces u_1, u_2, \dots, u_k . Let $i \in \{1, 2, \dots, k\}$ be fixed, and let $(Y_v^{(k)})_{v \in V}$ be the load of the vertices of G after the local search allocation places k balls with birthplaces $u_1, u_2, \dots, u_{i-1}, u'_i, u_{i+1}, u_{i+2}, \dots, u_k$, where $u'_i \in V$ is arbitrary. In other words, $(Y_v^{(k)})_{v \in V}$ is obtained from $(X_v^{(k)})_{v \in V}$ by changing the birthplace of the i th ball from u_i to u'_i . Assume that, for both processes, the local search allocation breaks ties either uniformly at random or via the permutations $\{\xi_v\}_{v \in V}$ described*

in Remark 1.2. Then, there exists a coupling such that

$$(2.1) \quad \sum_{v \in V} |X_v^{(k)} - Y_v^{(k)}| \leq 2.$$

Proof. We refer to the process defining the variables $X^{(k)}$ as the X process, and the process defining the variables $Y^{(k)}$ as the Y process. If ties are broken uniformly at random, then for each $v \in V$ and $i \geq 1$, we define $\xi_v^{(i)}$ to be an independent and uniformly random permutation of the neighbors of v . We use this permutation for both the X and Y processes to break ties when ball i is at vertex v . Then, since the first $i-1$ balls have the same birthplaces in both processes, we have that

$$(2.2) \quad X_v^{(i-1)} = Y_v^{(i-1)} \quad \text{for all } v \in V.$$

Now, when adding the i th ball, we let v_i be the vertex to which this ball is allocated in the X process and v'_i be the vertex to which this ball is allocated in the Y process. If $v_i = v'_i$, then $X_u^{(i)} = Y_u^{(i)}$ for all $u \in V$ and (2.1) holds. More generally, we have that

$$(2.3) \quad \begin{cases} X_{v_i}^{(i)} = Y_{v_i}^{(i)} + \mathbf{1}(v_i \neq v'_i) \\ Y_{v'_i}^{(i)} = X_{v'_i}^{(i)} + \mathbf{1}(v_i \neq v'_i) \\ X_u^{(i)} = Y_u^{(i)} \text{ for } u \in V \setminus \{v_i, v'_i\}. \end{cases}$$

If $i = k$, then this implies (2.1) and the lemma holds.

For the case $i < k$, we add ball $i+1$ and are going to define v_{i+1} and v'_{i+1} so that (2.3) holds with i replaced by $i+1$. Then the proof of the lemma is completed by induction. We assume that $v_i \neq v'_i$, otherwise (2.1) clearly holds. We note that v_{i+1} and v'_{i+1} will not be in the same way as v_i and v'_i . The role of v_{i+1} and v'_{i+1} is to be the only vertices whose loads in the X and Y processes are different. The definition of v_{i+1} and v'_{i+1} will vary depending on the situation. For this, let ball $i+1$ be born at u_{i+1} and define w to be the vertex on which ball $i+1$ is allocated in the X process and w' to be the vertex on which ball $i+1$ is allocated in the Y process. We can assume that $w \neq w'$, otherwise (2.3) holds with i replaced by $i+1$ by setting $v_{i+1} = v_i$ and $v'_{i+1} = v'_i$.

Now we analyze ball $i+1$. It is crucial to note that, during the local search of ball $i+1$, if it does not enter v_i in the Y process and does not enter v'_i in the X process, then ball $i+1$ follows the same path in both processes. Since we are in the case $w \neq w'$, we can assume without loss of generality that ball $i+1$ eventually visits v_i in the Y process. In this case, since the local search performed by ball i in the X process stops at vertex v_i , we have that v_i is a local

minimum for ball $i+1$ in the Y process, which implies that $w' = v_i$. (The case when ball $i+1$ visits v'_i in the X process follows by a symmetric argument.) So, since $w \neq w'$, we have $X_{v_i}^{(i+1)} = Y_{v_i}^{(i+1)}$. Then we let $v_{i+1} = w$. If $w = v'_i$, we set $v'_{i+1} = w$ and (2.3) holds since $X_u^{(i+1)} = Y_u^{(i+1)}$ for all $u \in V$. Otherwise we set $v'_{i+1} = v'_i$, and (2.3) holds as well. ■

The following is a consequence of Lemma 2.5.

LEMMA 2.6. (MONOTONICITY) *Let $k \geq 1$ be fixed and $u_1, u_2, \dots, u_k \in V$ be arbitrary. Let $(X_v^{(k)})_{v \in V}$ be the load of the vertices after k balls are allocated with birthplaces u_1, u_2, \dots, u_k . Let $i \in \{1, 2, \dots, k\}$ be fixed, and let $(Z_v^{(i,k)})_{v \in V}$ be the load of the vertices of G after $k-1$ balls are allocated with birthplaces $u_1, u_2, \dots, u_{i-1}, u_{i+1}, u_{i+2}, \dots, u_k$. In other words, $Z_v^{(i,k)}$ is obtained from $X_v^{(k)}$ by removing ball i . Assume that, for both processes, the local search allocation breaks ties either uniformly at random or via the variables $(\xi_v)_{v \in V}$ described in Remark 1.2. Then, there exists a coupling such that*

$$\sum_{v \in V} \left| X_v^{(k)} - Z_v^{(i,k)} \right| = 1.$$

Proof. Let G' be the graph obtained from G by adding an isolated node w ; i.e., G' has vertex set $V \cup \{w\}$ and the same edge set as G . Applying Lemma 2.5 to G' with the same choice of $u_1, \dots, u_k \in V$ and with $u'_i = w$ gives

$$\sum_{v \in V \cup \{w\}} \left| X_v^{(k)} - Y_v^{(k)} \right| = 2.$$

Since $Y_w^{(k)} = 1$, $X_w^{(k)} = 0$ and $Z_v^{(i,k)} = Y_v^{(k)}$ for any $v \in V$, we conclude that

$$\sum_{v \in V} \left| X_v^{(k)} - Z_v^{(i,k)} \right| = \sum_{v \in V} \left| X_v^{(k)} - Y_v^{(k)} \right| = 1.$$

We now use Lemma 2.6 to prove a type of subadditivity property. In a simpler statement, we show that, if for k independent copies of the local search allocation with m balls the maximum load is at most x , then the maximum load obtained after placing km balls via local search allocation is at most kx .

LEMMA 2.7. (SUBADDITIVITY) *For any $1 \leq z \leq n$, and any $x \geq 0$, it holds that*

$$\Pr \left[X_{\max}^{(n)} \geq \lceil n/z \rceil \cdot x \right] \leq \lceil n/z \rceil \cdot \Pr \left[X_{\max}^{(z)} \geq x \right].$$

Proof. Let U_1, U_2, \dots, U_n be independent uniform random samples from V . Then, $X_{\max}^{(n)}$ is the maximum load after n balls are added to G with birthplaces U_1, U_2, \dots, U_n . We define $k := \lceil n/z \rceil$ independent copies of the local search allocation, where in the first copy we allocate z balls according to the birthplaces U_1, U_2, \dots, U_z , in the second copy we allocate z balls according to the birthplaces $U_{z+1}, U_{z+2}, \dots, U_{2z}$ and so on and so forth. Hence, in total we allocate $\lceil n/z \rceil \cdot z \geq n$ balls. Let M_1, M_2, \dots, M_k be the maximum load of each copy after z balls are allocated. Then, we claim that $X_{\max}^{(n)} \leq X_{\max}^{(k \cdot z)} \leq \sum_{i=1}^k M_i$. This follows since, for each copy i , after allocating the z balls, we can successively add more balls in such a way that all vertices have load exactly M_i in copy i . Then, by taking the union of all copies, we obtain a balls-into-bins process with $\sum_{i=1}^k M_i$ balls, n of which have birthplaces U_1, U_2, \dots, U_n . Then, by Lemma 2.6, the maximum load in this process, which is $\sum_{i=1}^k M_i$, is at least $X_{\max}^{(n)}$. Therefore,

$$\begin{aligned} & \Pr \left[X_{\max}^{(n)} \geq \lceil n/z \rceil \cdot x \right] \\ & \leq \Pr \left[\sum_{i=1}^k M_i \geq \lceil n/z \rceil \cdot x \right] \\ & \leq \Pr \left[\bigcup_{i=1}^k \{M_i \geq x\} \right] \leq k \Pr \left[X_{\max}^{(z)} \geq x \right]. \end{aligned}$$

■

3 Expander graphs

In this section we give the proof of our main result, Theorem 1.1, which establishes an upper bound for the maximum load after n balls are allocated to the vertices of an expander graph. In fact, we can prove this theorem in a more general setting. As before, for each $u \in V$ and $r = 1, 2, \dots$, we define N_u^r to be the set of vertices of V whose distance to u is exactly r , and B_u^r to be the set of vertices of V whose distance to u is at most r ; in symbols,

■ $N_u^r = \{v \in V : \text{graph distance between } u \text{ and } v \text{ is } r\}$
and $B_u^r = \bigcup_{i=0}^r N_u^i$. We say that G has *exponential growth* if there exists a constant $\phi > 0$ so that

$$(3.1) \quad |B_u^r| \geq \min \left\{ e^{\phi r}, \frac{n}{2} \right\} \quad \text{for all } u \in V \text{ and } r \geq 0.$$

Note that any graph with exponential growth has a diameter of $\mathcal{O}(\log n)$. Moreover, if G is an expander, then it has exponential growth. Therefore, Theorem 1.1 follows from the theorem below.

THEOREM 3.1. *If G has the exponential growth property defined in (3.1) and constant maximum degree,*

then there exists a positive constant C so that, as $n \rightarrow \infty$,

$$\Pr \left[X_{\max}^{(n)} \geq C \log \log n \right] \leq n^{-\omega(1)}.$$

We devote the remainder of this section to prove Theorem 3.1. We start with a high level description of the proof. We claim that, for any vertex v and some properly chosen $r_0 = \mathcal{O}(\log \log n)$,

$$(3.2) \quad \Pr \left[\sum_{u \in B_v^{r_0}} X_u^{(n)} \geq C |B_v^{r_0}| \right] \leq n^{-\omega(1)};$$

i.e., with very high probability, the number of balls allocated to the vertices of $B_v^{r_0}$ is at most $C |B_v^{r_0}|$. Having established (3.2), the proof follows immediately by applying Lemma 2.3 and taking the union bound over all v . Now, in order to prove (3.2), we use that Lemma 2.5 establishes that the load of the vertices satisfies a Lipschitz condition; i.e., if the birthplace of one ball is changed, the load vector can only change in two vertices. Therefore, we can apply Azuma's inequality to bound the probability that $\sum_{u \in B_v^{r_0}} X_u^{(n)} \geq C |B_v^{r_0}|$; however, for $r_0 = \mathcal{O}(\log \log n)$, the probability bound obtained via Azuma's inequality is not small enough to take the union bound over $v \in V$ later. Nevertheless, Azuma's inequality gives a small enough bound when applied to any radii larger than some $R \gg r_0$. Then, the idea is to control the number of balls allocated to B_v^{R-1} by using the bounds obtained for all radii $r \geq R$, and then apply an inductive argument to finally establish (3.2).

The main intuition why this analysis works is because a ball can only be allocated inside B_v^r if the ball is either born inside B_v^r or it is born in a vertex u at distance j to B_v^r but whose load, at the moment the ball is born, is at least j . This is true because, at each time a ball moves from a vertex u to a vertex $u' \in N_u$ during the local search, the load of u must be strictly larger than the load of u' . In other words, the ball traverses a *load decreasing path* from its birthplace to the vertex on which the ball is placed. Therefore, for a ball allocated in B_v^r , the larger the distance between the birthplace of this ball and B_v^r , the smaller the number of possibilities for the birthplace of the ball since these vertices must have a large load at the moment the ball is born. This, in a high-level description, gives that if we change the birthplace of a ball to a uniformly random vertex, the load of the vertices inside B_v^r does not change with high probability. This allows us to control the variance of the Lipschitz condition and apply a more refined version of Azuma's inequality to move from

radius r to radius $r - 1$ inductively until we reach radius r_0 . We remark that we actually need to control not only the number of balls allocated to B_v^r for all $r \in [r_0, R]$, but also the number of balls allocated to nodes in B_v^r whose load is at least ℓ for all $r \in [r_0, R]$ and many values of ℓ .

Now we proceed to the rigorous argument. We start by showing that the load at any given vertex has an exponential tail.

LEMMA 3.2. *Let v be any given vertex of V , and let Δ be the maximum degree of G . Then, it holds for any $z \geq 8e\Delta$ that*

$$\Pr \left[X_v^{(n)} \geq z \right] \leq 2 \left(\frac{4e\Delta}{z} \right)^z.$$

Proof. We start defining a sequence of vertices w_0, w_1, \dots and time steps $t_0 > t_1 > \dots$ such that, for every $j \geq 1$, ball t_j is born at vertex w_j and allocated to w_{j-1} . We start by setting $w_0 = v$ and $t_0 = n$. Inductively for $j \geq 1$, we let t_j be the last ball allocated to w_{j-1} before time t_{j-1} and set w_j to be the vertex at which ball t_j is born. So, for $j = 1$, t_1 is the last ball allocated to w_0 and w_1 is the vertex where ball t_1 is born. Note that, whenever $X_v^{(n)} \geq z$, if $w_1 = w_0$, we know that $X_{w_1}^{(t_1)} \geq z - 1$. On the other hand, if $w_1 \neq w_0$, then we have that $X_{w_1}^{(t_1)} \geq z + d(w_0, w_1) - 1$, where $d(u, v)$ is the graph distance between u and v . More general, for all $\ell \geq 1$, we have that

$$X_{w_\ell}^{(t_\ell)} \geq z + \sum_{j=1}^{\ell} (d(w_{j-1}, w_j) - 1).$$

We continue this procedure until we find a value of ℓ such that $X_{w_\ell}^{(t_\ell)} = 0$. Note that, for each j , we have $X_{w_j}^{(t_j)} - X_{w_{j-1}}^{(t_{j-1})} \geq -1$; consequently, we can have $X_{w_\ell}^{(t_\ell)} = 0$ only for $\ell \geq z$. In order to obtain an upper bound for $\Pr \left[X_v^{(n)} \geq z \right]$, we apply the first-moment method over all possible sequences $(w_1, w_2, \dots, w_\ell) \in V^\ell$ and $t_1 > t_2 > \dots > t_\ell$, for every $\ell \geq 1$, such that $\sum_{j=1}^{\ell} (d(w_{j-1}, w_j) - 1) \leq -z$.

With this, we have

$$\begin{aligned}
& \Pr \left[X_v^{(n)} \geq z \right] \\
& \leq \sum_{\ell \geq z} \sum_{\substack{w_1, w_2, \dots, w_\ell \\ t_1 > t_2 > \dots > t_\ell}} \mathbf{1} \left(\sum_{j=1}^{\ell} (d(w_{j-1}, w_j) - 1) \leq -z \right) \\
& \quad \cdot \Pr \left[\bigcap_{j=1}^{\ell} \{\text{ball } t_j \text{ is born at } w_j\} \right] \\
(3.3) \quad & \leq \sum_{\ell \geq z} \sum_{\substack{w_1, w_2, \dots, w_\ell \\ t_1 > t_2 > \dots > t_\ell}} \mathbf{1} \left(\sum_{j=1}^{\ell} (d(w_{j-1}, w_j) - 1) \leq -z \right) \frac{1}{n^\ell}.
\end{aligned}$$

Let $\lambda_i = d(w_{j-1}, w_j) - 1 \in \{-1, 0, 1, \dots\}$. For any fixed ℓ , we can estimate the number of possible sequences by counting the number of possibilities to choose λ_j and t_j so that $\sum_{j=1}^{\ell} \lambda_j \leq -z$, and then counting the number of possibilities to choose w_j accordingly. Clearly, there are at most $\binom{n}{\ell}$ ways to choose t_j . Let $k \geq z$ be the number of values of j for which $\lambda_j = -1$. Then, the other $\ell - k$ values of λ_j are all non-negative and must sum to at most $k - z$. With this, we can bound above the number of choices for the t_j and λ_j by

$$\begin{aligned}
& \binom{n}{\ell} \sum_{k=z}^{\ell} \binom{\ell}{k} \binom{\ell-z}{k-z} \\
(3.4) \quad & \leq \left(\frac{e \cdot n}{\ell} \right)^\ell 2^{\ell-z} \sum_{k=z}^{\ell} \binom{\ell}{k} \leq \left(\frac{e \cdot n}{\ell} \right)^\ell 2^{2\ell-z}.
\end{aligned}$$

Once the λ_j are fixed, the number of choices for the w_j is at most

$$(3.5) \quad \prod_{j=1}^{\ell} \Delta^{\lambda_j+1} \leq \Delta^{\ell-z}.$$

Plugging the estimates in (3.4) and (3.5) into (3.3), we have

$$\begin{aligned}
& \Pr \left[X_v^{(n)} \geq z \right] \\
& \leq \sum_{\ell \geq z} \left(\frac{e \cdot n}{\ell} \right)^\ell \cdot \frac{2^{2\ell-z} \Delta^{\ell-z}}{n^\ell} \\
& \leq \sum_{\ell \geq z} \left(\frac{4e \cdot \Delta}{\ell} \right)^\ell \leq \sum_{\ell \geq z} \left(\frac{4e \cdot \Delta}{z} \right)^\ell \leq 2 \left(\frac{4e \cdot \Delta}{z} \right)^z,
\end{aligned}$$

where the last inequality uses the fact that $z \geq 8e \cdot \Delta$. \blacksquare

Throughout the section, we fix an arbitrary vertex v and bound the number of balls allocated to the

vertices of B_v^r for all $r_0 \leq r \leq R$, where

$$\begin{aligned}
r_0 & := \min \left\{ r : |B_v^r| \geq \log^{10} n \right\}, \quad \text{and} \\
(3.6) \quad R & := \min \left\{ r : |B_v^r| \geq \frac{n}{4\Delta} \right\}.
\end{aligned}$$

We will consider the balls $(B_v^r)_{r \geq r_0}$, and will bound the number of vertices in B_v^r with load at least ℓ for all integers $\ell \in [\ell_0, \ell_1]$, where

$$(3.7) \quad \ell_0 := 8e \cdot \Delta^2 \quad \text{and} \quad \ell_1 := \frac{\log n}{4 \log(2\Delta)}.$$

Then, for all r and ℓ , define

$$\Lambda_{r,\ell} = \left\{ u \in B_v^r : X_u^{(n)} \geq \ell \right\}.$$

In order to control $\Lambda_{r,\ell}$, we will need to estimate the probability that the n th ball changes the load of vertices in B_v^r . For this last value, we need to control the load of the vertices after $n-1$ balls have arrived. Then, we define

$$\Lambda_{r,\ell}^* := \left\{ u \in B_v^r : X_u^{(n-1)} \geq \ell \right\}.$$

Note that $\Lambda_{r,\ell}^* \subseteq \Lambda_{r,\ell}$. By Lemma 3.2, we have that, for $\ell \geq \ell_0$ and any $r \geq 1$,

$$\begin{aligned}
\mathbf{E} \left[\left| \bigcup_{k \geq 0} \Lambda_{r+k, \ell+k} \right| \right] & \leq \sum_{k \geq 0} |B_v^{r+k}| 2 \left(\frac{4e \cdot \Delta}{\ell+k} \right)^{\ell+k} \\
& \leq |B_v^r| \sum_{k \geq 0} \Delta^k 2 \left(\frac{1}{2\Delta} \right)^{\ell+k} \\
(3.8) \quad & \leq 4(2\Delta)^{-\ell} |B_v^r|.
\end{aligned}$$

Next define

$$\Lambda_\ell^* := \left\{ u \in V : X_u^{(n-1)} \geq \ell \right\},$$

hence, $\Lambda_\ell^* = \bigcup_{r=1}^{\infty} \Lambda_{r,\ell}^*$. Define the event

$$(3.9) \quad L_R^* = \bigcap_{\ell=\ell_0}^{6\ell_1} \left\{ |\Lambda_\ell^*| \leq \frac{n}{4\Delta} \cdot (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell} \right\}.$$

From now on let \mathcal{F}_i be the σ -algebra induced by the configuration obtained after i balls are placed. More formally, if U_1, U_2, \dots, U_i are the birthplaces of the first i balls and, for each $v \in V$ and $j = 1, 2, \dots, i$, we define $\xi_v^{(j)}$ to be an independent uniformly random permutation of the neighbors of v , where $\{\xi_v^{(j)}\}_{v \in V}$ are the permutations used to break ties uniformly at random for ball j , then \mathcal{F}_i is the σ -algebra induced by U_1, \dots, U_i and $\{\xi_v^{(j)}\}_{v \in V, 1 \leq j \leq i}$.

LEMMA 3.3. Let $v \in V$ be fixed, and let R and L_R^* be as defined in (3.6) and (3.9), respectively. Then, there exist n_0 so that, for all $n \geq n_0$, we have

$$\Pr[L_R^*] \geq 1 - 2n^{-\log^5 n}.$$

Proof. Recall that by Lemma 2.1, there is a coupling such that with probability 1, $\bar{X}^{(n)}$ majorizes $X^{(n)}$. Hence the claim follows directly by Lemma A.4 and a union bound over all ℓ with $\ell_0 \leq \ell \leq \ell_1$. ■

For any r with $1 \leq r < R$ and any $\ell \geq \ell_0$, in order to bound the number of vertices in $\Lambda_{r,\ell}$, we will look at the probability that the n th ball affects $\Lambda_{r,\ell}$. In other words, we control the probability that $\Lambda_{r,\ell}$ is different from $\Lambda_{r,\ell}^*$. Note that it is only possible that $\Lambda_{r,\ell} \neq \Lambda_{r,\ell}^*$ if the n th ball is born at a vertex of $\Lambda_{r,\ell-1}^*$ or if it is born at a vertex of N_v^{r+k} with load at least $\ell - 1 + k$ for some $k \geq 1$; we shall bound this last set of vertices by $\Lambda_{r+k,\ell-1+k}$. We define inductively for $r < R$ that

$$L_{r,\ell_0-1}^* = L_{r+1,\ell_1}^* \quad \text{and} \quad L_{R-1,\ell_0-1}^* = L_R^*$$

and, for $\ell \geq \ell_0$,

$$(3.10) \quad L_{r,\ell}^* = L_{r,\ell-1}^* \cap \left\{ |\Lambda_{r,\ell}^*| \leq \frac{8 \cdot |B_v^r|}{(2\Delta)^\ell} + \frac{\log^7 n}{\ell} \right\}.$$

The next lemma establishes that, with high probability, the last ball cannot affect the load of a small set of vertices.

LEMMA 3.4. Let v be any fixed vertex, $r \geq 1$, and $\ell \geq \ell_0$. Then, there exists a positive constant $c = c(\Delta)$ such that

$$\begin{aligned} & \Pr \left[\bigcup_{u \in \Lambda_{r,\ell-1}^*} \{X_u^{(n-1)} \neq X_u^{(n)}\} \mid L_{r,\ell-1}^* \right] \\ & \leq \frac{c|B_v^r|(2\Delta)^{-\ell}}{n} + \frac{3 \log^7 n \log \log n}{n}. \end{aligned}$$

Proof. Note that the n th ball can only change the load of a vertex in $\Lambda_{r,\ell-1}^*$ if it is born at a vertex of B_v^r of load at least $\ell - 1$ or if it is born at a vertex $u \in N_v^{r+k}$ with $X_u^{(n-1)} \geq \ell - 1 + k$ for some $k \geq 1$. Fixing any realization for the birthplaces of the first $n - 1$ balls, and thereby fixing the sets $\Lambda_{r,\ell}^*$ for all r and ℓ , we have the following upper bound

$$\begin{aligned} & \Pr \left[\bigcup_{u \in \Lambda_{r,\ell-1}^*} \{X_u^{(n-1)} \neq X_u^{(n)}\} \mid \mathcal{F}_{n-1} \right] \\ & \leq \frac{|\Lambda_{r,\ell-1}^*|}{n} + \frac{|\bigcup_{k=1}^{\infty} \Lambda_{r+k,\ell-1+k}^*|}{n}, \end{aligned}$$

where the probability above is taken over the choice of U_n only. Note that $L_{r,\ell-1}^*$ is measurable with

respect to \mathcal{F}_{n-1} since the birthplace of the n th ball is independent of any event in \mathcal{F}_{n-1} . Then, for $\ell > \ell_0$, we obtain that

$$\begin{aligned} & \Pr \left[\bigcup_{u \in \Lambda_{r,\ell-1}^*} \{X_u^{(n-1)} \neq X_u^{(n)}\} \mid L_{r,\ell-1}^* \right] \\ & \leq \frac{8|B_v^r|(2\Delta)^{-\ell+1} + \frac{\log^7 n}{\ell-1}}{n} \\ & \quad + \sum_{k=1}^{R-r} \frac{8|B_v^{r+k}|(2\Delta)^{-\ell+1-k} + \frac{\log^7 n}{\ell-1+k}}{n} \\ & \quad + \frac{8|B_v^R|(2\Delta)^{-\ell-R+r} + \frac{\log^7 n}{\ell-1}}{n}, \end{aligned}$$

where the last term comes from L_R^* . Using the bounds $|B_v^{r+k}| \leq |B_v^r| \Delta^k$ and $|B_v^R| \leq |B_v^r| \Delta^{R-r}$, we obtain that

$$\begin{aligned} & \Pr \left[\bigcup_{u \in \Lambda_{r,\ell}^*} \{X_u^{(n-1)} \neq X_u^{(n)}\} \mid L_{r,\ell-1}^* \right] \\ & \leq \frac{8|B_v^r|(2\Delta)^{-\ell+1} + \frac{\log^7 n}{\ell-1}}{n} \\ & \quad + \sum_{k=1}^{R-r+1} \frac{8|B_v^r| \Delta^k (2\Delta)^{-\ell+1-k}}{n} + \sum_{k=0}^{R-r} \frac{\log^7 n}{(\ell-1+k)n} \\ & \leq \frac{8|B_v^r|(2\Delta)^{-\ell+1} + \frac{\log^7 n}{\ell-1}}{n} + \frac{16|B_v^r|(2\Delta)^{-\ell+1}}{n} \\ & \quad + \frac{2 \log^7 n \log \log n}{n} \\ & \leq \frac{24|B_v^r|(2\Delta)^{-\ell+1}}{n} + \frac{3 \log^7 n \log \log n}{n}. \end{aligned}$$

For $\ell = \ell_0$, we simply bound $|\Lambda_{r,\ell_0-1}^*|$ by $|B_v^r|$, which gives that

$$\begin{aligned} & \Pr \left[\bigcup_{u \in \Lambda_{r,\ell_0}^*} \{X_u^{(n-1)} \neq X_u^{(n)}\} \mid L_{r,\ell_0}^* \right] \\ & \leq \frac{|B_v^r|}{n} + \sum_{k=1}^{R-r+1} \frac{8|B_v^r| \Delta^k (2\Delta)^{-\ell_0+1-k}}{n} \\ & \quad + \sum_{k=0}^{R-r} \frac{\log^7 n}{(\ell_0-1+k)n} \\ & \leq \frac{|B_v^r|}{n} + \frac{16|B_v^r|(2\Delta)^{-\ell_0+1}}{n} + \frac{2 \log^7 n \log \log n}{n}. \end{aligned}$$

Now, for any $i \geq 1$, $r \in [r_0, R]$ and $\ell \in [\ell_0, \ell_1]$, we define $\Upsilon_{r,\ell}^i \subseteq V^i$ as the set

$$(3.11) \quad \Upsilon_{r,\ell}^i = \left\{ (u_1, \dots, u_i) \in V^i : \forall 1 \leq j_0 \leq i: \right.$$

$$\left. \Pr \left[L_{r,\ell}^* \mid \bigcap_{j=1}^{j_0} \{U_j = u_j\} \right] \geq 1 - \frac{1}{n^2} \right\}.$$

Intuitively, for any given r, ℓ and i , the set $\Upsilon_{r,\ell}^i$ contains the *good* birthplace for the first i balls so that the event $L_{r,\ell}^*$ is likely to occur, conditioning on any prefix of the birthplaces.

While Lemma 3.4 considered the effect of the last ball n , the following lemma studies the effect of replacing the birthplace of ball i by a randomly chosen bin.

LEMMA 3.5. *Let i, ℓ, r and v be fixed. Let $f: V^n \rightarrow \mathbb{Z}$ be an increasing function that depends only on $(X_{v'}^{(n)})_{v' \in \Lambda_{r,\ell}}$ and is 1-Lipschitz. Let W_0, W_1, \dots, W_{n-i} and \widehat{W}_0 be i.i.d. random variables chosen uniformly from V . Then, if $(u_1, u_2, \dots, u_{i-1}) \in \Upsilon_{r,\ell}^{i-1}$, we have*

$$\begin{aligned} & \Pr \left[\begin{array}{l} f(u_1, \dots, u_{i-1}, W_0, W_1, \dots, W_{n-i}) \\ \neq f(u_1, \dots, u_{i-1}, \widehat{W}_0, W_1, \dots, W_{n-i}) \end{array} \right] \\ & \leq \frac{6c|B_v^r|(2\Delta)^{-\ell+1}}{n} + \frac{18 \log^7 n \log \log n}{n} + \frac{5}{n^2}, \end{aligned}$$

where c is the constant from Lemma 3.4.

Proof. In this proof all the probabilities are taken conditional on $U_j = u_j$ for all $j = 1, 2, \dots, i-1$, but we will omit this dependence from the notation. The idea is to relate the probability above to

$$\Pr \left[\begin{array}{l} f(u_1, \dots, u_{i-1}, W_1, \dots, W_{n-i}, W_0) \\ \neq f(u_1, \dots, u_{i-1}, W_1, \dots, W_{n-i}, \widehat{W}_0) \end{array} \right],$$

which corresponds to changing the n th ball instead of the i th ball; this will allow us to apply Lemma 3.4.

Consider the three events below:

$$\begin{aligned} E_1 &= \left\{ \begin{array}{l} f(u_1, \dots, u_{i-1}, W_0, W_1, \dots, W_{n-i}) \\ \neq f(u_1, \dots, u_{i-1}, W_1, \dots, W_{n-i}, W_0) \end{array} \right\} \\ \widehat{E}_1 &= \left\{ \begin{array}{l} f(u_1, \dots, u_{i-1}, \widehat{W}_0, W_1, \dots, W_{n-i}) \\ \neq f(u_1, \dots, u_{i-1}, W_1, \dots, W_{n-i}, \widehat{W}_0) \end{array} \right\} \\ E_{\text{end}} &= \left\{ \begin{array}{l} f(u_1, \dots, u_{i-1}, W_1, \dots, W_{n-i}, W_0) \\ \neq f(u_1, \dots, u_{i-1}, W_1, \dots, W_{n-i}, \widehat{W}_0) \end{array} \right\}. \end{aligned}$$

Clearly, $f(u_1, \dots, u_{i-1}, W_0, W_1, \dots, W_{n-i})$ and $f(u_1, \dots, u_{i-1}, \widehat{W}_0, W_1, \dots, W_{n-i})$ can only be different if at least one of E_1, \widehat{E}_1 or E_{end} happens. Therefore, we can write

$$\begin{aligned} & \Pr \left[\begin{array}{l} f(u_1, \dots, u_{i-1}, W_0, W_1, \dots, W_{n-i}) \\ \neq f(u_1, \dots, u_{i-1}, \widehat{W}_0, W_1, \dots, W_{n-i}) \end{array} \right] \\ (3.12) \quad & \leq \Pr[E_1] + \Pr[\widehat{E}_1] + \Pr[E_{\text{end}}]. \end{aligned}$$

We start with the term $\Pr[E_{\text{end}}]$. Let I_{end} be the event that $L_{r,\ell-1}^*$ happens given that the birthplaces of the first $n-1$ balls are according to the sequence $(u_1, \dots, u_{i-1}, W_1, \dots, W_{n-i})$. If $Y_u^{(n)}$ is the load of vertex u when n balls are added with birthplaces $u_1, u_2, \dots, u_{i-1}, W_1, W_2, \dots, W_{n-i}, W_0$

and $\widehat{Y}_u^{(n)}$ is the load of vertex u when n balls are added with birthplaces $u_1, u_2, \dots, u_{i-1}, W_1, W_2, \dots, W_{n-i}, \widehat{W}_0$, we have that

(3.13)

$$\begin{aligned} & \Pr[E_{\text{end}}] \leq \Pr[I_{\text{end}}^c] \\ & + \Pr \left[\bigcup_{u \in B_v^r} \left(\{Y_u^{(n-1)} \neq Y_u^{(n)}\} \cap \{Y_u^{(n)} \geq \ell\} \right) \cap I_{\text{end}} \right] \\ & + \Pr \left[\bigcup_{u \in B_v^r} \left(\{\widehat{Y}_u^{(n-1)} \neq \widehat{Y}_u^{(n)}\} \cap \{\widehat{Y}_u^{(n)} \geq \ell\} \right) \cap I_{\text{end}} \right], \end{aligned}$$

where $\Pr[I_{\text{end}}^c] \leq n^{-2}$ by (3.11) since $(u_1, \dots, u_{i-1}) \in \Upsilon_{r,\ell}^{i-1}$ and W_1, \dots, W_{n-i} are i.i.d. uniform samples from V . The other two terms in (3.13) can be bounded by Lemma 3.4.

Now it remains to bound $\Pr[E_1]$, since by symmetry we have $\Pr[E_1] = \Pr[\widehat{E}_1]$. In order to bound $\Pr[E_1]$, we consider all cyclic permutations of $(u_1, \dots, u_{i-1}, W_0, W_1, \dots, W_{n-i})$. More specifically, we first compare

$$f(u_1, \dots, u_{i-1}, W_0, W_1, \dots, W_{n-i}) \quad \text{with}$$

$$f(u_1, \dots, u_{i-1}, W_1, W_2, \dots, W_{n-i}, W_0),$$

then we compare

$$f(u_1, \dots, u_{i-1}, W_1, W_2, \dots, W_{n-i}, W_0) \quad \text{with}$$

$$f(u_1, \dots, u_{i-1}, W_2, W_3, \dots, W_{n-i}, W_0, W_1),$$

and so on and so forth until we compare

$$f(u_1, \dots, u_{i-1}, W_{n-i}, W_0, W_1, \dots, W_{n-i-1}) \quad \text{with}$$

$$f(u_1, \dots, u_{i-1}, W_0, W_1, \dots, W_{n-i}).$$

In order to do this, we define a graph H whose vertex set is $\{0, 1, \dots, n-i\} \times V^{n-i+1}$; so the vertices of H have the form $(j, z_0, z_1, \dots, z_{n-i})$. We let each vertex of H have exactly one outgoing edge and one incoming edge by having a directed edge from each vertex $(j, z_0, z_1, \dots, z_{n-i})$ to $(j', z_1, z_2, \dots, z_{n-i}, z_0)$ where $j' = j+1 \pmod{n-i+1}$. With this, the sequence

$$\begin{aligned} & (0, z_0, z_1, \dots, z_{n-i}), (1, z_1, \dots, z_{n-i}, z_0), \dots, \\ & (j, z_j, \dots, z_{n-i}, z_0, z_1, \dots, z_{j-1}), \dots, \\ & (n-i, z_{n-i}, z_0, \dots, z_{n-i-1}), (0, z_0, z_1, \dots, z_{n-i}) \end{aligned}$$

forms a directed cycle in H . Since each vertex has in-degree and out-degree equal to one, we have that each connected component of H is a directed

cycle with $n - i + 1$ vertices. Now, for any vertex $(j, z_0, z_1, \dots, z_{n-i})$ of H , we say that the edge from $(j, z_0, z_1, \dots, z_{n-i})$ to $(j + 1, z_1, z_2, \dots, z_{n-i}, z_0)$ is f -increasing if

$$\begin{aligned} & f(u_1, u_2, \dots, u_{i-1}, z_0, z_1, \dots, z_{n-i}) \\ & < f(u_1, u_2, \dots, u_{i-1}, z_1, z_2, \dots, z_{n-i}, z_0) . \end{aligned}$$

Similarly we say that the edge is f -decreasing if the opposite inequality holds:

$$\begin{aligned} & f(u_1, u_2, \dots, u_{i-1}, z_0, z_1, \dots, z_{n-i}) \\ & > f(u_1, u_2, \dots, u_{i-1}, z_1, z_2, \dots, z_{n-i}, z_0) . \end{aligned}$$

Since f is 1-Lipschitz and integer-valued, for any directed cycle, the number of f -increasing edges is the same as the number of f -decreasing edges, and therefore the number of f -increasing edges in H is the same as the number of f -decreasing edges.

Note that choosing a sequence $(W_0, W_1, \dots, W_{n-i})$ uniformly at random from V^{n-i+1} and a number k uniformly at random from $\{0, 1, \dots, n - i\}$ gives a uniformly random vertex of H . Similarly, if we choose a vertex $(j, z_0, z_1, \dots, z_{n-i})$ uniformly at random from the vertex set of H , then $(z_0, z_1, \dots, z_{n-i})$ is a uniform sample from V^{n-i+1} . Since each vertex of H has a unique outgoing edge, one can choose a uniformly random edge of H by choosing the outgoing edge of a uniformly random vertex of H .

We fix a uniformly random vertex $(j, z_0, z_1, \dots, z_{n-i})$ of H and let e be its unique outgoing edge; i.e., e is the edge from $(j, z_0, z_1, \dots, z_{n-i})$ to $(j + 1, z_1, \dots, z_{n-i}, z_0)$. Therefore,

$$\begin{aligned} \Pr[E_1] & \leq \Pr[e \text{ is } f\text{-increasing}] \\ & \quad + \Pr[e \text{ is } f\text{-decreasing}] \\ (3.14) \quad & = 2 \Pr[e \text{ is } f\text{-increasing}] . \end{aligned}$$

Let $Y_u^{(k)}$ be the load of vertex u after k balls are added with birthplaces given by the sequence $(u_1, \dots, u_{i-1}, z_0, \dots, z_{n-i})$, and let $\widehat{Y}_u^{(k)}$ be the load of vertex u after k balls are added with birthplaces given by the sequence $(u_1, \dots, u_{i-1}, z_1, \dots, z_{n-i}, z_0)$.

Note that, by Lemma 2.6, removing a ball cannot increase the load of any vertex; this gives that

$$Y_u^{(n)} \geq \widehat{Y}_u^{(n-1)} \quad \text{for all } u \in V .$$

Therefore, e can only be f -increasing if $\widehat{Y}_u^{(n)} \neq \widehat{Y}_u^{(n-1)}$ and $\widehat{Y}_u^{(n)} \geq \ell$ for some $u \in B_v^r$. Let I_e be the indicator that $L_{r,\ell}^*$ holds given that the birthplaces of the balls are according to the sequence $(u_1, \dots, u_{i-1}, z_1, \dots, z_{n-i}, z_0)$. Hence, using (3.14),

we have that

$$\begin{aligned} & \Pr[E_1] \\ & \leq 2 \Pr \left[\bigcup_{u \in B_v^r} \left(\{\widehat{Y}_u^{(n-1)} \neq \widehat{Y}_u^{(n)}\} \cap \{\widehat{Y}_u^{(n)} \geq \ell\} \right) \mid I_e \right] \\ & \quad + 2 \Pr[I_e^c] . \end{aligned}$$

By (3.11), $\Pr[I_e^c] \leq n^{-2}$ and the other term can be bounded by Lemma 3.4. Then, we put this and (3.13) into (3.12) to complete the proof. \blacksquare

Now we will use an inductive argument to bound the probability that $L_{r,\ell}^*$ happens for all r and ℓ .

LEMMA 3.6. *There exists a positive constant c' such that, for any $v \in V$, $r < R$ and any integer $\ell \in [\ell_0, \ell_1]$, we have*

$$\begin{aligned} & \Pr \left[|\Lambda_{r,\ell}| > 8|B_v^r| (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell} \right] \\ & \leq e^{-\log^4 n} + n^2 \cdot \Pr \left[(L_{r,\ell-1}^*)^c \right] \end{aligned}$$

for all large enough n .

Proof. Recall that, by Lemma 3.2, we have

$$\mathbf{E}[|\Lambda_{r,\ell}|] \leq 2|B_v^r| \left(\frac{4e \cdot \Delta}{\ell} \right)^\ell \leq 4|B_v^r| (2\Delta)^{-\ell} .$$

Define the stopping time

$$\tau = n \wedge \min \left\{ j : (U_1, U_2, \dots, U_j) \notin \Upsilon_{r,\ell}^j \right\} ,$$

where for any two numbers $a \wedge b := \min\{a, b\}$. Let $X_i = \mathbf{E}[|\Lambda_{r,\ell}| \mid \mathcal{F}_i]$ and $Z_i = X_{i \wedge \tau}$. Note that Z_i is a martingale since τ is a stopping time. Moreover, τ is a *bounded* stopping time and, using the optional stopping theorem, we have that $\mathbf{E}[Z_\tau] = \mathbf{E}[Z_0] = \mathbf{E}[X_0]$. We want to bound the conditional variance of $Z_i - Z_{i-1}$ uniformly over all i , which is given by

$$\begin{aligned} & \mathbf{Var}_{U_i} [Z_i - Z_{i-1} \mid \mathcal{F}_{i-1}] \\ & = \mathbf{E}_{U_i} \left[(Z_i - Z_{i-1})^2 \mid \mathcal{F}_{i-1} \right] \\ & \quad - (\mathbf{E}_{U_i} [Z_i - Z_{i-1} \mid \mathcal{F}_{i-1}])^2 \\ & = \mathbf{E}_{U_i} \left[(Z_i - Z_{i-1})^2 \mid \mathcal{F}_{i-1} \right] , \end{aligned}$$

where the variance and expectation are taken over the choice of U_i only. Now, we write

$$\begin{aligned} & \mathbf{E}_{U_i} \left[(Z_i - Z_{i-1})^2 \mid \mathcal{F}_{i-1} \right] \\ & = \mathbf{E}_{U_i} \left[(Z_i - Z_{i-1})^2 \mathbf{1}(\tau \geq i) \mid \mathcal{F}_{i-1} \right] \\ & \quad + \mathbf{E}_{U_i} \left[(Z_i - Z_{i-1})^2 \mathbf{1}(\tau < i) \mid \mathcal{F}_{i-1} \right] \\ & = \mathbf{E}_{U_i} \left[(Z_i - Z_{i-1})^2 \mathbf{1}(\tau \geq i) \mid \mathcal{F}_{i-1} \right] , \end{aligned}$$

since $Z_i = Z_{i-1}$ whenever $\tau < i$. We want to bound

$$\mathbf{E}_{U_i} \left[(Z_i - Z_{i-1})^2 \left| \bigcap_{j=1}^{i-1} \{U_j = u_j\} \right. \right],$$

uniformly over all $(u_1, u_2, \dots, u_{i-1}) \in \Upsilon_{r,\ell}^{i-1}$. Let ζ_u be the value of Z_i when $U_i = u$ and let $\bar{\zeta} = \frac{1}{n} \sum_{u \in V} \zeta_u$. Then we have

$$\begin{aligned} \mathbf{E}_{U_i} & \left[(Z_i - Z_{i-1})^2 \left| \bigcap_{j=1}^{i-1} \{U_j = u_j\} \right. \right] \\ & = \frac{1}{n} \sum_{u \in V} (\zeta_u - \bar{\zeta})^2. \end{aligned}$$

Since X_i is 1-Lipschitz we have $|\zeta_u - \zeta_{u'}| \leq 1$ for all $u, u' \in V$. With this, we can write

$$\begin{aligned} \frac{1}{n} \sum_{u \in V} (\zeta_u - \bar{\zeta})^2 & \leq \frac{1}{n} \sum_{u \in V} |\zeta_u - \bar{\zeta}| \\ & = \frac{1}{n} \sum_{u \in V} \left| \sum_{u' \in V} \frac{1}{n} (\zeta_u - \zeta_{u'}) \right| \\ & \leq \frac{1}{n^2} \sum_{u \in V} \sum_{u' \in V} |\zeta_u - \zeta_{u'}|. \end{aligned}$$

Now, using Lemma 3.5, we have that

$$\begin{aligned} & \frac{1}{n^2} \sum_{u \in V} \sum_{u' \in V} |\zeta_u - \zeta_{u'}| \\ & \leq \frac{6c|B_v^r|(2\Delta)^{-\ell+1}}{n} + \frac{18 \log^7 n \log \log n}{n} + \frac{5}{n^2}, \end{aligned}$$

which gives that

$$\begin{aligned} & \mathbf{E}_{U_i} [(Z_i - Z_{i-1})^2 | \mathcal{F}_{i-1}] \\ & \leq \frac{6c|B_v^r|(2\Delta)^{-\ell+1}}{n} + \frac{18 \log^7 n \log \log n}{n} + \frac{5}{n^2}, \end{aligned}$$

uniformly over i . Now, note that, for any $\lambda > 0$,

$$\begin{aligned} & \mathbf{Pr}[|X_n - X_0| > \lambda] \\ & \leq \mathbf{Pr}[\{|X_n - X_0| > \lambda\} \cap \{\tau \geq n\}] + \mathbf{Pr}[\tau < n] \\ & = \mathbf{Pr}[\{|Z_n - Z_0| > \lambda\} \cap \{\tau \geq n\}] + \mathbf{Pr}[\tau < n]. \end{aligned}$$

Also, we have that for any $1 \leq i \leq n$,

$$\begin{aligned} & \mathbf{Pr}[(L_{r,\ell-1}^*)^c] \\ & = \sum_{j_0=1}^i \mathbf{Pr} \left[\begin{aligned} & (L_{r,\ell-1})^c \cap \{(U_1, \dots, U_{j_0}) \notin \Upsilon_{r,\ell-1}^{j_0}\} \\ & \cap (\bigcap_{k=1}^{j_0-1} \{(U_1, \dots, U_k) \in \Upsilon_{r,\ell-1}^k\}) \end{aligned} \right] + \\ & \mathbf{Pr} \left[(L_{r,\ell-1}^*)^c \cap \left(\bigcap_{k=1}^i \{(U_1, \dots, U_k) \in \Upsilon_{r,\ell-1}^k\} \right) \right] \\ & \geq \mathbf{Pr}[(U_1, \dots, U_i) \notin \Upsilon_{r,\ell-1}^i] \cdot \frac{1}{n^2}. \end{aligned}$$

This gives that

$$\begin{aligned} \mathbf{Pr}[\tau < n] & \leq \mathbf{Pr}[(U_1, \dots, U_{n-1}) \notin \Upsilon_{r,\ell-1}^{n-1}] \\ & \leq n^2 \cdot \mathbf{Pr}[(L_{r,\ell-1}^*)^c]. \end{aligned}$$

Then, applying the version of Azuma's inequality from Lemma A.2, we have that

$$\begin{aligned} & \mathbf{Pr} \left[|\Lambda_{r,\ell}| > 8|B_v^r|(2\Delta)^{-\ell} + \frac{\log^7 n}{\ell} \right] \\ & \leq \mathbf{Pr} \left[|Z_n - Z_0| > 4|B_v^r|(2\Delta)^{-\ell} + \frac{\log^7 n}{\ell} \right] \\ & \quad + n^2 \cdot \mathbf{Pr}[(L_{r,\ell-1}^*)^c] \\ & \leq \exp \left(-\frac{(4|B_v^r|(2\Delta)^{-\ell} + \frac{\log^7 n}{\ell})^2}{14c|B_v^r|(2\Delta)^{-\ell} + 37 \log^7 n \log \log n} \right) \\ & \quad + n^2 \cdot \mathbf{Pr}[(L_{r,\ell-1}^*)^c]. \end{aligned}$$

If $|B_v^r|(2\Delta)^{-\ell} \geq \log^7 n \log \log n$, the exponential term above is at most

$$\begin{aligned} & \exp \left(-\frac{(4|B_v^r|(2\Delta)^{-\ell})^2}{(14c+37)|B_v^r|(2\Delta)^{-\ell}} \right) \\ & \leq \exp \left(-\frac{16|B_v^r|(2\Delta)^{-\ell}}{14c+37} \right) \leq \exp(-\log^4 n); \end{aligned}$$

otherwise we bound above the exponential term by

$$\begin{aligned} & \exp \left(-\frac{(\frac{\log^7 n}{\ell})^2}{(14c+37) \log^7 n \log \log n} \right) \\ & \leq \exp \left(-\frac{\log^7 n}{(14c+37)\ell^2 \log \log n} \right) \leq \exp(-\log^4 n), \end{aligned}$$

where the last inequality holds for all large enough n since $\ell \leq \ell_1 = O(\log n)$. \blacksquare

Proof. [of Theorem 3.1] First note that the lower bound on the maximum load is established by Theorem 1.6. So we now prove the upper bound. Let v be an arbitrary vertex of V . We start the proof by showing that there exist positive constants C and c such that, for all large enough n , it holds that

$$(3.15) \quad \mathbf{Pr} \left[\sum_{u \in B_v^{r_0}} X_u^{(n)} > C|B_v^{r_0}| \right] \leq e^{-c \log^3 n}.$$

Then, it follows by Lemma 2.3 that

$$X_v^{(n)} \leq C + \sum_{i=0}^{r_0} \frac{|N_v^i|}{|B_v^{r_0}|}.$$

Note that $\sum_{i=0}^r \frac{|N_v^i|}{|B_v^{r_0}|}$ is the average distance between v and a vertex in $B_v^{r_0}$, which is at most r_0 . Therefore,

we obtain $X_v^{(n)} \leq C + r_0$. By the definition of r_0 , we have $|B_v^{r_0}| \leq \Delta \log^{10} n$. Combing this with the exponential growth property of G yields

$$r_0 \leq \frac{1}{\phi} \log |B_v^{r_0}| = \mathcal{O}(\log \log n),$$

where ϕ is the parameter defined by (3.1).

It remains to establish (3.15). First note that

$$\sum_{u \in B_v^{r_0}} X_u^{(n)} = \sum_{\ell=1}^{\infty} |\Lambda_{r_0, \ell}| \leq |B_v^{r_0}| \ell_0 + \sum_{\ell=\ell_0+1}^{\infty} |\Lambda_{r_0, \ell}|.$$

Therefore, we have that

$$\begin{aligned} & \Pr \left[\sum_{u \in B_v^{r_0}} X_u^{(n)} > C |B_v^{r_0}| \right] \\ & \leq \Pr \left[\sum_{\ell=\ell_0}^{\infty} |\Lambda_{r_0, \ell}| > (C - \ell_0) |B_v^{r_0}| \right]. \end{aligned}$$

Now, if $|\Lambda_{r_0, \ell}| \leq 8 |B_v^{r_0}| (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell}$ for all $\ell = \ell_0, \ell_0 + 1, \dots, \ell_1$, then

$$\sum_{\ell=\ell_0}^{\infty} |\Lambda_{r_0, \ell}| \leq (C - \ell_0) |B_v^{r_0}|$$

for some large enough C since $|B_v^{r_0}| \geq \log^{10} n$. For any given ℓ , using Lemma 3.6, we have

$$\begin{aligned} & \Pr \left[|\Lambda_{r_0, \ell}| > 8 |B_v^{r_0}| (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell} \right] \\ (3.16) \quad & \leq \exp(-\log^4 n) + n^2 \cdot \Pr \left[(L_{r_0, \ell-1}^*)^c \right]. \end{aligned}$$

By definition of $L_{r, \ell}^*$ (cf. (3.10)), we have that

$$\begin{aligned} & \Pr \left[(L_{r_0, \ell-1}^*)^c \right] \\ & \leq \Pr \left[|\Lambda_{r_0, \ell-1}^*| > 8 |B_v^{r_0}| (2\Delta)^{-\ell+1} + \frac{\log^7 n}{\ell-1} \right] \\ & \quad + \Pr \left[(L_{r_0, \ell-2}^*)^c \right] \\ & \leq \Pr \left[|\Lambda_{r_0, \ell-1}| > 8 |B_v^{r_0}| (2\Delta)^{-\ell+1} + \frac{\log^7 n}{\ell-1} \right] \\ & \quad + \Pr \left[(L_{r_0, \ell-2}^*)^c \right] \\ & \leq \exp(-\log^4 n) + (n^2 + 1) \Pr \left[(L_{r_0, \ell-2}^*)^c \right], \end{aligned}$$

where the second inequality follows since $\Lambda_{r_0, \ell-1}^* \subseteq \Lambda_{r_0, \ell-1}$ and last inequality follows from Lemma 3.6.

Applying this into (3.16), we have

$$\begin{aligned} & \Pr \left[|\Lambda_{r_0, \ell}| > 8 |B_v^{r_0}| (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell} \right] \\ & \leq \exp(-\log^4 n) + (n^2 + 1) \exp(-\log^4 n) \\ & \quad + (n^2 + 1)^2 \Pr \left[(L_{r_0, \ell-2}^*)^c \right] \\ & \leq \sum_{k=0}^{\ell-\ell_0} (n^2 + 1)^k \exp(-\log^4 n) \\ & \quad + (n^2 + 1)^{\ell-\ell_0+1} \Pr \left[(L_{r_0+1, \ell_1}^*)^c \right]. \end{aligned}$$

Using the same argument, we obtain for any r that

$$\begin{aligned} & \Pr \left[(L_{r, \ell_1}^*)^c \right] \\ & \leq \sum_{k=0}^{\ell_1-\ell_0} (n^2 + 1)^k \exp(-\log^4 n) \\ & \quad + (n^2 + 1)^{\ell_1-\ell_0+1} \Pr \left[(L_{r+1, \ell_1}^*)^c \right] \\ & \leq \sum_{j=0}^{R-r-1} \sum_{k=0}^{\ell_1-\ell_0} (n^2 + 1)^{\ell_1 j + k} \exp(-\log^4 n) \\ & \quad + (n^2 + 1)^{\ell_1(R-r)} \Pr \left[(L_R^*)^c \right]. \end{aligned}$$

Using Lemma 3.3,

$$\Pr \left[(L_R^*)^c \right] \leq 2n^{-\log^5 n},$$

and plugging this into (3.16), and using the union bound over ℓ , we obtain that

$$\begin{aligned} & \Pr \left[\bigcup_{\ell=\ell_0}^{\ell_1} \left\{ |\Lambda_{r_0, \ell}| > 8 |B_v^{r_0}| (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell} \right\} \right] \\ & \leq \exp(-c \log^3 n) \end{aligned}$$

for some positive constant c , which establishes (3.15). \blacksquare

4 Grid graphs

In this section, we analyze the maximum load of the local allocation process on any d -dimensional grid, where d is an arbitrary constant. We show that the maximum load is $\Theta\left(\left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}\right)$. Interestingly, the analysis on the grid turns out to be much easier than the analysis on expander graphs. Roughly speaking, this is due to the fact that on grid graphs the number of paths the local search could take is much smaller than on the expansion. Formally, we define the d -dimensional grid by the vertex set

$$V = \left\{ u : u = (u_1, \dots, u_d), u_i = 0, \dots, n^{1/d} - 1 \right\}$$

and edge set $E = \{\{u, v\} : \text{dist}(u, v) = 1\}$, where

$$\text{dist}(u, v) = \sum_{i=1}^d \text{dist}(u_i, v_i)$$

and

$$\text{dist}(u_i, v_i) = \min \left\{ |u_i - v_i|, n^{1/d} - |u_i - v_i| \right\}.$$

Proof. [of Theorem 1.3] We start with the *upper bound*. Within this proof, we use the following notation:

$$\tilde{B}_u^r := \{v \in V : \text{dist}(u_i, v_i) \leq r, \forall i = 1, \dots, d\}.$$

Roughly speaking, \tilde{B}_u^r can be seen as the “ ℓ_∞ -version” of the set B_u^r used in Section 3. Note that for any $r \leq n^{1/d}/2 - 1/2$, $|\tilde{B}_u^r| = (2r + 1)^d$. We first define an event that gives an upper bound for the number of balls born in \tilde{B}_u^r for various u and r :

$$\mathcal{E} := \bigcap_{u \in V} \bigcap_{r = (4d)^d \left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}}^{n^{1/d}/2} \left\{ \sum_{v \in \tilde{B}_u^r} Z_v^{(n)} \leq \rho(r) \right\},$$

where $\rho(r) := 4e \cdot (d + 1) \left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}} \cdot (3r)^d$ and $Z_v^{(n)} = \sum_{i=1}^n \mathbf{1}(U_i = v)$ is the number of balls born on v during the first n rounds.

To prove that \mathcal{E} holds with high probability, fix any vertex $u \in V$ and $r \geq (4d)^d \left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}$. We have that

$$\begin{aligned} & \Pr \left[\sum_{v \in \tilde{B}_u^r} Z_v^{(n)} \geq \rho(r) \right] \\ & \leq \binom{n}{\rho(r)} \left(\frac{|\tilde{B}_u^r|}{n}\right)^{\rho(r)} \leq \left(\frac{e \cdot n}{\rho(r)}\right)^{\rho(r)} \left(\frac{(2r + 1)^d}{n}\right)^{\rho(r)} \\ & \leq \left(\frac{e \cdot (3r)^d}{\rho(r)}\right)^{\rho(r)} \\ & = \left(\frac{e}{4e \cdot (d + 1) \left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}}\right)^{4e \cdot (d + 1) \left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}} \cdot (3r)^d} \\ & \leq n^{-3}. \end{aligned}$$

Taking the union bound over the n vertices and at most $n/2$ possible values for r yields $\Pr[\mathcal{E}] \geq 1 - n^{-1}$.

Assuming that \mathcal{E} occurs, we now infer the upper bound on the maximum load. Assume for the sake of contradiction that the maximum load is in the interval $[\alpha/2, \alpha]$ where α is any value larger than $16e \cdot (16d)^d (d + 1) \left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}$. Let $u \in V$ be a vertex with $X_u^{(n)} \in [\alpha/2, \alpha]$. Since the maximum load is at most α , only balls that are born in $\tilde{B}_u^{2\alpha}$ can reach \tilde{B}_u^α . Since \mathcal{E} occurs, we know for $r = 2\alpha$ that

$$(4.1) \quad \sum_{v \in \tilde{B}_u^{2\alpha}} Z_v^{(n)} \leq 4e(d + 1)(6\alpha)^d \left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}.$$

On the other hand, if u has load at least $\alpha/2$, then

$$\begin{aligned} \sum_{v \in \tilde{B}_u^{2\alpha}} Z_v^{(n)} & \geq |\tilde{B}_u^{\alpha/(4d)}| \cdot \left(\frac{\alpha}{2} - d \cdot \frac{\alpha}{4d}\right) \\ & \geq \left(\frac{\alpha}{2d}\right)^d \cdot \frac{\alpha}{4} = \frac{1}{4} \cdot \frac{\alpha}{(16d)^d} \cdot (8\alpha)^d \\ (4.2) \quad & \geq 4e \cdot (d + 1) \cdot \left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}} \cdot (8\alpha)^d, \end{aligned}$$

where the last step used our lower bound on α . The desired contradiction follows now from (4.1) and (4.2), and the proof of the upper bound is complete.

Now we proceed to establish the *lower bound*. It is a well-known fact (cf. [13, Lemma 5.12]) that with probability at least $1 - n^{-1}$, there is a vertex $u \in V$ on which at least $\frac{\log n}{\log \log n}$ balls are born. Applying Lemma 2.4 with $S = \{u\}$, $\Phi_S = \frac{\log n}{\log \log n}$ implies that the maximum load $\beta := X_{\max}^{(n)}$ satisfies

$$\beta \cdot |B_u^\beta| \geq \frac{\log n}{\log \log n}.$$

Hence, as $|B_u^\beta| \leq |\tilde{B}_u^\beta| \leq (2\beta + 1)^d$ and d is a constant, we obtain that $\beta = \Omega\left(\left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}\right)$. \blacksquare

5 Dense graphs

In this section, we analyze *dense* graphs which we define as graphs where the minimum degree is $\Omega(\log n)$ and the ratio between the maximum and minimum degrees is constant. This includes, for instance, the $\log n$ -dimensional hypercube and Erdős-Rényi random graphs with average degree $(1 + \varepsilon) \log n$, for any $\varepsilon > 0$. The key idea of the analysis is that as long as less than α balls are allocated, where $\alpha < n$, every vertex has a constant fraction of neighbors which have received no ball, and hence, the maximum load is bounded by 1. Then Lemma 2.7 implies that after n balls are allocated, the maximum load is at most n/α . To make the analysis work, we need to assume that ties are broken uniformly at random; i.e., whenever a ball has more than one vertex to be forwarded to, the vertex is chosen independently and uniformly at random among the set of possible vertices.

Proof. [of Theorem 1.4] We divide the process of allocating the n balls into different phases, where each phase allocates a batch of consecutive $\alpha \leq n$ balls (hence the number of phases is $\lceil n/\alpha \rceil$). Then, by subadditivity (cf. Lemma 2.7), the maximum load at the end is at most $\lceil n/\alpha \rceil$ times the maximum load of a single phase.

Let G be an almost regular graph with minimum degree $\delta = c \cdot \log n$ and maximum degree $\Delta \leq C \cdot \delta$,

where $c > 0$ is any value bounded below by a constant and $C \geq 1$ is a constant. For any load assignment of the vertices $(x_u)_{u \in V}$, we define the exponential potential as

$$\Phi((x_u)_{u \in V}) := \sum_{u \in V} \exp \left(\sigma \cdot \sum_{v \in N_u} x_v \right),$$

where $\sigma := \max\{4 \log(n)/\delta, 1\} = \mathcal{O}(1)$. We also define for any $1 \leq t \leq n$,

$$\Phi^{(t)} := \Phi((X_u^{(t)})_{u \in V}) = \sum_{u \in V} \exp \left(\sigma \cdot \sum_{v \in N_u} X_v^{(t)} \right),$$

hence $\Phi^{(0)} = n$. Our goal is to bound the expected multiplicative increase in $\Phi^{(t+1)}$ compared to $\Phi^{(t)}$. In order to do that, we will actually also assume that $\Phi^{(t)}$ is small.

Specifically, assume that $(x_u)_{u \in V}$ be any vector in $(\mathbb{N} \cup \{0\})^n$ such that $\Phi((x_u)_{u \in V}) \leq n \cdot e^{\delta/2}$ and suppose that the load vector at the end of round t is $(x_u)_{u \in V}$, i.e., $X^{(t)} = x$. Then this implies for every vertex $u \in V$ that

$$\exp \left(\sigma \cdot \sum_{v \in N_u} X_v^{(t)} \right) \leq n \cdot e^{\delta/2} = e^{\log n + \delta/2},$$

and consequently, $\sum_{v \in N_u} X_v^{(t)} \leq (1/\sigma) \cdot (\log n + \delta/2) \leq \delta/4 + \delta/2 = (3/4)\delta$. Hence, there are at least $\deg(u) - (3/4)\delta \geq (1/4)\delta$ neighbors of u which have no ball, where $\deg(u)$ is the degree of vertex u . In particular, this implies that the next ball $t+1$ will be allocated either on its birthplace or at a direct neighbor. Therefore, with $Z_u^{(t)} := \sum_{v \in N_u} X_v^{(t)}$,

$$\begin{aligned} & \Pr \left[Z_u^{(t+1)} = Z_u^{(t)} + 1 \mid X^{(t)} = x \right] \\ & \leq \sum_{\substack{v \in N_u \\ w \in N_v \cup \{v\}}} \Pr \left[\text{ball } t+1 \text{ born at } w \text{ and allocated on } v \mid X^{(t)} = x \right] \\ (5.1) \quad & \leq \sum_{v \in N_u} \left(\sum_{w \in N_v} \left(\frac{1}{n} \cdot \frac{4}{\delta} \right) + \frac{1}{n} \right) \leq \frac{5C\Delta}{n}. \end{aligned}$$

This yields that

$$\begin{aligned} & \mathbf{E} \left[\Phi^{(t+1)} \mid X^{(t)} = x \right] \\ & \leq \sum_{u \in V} \left(\Pr \left[Z_u^{(t+1)} = Z_u^{(t)} + 1 \mid X^{(t)} = x \right] \right. \\ & \quad \cdot \exp \left(\sigma \cdot (Z_u^{(t)} + 1) \right) \\ & \quad \left. + \Pr \left[Z_u^{(t+1)} = Z_u^{(t)} \mid X^{(t)} = x \right] \right. \\ & \quad \left. \cdot \exp \left(\sigma \cdot Z_u^{(t)} \right) \right) \\ & = \sum_{u \in V} \left(e^\sigma \cdot \Pr \left[Z_u^{(t+1)} = Z_u^{(t)} + 1 \mid X^{(t)} = x \right] + 1 \right) \\ & \quad \cdot \exp \left(\sigma \cdot Z_u^{(t)} \right) \\ (5.2) \quad & \leq \left(1 + e^\sigma \cdot \frac{5C\Delta}{n} \right) \cdot \Phi^{(t)}, \end{aligned}$$

where (5.2) follows from (5.1). Note that if we only consider the allocation of $\alpha := n/(e^\sigma \cdot 25C^2) = \Theta(n)$ balls, then we have that

$$\left(1 + e^\sigma \cdot \frac{5C\Delta}{n} \right)^\alpha \leq e^{\Delta/(5C)} \leq e^{\delta/4}.$$

Define $\Psi^{(t)} := \min\{\Phi^{(t)}, n \cdot e^{\delta/2}\}$. Then, since $\Phi^{(t)}$ is increasing in t , (5.2) yields that

$$\mathbf{E} \left[\Psi^{(t+1)} \right] \leq \left(1 + e^\sigma \cdot \frac{5C\Delta}{n} \right) \cdot \Psi^{(t)}$$

and thus inductively, $\mathbf{E} \left[\Psi^{(\alpha)} \right] \leq \left(1 + e^\sigma \cdot \frac{5C\Delta}{n} \right)^\alpha \cdot \Psi^{(0)} \leq e^{\delta/4} \cdot n$. Hence, applying Markov's inequality gives $\Pr \left[\Psi^{(\alpha)} < e^{\delta/2} \cdot n \right] \geq 1 - e^{-\delta/4}$. By definition of $\Psi^{(\alpha)}$, if $\Psi^{(\alpha)} < e^{\delta/2} \cdot n$, then $\Psi^{(\alpha)} = \Phi^{(\alpha)}$. Hence, $\Pr \left[\Phi^{(\alpha)} < e^{\delta/2} \cdot n \right] \geq 1 - e^{-\delta/4}$, as required. If $\Phi^{(\alpha)} < e^{\delta/2} \cdot n$ occurs, then since every vertex has at least one neighbor with load zero, the maximum load after the allocation of α balls is 1. Then, we use subadditivity (cf. Lemma 2.7) to conclude that the maximum load after all $\lceil n/\alpha \rceil \cdot \alpha$ balls are allocated is at most $1 \cdot \lceil n/\alpha \rceil$ with probability at least $1 - \lceil n/\alpha \rceil \cdot o(1) = 1 - o(1)$. ■

6 Impact of tie-breaking rules

Proof. [of Theorem 1.5] We now analyze the effect of employing different tie-breaking rules. We first describe the construction of the graph G , which will be a d -regular graph, where $d = \omega(1)$ as $n \rightarrow \infty$. Additionally, we may assume that $d \leq \sqrt{n}$, since otherwise the claimed lower bound is trivial. For the

construction of G , we assume that there is an integer k such that $1 + \sum_{i=0}^{k-1} d \cdot (d-1)^i = n$. Note that $k = \Theta(\log n / \log d)$. Then, let G be a balanced tree with root s so that

$$|N_s^i| = d \cdot (d-1)^{i-1} \quad \text{for any } i \geq 1 .$$

Hence G is a tree where all vertices except for the leaves and the root have $d-1$ successors; the root has d successors, and the leaves have no successor. Hence, the root has degree d , the inner vertices have degree d as well and the leaves have degree 1. Further, note that the number of leaves is $d \cdot (d-1)^{k-1}$. In order to make the graph d -regular, we simply add edges among the leaves in G so that, after all edges have been added, every leaf has degree d (this is possible, since the number of edges to add is smaller than the total number of leaves).

We choose $\alpha := \min\{(d-1)^{1/4}, k-2\}$. The process of allocating the n balls will be divided into α phases and in each phase we consider the allocation of n/α balls. To prove the desired lower bound, we focus on the vertices in $B_s^{2\alpha}$.

Next, we define an event that essentially shows that there are always enough balls so that the tie-breaking rule can send balls towards the root:

$$\mathcal{E} := \bigcap_{p=1}^{\alpha} \bigcap_{\ell=1}^{2\alpha} \left\{ \forall u \in N_s^\ell \exists v \in N_u \cap N_s^{\ell+1} : Z_p(v) \geq 2 \right\} ,$$

where $Z_p(v) := \sum_{t=(p-1) \cdot (n/\alpha) + 1}^{p \cdot (n/\alpha)} \mathbf{1}(U_t = v)$ is the number of balls born on v in phase p . Hence the event \mathcal{E} means that for each vertex in N_s^ℓ , there is in each phase at least one neighbor in $N_s^{\ell+1}$ on which two balls are born; hence, for at least one ball we may be able to use the tie-breaking rule and forward the ball towards the root s .

Let us estimate the probability that the event \mathcal{E} occurs. First, for any fixed $u \in N_s^\ell$ and $v \in N_u \cap N_s^{\ell+1}$

$$\begin{aligned} \Pr[Z_p(v) \geq 2] &\geq \binom{n/\alpha}{2} \cdot \frac{1}{n^2} \cdot \left(1 - \frac{1}{n}\right)^{n/\alpha-2} \\ &\geq \frac{1}{8\alpha^2} . \end{aligned}$$

Since the events $\{Z_p(v) \geq 2\}_{v \in N_u}$ are negatively correlated, we have that, for any fixed $u \in N_s^\ell$,

$$\begin{aligned} \Pr[\exists v \in N_u \cap N_s^{\ell+1} : Z_p(v) \geq 2] \\ \geq 1 - \left(1 - \frac{1}{8\alpha^2}\right)^{d-1} \geq 1 - \exp\left(-\frac{d-1}{8\alpha^2}\right) . \end{aligned}$$

Hence,

$$\begin{aligned} \Pr[\mathcal{E}] &\geq 1 - \sum_{\substack{1 \leq p \leq \alpha \\ 1 \leq \ell \leq 2\alpha}} \sum_{u \in N_s^\ell} \Pr[\neg \exists v \in N_u \cap N_s^{\ell+1} : Z_p(v) \geq 2] \\ &\geq 1 - \alpha \cdot \sum_{1 \leq \ell \leq 2\alpha} |N_s^\ell| \cdot \exp\left(-\frac{d-1}{8\alpha^2}\right) \\ (6.1) \quad &\geq 1 - \exp\left(-\frac{d-1}{8\alpha^2}\right) \cdot \alpha \cdot 2d(d-1)^{2\alpha-1} . \end{aligned}$$

We now claim that the last term in (6.1) is $1 - o(1)$. To this end, recall the choice of α and $d = \omega(1)$. First, since $\alpha \leq (d-1)^{1/4}$, $\exp\left(-\frac{d-1}{8\alpha^2}\right) \leq \exp\left(-\frac{(d-1)^{1/2}}{8}\right)$, whereas $\alpha \cdot 2d \cdot (d-1)^{2\alpha-1} \leq 2d^{3\alpha} \leq 2d^{3(d-1)^{1/4}} = 2 \cdot \exp(\log d \cdot 3(d-1)^{1/4})$. Hence, as $d \rightarrow \infty$, the probability on the right-hand side in (6.1) is $1 - o(1)$; i.e., we have shown that

$$\Pr[\mathcal{E}] \geq 1 - o(1) .$$

It remains to show that if the event \mathcal{E} occurs, then it is possible to break the ties so that we end up with a maximum load of α after the allocation of n balls (see Figure 2 for an illustration). Our tie-breaking rule follows the simple strategy that, whenever possible, a ball is forwarded to a vertex closer to the root s .

By Lemma 2.6, we may assume that balls are only generated in the set B_s^α and, in each phase, every vertex in this set generates at most two balls. Then after the first phase is completed, every vertex in $B_s^{2\alpha}$ contains at least one ball, since for each such vertex there is at least one neighbor in $B_s^{2\alpha+1}$ on which two balls are born.




Figure 2: Illustration of the $\alpha = 3$ phases after which the root vertex s has a load of at least 3. The red color indicates the ball which are placed on the vertex in the recent phase. The existence of a red ball on a vertex, say, u , follows, since there exists at least one successor of u , say, $v \in N_u$, on which at least two balls are born in that phase.

Moreover, in the second phase, every vertex u in $B_s^{2\alpha-2}$ has at least one neighbor $v \in N_u$ in the next level such that: (i) two balls are generated on v in the second phase and (ii) all neighbors of v and v itself have at least one ball at the beginning of phase two. Hence if ties are always broken in the direction towards the root, every vertex in $B_s^{2\alpha-2}$ will have at least two balls at the end of phase two. Completing the induction, we conclude that at the end of phase α , the vertex s will have at least α balls. Overall, we conclude that the maximum load is at least α whenever the event \mathcal{E} occurs. Since \mathcal{E} occurs with

probability $1 - o(1)$, the proof is complete. \blacksquare

7 Lower bounds for sparse graphs

Proof. [of Theorem 1.6] In the first part, we show that the maximum load after n balls are allocated is $\Omega\left(\frac{\log \log n}{\log \Delta}\right)$, where G is any graph with maximum degree Δ . Our arguments are almost the same as in the proof of the lower bound of Theorem 1.3. We are using again the fact that with probability at least $1 - n^{-1}$, there is a vertex $u \in V$ on which at least $\frac{\log n}{\log \log n}$ balls are born (cf. [13, Lemma 5.12]). Then, applying Lemma 2.4 with $S = \{u\}$, we obtain that the maximal load $\beta := X_{\max}^{(n)}$ satisfies

$$\beta \cdot |B_u^\beta| \geq \frac{\log n}{\log \log n}.$$

Since $|B_u^\beta| \leq \Delta^\beta$, the above inequality implies that

$$\beta \cdot \Delta^\beta \geq \frac{\log n}{\log \log n},$$

which in turn implies that

$$\beta = \Omega\left(\frac{\log\left(\frac{\log n}{\log \log n}\right)}{\log \Delta} - \log(\beta)\right),$$

i.e., $\beta = \Omega\left(\frac{\log \log n}{\log \Delta}\right)$.

Now, for the second part, we show that there exists a d -regular graph for which the maximum load after n balls are allocated is $\Omega\left(\sqrt{\frac{\log n}{d \cdot \log\left(\frac{\log n}{d}\right)}}\right)$. We

first describe the construction of the d -regular graph G . First take $n/(d-1)$ disjoint cliques of size $d-1$ and arrange them in a cycle. Then connect two cliques which are next to each other in the cycle by $d-1$ vertex-disjoint edges. This way we obtain a d -regular graph, which can be also defined as the Cartesian product of a cycle of length $n/(d-1)$ and a clique of size $d-1$. See Figure 3 below.




Figure 3: Illustration of the construction of the graph G , where $d-1 = 4$ and $n/(d-1) = 9$, so $n = 36$.

Let us now consider the number of balls that are born in each clique. This can be seen as the 1-choice process where n balls are randomly placed into

$n/(d-1)$ bins. By Raab and Steger [17, Theorem 1, Cases 1 & 2] it follows that, with probability $1 - o(1)$, there exists a clique with vertex set S , $|S| = d-1$, so that the number of balls born in S is at least

$$\Phi_S := C \cdot \left(\frac{\log n}{\log \left(\frac{\log(n/(d-1))}{d-1} \right)} \right),$$

for some constant $C > 0$. Now we use Lemma 2.4 with the Φ_S above to conclude that the maximum load $\beta = X_{\max}^{(n)}$ satisfies

$$\beta \cdot |B_S^\beta| \geq C \cdot \left(\frac{\log n}{\log \left(\frac{\log(n/(d-1))}{d-1} \right)} \right).$$

Since $|B_S^\beta| \leq (2\beta+1) \cdot d$ and $\log(n/(d-1)) = \Theta(\log n)$, this implies $\beta = \Omega \left(\sqrt{\frac{\log n}{d \cdot \log \left(\frac{\log n}{d} \right)}} \right)$. ■

References

- [1] M. Adler, E. Halperin, R. M. Karp, and V. V. Vazirani. A stochastic process on the hypercube with applications to peer-to-peer networks. In *Proc. 35th Symp. on Theory of Computing (STOC)*, pages 575–584, 2003.
- [2] N. Alon and J. Spencer. *The probabilistic method*. John Wiley & Sons, 3rd edition, 2008.
- [3] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, 1999.
- [4] P. Berenbrink, A. Czumaj, A. Steger, and B. Vöcking. Balanced allocations: The heavily loaded case. *SIAM J. Comput.*, 35(6):1350–1385, 2006.
- [5] P. Berenbrink, T. Friedetzky, L. A. Goldberg, P. W. Goldberg, Z. Hu, and R. A. Martin. Distributed selfish load balancing. *SIAM J. Comput.*, 37(4):1163–1181, 2007.
- [6] A. Z. Broder and M. Mitzenmacher. Multidimensional balanced allocations. In *37th Annual ACM Symposium on Theory of Computing (STOC'05)*, pages 195–196, 2005.
- [7] F. Chung and L. Lu. Concentration inequalities and Martingale inequalities: a survey. *Internet Mathematics*, 3:79–127, 2006.
- [8] E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to nash equilibrium in load balancing. *ACM Transactions on Algorithms*, 3(3), 2007.
- [9] B. Godfrey. Balls and bins with structure: balanced allocations on hypergraphs. In *19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'08)*, pages 511–517, 2008.
- [10] P. W. Goldberg. Bounds for the convergence rate of randomized local search in a multiplayer load-balancing game. In *23rd Annual ACM-SIGOPT Principles of Distributed Computing (PODC'04)*, pages 131–140, 2004.
- [11] R. M. Karp, M. Luby, and F. Meyer auf der Heide. Efficient PRAM simulation on a distributed memory machine. *Algorithmica*, 16(4/5):517–542, 1996.
- [12] K. Kenthapadi and R. Panigrahy. Balanced allocation on graphs. In *17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'06)*, pages 434–443, 2006.
- [13] M. Mitzenmacher and E. Upfal. *Probability and Computing: randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.
- [14] M. Mitzenmacher, A. Richa, and R. Sitaraman. The power of two random choices: A survey of techniques and results. In *Handbook of Randomized Computing*, volume 1, pages 255–312. 2001.
- [15] M. Mitzenmacher, B. Prabhakar, and D. Shah. Load balancing with memory. In *43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02)*, pages 799–808, 2002.
- [16] Y. Peres, K. Talwar, and U. Wieder. The $(1 + \beta)$ -choice process and weighted balls-into-bins. In *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'10)*, pages 1613–1619, 2010.
- [17] M. Raab and A. Steger. Balls into bins - a simple and tight analysis. In *2nd International Workshop on Randomization and Computation (RANDOM'98)*, pages 159–170, 1998.
- [18] B. Vöcking. How asymmetry helps load balancing. *J. ACM*, 50(4), 2003.

A Standard technical results

LEMMA A.1. ([2, THEOREM 7.2.1]) *Let X_0, X_1, \dots, X_m be a martingale such that there exists a fixed positive c for which $|X_i - X_{i-1}| \leq c$ for all i . Then,*

$$\Pr[|X_m - X_0| \geq \lambda] \leq \exp\left(-\frac{\lambda^2}{2c^2m}\right).$$

LEMMA A.2. ([7, THEOREM 6.1]) *Let X_0, X_1, \dots, X_m be a martingale adapted to the*

filtration \mathcal{F}_i . Suppose that there exists a fixed positive c for which $|X_i - X_{i-1}| \leq c$ for all i and there exists c' such that $\mathbf{E}[(X_i - X_{i-1})^2 | \mathcal{F}_{i-1}] \leq c'$ for all i . Then,

$$\Pr[|X_m - X_0| \geq \lambda] \leq \exp\left(-\frac{\lambda^2}{2c'm + c\lambda/3}\right).$$

For the special case where X_1, \dots, X_m are independent Bernoulli random variables, we can apply the above lemma to the random variables $(X_i - \mathbf{E}[X_i])_i$ with $c' = \mathbf{E}[X_1]$ and $c = 1$ to obtain the inequality below.

LEMMA A.3. Let X_1, \dots, X_m be m independent, identically distributed Bernoulli random variables. Let $X := \sum_{i=1}^m X_i$. Then, for any $\lambda > 0$,

$$\Pr[|X - \mathbf{E}[X]| \geq \lambda] \leq \exp\left(-\frac{\lambda^2}{2\mathbf{E}[X] + \lambda/3}\right).$$

LEMMA A.4. Consider the 1-choice process $\{\bar{X}_v^{(n)}\}_{v \in V}$ where n balls are allocated to n bins chosen independently and uniformly at random. Let $\ell_0 := 8e\Delta^2$, and let $\Lambda_\ell := \{u \in V: X_u^{(n)} \geq \ell\}$. Then, for any ℓ with $\ell_0 \leq \ell = o(\log^2 n)$,

$$\Pr\left[|\Lambda_\ell| \geq \frac{n}{4\Delta} \cdot (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell}\right] \leq 2n^{-\log^5 n}.$$

Proof. Fix any ℓ with $\ell \geq \ell_0$. Let $\{\tilde{X}_v\}_{v \in V}$ be n independent poisson random variables with mean 1. Define $\tilde{\Lambda}_\ell := \{u \in V: \tilde{X}_u^{(n)} \geq \ell\}$. Since $\left\{|\tilde{\Lambda}_\ell| \geq \frac{|V|}{4\Delta} \cdot (2\Delta)^{-\ell} + \frac{\log^7 |V|}{\ell}\right\}$ is a monotone event in the number of balls $1 \leq n \leq |V|$, it follows by a standard ‘‘Poissonization argument’’ (see, e.g., [13, Corollary 5.11]) that

$$\begin{aligned} \Pr\left[|\Lambda_\ell| \geq \frac{n}{4\Delta} \cdot (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell}\right] \\ \leq 2 \cdot \Pr\left[\tilde{\Lambda}_\ell \geq \frac{n}{4\Delta} \cdot (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell}\right]. \end{aligned}$$

To bound the latter probability, let us first estimate $\mathbf{E}\left[|\tilde{\Lambda}_\ell|\right]$. First, if P is a Poisson random variable with parameter 1, then we have the following Chernoff-type inequality ([2, Theorem A.1.15]): for any $\varepsilon > 0$,

$$\Pr[P \geq (1 + \varepsilon)] \leq e^\varepsilon (1 + \varepsilon)^{-(1 + \varepsilon)}.$$

As long as $\varepsilon \geq 8e\Delta^2 - 1$, we can write

$$\Pr[P \geq 1 + \varepsilon] \leq e^{1 + \varepsilon} (8e\Delta^2)^{-(1 + \varepsilon)} \leq (8\Delta^2)^{-(1 + \varepsilon)},$$

and hence, replacing $1 + \varepsilon$ by ℓ gives

$$\Pr[P \geq \ell] \leq (8\Delta^2)^{-\ell}.$$

Now observe that $|\tilde{\Lambda}_\ell|$ is stochastically smaller than the sum of n independent Bernoulli random variables Z_1, \dots, Z_n , each with parameter $(8\Delta^2)^{-\ell}$. Hence, if we denote $Z := \sum_{i=1}^n Z_i$, then

$$\begin{aligned} \Pr\left[|\tilde{\Lambda}_\ell| \geq \frac{n}{4\Delta} \cdot (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell}\right] \\ \leq \Pr\left[Z \geq \frac{n}{4\Delta} \cdot (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell}\right]. \end{aligned}$$

Note that $\mathbf{E}[Z] = n \cdot (8\Delta^2)^{-\ell}$. Hence by Lemma A.3,

$$\begin{aligned} \Pr\left[Z \geq \frac{n}{4\Delta} \cdot (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell}\right] \\ \leq \Pr\left[Z \geq \mathbf{E}[Z] + \frac{\log^7 n}{\ell} + \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}\right] \\ \leq \exp\left(-\frac{\left(\frac{\log^7 n}{\ell} + \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}\right)^2}{2\mathbf{E}[Z] + \frac{\log^7 n}{\ell} + \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}}\right). \end{aligned}$$

To bound the last term, we proceed by a case distinction. The first case is when $\log^7 n/\ell \geq \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}$. Then also $\mathbf{E}[Z] \leq \log^7 n/\ell$ and hence

$$\begin{aligned} \exp\left(-\frac{\left(\frac{\log^7 n}{\ell} + \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}\right)^2}{2\mathbf{E}[Z] + \frac{\log^7 n}{\ell} + \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}}\right) \\ \leq \exp\left(-\frac{\left(\frac{\log^7 n}{\ell}\right)^2}{4 \cdot \frac{\log^7 n}{\ell}}\right) \leq n^{-\log^5 n}. \end{aligned}$$

Otherwise, $\log^7 n/\ell < \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}$. Then,

$$\begin{aligned} \exp\left(-\frac{\left(\frac{\log^7 n}{\ell} + \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}\right)^2}{2\mathbf{E}[Z] + \frac{\log^7 n}{\ell} + \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}}\right) \\ \leq \exp\left(-\frac{\left(\frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}\right)^2}{4 \cdot \frac{n}{8\Delta} \cdot (2\Delta)^{-\ell}}\right) \leq n^{-\log^5 n}. \end{aligned}$$

Hence,

$$\Pr\left[|\Lambda_\ell| \geq \frac{n}{4\Delta} \cdot (2\Delta)^{-\ell} + \frac{\log^7 n}{\ell}\right] \leq 2 \cdot n^{-\log^5 n}.$$

as desired. \blacksquare