

Balls into bins via local search: cover time and maximum load*

Karl Bringmann¹, Thomas Sauerwald², Alexandre Stauffer³, and He Sun¹

1 Max Planck Institute for Informatics, Saarbrücken, Germany

2 University of Cambridge, UK

3 University of Bath, UK

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 $m \geq n$.

1998 ACM Subject Classification G.3 Mathematics of Computing: Probability and Statistics

Keywords and phrases Balls and Bins, Stochastic Process, Randomized Algorithm


Digital Object Identifier 10.4230/LIPIcs.STACS.2014.187

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. [2] 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. [3] 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

* Karl Bringmann is a recipient of the *Google Europe Fellowship in Randomized Algorithms*, and this research is supported in part by this Google Fellowship. The research of Alexandre Stauffer is supported in part by a Marie Curie Career Integration Grant PCIG13-GA-2013-618588 DSRELIS. The research of He Sun has partially been funded by the Cluster of Excellence “Multimodal Computing and Interaction” within the Excellence Initiative of the German Federal Government.



■ **Figure 1** Illustration of the local search allocation. Black circles represent the vertices 1–6 arranged as a path, and 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. 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. Figures (c) and (d) show the placement of balls $i + 1$ and $i + 2$.

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 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. [4] 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 time 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 the ball repeatedly moves to the adjacent vertex with the smallest load, provided that this load is strictly smaller than the load of its current vertex. 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 [4] 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 [4], 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 new upper and lower bounds for the maximum load and 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 :

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

and

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

where B_u^r denotes the set of vertices within distance r from vertex u . Note that $R_1 \leq R_2$ for all G . For the sake of clarity, we state our results here for *vertex-transitive* graphs only. In later sections we state our results in full 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 [4, 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 3.1, 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 3.1 in Section 3.

Note that 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 [4] are special cases of Theorems 1.1 and 3.1. Furthermore, the proof of Theorems 1.1 and 3.1 uses different techniques (it follows by a subtle coupling with the 1-choice process) and is substantially shorter than the proofs in [4].

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. 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 [3].

► **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 $\mathcal{O}\left(\frac{m}{n} + R_2\right)$.*

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 $n \log n + \Theta(n)$ [9, Section 2.4.1]. 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. 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 for any $d = \Theta(1)$. In particular, we have $R_2 = \Theta(\log \log n)$ for bounded-degree expanders, and $R_2 = \Theta\left((\log n)^{\frac{1}{d+1}}\right)$ for d -dimensional grids.

► **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 4.2, which we state and prove in Section 4.

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 4.3, which gives an upper bound on the cover time for all regular graphs. We state and prove Theorem 4.3 in Section 4.

► **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 $\Theta(n)$.*

Due to space limitations, we skip some proofs. The full version can be found in [5].

2 Key technical argument

Aside from Theorem 1.4, we assume throughout this paper that G has bounded degrees; i.e., the maximum degree Δ is bounded above by a constant independent of n . We also assume that, in the local search allocation, ties are broken independently and uniformly at random.

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, $T_{\text{cov}} = \min\{m \geq 0: X_v^{(m)} \geq 1 \text{ for all } v \in V\}$.

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,

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

We now prove a key technical result (Lemma 2.2 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:

$$\text{for any two neighbors } u, v \in V, \text{ we have } |\mu(u) - \mu(v)| \leq 1. \quad (2.2)$$

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

$$W_v^{(m)} = X_v^{(m)} + \mu(v). \quad (2.3)$$

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 2.1.** *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 $W_{v'}^{(m-1)} \leq W_v^{(m-1)}$. ◀

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

$$\sum_{i=1}^{\kappa} a_{j_i} \geq \sum_{i=1}^{\kappa} a'_{j'_i} \quad \text{for all } \kappa = 1, 2, \dots, n. \quad (2.4)$$

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 establishes that $\overline{W}^{(m)}$ majorizes $W^{(m)}$ for any m .

► **Lemma 2.2.** *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 [2].

► **Lemma 2.3** ([2, 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 2.2. 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 vector which is obtained from $W^{(m-1)}$ by allocating ball m to vertex i_ℓ (the birthplace of ball m). Applying Lemma 2.3 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

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

so we obtain the vector $W^{(m)}$ from $\widetilde{W}^{(m)}$ by removing one ball from vertex i_ℓ and adding one ball to vertex i_i . By Lemma 2.1, we have $W_{i_i}^{(m-1)} \leq W_{i_\ell}^{(m-1)}$. This implies $\widetilde{W}_{i_\ell}^{(m)} = W_{i_\ell}^{(m-1)} + 1 \geq W_{i_i}^{(m-1)} + 1$ 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. \blacktriangleleft

Now we illustrate the usefulness of the above result by relating the probability of a vertex to have a certain load to the probability that balls are born in a neighborhood around a vertex. For any two vertices $u, v \in V$, we denote by $d_G(u, v)$ their distance on G .

► **Lemma 2.4.** *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 $\overline{\mathcal{A}}^{(m)} = \{\min_{u \in V} \overline{W}_u^{(m)} \geq \ell\}$. By Lemma 2.2, we have that $\Pr[\mathcal{A}^{(m)}] \geq \Pr[\overline{\mathcal{A}}^{(m)}]$.

Clearly, we have that $\mathcal{A}^{(m)}$ implies $\{X_v^{(m)} \geq \ell\}$, but the two events are in fact equal since, by the smoothness of the load vector ([4, Lemma 2.2]), $\{X_v^{(m)} \geq \ell\}$ implies $\mathcal{A}^{(m)}$. The proof is then complete since $\overline{\mathcal{A}}^{(m)} = \bigcap_{w \in B_v^\ell} \left\{ \overline{X}_w^{(m)} \geq \ell - d_G(v, w) \right\}$.

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 2.2 we obtain that $\Pr[\mathcal{B}^{(m)}] \leq \Pr[\overline{\mathcal{B}}^{(m)}] = \Pr \left[\bigcup_{w \in V} \left\{ \overline{X}_w^{(m)} \geq \ell + d_G(v, w) \right\} \right]$. \blacktriangleleft

► **Remark.** 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 2.4 to derive an upper bound on the load of a subset of vertices. The proof of this proposition can be found in the full version [5].

► **Proposition 2.5.** 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 $\Pr \left[\sum_{v \in S} X_v^{(m)} \geq \ell |S| \right] \leq 4 \exp \left(-\frac{|S|\ell}{14} \log \left(\frac{\ell n}{m} \right) \right) + \exp \left(-\frac{m}{4} \right)$. Moreover, for any given $u \in V$, it holds that $\Pr \left[X_u^{(m)} \geq 2\ell \right] \leq 4 \exp \left(-\frac{|B_u^\ell| \ell}{14} \log \left(\frac{\ell n}{m} \right) \right) + \exp \left(-\frac{m}{4} \right)$.

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 2.6.** *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 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

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

Note that $R_1^{(\gamma)}$ is non-increasing with γ . 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 $u \in V$. 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 when $m = n$ is $\Theta(R_1)$. In the following, $\omega(1)$ stands for a term that goes to ∞ as $n \rightarrow \infty$.

► **Theorem 3.1** (General version of Theorem 1.1). *Let G be any graph with bounded degrees. For any $\gamma \in (0, 1/2]$ and $\alpha \geq 1$, we have*

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

Proof. We start establishing a lower bound for $X_{\max}^{(n)}$. Let A be a Poisson random variable with mean 1. 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 2.6). For any $v \in V$ and any $\ell > 0$, Lemma 2.4 gives that

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

where N_v^r is the set of vertices at distance r from v so that $B_v^\ell = \bigcup_{r=0}^{\ell} N_v^r$. Hence,

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

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

$$\Pr \left[X_v^{(n)} \geq \frac{\gamma R_1^{(\gamma)}}{4} \right] \geq \exp \left(-\frac{\gamma R_1^{(\gamma)} |B_v^{R_1^{(\gamma)}}| \log(R_1^{(\gamma)})}{2} \right) \geq n^{-\gamma/2} \quad \text{for all } v \in S. \quad (3.1)$$

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 standard tail bounds, $\Pr[K > 2en] \leq 2^{1-2ne}$. 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 [4, Lemma 2.5], and (3.1) implies that

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

With this, we apply the method of bounded differences [8, Lemma 1.2] to obtain

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

This result can then be translated to the non-Poissonized model via Lemma 2.6.

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

Proof of Theorem 1.2. Applying Proposition 2.5 with $\ell = \left(\frac{m}{n} + R_2\right)c$ for any constant $c \geq 300\Delta$, we obtain

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

where $B_u^{R_2}$ denotes the set of vertices within distance R_2 from u . By setting $c > 0$ 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, 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 a union bound over all $u \in V$ completes the proof. \blacktriangleleft

4 Cover time

The proposition below gives an upper bound for the cover time.

► **Proposition 4.1.** Let G be a graph with bounded degrees. Then for any $\alpha > 1$ there exists a $C = C(\alpha) > 0$ such that for all $m \geq CR_2n$ 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 2. Define $\mu(v) = d_G(u, v)$ and $W_v^{(m)} = X_v^{(m)} + \mu(v)$. Similarly, for the 1-choice process, define $\bar{W}_v^{(m)} = \bar{X}_v^{(m)} + \mu(v)$. Let $Y := \min_{v \in V} \bar{W}_v^{(m)}$ be the minimum weight of all vertices in V in the 1-choice process. Let $\ell = \frac{m}{28n \log \Delta}$ and recall that B_u^r is the set of vertices within distance r from u . We have

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

Using a variant of Hoeffding’s inequality, we obtain

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

where the last inequality holds since $m/n \geq CR_2 = \omega(1)$ for bounded degree graphs. Now define \bar{Z} as the sum of the $|B_u^{R_2}|$ smallest values of $\{\bar{W}_v^{(m)} : v \in V\}$ and Z as the sum of the $|B_u^{R_2}|$ smallest values of $\{W_v^{(m)} : v \in V\}$. By Lemma 2.2, we can couple $W^{(m)}$ and $\bar{W}^{(m)}$ so that, with probability 1, $Z \geq \bar{Z}$. Further, $\mathbf{E}[\bar{Z}] \geq \frac{\ell |B_u^{R_2}|}{2}$. We now apply Azuma’s inequality [6, Theorem 6.1] in order to show that \bar{Z} is likely to be at least $\frac{\ell |B_u^{R_2}|}{4}$. Let A_1, A_2, \dots, A_m be the martingale adapted to the filtration \mathcal{F}_i generated by U_1, U_2, \dots, U_i ; i.e., $A_i = \mathbf{E}[\bar{Z} \mid \mathcal{F}_i]$. Since changing the birthplace of ball i (and keeping all other birthplaces the same) can change Z by at most one [4, Lemma 2.5], we have that $\mathbf{E}[A_i - A_{i-1} \mid \mathcal{F}_{i-1}] \leq 1$.

Now fix i . 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

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

Note that, for any realization of $U_1, U_2, \dots, U_{i-1}, U_{i+1}, \dots, U_m$, ζ_u and $\zeta_{u'}$ only differ if exactly one of u or u' is among the $|B_u^{R_2}|$ smallest loads. 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, Azuma’s

inequality [6, Theorem 6.1] gives

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

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 and m , 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 [4, 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 $C \geq 672 \log \Delta$. The result follows by taking the union bound over all $u \in V$, which yields 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 completed by setting C large enough with respect to α so that $\frac{C}{68(28 \log \Delta)^2} - 1 \geq \alpha$. \blacktriangleleft

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

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

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 $\gamma > 0$, 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 4.2 (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 $C = C(\alpha, \Delta)$ such that*

$$\Pr \left[T_{\text{cov}} < \frac{\gamma R_2^{(\gamma)} n}{8\Delta} \right] \leq n^{-\omega(1)} \quad \text{and} \quad \Pr [T_{\text{cov}} \geq CR_2 n] \leq n^{-\alpha}.$$

Proof. The second inequality is established by Proposition 4.1. 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 $m = \frac{\gamma R_2^{(\gamma)} n}{8\Delta}$. 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 $X_u^{(m)} = 0$. By the second part of Lemma 2.4, we have that

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

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

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

where the last inequality follows by a Chernoff bound [1, Theorem A.1.15]. Using the simple bound $|N_u^i| \leq \Delta^i$ and the fact 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

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

Now let Y be the random variable defined as the number of vertices v 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 [4, Lemma 2.5] 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 $m = \mathcal{O}(n \log n)$ for any G . With this, we apply the method of bounded differences [8, Lemma 1.2] and a standard tail bound to obtain

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

This result can then be translated to the non-Poissonized process using Lemma 2.6 and the fact that $m = \frac{\gamma R_2^{(\gamma)} n}{4} = \mathcal{O}(n \log n)$. \blacktriangleleft

We now state a stronger version of Theorem 1.4. The proof is in the full version [5].

► Theorem 4.3 (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*

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

5 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 4.1 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 $R_2(G)$ are 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.

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 $m = \omega(nR_2)$, 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 $\delta > 1$? Also, 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 $E \subset E'$. Is the maximum load on G stochastically dominated by the maximum load on G' for any m ?

References

- 1 N. Alon and J.H. Spencer. *The probabilistic method*. John Wiley & Sons, 3rd edition, 2008.
- 2 Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, 1999.
- 3 Petra Berenbrink, Artur Czumaj, Angelika Steger, and Berthold Vöcking. Balanced allocations: The heavily loaded case. *SIAM J. Comput.*, 35(6):1350–1385, 2006.
- 4 P. Bogdan, T. Sauerwald, A. Stauffer, and H. Sun. Balls into bins via local search. In *Proc. of the 24th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 16–34, 2013.
- 5 K. Bringmann, T. Sauerwald, H. Sun, and A. Stauffer. Balls into bins via local search: cover time and maximum load, 2013. Preprint at arXiv:1310.0801.
- 6 F. Chung and L. Lu. Concentration inequalities and Martingale inequalities: a survey. *Internet Mathematics*, 3:79–127, 2006.
- 7 Richard M. Karp, Michael Luby, and Friedhelm 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 Martin Raab and Angelika Steger. Balls into bins – a simple and tight analysis. In *2nd Int'l Workshop on Randomization and Computation (RANDOM'98)*, pages 159–170, 1998.
- 11 Berthold Vöcking. How asymmetry helps load balancing. *J. ACM*, 50(4), 2003.