Citation for published version:

Bringmann, K, Sauerwald, T, Stauffer, A & Sun, H 2016, 'Balls into bins via local search: cover time and maximum load', *Random Structures and Algorithms*, vol. 48, no. 4, pp. 681-702.
<https://doi.org/10.1002/rsa.20602>

DOI:

[10.1002/rsa.20602](https://doi.org/10.1002/rsa.20602)

Publication date:

2016

Document Version

Peer reviewed version

[Link to publication](#)

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 via local search: cover time and maximum load^{*}

Karl Bringmann[†]   Thomas Sauerwald[‡]   Alexandre Stauffer[§]   He Sun[¶]

Abstract

We study a natural process for allocating m balls into n bins that are organized as the vertices of an undirected graph G . Balls arrive one at a time. When a ball arrives, it 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. Then the next ball arrives and this procedure is repeated. For the case $m = n$, we give an upper bound for the maximum load on graphs with bounded degrees. We also propose the study of the *cover time* of this process, which is defined as the smallest m so that every bin has at least one ball allocated to it. We establish an upper bound for the cover time on graphs with bounded degrees. Our bounds for the maximum load and the cover time are tight when the graph is vertex transitive or sufficiently homogeneous. We also give upper bounds for the maximum load when $m \geq n$.

Keywords and phrases. balls-into-bins, load balancing, stochastic process, local search.

1 Introduction

A very simple procedure for allocating m balls into n bins is to place each ball into a bin chosen independently and uniformly at random. We refer to this process as *1-choice process*. It is well known that, when $m = n$, the maximum load for the 1-choice process (i.e., the maximum number of balls allocated to any single bin) is $\Theta\left(\frac{\log n}{\log \log n}\right)$ [10]. Alternatively, in the d -choice process, balls arrive sequentially one after the other, and when a ball arrives, it chooses d bins independently and uniformly at random, and places itself in the bin that currently has the smallest load among the d bins (ties are broken uniformly at random). It was shown by Azar et al. [3] and Karp et al. [7] that the maximum load for the d -choice process with $m = n$ and $d \geq 2$ is $\Theta\left(\frac{\log \log n}{\log d}\right)$. The constants omitted in the Θ are known and, as shown by Vöcking [11], they can be reduced with a slight modification of the d -choice process. Berenbrink et al. [4] extended these results to the case $m \gg n$.

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

---

^{*}A preliminary version of this paper appeared in the 31st Symposium on Theoretical Aspects of Computer Science (STACS'14).

[†]Institute for Theoretical Computer Science, ETH Zurich, Switzerland. Email: karlb@inf.ethz.ch. Supported by the ETH Zurich Postdoctoral Fellowship Program.

[‡]Computer Laboratory, University of Cambridge, UK. Email: thomas.sauerwald@cl.cam.ac.uk

[§]Department of Mathematical Sciences, University of Bath, UK. Email: a.stauffer@bath.ac.uk. Supported in part by a Marie Curie Career Integration Grant PCIG13-GA-2013-618588 DSRELIS.

[¶]Department of Computer Science, University of Bristol, UK. Email: hsun@mpi-inf.mpg.de

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 $i + 2$.

processors to achieve a more balanced load distribution. In such settings, allocating balls close to the processor that created them is certainly very desirable as it reduces the costs of probing the load of a processor and allocating the task.

With this motivation in mind, Bogdan et al. [5] introduced a natural allocation process called *local search allocation*. Consider that the bins are organized as the vertices of a graph $G = (V, E)$ with $n = |V|$. At each step a ball is “born” at a vertex chosen independently and uniformly at random from V , which we call the *birthplace* of the ball. Then, starting from its birthplace, the ball performs a local search in G , where in each step the ball moves to the adjacent vertex with the smallest load, provided that the load is strictly smaller than the load of the vertex the ball is currently in. We assume that ties are broken independently and 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 the procedure above is repeated. See Figure 1 for an illustration.

The main result in [5] establishes that when G is an expander graph with bounded maximum degree, the maximum load after n balls have been allocated is $\Theta(\log \log n)$. Hence, local search allocation on bounded-degree expanders achieves the same maximum load (up to constants) as in the d -choice process, but has the extra benefit of requiring only local information during the allocation. In [5], it was also established that the maximum load is $\Theta\left(\left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}\right)$ on d -dimensional grids, and $\Theta(1)$ on regular graphs of degrees $\Omega(\log n)$.

1.1 Results

In this paper we derive upper and lower bounds for the maximum load that hold for all bounded-degree graphs. In addition, we propose the study of another natural quantity, which we refer to as the *cover time*. In order to state our results, we need to introduce the following two quantities that are related to the local neighborhood growth of G :

$$R_1 = R_1(G) = \min\{r \in \mathbb{N} : r|B_u^r| \log r \geq \log n \text{ for all } u \in V\} \quad (1.1)$$

and

$$R_2 = R_2(G) = \min\{r \in \mathbb{N} : r|B_u^r| \geq \log n \text{ for all } u \in V\}, \quad (1.2)$$

where B_u^r denotes the set of vertices within distance at most r from vertex u . Note that $R_1 \leq R_2$ for all G . For bounded-degree expander graphs, we have $|B_u^r| = e^{\Theta(r)}$ which implies that R_1

and R_2 are of order $\log \log n$, whereas for a d -dimensional grids we have $|B_u^r| = \Theta(r^d)$ which gives that $R_1 = \Theta\left(\left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}\right)$ and $R_2 = \Theta\left((\log n)^{\frac{1}{d+1}}\right)$.

For the sake of clarity, we state our results here for *vertex-transitive* graphs only. In later sections we state our results in fullest generality, which will require a more refined definition of R_1 and R_2 . We also highlight that for all the results below (and throughout this paper) we assume that ties are broken independently and uniformly at random; the impact of tie-breaking procedures in local search allocation was investigated in [5, Theorem 1.5].

Maximum load

We derive an upper bound for the maximum load after n balls have been allocated. Our bound holds for *all* bounded-degree graphs, and is tight for vertex-transitive graphs (and, more generally, for graphs where the neighborhood growth is sufficiently homogeneous across different vertices).

Theorem 1.1 (Maximum load when $m = n$). *Let G be any vertex-transitive graph with bounded degrees. Then, with probability at least $1 - n^{-1}$, the maximum load after n balls have been allocated is $\Theta(R_1)$.*

Theorem 1.1 is a special case of Theorem 4.1, which gives a more precise version of the result above and generalizes it to non-transitive graphs; in particular, we obtain that for any graph with bounded degrees the maximum load is $\mathcal{O}(R_1)$ with high probability. We state and prove Theorem 4.1 in Section 4.

As mentioned above, for bounded-degree expanders we have $R_1 = \Theta(\log \log n)$, and for d -dimensional grids we have $R_1 = \Theta\left(\left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}\right)$. Hence the results for bounded-degree graphs in [5] are special cases of Theorems 1.1 and 4.1. Furthermore, the proof of Theorems 1.1 and 4.1 uses novel techniques, and is substantially shorter than the proofs in [5]. In particular, in [5], different proofs are needed to handle the cases of expander graphs and grids, and the proof for expander graphs uses a rather involved recursive application of Azuma's inequality. Our proof here, on the other hand, is applicable to all bounded-degree graphs and is quite elegant. We construct a subtle coupling between local search allocation and the 1-choice process, which allows us to express the probability that a vertex has a large load in terms of the probability for the number of balls being born at a given vertex, which are much simpler to control (see Lemma 3.4 below).

Our second result establishes an upper bound for the maximum load when $m \geq n$. We point out that all other results known so far were limited to the case $m = n$. We establish that, when $m = \Omega(R_2 n)$, the maximum load is of order $\Theta(m/n)$ (i.e., the same order as the average load). We note that the difference between the maximum load and the average load for the local search allocation is always bounded above by the diameter of the graph (see Lemma 2.2 below). This is in some sense similar to the d -choice process, where the difference between the maximum load and the average load does not depend on m [4].

Theorem 1.2 (Maximum load when $m \geq n$). *Let G be any graph with bounded degrees. Then for any $m \geq n$, with probability at least $1 - n^{-1}$, the maximum load after m balls have been allocated is $\mathcal{O}(\frac{m}{n} + R_2)$.*

Cover time

We propose to study the following natural quantity related to any process based on allocating balls into bins. Define the *cover time* as the first time at which all bins have at least one ball allocated to them. This is in analogy with cover time of random walks on graphs, which is the

first time at which the random walk has visited all vertices of the graph. Note that for the 1-choice process, the cover time corresponds to the time of a *coupon collector* problem, which is known to be of order $n \log n$ with high probability [9, Theorem 5.13]. For the d -choice process with $d = \Theta(1)$, we obtain that the cover time is also of order $n \log n$. We show that for the local search allocation the cover time can be much smaller than $n \log n$.

Our next theorem establishes that the cover time for vertex-transitive bounded-degree graphs is $\Theta(R_2 n)$ with high probability.

Theorem 1.3 (Cover time for bounded-degree graphs). *Let G be any vertex-transitive graph with bounded degrees. Then, with probability at least $1 - n^{-1}$, the cover time of local search allocation on G is $\Theta(R_2 n)$.*

The theorem above is a special case of Theorem 5.2, which we state and prove in Section 5; in particular, we establish there that the upper bound of $\mathcal{O}(R_2 n)$ holds for all bounded-degree graphs. Since $R_2 = \mathcal{O}(\sqrt{\log n})$ for all connected graphs, it follows that the cover time for any connected, bounded-degree graph is at most $\mathcal{O}(n\sqrt{\log n})$, which is significantly smaller than the cover time of the d -choice process, which is $\Theta(n \log n)$ for any $d = \Theta(1)$. In particular, we obtain $R_2 = \Theta(\log \log n)$ for bounded-degree expanders, and $R_2 = \Theta((\log n)^{\frac{1}{d+1}})$ for d -dimensional grids.

Our final result provides a general upper bound on the cover time for dense graphs. Theorem 1.4 below is a special case of Theorem 5.3, which gives an upper bound on the cover time for all regular graphs. We state and prove Theorem 5.3 in Section 5.

Theorem 1.4 (Cover time for dense graphs). *Let G be any d -regular graph with $d = \Omega(\log n \log \log n)$. Then, with probability at least $1 - n^{-1}$, the cover time is $\Theta(n)$.*

2 Background and notation

In this section we recall some basic properties of the local search allocation that will be useful in our proofs.

Let $G = (V, E)$ be an undirected, not necessarily connected, graph with n vertices, and let Δ be the maximum degree of G . We assume that, in the local search allocation, ties are broken independently and uniformly at random.

We denote by $d_G(u, v)$ the distance between u and v in G and define $d_G(v, S) = \min_{v' \in S} d_G(v, v')$ for any non-empty subset $S \subseteq V$. Further, we define $B_u^r = \{v \in V : d_G(v, u) \leq r\}$ and for any non-empty set $S \subseteq V$, $B_S^r = \{v \in V : d_G(v, S) \leq r\}$.

For each $m \geq 0$ and vertex $v \in V$, let $X_v^{(m)}$ denote the load of v (i.e., the number of balls allocated to v) after m balls have been allocated. Initially we have $X_v^{(0)} = 0$ for all $v \in V$ and, for any $m \geq 0$, we have $\sum_{v \in V} X_v^{(m)} = m$. Denote by $X_{\max}^{(m)}$ the maximum load after m balls have been allocated; i.e.,

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

Also, denote by $T_{\text{cov}} = T_{\text{cov}}(G)$ the *cover time* of G , which we define as the first time at which all vertices have load at least 1. More formally,

$$T_{\text{cov}} = \min\{m \geq n : X_v^{(m)} \geq 1 \text{ for all } v \in V\}.$$

Let $U_i \in V$ denote the birthplace of ball i , and for each $m \geq 0$ and $v \in V$, let $\bar{X}_v^{(m)}$ denote the load of v after m balls have been allocated according to the 1-choice process. Let $\bar{X}_{\max}^{(m)}$ denote the maximum load for the 1-choice process. More formally, for any $m \geq 0$ and any $v \in V$,

$$\bar{X}_v^{(m)} = \sum_{i=1}^m \mathbf{1}(U_i = v) \quad \text{and} \quad \bar{X}_{\max}^{(m)} = \max_{v \in V} \bar{X}_v^{(m)}, \quad (2.1)$$

where $\mathbf{1}(\cdot)$ denotes the indicator function.

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 integers from $\{1, 2, \dots, n\}$ 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 integers from $\{1, 2, \dots, n\}$ such that $a'_{j'_1} \geq a'_{j'_2} \geq \dots \geq a'_{j'_n}$, then A majorizes A' if

$$\sum_{i=1}^{\kappa} a_{j_i} \geq \sum_{i=1}^{\kappa} a'_{j'_i} \quad \text{for all } \kappa = 1, 2, \dots, n. \quad (2.2)$$

For two random variables Y and Z on $\mathbb{R}$, we say that Y *stochastically dominates* Z if we can couple the probability distributions of Y and Z so that $Y \geq Z$ with probability 1.

The lemma below establishes that the load vector obtained by the 1-choice process majorizes the load vector obtained by the local search allocation. This implies that the maximum load in the 1-choice process stochastically dominates the maximum load obtained by the local search allocation. As a consequence, we have that $X_{\max}^{(n)} = \mathcal{O}\left(\frac{\log n}{\log \log n}\right)$ and $T_{\text{cov}} = \mathcal{O}(n \log n)$ for all G . Later, in Section 3, we state and prove Lemma 3.2, which is a generalization of Lemma 2.1.

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 $X_{\max}^{(k)}$.*

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. $\square$

The next lemmas establish that the load vector $X^{(n)}$ satisfies a Lipschitz and monotonicity condition.

Lemma 2.3 (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 . Then, there exists a coupling such that, with probability 1,*

$$\sum_{v \in V} |X_v^{(k)} - Y_v^{(k)}| \leq 2. \quad (2.3)$$

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. 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 . Since the first $i - 1$ balls have the same birthplaces in both processes, we have that

$$X_v^{(i-1)} = Y_v^{(i-1)} \quad \text{for all } v \in V. \quad (2.4)$$

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 (2.3) holds. More generally, we have that

$$\begin{cases} X_{v_i}^{(i)} = Y_{v_i}^{(i)} + 1 & (v_i \neq v'_i) \\ Y_{v'_i}^{(i)} = X_{v'_i}^{(i)} + 1 & (v_i \neq v'_i) \\ X_u^{(i)} = Y_u^{(i)} & \text{for } u \in V \setminus \{v_i, v'_i\}. \end{cases} \quad (2.5)$$

If $i = k$, then this implies (2.3) 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 (2.5) 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 (2.3) clearly holds. We note that v_{i+1} and v'_{i+1} will not necessarily be the vertices to which ball $i + 1$ is allocated in the X and Y processes. 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 (2.5) holds with i replaced by $i + 1$ by setting $v_{i+1} = v_i$ and $v'_{i+1} = v'_i$.

Now we analyze ball $i + 1$. First, 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. In order to see this, consider that during the local search of ball $i + 1$, the ball is on a vertex $z \in V \setminus \{v_i, v'_i\}$ and has so far performed the same steps in both the X and Y processes. It is enough to show that, in the next step of the local search, if ball $i + 1$ does not go to v_i or v'_i , then it goes to the same vertex in both the X and Y processes. This establishes the above observation by induction. There are two cases, the first being if z is not a neighbor of v_i or v'_i . Then the neighbors of z have the same load in processes X and Y , so ball $i + 1$ does the same step in each process. On the other hand, if z is a neighbor of v_i (the same reasoning will apply with v_i replaced by v'_i), given that ball $i + 1$ never enters v_i , z must have a neighbor $z' \neq v_i$ so that z' has the smallest load among the neighbors of z in both X and Y , and in case of ties, z' is the neighbor of smallest load appearing first in $\xi_z^{(i+1)}$. Since the local searches of ball $i + 1$ in the X and Y processes break ties according to the same permutation $\xi_z^{(i+1)}$, we obtain that ball $i + 1$ goes to z' in both X and Y .

As a consequence of the above reasoning, 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 (2.5) holds since $X_u^{(i+1)} = Y_u^{(i+1)}$ for all $u \in V$. Otherwise we set $v'_{i+1} = v'_i$, and (2.5) holds as well. $\square$

Lemma 2.4 (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 . There exists a coupling such that, with probability 1,

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

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 2.3 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

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

In many of our proofs we analyze a continuous-time variant where the number of balls is *not* fixed, but is given by a Poisson random variable with mean m . Equivalently, in this variant balls are born at each vertex according to a Poisson process of rate $1/n$. We refer to this as the *Poissonized* version. We will use the Poissonized versions of both the local search allocation and the 1-choice process in our proofs. Since the probability that a mean- m Poisson random variable takes the value m is of order $\Theta(m^{-1/2})$ we obtain the following relation.

Lemma 2.5. *Let $\mathcal{A}$ be an event that holds for the Poissonized version of the local search allocation (respectively, 1-choice process) with probability $1 - \varepsilon$ for some $\varepsilon \in (0, 1)$. Then, the probability that $\mathcal{A}$ holds for the non-Poissonized version of the local search allocation (respectively, 1-choice process) is at least $1 - \mathcal{O}(\varepsilon\sqrt{m})$.*

3 Key technical argument

In this section we prove a key technical result (Lemma 3.2 below) that will play a central role in our proofs later.

Let $\mu: V \rightarrow \mathbb{Z}$ be any integer function on the vertices of G that satisfies the following property:

$$\text{for any two neighbors } u, v \in V, \text{ we have } |\mu(u) - \mu(v)| \leq 1. \quad (3.1)$$

We see μ as an initial attribution of weights to the vertices of G . Then, for any $m \geq 1$, after m balls are allocated, we define the weight of vertex v by

$$W_v^{(m)} = X_v^{(m)} + \mu(v). \quad (3.2)$$

Note that for any $m \geq 1$ and $v \in V$, we have that W_v can increase by at most one after each step; i.e., $W_v^{(m)} \in \{W_v^{(m-1)}, W_v^{(m-1)} + 1\}$. The lemma below establishes that a ball cannot be allocated to a vertex with larger weight than the vertex where the ball is born.

Lemma 3.1. *Let $m \geq 1$ and denote by v the vertex where ball m is born (i.e., $v = U_m$). Let v' be the vertex where ball m is allocated. Then, $W_{v'}^{(m-1)} \leq W_v^{(m-1)}$.*

Proof. Assume that $v \neq v'$, thus the local search of ball m visits at least two vertices. Let w be the second vertex visited during the local search. Since v and w are neighbors in G , we have

$$W_w^{(m-1)} = X_w^{(m-1)} + \mu(w) = X_v^{(m-1)} - 1 + \mu(w) \leq X_v^{(m-1)} + \mu(v) = W_v^{(m-1)}.$$

Proceeding inductively for each step of the local search we obtain $W_{v'}^{(m-1)} \leq W_v^{(m-1)}$. $\square$

Now we use the definition of majorization from (2.2). Let $\overline{W}_v^{(m)}$ be the weight of vertex v after m balls are allocated according to the 1-choice process; i.e., $\overline{W}_v^{(m)} = \overline{X}_v^{(m)} + \mu(v)$ for all $v \in V$. The lemma below extends the result of Lemma 2.1 to the weights of the vertices; Lemma 2.1 can be obtained from Lemma 3.2 by setting $\mu(v) = 0$ for all $v \in V$.

Lemma 3.2. *For any fixed $m \geq 0$, we can couple $W^{(m)}$ and $\overline{W}^{(m)}$ so that, with probability 1, $\overline{W}^{(m)}$ majorizes $W^{(m)}$.*

For the proof of this lemma, we need the following result from [3].

Lemma 3.3 ([3, Lemma 3.4]). *Let $v = (v_1, v_2, \dots, v_n)$, $u = (u_1, u_2, \dots, u_n)$ be two vectors such that $v_1 \geq v_2 \geq \dots \geq v_n$ and $u_1 \geq u_2 \geq \dots \geq u_n$. If v majorizes u , then also $v + e_i$ majorizes $u + e_i$, where e_i is the i th unit vector.*

Proof of Lemma 3.2. The proof is by induction on m . Clearly, for $m = 0$, we have $W_v^{(0)} = \overline{W}_v^{(0)} = \mu(v)$ for all $v \in V$. Now, assume that we can couple $W^{(m-1)}$ with $\overline{W}^{(m-1)}$ so that $\overline{W}^{(m-1)}$ majorizes $W^{(m-1)}$. Let $i_1, i_2, \dots, i_n$ be distinct elements of V so that

$$W_{i_1}^{(m-1)} \geq W_{i_2}^{(m-1)} \geq \dots \geq W_{i_n}^{(m-1)}.$$

Similarly, let $j_1, j_2, \dots, j_n$ be distinct elements of V so that

$$\overline{W}_{j_1}^{(m-1)} \geq \overline{W}_{j_2}^{(m-1)} \geq \dots \geq \overline{W}_{j_n}^{(m-1)}.$$

Let ℓ be a uniformly random integer from 1 to n . Then, for the process $(W_v^{(m)})_{v \in V}$, let the birthplace of ball m be vertex i_ℓ and for the process $(\overline{W}_v^{(m)})_{v \in V}$, let the birthplace of ball m be j_ℓ . For the process $(W_v^{(m)})_{v \in V}$, ball m may not necessarily be allocated at vertex i_ℓ , so let us define ι as the integer so that i_ι is the vertex to which ball m is allocated.

In order to prove that $\overline{W}^{(m)}$ majorizes $W^{(m)}$, let us define by $\widetilde{W}^{(m)}$ the load vector which is obtained from $W^{(m-1)}$ by allocating ball m to vertex i_ℓ (the birthplace of ball m). Applying Lemma 3.3 gives that $\overline{W}^{(m)}$ majorizes $\widetilde{W}^{(m)}$, since by the induction hypothesis $\overline{W}^{(m-1)}$ majorizes $W^{(m-1)}$. Next observe that

$$W^{(m)} = \widetilde{W}^{(m)} - e_{i_\ell} + e_{i_\iota},$$

so we obtain the vector $W^{(m)}$ from $\widetilde{W}^{(m)}$ by removing one from vertex i_ℓ and adding one to vertex i_ι . By Lemma 3.1, we have $W_{i_\iota}^{(m-1)} \leq W_{i_\ell}^{(m-1)}$. This implies $\widetilde{W}_{i_\ell}^{(m)} = W_{i_\ell}^{(m-1)} + 1 \geq W_{i_\iota}^{(m-1)} + 1$ and in turn that $\widetilde{W}^{(m)}$ majorizes $W^{(m)}$. Combining this with the insight that $\overline{W}^{(m)}$ majorizes $\widetilde{W}^{(m)}$ implies that $\overline{W}^{(m)}$ majorizes $W^{(m)}$. This completes the induction and the proof. $\square$

Now we illustrate the usefulness of the above result by relating the probability of a vertex to have a certain load with the probability that balls are born in a neighborhood around a vertex. Recall that the load vector is smooth (cf. Lemma 2.2), which means that if a vertex v has load ℓ , then a vertex at distance r from v has load at least $\ell - r$ and at most $\ell + r$.

Lemma 3.4. *For any $v \in V$, and any $\ell, m \geq 1$, we have*

$$\Pr \left[X_v^{(m)} \geq \ell \right] \geq \Pr \left[\bigcap_{w \in B_v^{\ell-1}} \left\{ \overline{X}_w^{(m)} \geq \ell - d_G(v, w) \right\} \right]$$

and

$$\Pr \left[X_v^{(m)} \geq \ell \right] \leq \Pr \left[\bigcup_{w \in V} \left\{ \overline{X}_w^{(m)} \geq \ell + d_G(v, w) \right\} \right].$$

Proof. For the first inequality, set $\mu(w) = d_G(v, w)$ for all $w \in V$. Let $\mathcal{A}^{(m)}$ be the event that all vertices have weight at least ℓ after m balls are allocated, and let $\overline{\mathcal{A}}^{(m)}$ be the same event for the 1-choice process. In symbols, $\mathcal{A}^{(m)} = \{\min_{u \in V} W_u^{(m)} \geq \ell\}$ and $\overline{\mathcal{A}}^{(m)} = \{\min_{u \in V} \overline{W}_u^{(m)} \geq \ell\}$. Note that $\mathcal{A}^{(m)}$ and $\overline{\mathcal{A}}^{(m)}$ are events defined on the *weights* of the vertices while the events in the statement of the lemma are defined on the *loads* of the vertices. By Lemma 3.2, we have that $\Pr[\mathcal{A}^{(m)}] \geq \Pr[\overline{\mathcal{A}}^{(m)}]$. Since $W_v^{(m)} = X_v^{(m)}$ for all m , we have $\mathcal{A}^{(m)} \subseteq \{X_v^{(m)} \geq \ell\}$, but the two events are in fact equal since, by the smoothness of the load vector (cf. Lemma 2.2), $\{X_v^{(m)} \geq \ell\}$ implies $\mathcal{A}^{(m)}$. The proof is then complete since $\overline{\mathcal{A}}^{(m)} = \bigcap_{w \in B_v^{\ell-1}} \{\overline{X}_w^{(m)} \geq \ell - d_G(v, w)\}$.

For the second inequality, set $\mu(w) = -d_G(v, w)$ for all $w \in V$. Then define $\mathcal{B}^{(m)}$ to be the event that there exists at least one vertex with weight at least ℓ after m balls are allocated, and let $\overline{\mathcal{B}}^{(m)}$ be the corresponding event for the 1-choice process. Thus, $\mathcal{B}^{(m)} = \{\max_{u \in V} W_u^{(m)} \geq \ell\}$ and $\overline{\mathcal{B}}^{(m)} = \{\max_{u \in V} \overline{W}_u^{(m)} \geq \ell\}$. Similarly as for the event $\mathcal{A}^{(m)}$, we have that the events $\{X_v^{(m)} \geq \ell\}$ and $\mathcal{B}^{(m)}$ are identical. Applying Lemma 3.2 we obtain that $\Pr[\mathcal{B}^{(m)}] \leq \Pr[\overline{\mathcal{B}}^{(m)}] = \Pr\left[\bigcup_{w \in V} \{\overline{X}_w^{(m)} \geq \ell + d_G(v, w)\}\right]$. $\square$

Remark 3.5. The lemma above states that one can *couple* $\{X_v^{(m)}\}_{v \in V}$ and $\{\overline{X}_v^{(m)}\}_{v \in V}$ so that if $\overline{X}_w^{(m)} \geq \ell - d_G(v, w)$ for all $w \in B_v^{\ell-1}$, then $X_v^{(m)} \geq \ell$. However, this is not necessarily achieved with the “trivial” coupling where each ball is born at the same vertex for both processes $\{X_v^{(m)}\}_{v \in V}$ and $\{\overline{X}_v^{(m)}\}_{v \in V}$. In other words, knowing that the number of balls born at vertex w is at least $\ell - d_G(v, w)$ for all $w \in B_v^\ell$ does *not* imply that $X_v^{(m)} \geq \ell$.

Now we extend the proof of Lemma 3.4 to derive an upper bound on the load of a subset of vertices.

Proposition 3.6. *Let $S \subset V$ be fixed and Δ be the maximum degree in G . Then, for all $m \geq n$ and $\ell \geq \frac{300\Delta m}{n}$ we have*

$$\Pr\left[\sum_{v \in S} X_v^{(m)} \geq \ell|S|\right] \leq 4 \exp\left(-\frac{|S|\ell}{14} \log\left(\frac{\ell n}{m}\right)\right) + \exp\left(-\frac{m}{4}\right).$$

The above inequality implies that, for any given $u \in V$,

$$\Pr\left[X_u^{(m)} \geq 2\ell\right] \leq 4 \exp\left(-\frac{|B_u^\ell|\ell}{14} \log\left(\frac{\ell n}{m}\right)\right) + \exp\left(-\frac{m}{4}\right).$$

Proof. For any $v \in V$, define $\mu(v) = -d_G(v, S)$ and (cf. (3.2)) $W_v^{(m)} = X_v^{(m)} + \mu(v)$. Let $K_S^{(m)}$ be the sum of the weights of the $|S|$ vertices with largest weights after m balls are allocated, and $\overline{K}_S^{(m)}$ be the corresponding value for the 1-choice process. Then,

$$\sum_{v \in S} X_v^{(m)} = \sum_{v \in S} W_v^{(m)} \leq K_S^{(m)} \leq \overline{K}_S^{(m)},$$

where the last step follows by majorization (cf. Lemma 3.2). Let $\widehat{W}_v^{(k)}$ be the weight of vertex v for the Poissonized version of the 1-choice process with expected number of balls equal to k , and $\widehat{K}_S^{(k)}$ be the sum of the weights of the $|S|$ vertices with largest weight for this Poissonized version. If the Poissonized version with $k = 2m$ allocates at least m balls, then we can couple the allocations of the first m balls with the allocation in the non-Poissonized version of the 1-choice process, and it holds that $\widehat{K}_S^{(2m)} \geq \overline{K}_S^{(m)}$. Hence by the first statement of Lemma A.4 we have that

$$\Pr\left[\widehat{K}_S^{(2m)} \geq \overline{K}_S^{(m)}\right] \geq 1 - \exp\left(-\frac{m}{4}\right). \quad (3.3)$$

From now on, we consider only the Poissonized version. Let $\tilde{K}^{(2m)}$ be the sum of the weights of the vertices with weight at least $\ell/16$. More formally, $\tilde{K}^{(2m)} = \sum_{v \in V: \widehat{W}_v^{(2m)} \geq \ell/16} \widehat{W}_v^{(2m)}$. Then, we have that, on the event $\widehat{K}_S^{(2m)} \geq \overline{K}_S^{(m)}$,

$$\sum_{v \in S} X_v^{(m)} \leq \widehat{K}_S^{(2m)} \leq \frac{\ell}{16}|S| + \tilde{K}^{(2m)}.$$

We can estimate the weight of vertices that reach weight $\ell/16$ as follows. For each vertex, let balls arrive according to a rate-1 Poisson point process up to time $2m/n$ or until the vertex reaches weight $\ell/16$, whatever happens first. Then, if the vertex reaches weight $\ell/16$, continue adding balls for an additional time interval of length $2m/n$. This construction stochastically dominates the weight of the vertices by the memoryless property of Poisson random variables. The probability that a vertex v with $\mu(v) = -k$ reaches weight $\ell/16$ is

$$\sum_{x=\ell/16+k}^{\infty} \frac{e^{-2m/n}(2m/n)^x}{x!} \leq \sum_{x=\ell/16+k}^{\infty} \left(\frac{2me}{nx}\right)^x \leq 2 \left(\frac{2me}{n(\ell/16+k)}\right)^{\ell/16+k},$$

since $\frac{2me}{n(\ell/16+k)} \leq \frac{1}{2}$ for all $k \geq 0$ and $x! \geq (x/e)^x$ for any integer x . Now any Bernoulli random variable with mean $p \leq 1/2$ is stochastically dominated by a Poisson random variable with mean $2p$, which follows from the fact that $e^{-2p} \leq 1 - p$ for $0 \leq p \leq 1/2$. Using this, and denoting by N_S^k the set of vertices at distance k from S , we have that the number of vertices reaching weight $\ell/16$ is a Poisson random variable of mean

$$\begin{aligned} \sum_{k \geq 0} |N_S^k| 4 \left(\frac{2me}{n(\ell/16+k)}\right)^{\ell/16+k} &\leq 4|S| \left(\frac{32me}{n\ell}\right)^{\ell/16} \sum_{k \geq 0} \Delta^k \left(\frac{2me}{n(\ell/16+k)}\right)^k \\ &\leq 8|S| \left(\frac{32me}{n\ell}\right)^{\ell/16}, \end{aligned}$$

for large enough $\ell \geq 300\Delta m/n$. Then the probability that the number of vertices reaching weight $\ell/16$ is larger than $8|S|$ is at most

$$\begin{aligned} \sum_{k \geq 8|S|} \left(\frac{8e|S| \left(\frac{32me}{n\ell}\right)^{\ell/16}}{k}\right)^k &\leq 2 \left(\frac{8e|S| \left(\frac{32me}{n\ell}\right)^{\ell/16}}{8|S|}\right)^{8|S|} \\ &= 2 \exp \left(-8|S| \left(\frac{\ell}{16} \log \left(\frac{\ell n}{32me}\right) - 1\right)\right), \end{aligned}$$

since $\frac{8e|S| \left(\frac{32me}{n\ell}\right)^{\ell/16}}{8|S|} = e \left(\frac{32me}{n\ell}\right)^{\ell/16} \leq \frac{1}{2}$. Using that $\ell \geq 300\Delta m/n$, we have

$$\frac{\ell}{16} \log \left(\frac{\ell n}{32me}\right) - 1 \geq \frac{\ell}{32} \log \left(\frac{\ell n}{m}\right) - 1 \geq \frac{\ell}{96} \log \left(\frac{\ell n}{m}\right).$$

Putting the last two equations together, we obtain that

$$\mathbf{Pr} \left[\text{more than } 8|S| \text{ vertices reach weight } \frac{\ell}{16} \right] \leq 2 \exp \left(-\frac{\ell|S|}{12} \log \left(\frac{\ell n}{m}\right)\right). \quad (3.4)$$

If the event above occurs, then $\tilde{K}^{(2m)}$ is stochastically dominated by $8|S| \cdot \frac{\ell}{16} = \frac{\ell|S|}{2}$ plus a Poisson random variable of mean $8|S| \cdot \frac{2m}{n} = 16|S| \frac{m}{n}$, which is larger than $\frac{7\ell|S|}{16}$ with probability

at most

$$\begin{aligned} \sum_{k=\frac{7\ell|S|}{16}}^{\infty} \left(\frac{16e|S|m}{nk} \right)^k &\leq 2 \left(\frac{16 \cdot 16e|S|m}{7n\ell|S|} \right)^{\frac{7\ell|S|}{16}} \\ &= 2 \exp \left(-\frac{7\ell|S|}{16} \log \left(\frac{7\ell n}{256em} \right) \right) \leq 2 \exp \left(-\frac{7\ell|S|}{96} \log \left(\frac{\ell n}{m} \right) \right). \end{aligned} \quad (3.5)$$

Therefore, by summing the right-hand sides of (3.4) and (3.5), with probability at least $1 - 4 \exp \left(-\frac{\ell|S|}{14} \log \left(\frac{\ell n}{m} \right) \right)$, we have $\tilde{K}^{(2m)} \leq \frac{\ell|S|}{2} + \frac{7\ell|S|}{16} \leq \frac{15\ell|S|}{16}$. This and the fact that $\hat{K}_S^{(2m)} \leq \frac{\ell|S|}{16} + \tilde{K}^{(2m)}$, together with (3.3), establish the first part of the lemma.

The second part of Proposition 3.6 holds by setting $S = B_u^\ell$. If u has load $k > \ell$, then the total number of balls allocated to B_u^ℓ is at least

$$\sum_{i=0}^{\ell} (k-i) |N_u^i| \geq (k-\ell) |B_u^\ell|.$$

Then setting $k = 2\ell$ and applying the first part of the proposition yields the result. $\square$

4 Maximum Load

We start stating a stronger version of Theorem 1.1 which also holds for non-transitive graphs. For $\gamma \in (0, 1/2]$, let

$$R_1^{(\gamma)} = R_1^{(\gamma)}(G) = \max \left\{ r \in \mathbb{N} : \text{there exists } S \subseteq V \text{ with } |S| \geq n^{\frac{1}{2}+\gamma} \right. \\ \left. \text{such that } r|B_u^r| \log r < \log n \text{ for all } u \in S \right\}.$$

Note that $R_1^{(\gamma)}$ is non-increasing with γ , and R_1 as defined in (1.1) is larger than $R_1^{(\gamma)}$ for any $\gamma \geq 0$. Also, when G is vertex transitive, we have $R_1 = R_1^{(\gamma)} + 1$ for all $\gamma \in (0, 1/2]$, because in this case, for any given r , the size of B_u^r is the same for all $u \in V$. Intuitively, $R_1^{(\gamma)}$ can be seen as the “lower bound counterpart” of R_1 which, as mentioned earlier, relates to the (worst-case) local neighborhood growth in G . The condition $|S| \geq n^{1/2+\gamma}$ ensures that there are sufficiently many vertices for which the neighborhood growth can be bounded *uniformly*. Lower bounding the size of S is necessary, since otherwise G could consist of a small part which has low local neighborhood growth, whereas the remaining part has high local neighborhood growth.

The theorem below establishes that, for any bounded-degree graph, if there exists a $\gamma \in (0, 1/2]$ for which $R_1^{(\gamma)} = \Theta(R_1)$, then the maximum load is $\Theta(R_1)$, when $m = n$.

Theorem 4.1 (General version of Theorem 1.1). *Let G be any graph with bounded degrees. For any constants $\gamma \in (0, 1/2]$ and $\alpha \geq 1$, we have*

$$\mathbf{Pr} \left[X_{\max}^{(n)} < \frac{\gamma R_1^{(\gamma)}}{4} \right] \leq n^{-\omega(1)} \quad \text{and} \quad \mathbf{Pr} \left[X_{\max}^{(n)} \geq 56\alpha R_1 \right] \leq 5n^{-\alpha}.$$

Proof. We start establishing a lower bound for $X_{\max}^{(n)}$. We first consider the Poissonized versions of the local search allocation and the 1-choice process (recall the definition of these variants from the paragraph preceding Lemma 2.5). We abuse notation slightly, and let $X_v^{(m)}$ and $\bar{X}_v^{(m)}$ denote the load of v for the Poissonized version of the local search allocation and 1-choice process, respectively, when the expected number of balls allocated in total is m . Note that since Lemma 3.4 works for all m , it also works for the Poissonized version by conditioning on

the total number of balls. Let A be a Poisson random variable with mean 1. For any positive integer z , we have that

$$\Pr[A \geq z] = \sum_{y=z}^{\infty} e^{-1} \frac{1}{y!} \geq \sum_{y=z}^{\infty} e^{-1} y^{-y} \geq e^{-1} z^{-z}.$$

For any $v \in V$ and any $\ell > 0$, Lemma 3.4 and the above inequality give that

$$\Pr[X_v^{(n)} \geq \ell] \geq \prod_{r=0}^{\ell-1} (\Pr[A \geq \ell - r])^{|N_v^r|} \geq \prod_{r=0}^{\ell-1} (e^{-1}(\ell - r)^{-\ell+r})^{|N_v^r|} \geq (e^{-1}\ell^{-\ell})^{\sum_{r=0}^{\ell-1} |N_v^r|},$$

where N_v^r is the set of vertices at distance r from v , and the last inequality follows since the function z^{-z} is decreasing for all $z \geq 1$. Since $B_v^\ell = \bigcup_{r=0}^{\ell} N_v^r$, **Did we define N_v^r before?**

$$\Pr[X_v^{(n)} \geq \ell] \geq \exp(-|B_v^\ell| - \ell|B_v^\ell| \log(\ell)) \geq \exp(-2\ell|B_v^\ell| \log(\ell)),$$

where the last step follows for all $\ell \geq 2$. Given $\gamma > 0$, set $\ell = \frac{\gamma R_1^{(\gamma)}}{4}$. Hence, since $|B_v^r| \log r$ is increasing with r , we have that there exists a set S with $|S| = \lceil n^{\frac{1}{2}+\gamma} \rceil$ such that

$$\Pr\left[X_v^{(n)} \geq \frac{\gamma R_1^{(\gamma)}}{4}\right] \geq \exp\left(-\frac{\gamma R_1^{(\gamma)} |B_v^{R_1^{(\gamma)}}| \log(R_1^{(\gamma)})}{2}\right) \geq n^{-\gamma/2} \quad \text{for all } v \in S. \quad (4.1)$$

Let $Y = Y(\gamma)$ be the random variable defined as the number of vertices v satisfying $X_v^{(n)} \geq \frac{\gamma R_1^{(\gamma)}}{4}$. Let K be the total number of balls allocated in the Poissonized version of the local search allocation. Note that $\mathbf{E}[K] = n$ and by the last statement of Lemma A.4, $\Pr[K > 2en] \leq 2^{1-2ne}$. Regard Y as a function of the K independently chosen birthplaces $U_1, U_2, \dots, U_K$. Then, for any given K , Y is 1-Lipschitz by Lemma 2.3

$$\begin{aligned} \mathbf{E}[Y \mid K \leq 2en] &\geq \sum_{v \in S} \Pr\left[X_v^{(n)} \geq \frac{\gamma R_1^{(\gamma)}}{4} \mid K \leq 2en\right] \\ &\geq \sum_{v \in S} \frac{\Pr\left[X_v^{(n)} \geq \frac{\gamma R_1^{(\gamma)}}{4}\right] - \Pr[K > 2en]}{\Pr[K \leq 2en]}. \end{aligned}$$

Applying (4.1) to the above inequality yields

$$\mathbf{E}[Y \mid K \leq 2en] \geq |S| \cdot \left(\frac{n^{-\gamma/2} - \Pr[K > 2en]}{\Pr[K \leq 2en]} \right) \geq \frac{n^{\frac{1}{2}+\gamma}}{2}.$$

With this, we apply Lemma A.1 to obtain

$$\begin{aligned} &\Pr\left[X_{\max}^{(n)} < \frac{\gamma R_1^{(\gamma)}}{4}\right] \\ &= \Pr[Y = 0] \\ &\leq \Pr[Y = 0 \mid K \leq 2en] + \Pr[K > 2en] \\ &\leq \Pr\left[|Y - \mathbf{E}[Y \mid K \leq 2en]| \geq \frac{1}{2} \mathbf{E}[Y \mid K \leq 2en] \mid K \leq 2en\right] + \Pr[K > 2en] \\ &\leq 2 \exp\left(-\frac{n^{1+\gamma}}{2 \cdot 2en}\right) + 2^{1-2ne} = n^{-\omega(1)}. \end{aligned}$$

This result can then be translated to the non-Poissonized model via Lemma 2.5.

Now we establish the upper bound, where we consider the non-Poissonized process. For any fixed $u \in V$, we have from the second part of Proposition 3.6 (with $m = n$) that

$$\begin{aligned} \Pr \left[X_u^{(n)} \geq 56\alpha R_1 \right] &\leq 4 \exp \left(-\frac{28\alpha R_1 |B_u^{28\alpha R_1}|}{14} \log(28\alpha R_1) \right) + \exp \left(-\frac{n}{4} \right) \\ &\leq 4 \exp \left(-2\alpha R_1 |B_u^{R_1}| \log R_1 \right) + \exp \left(-\frac{n}{4} \right) \leq 5n^{-2\alpha}. \end{aligned}$$

Taking the union bound over u we obtain that

$$\Pr \left[X_{\max}^{(n)} \geq 56\alpha R_1 \right] \leq 5n^{-2\alpha+1} \leq 5n^{-\alpha}.$$

□

Proof of Theorem 1.2. Applying Proposition 3.6 with $\ell = \left(\frac{m}{n} + R_2\right) c$ for any constant $c \geq 300\Delta$, we obtain

$$\begin{aligned} \Pr \left[\sum_{u \in B_u^{R_2}} X_u^{(m)} \geq \left(\frac{m}{n} + R_2\right) c \cdot |B_u^{R_2}| \right] \\ \leq 4 \exp \left(-\left(\frac{m}{n} + R_2\right) \frac{c |B_u^{R_2}|}{14} \log c \right) + \exp \left(-\frac{m}{4} \right) \\ \leq 4 \exp \left(-\frac{c R_2 |B_u^{R_2}|}{14} \log c \right) + \exp \left(-\frac{m}{4} \right). \end{aligned}$$

By setting $c \geq 1$ sufficiently large, the right-hand side above can be made smaller than n^{-2} . If u has load k , then the number of balls allocated to vertices in $B_u^{R_2}$ is at least

$$\sum_{i=0}^{R_2} (k-i) |N_u^i| \geq (k-R_2) |B_u^{R_2}|.$$

Therefore we obtain that, on the event $\sum_{u \in B_u^{R_2}} X_u^{(m)} \leq \left(\frac{m}{n} + R_2\right) c |B_u^{R_2}|$, we have $X_u^{(m)} \leq c \left(\frac{m}{n} + R_2\right) + R_2 \leq 2c \left(\frac{m}{n} + R_2\right)$. Taking the union bound over all u completes the proof. □

5 Cover time

Our next result will be a lower bound on the minimum load, which in turn gives an upper bound for the cover time. Given that for the cover time we only need to require all vertices to have a nonzero load, it may seem a bit crude that in the proposition below we require all vertices to have load at least $\frac{m}{224n \log \Delta}$ after m balls have been allocated. However, we believe that the cover time and the blanket time (which is the first time until all vertices have load of order m/n) are in fact of the same order for many natural graph classes including vertex-transitive graphs (cf. the discussion in Section 6).

Proposition 5.1. *Let G be a graph with bounded degrees. Then for any $\alpha > 1$ there exists a $C = C(\alpha, \Delta) > 0$ such that for all $m \geq CR_2 n$ we have*

$$\Pr \left[X_{\min}^{(m)} < \frac{m}{224n \log \Delta} \right] \leq n^{-\alpha},$$

where $X_{\min}^{(m)} = \min_{v \in V} X_v^{(m)}$.

Proof. Fix an arbitrary vertex $u \in V$. We will use the concept of weights defined in Section 3. Define $\mu(v) = d_G(u, v)$ and $W_v^{(m)} = X_v^{(m)} + \mu(v)$. Similarly, for the 1-choice process, define $\overline{W}_v^{(m)} = \overline{X}_v^{(m)} + \mu(v)$. Let $Y := \min_{v \in V} \overline{W}_v^{(m)}$ be the minimum weight of all vertices in V in the 1-choice process. Let

$$\ell = \frac{m}{28n \log \Delta}. \quad (5.1)$$

We have

$$\begin{aligned} \Pr[Y < \ell] &= \Pr \left[\bigcup_{v \in B_u^{\ell-1}} \left\{ \overline{W}_v^{(m)} < \ell \right\} \right] \\ &\leq |B_u^{\ell-1}| \Pr \left[\overline{X}_u^{(m)} < \ell \right] \\ &\leq |B_u^{\ell-1}| \Pr \left[\left| \overline{X}_u^{(m)} - \mathbf{E} \left[\overline{X}_u^{(m)} \right] \right| > \frac{m}{n} \left(1 - \frac{1}{28 \log \Delta} \right) \right]. \end{aligned}$$

Using Lemma A.3, we obtain

$$\Pr[Y < \ell] \leq |B_u^\ell| \exp \left(-\frac{\frac{m^2}{n^2} \left(1 - \frac{1}{28 \log \Delta} \right)^2}{\frac{7m}{3n}} \right) \leq |B_u^\ell| \exp \left(-\frac{3m}{28n} \right) \leq \exp \left(\frac{m}{28n} - \frac{3m}{28n} \right) \leq \frac{1}{2},$$

where the last inequality holds since $m/n \geq CR_2 = \omega(1)$ for bounded degree graphs. Now define $\overline{Z}$ as the sum of the $|B_u^{R_2}|$ smallest values of $\{\overline{W}_v^{(m)} : v \in V\}$ and Z as the sum of the $|B_u^{R_2}|$ smallest values of $\{W_v^{(m)} : v \in V\}$. By Lemma 3.2, we can couple $W^{(m)}$ and $\overline{W}^{(m)}$ so that, with probability 1, $Z \geq \overline{Z}$. Further,

$$\mathbf{E}[\overline{Z}] \geq \Pr[Y \geq \ell] \cdot \ell |B_u^{R_2}| \geq \frac{\ell |B_u^{R_2}|}{2}.$$

We now apply a variant of Azuma's inequality for martingales (Lemma A.2) in order to show that $\overline{Z}$ is likely to be at least $\frac{\ell |B_u^{R_2}|}{4}$. Let $A_0 = \mathbf{E}[\overline{Z}]$ and define $A_1, A_2, \dots, A_m$ to be the martingale adapted to the filtration $\mathcal{F}_i$ generated by $U_1, U_2, \dots, U_i$; i.e., $A_i = \mathbf{E}[\overline{Z} \mid \mathcal{F}_i]$.

We first establish that the martingale A_i satisfies the 1-Lipschitz condition. Changing the birthplace of ball i and keeping all other birthplaces the same can change Z by at most one, cf. Lemma 2.3. In symbols, this means

$$\mathbf{E}[|A_i - A_{i-1}| \mid \mathcal{F}_{i-1}] \leq 1.$$

In the following we will derive an upper bound on the second moment of $A_i - A_{i-1}$ exploiting the simple fact that the set of the $|B_u^{R_2}|$ smallest loads can only change if the uniformly chosen bin U_i is in $B_u^{R_2}$. Now fix $i \in \{1, 2, \dots, m\}$ and fix $U_1, U_2, \dots, U_{i-1}$. Let ζ_u be the value of A_i when $U_i = u$ and let $\bar{\zeta} = \frac{1}{n} \sum_{u \in V} \zeta_u$. Then we have

$$\mathbf{E}_{U_i} \left[(A_i - A_{i-1})^2 \mid \bigcap_{j=1}^{i-1} \{U_j = u_j\} \right] = \frac{1}{n} \sum_{u \in V} (\zeta_u - \bar{\zeta})^2,$$

where the expectation above is taken with respect to U_i . Since $|\zeta_u - \zeta_{u'}| \leq 1$ for all $u, u' \in V$, we can write

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

Now consider a given realization of $U_1, U_2, \dots, U_{i-1}, U_{i+1}, \dots, U_m$, and let $\Gamma \subset V$ be the set of $|B_u^{R_2}|$ vertices with smallest loads (note that we skip the i th ball in the definition of Γ). Then, by adding the i th ball, ζ_u and $\zeta_{u'}$ only differ if at least one of u or u' is in Γ . Hence, $\sum_{u \in V} \sum_{u' \in V} |\zeta_u - \zeta_{u'}| \leq 2|B_u^{R_2}|n$. Consequently,

$$\mathbf{E}_{U_i} \left[(A_i - A_{i-1})^2 \mid \bigcap_{j=1}^{i-1} \{U_j = u_j\} \right] \leq \frac{2|B_u^{R_2}|}{n}.$$

Now, Lemma A.2 gives

$$\Pr \left[\bar{Z} < \frac{\ell|B_u^{R_2}|}{4} \right] \leq \Pr \left[|\bar{Z} - \mathbf{E}[\bar{Z}]| \geq \frac{1}{2} \mathbf{E}[\bar{Z}] \right] \leq \exp \left(- \frac{(\frac{1}{2} \mathbf{E}[\bar{Z}])^2}{4 \cdot \frac{|B_u^{R_2}|}{n} \cdot m + \frac{1}{6} \mathbf{E}[\bar{Z}]} \right).$$

Clearly, $\mathbf{E}[\bar{Z}] \leq \frac{m|B_u^{R_2}|}{n}$, which gives that

$$\Pr \left[\bar{Z} < \frac{\ell|B_u^{R_2}|}{4} \right] \leq \exp \left(- \frac{\mathbf{E}[\bar{Z}]^2}{16 \cdot \frac{|B_u^{R_2}|}{n} \cdot m + \frac{2m|B_u^{R_2}|}{3n}} \right) \leq \exp \left(- \frac{\ell^2|B_u^{R_2}|/4}{17m/n} \right).$$

Using the value of ℓ from (5.1) and using that $m \geq CR_2n$ from the condition in the statement of the proposition, we have

$$\Pr \left[\bar{Z} < \frac{\ell|B_u^{R_2}|}{4} \right] \leq \exp \left(- \frac{\frac{m}{n}|B_u^{R_2}|}{68(28 \log \Delta)^2} \right) \leq \exp \left(- \frac{CR_2|B_u^{R_2}|}{68(28 \log \Delta)^2} \right) \leq n^{-\frac{C}{68(28 \log \Delta)^2}}.$$

Due to our coupling which gives $Z \geq \bar{Z}$ we conclude that with probability at least $1 - n^{-\frac{C}{68(28 \log \Delta)^2}}$ there exists a vertex $v \in B_u^{R_2}$ with $W_v^{(m)} \geq \frac{\ell}{4}$ and thus $X_v^{(m)} \geq \frac{\ell}{4} - R_2$. Then, by smoothness of the load vector (cf. Lemma 2.2), we have that with probability at least $1 - n^{-\frac{C}{68(28 \log \Delta)^2}}$, every vertex in $B_u^{R_2}$ has load at least $\frac{\ell}{4} - 3R_2 \geq \frac{m}{224n \log \Delta}$, where the last step follows for all $C \geq 672 \log \Delta$. Then the result follows by taking the union bound over all $u \in V$, which gives that with probability at least $1 - n^{-\frac{C}{68(28 \log \Delta)^2} + 1}$, all vertices have load at least $\frac{m}{224n \log \Delta}$. The proof is then completed by setting C large enough with respect to α so that $\frac{C}{68(28 \log \Delta)^2} - 1 \geq \alpha$. $\square$

We prove a stronger version of Theorem 1.3, which holds also for non-transitive graphs. For $\gamma \in (0, 1/2]$, let

$$R_2^{(\gamma)} = R_2^{(\gamma)}(G) = \max \left\{ r \in \mathbb{N} : \text{there exists } S \subseteq V \text{ with } |S| \geq n^{\frac{1}{2} + \gamma} \right. \\ \left. \text{such that } r|B_u^r| < \log n \text{ for all } u \in S \right\}.$$

Note that $R_2^{(\gamma)}$ is non-increasing with γ . Also, when G is vertex transitive, we have $R_2 = R_2^{(\gamma)} + 1$ for all $\gamma \in (0, 1/2]$, because in this case, for any given r , the size of B_u^r is the same for all $u \in V$. The theorem below establishes that, for any bounded-degree graph, if there exists a $\gamma \in (0, 1/2]$ for which $R_2^{(\gamma)} = \Theta(R_2)$, then the cover time is $\Theta(R_2)$.

Theorem 5.2 (General version of Theorem 1.3). *Let G be any graph with bounded degrees. For any $\gamma \in (0, 1/2]$ and $\alpha \geq 1$, there exists $C = C(\alpha, \Delta)$ so that*

$$\Pr \left[T_{\text{cov}} < \frac{\gamma R_2^{(\gamma)} n}{8\Delta} \right] \leq n^{-\omega(1)} \quad \text{and} \quad \Pr [T_{\text{cov}} \geq CR_2n] \leq n^{-\alpha}.$$

Proof. The second inequality is established by Proposition 5.1. For the first inequality, let S be a set of $n^{\frac{1}{2}+\gamma}$ vertices u for which $R_2^{(\gamma)} \cdot |B_u^{R_2^{(\gamma)}}| < \log n$. Let

$$m = \frac{\gamma R_2^{(\gamma)} n}{8\Delta}. \quad (5.2)$$

We consider the Poissonized version of the local search allocation and the 1-choice process. We abuse notation slightly and let $X_v^{(m)}$ and $\bar{X}_v^{(m)}$ denote the load of v for the Poissonized version of the local search allocation and 1-choice process, respectively, when the expected number of balls allocated in total is m . For any $u \in S$, we will bound the probability that $X_u^{(m)} = 0$. First note that since Lemma 3.4 works for all m , it also works for the Poissonized version by conditioning on the total number of balls. Applying the second part of Lemma 3.4, we have that

$$\Pr[X_u^{(m)} = 0] \geq \Pr\left[\bigcap_{w \in V} \{\bar{X}_w^{(m)} \leq d_G(u, w)\}\right].$$

Recall that N_u^r is the set of vertices at distance r from u and $B_u^\ell = \bigcup_{r=0}^\ell N_u^r$. By independence of the Poissonized model, we can write

$$\begin{aligned} \Pr[X_u^{(m)} = 0] &\geq \Pr\left[\bigcap_{w \in B_u^{R_2^{(\gamma)}}} \{\bar{X}_w^{(m)} = 0\}\right] \Pr\left[\bigcap_{i > R_2^{(\gamma)}} \bigcap_{w \in N_u^i} \{\bar{X}_w^{(m)} \leq i\}\right] \\ &\geq \exp\left(-\frac{m|B_u^{R_2^{(\gamma)}}|}{n}\right) \left(1 - \sum_{i > R_2^{(\gamma)}} \sum_{w \in N_u^i} \Pr[\bar{X}_w^{(m)} > i]\right) \\ &\geq \exp\left(-\frac{m|B_u^{R_2^{(\gamma)}}|}{n}\right) \left(1 - 2 \sum_{i > R_2^{(\gamma)}} \sum_{w \in N_u^i} \left(\frac{me}{ni}\right)^i\right), \end{aligned}$$

where the last inequality follows by the last statement of Lemma A.4. Using the simple bound $|N_u^i| \leq \Delta^i$ and the value of m from (5.2), which implies that $\frac{me\Delta}{ni} \leq \frac{1}{2}$ for all $i \geq R_2^{(\gamma)}$ (as $\Delta/R_2^{(\gamma)} = o(1)$ since $\Delta = \mathcal{O}(1)$), we have

$$\Pr[X_u^{(m)} = 0] \geq \exp\left(-\frac{m|B_u^{R_2^{(\gamma)}}|}{n}\right) \left(1 - 4 \left(\frac{me\Delta}{nR_2^{(\gamma)}}\right)^{R_2^{(\gamma)}}\right) \geq n^{-\frac{\gamma}{8\Delta}} \cdot \frac{1}{2} \geq n^{-\frac{\gamma}{8}} \cdot \frac{1}{2}.$$

Now let Y be the random variable defined as the number of vertices $v \in S$ satisfying $X_v^{(m)} = 0$. Let K be the random variable for the total number of balls allocated and regard Y as a function of the K independently chosen birthplaces $U_1, U_2, \dots, U_K$. Then, Y is 1-Lipschitz by Lemma 2.3 for any given K . The calculations above give that

$$\mathbf{E}[Y \mid K \leq 2em] \geq \mathbf{E}[Y] \geq \frac{n^{\frac{1}{2}+\frac{7\gamma}{8}}}{2}.$$

Note that $m = \frac{\gamma R_2^{(\gamma)} n}{8\Delta} = \mathcal{O}(n \log n)$ for any G . With this, we apply Lemma A.1 and the last statement of Lemma A.4 to obtain

$$\begin{aligned} \Pr[X_{\min}^{(n)} = 0] &\leq \Pr\left[\{|Y - \mathbf{E}[Y \mid K \leq 2em]| \geq \frac{1}{2} \mathbf{E}[Y \mid K \leq 2em]\} \mid \{K \leq 2em\}\right] + \Pr[K > 2em] \\ &\leq 2 \exp\left(-\frac{n^{1+14\gamma/8}}{8(2em)}\right) + 2^{1-2me} = n^{-\omega(1)}. \end{aligned}$$

This result can then be translated to the non-Poissonized process using Lemma 2.5 and the fact that $m = \mathcal{O}(n \log n)$. $\square$

We now state and prove a stronger version of Theorem 1.4.

Theorem 5.3 (General version of Theorem 1.4). *Let G be any d -regular graph. Then, for any $\alpha > 1$ there exists $C = C(\alpha) > 0$ such that*

$$\Pr \left[T_{\text{cov}} \geq C \cdot \left(n \left(1 + \frac{\log n \cdot \log d}{d} \right) \right) \right] \leq n^{-\alpha}.$$

Proof. The result is shown by a coupling with the following stochastic process, introduced in [1], which we call *coupon collector process*. Initially, every node of G is uncovered. Then in each round i , a node $\tilde{U}_i$ is chosen independently and uniformly at random. If node $\tilde{U}_i$ is uncovered, then it becomes covered. Otherwise, if $\tilde{U}_i$ has any uncovered neighbor, then a random node among this set becomes covered. For this process, let us denote by $\tilde{C}^{(i)}$ the set of covered nodes after round i . We shall prove that there is a coupling so that for every round i , $\tilde{C}^{(i)} \subseteq \{v \in V : X_v^{(i)} \geq 1\}$; in other words, every node which is covered by the coupon collector process in round i is also covered by the local search allocation after the allocation of ball i .

The coupling is shown by induction. Clearly, the claim holds for $i = 1$. Consider now the execution of any round $i + 1$, assuming that the induction hypothesis holds for round i . In our coupling, we choose the same node v for $\tilde{U}_{i+1}$ and U_{i+1} .

In the first case, we assume that v is uncovered in the coupon collector process. Then the coupon collector process will cover node v in round $i + 1$. If v has not been covered by the local search allocation, then we have $X_v^{(i)} = 0$ and hence ball $i + 1$ will be allocated on node v in round $i + 1$. Otherwise, v has been covered previously. In either case, we conclude that node v is covered after round $i + 1$ in the local search allocation.

For the second case, suppose that v is covered in the coupon collector process. Then the coupon collector process will try to cover an uncovered neighbor of v if there exists one. This uncovered neighbor is chosen uniformly at random from all uncovered neighbors of v . This random experiment can be described by first choosing a random ranking of all $\deg(v)$ neighbors and then picking the uncovered neighbor with the highest rank, say node u . In our coupling, we assume that the local search allocation chooses the same ranking of all $\deg(v)$ neighbors. This, together, with the induction hypothesis, guarantees that if there is node u which becomes covered by the coupon collector process, then this node u also becomes covered by the local search allocation if it has not been covered in an earlier round.

Combining the two cases, we have shown that there is a coupling such that $\tilde{C}^{(i)} \subseteq \{v \in V : X_v^{(i)} \geq 1\}$ for any integer $i \geq 1$. Since it was shown for the coupon collector process in [1] that with probability $1 - n^{-c}$ for some constant $c > 0$, $O(n(1 + \frac{\log n \cdot \log d}{d}))$ rounds suffice to cover all nodes, the theorem follows. $\square$

6 Remarks and open questions

Blanket time

In analogy with the cover time for random walks, for each $\delta > 1$, we can define the blanket time as the first time at which the load of each vertex is in the interval $(\frac{1}{\delta} \cdot \frac{m}{n}, \delta \cdot \frac{m}{n})$. It follows from Theorem 1.2 and Proposition 5.1 that, for bounded-degree vertex-transitive graphs, the blanket time is $\Theta(nR_2)$ for all large enough δ .

Extreme graphs

Note that for any connected graph G , we have $R_1(G) \leq \sqrt{\frac{\log n}{\log \log n}}$ and $R_2(G) \leq \sqrt{\log n}$. Thus, the cycle is the graph with the largest possible maximum load (when $m = n$) and largest possible cover time among all bounded-degree graphs up to constant factors. Also, for any graph G

with bounded degrees, we have $R_1(G)$ and $R_2(G)$ of order $\Omega(\log \log n)$. Thus bounded-degree expanders are the graphs with the smallest maximum load (when $m = n$) and smallest cover time among all bounded-degree graphs up to constant factors.

Maximum load and birthplace

Consider the case $m = n$ and let $u_\star$ be the vertex at which the maximum number of balls are born. A simple lower bound for the maximum load can be obtained by considering only the balls that are born at $u_\star$. Then, by monotonicity (cf. Lemma 2.4), we may consider the scenario where $\frac{\log n}{\log \log n}$ balls are born at $u_\star$ and no other vertex is a birthplace of a ball. Since a ball born at a vertex with load β will be allocated within distance β , it follows that

$$X_{\max}^{(n)} \geq \min \left\{ \beta \geq 0 : \beta |B_{u_\star}^\beta| \geq \frac{\log n}{\log \log n} \right\}. \quad (6.1)$$

This is the strategy employed in [5, Lemma 2.4]. Below we show that this is not tight even for vertex-transitive graphs; in other words, the balls born at the single vertex $u_\star$ do not determine the maximum load.

Consider the graph G obtained via the Cartesian product of a vertex-transitive expander of size $n_0 = \frac{\log n}{(\log \log n)^3}$ and a cycle of size n/n_0 . Let r_0 be the diameter of the expander; we have $r_0 = \Theta(\log \log n)$. Then, for any vertex u of G and any $r \geq 1$, we have

$$(r - r_0)n_0 \leq |B_u^r| \leq (2r + 1)n_0. \quad (6.2)$$

Therefore, the bound in (6.1) gives that $X_{\max}^{(n)} = \Omega(\log \log n)$.

Now we estimate the maximum number of balls that are born in each expander graph of G . This can be done via a 1-choice process with $n' = \frac{n}{n_0}$ bins and $m' = n = \Theta\left(\frac{n' \log n'}{(\log \log n')^3}\right)$ balls, which using [10, Theorem 1] has maximum load

$$\Theta\left(\frac{\log n'}{\log \frac{n' \log n'}{m'}}\right) \geq \frac{\varepsilon \log n}{\log \log \log n},$$

for some constant $\varepsilon > 0$ and all large enough n . Let S denote the expander with the maximum number of balls born there. Then, using the same reasoning as in (6.1), and denoting by B_S^β the set of vertices within distance β from S , we have

$$\begin{aligned} X_{\max}^{(n)} &\geq \min \left\{ \beta \geq 0 : \beta |B_S^\beta| \geq \frac{\varepsilon \log n}{\log \log \log n} \right\} \\ &\geq \min \left\{ \beta \geq 0 : \beta(2\beta + 1) \frac{\log n}{(\log \log n)^3} \geq \frac{\varepsilon \log n}{\log \log \log n} \right\} \\ &= \Omega\left(\frac{(\log \log n)^{1.5}}{(\log \log \log n)^{0.5}}\right), \end{aligned}$$

where the first inequality follows from (6.2).

Open questions

1. For any vertex-transitive graph (not necessarily of bounded degrees), does it hold that $X_{\max}^{(n)} = \Theta(R_1)$ and $T_{\text{cov}} = \Theta(R_2 n)$ with high probability?
2. For any vertex-transitive graph (not necessarily of bounded degrees) and any $m = \omega(n R_2)$, does it hold that $X_{\max}^{(m)} = \frac{m}{n} + \Theta(R_2)$ with high probability?

3. For any vertex-transitive graph, is the blanket time of order nR_2 for all $\varepsilon \in (0, 1)$? In particular, is the blanket time of the same order as the cover time for all vertex-transitive graphs?
4. Let $G = (V, E)$ and $G' = (V, E')$ be two graphs such that $E \subset E'$. For any m , does it hold that the maximum load on G' is stochastically dominated by the maximum load on G ?

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. Bogdan, T. Sauerwald, A. Stauffer, and H. Sun. Balls into bins via local search. In *Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 16–34, 2013.
- [6] F. Chung and L. Lu. Concentration inequalities and Martingale inequalities: a survey. *Internet Mathematics*, 3:79–127, 2006.
- [7] 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.
- [8] C. McDiarmid. On the method of bounded differences. *Surveys in Combinatorics*, 141:148–188, 1989.
- [9] M. Mitzenmacher and E. Upfal. *Probability and Computing: randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.
- [10] 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.
- [11] B. Vöcking. How asymmetry helps load balancing. *J. ACM*, 50(4), 2003.

A Standard technical results

Lemma A.1 ([8, Lemma 1.2]). *Let $X_1, X_2, \dots, X_n$ be independent random variables with X_k taking values in a set Λ_k for each k . Suppose that the measurable function $f : \prod_{k=1}^n \Lambda_k \rightarrow \mathbb{R}$ satisfies for every k that*

$$|f(x) - f(x')| \leq c_k,$$

whenever the vectors x and x' differ only in the k th coordinate. Then for any $\lambda > 0$,

$$\Pr[|f - \mathbf{E}[f]| \geq \lambda] \leq 2 \cdot \exp\left(-\frac{2\lambda^2}{\sum_{k=1}^n c_k^2}\right).$$

Lemma A.2 ([6, 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,*

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

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 $\lambda > 0$,*

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

Lemma A.4 ([2, Theorem A.1.15]). *Let X have Poisson distribution. Then for any $0 < \varepsilon < 1$,*

$$\mathbf{Pr}[X \leq (1 - \varepsilon)\mathbf{E}[X]] \leq \exp\left(-\frac{\varepsilon^2 \mathbf{E}[X]}{2}\right).$$

Also, for any $x \geq 2e\mathbf{E}[X]$, it follows by Stirling's approximation that

$$\mathbf{Pr}[X \geq x] \leq 2 \left(\frac{\mathbf{E}[X]e}{x}\right)^x.$$