




Dynamic Averaging Load Balancing on Cycles

Dan Alistarh¹ · Giorgi Nadiradze¹  · Amirmojtaba Sabour¹

Received: 17 September 2020 / Accepted: 19 November 2021
© The Author(s) 2021

Abstract

We consider the following dynamic load-balancing process: given an underlying graph G with n nodes, in each step $t \geq 0$, a random edge is chosen, one unit of load is created, and placed at one of the endpoints. In the same step, assuming that loads are arbitrarily divisible, the two nodes *balance* their loads by averaging them. We are interested in the expected gap between the minimum and maximum loads at nodes as the process progresses, and its dependence on n and on the graph structure. Peres et al. (Random Struct Algorithms 47(4):760–775, 2015) studied the variant of this process, where the unit of load is placed in the least loaded endpoint of the chosen edge, and the averaging is not performed. In the case of dynamic load balancing on the cycle of length n the only known upper bound on the expected gap is of order $\mathcal{O}(n \log n)$, following from the majorization argument due to the same work. In this paper, we leverage the power of averaging and provide an improved upper bound of $\mathcal{O}(\sqrt{n} \log n)$. We introduce a new potential analysis technique, which enables us to bound the difference in load between k -hop neighbors on the cycle, for any $k \leq n/2$. We complement this with a “gap covering” argument, which bounds the maximum value of the gap by bounding its value across all possible subsets of a certain structure, and recursively bounding the gaps within each subset. We also show that our analysis can be extended to the specific instance of Harary graphs. On the other hand, we prove that the expected second moment of the gap is lower bounded by $\Omega(n)$. Additionally, we provide experimental evidence that our upper bound on the gap is tight up to a logarithmic factor.

Keywords Algorithms · Load balancing · Heavily loaded case · Randomized algorithms · Graph algorithms

Special Issue on Combinatorial Algorithms (IWOC A 2020).

Extended author information available on the last page of the article

Published online: 24 December 2021

1 Introduction

This paper considers balls-into-bins processes where a sequence of m weights are placed into n bins via some randomized procedure, with the goal of minimizing the load imbalance between the most loaded and the least loaded bin. This family of randomized processes has been used to model several practical allocation problems, such as load-balancing [3,14,20], hashing [9], or even relaxed data structures [1,2].

The classic formulation of this problem is known as d -choice process, in each step, a new weight is generated, and is placed in the least loaded of d randomly chosen bins. If $d = 1$, then we have the uniform random choice scheme, whose properties are well understood, e.g. [18]. In particular, if we place $m = n$ unit weights into the bins, then it is known that the most loaded bin will have expected $\Theta(\log n / \log \log n)$ load, whereas if $m = \Omega(n \log n)$ we have that the expected maximum load is $m/n + \Theta(\sqrt{m \log n/n})$. Seminal work by Azar et al. [3] showed that, if we place n unit weights into n bins by the d -choice process with $d \geq 2$, then, surprisingly, the maximum load is reduced to $\Theta(\log \log n / \log d)$. A technical tour-de-force by Berenbrink, Czumaj, Steger, and Vöcking [4] extended this result to the “heavily-loaded” case where $m \gg n$, showing that in this case the maximum load is $m/n + \log \log n / \log d + O(1)$ with failure probability at most $1/\text{polyn}$. An elegant alternative proof for a slightly weaker version of this result was later provided by Talwar and Wieder [23].

More recently, Peres et al. [17] considered the *graphical* version of this process, where the bins are the vertices of a graph, an *edge* is chosen at every step, and the weight is placed at the less loaded endpoint of the edge, breaking ties arbitrarily. (The reader will notice that the classic 2-choice process corresponds to the case where the graph is a clique.) The authors focus on the evolution of the gap between the highest and lowest loaded bins, showing that, for graphs of β -edge-expansion [17], this gap is $O(\log n/\beta)$, with probability $1 - 1/\text{polyn}$.

Another closely related line of work considers *static load-balancing* processes, where each node in a graph starts with an arbitrary initial load, and the endpoints average their current loads at each step. Note that load balancing schemes which are commonly used in this setting are usually more involved than simply averaging the loads of the endpoints of the randomly chosen node, but the tools used in their analysis are still applicable to the static version of our process. To analyze such processes, it is common to map the process to a Markov chain and analyse its convergence [6,10,12], or derive an upper bound on a potential function which captures the discrepancy between the loads [5]. In both cases, the gap between the highest and lowest loaded bins can be characterized by the spectral gap of the graph [7,11,19,21,22].

By contrast to these two lines of previous work, in this paper we consider a graphical load balancing in the *dynamic* case, where weights arrive at each step rather than being statically allocated initially, but we allow balancing via *continuous averaging*, i.e. the resulting weights after the balancing step equal the average of the sum of the weights of the two nodes prior to balancing. Thus, our averaging step is more powerful relative to d -choice or static averaging, but it is applied in the more challenging *dynamic* scenario.

We will focus on the gap in the dynamic case on graphs of low expansion, specifically on cycles. In [17], it is shown that, in this case (but without averaging), the gap

is $O(n \log n)$ both in expectation and with high probability. The techniques used in [17] imply that the averaging of the loads does not worsen the gap in expectation, our aim is to show that it actually helps to reduce it. Also, directly applying the tools from the static process to the dynamic one results in the upper bound which is larger than $O(n \log n)$. Upper bounding the gap for cycle graphs is known to be a challenging open problem [16]. As suggested in [17], to deal with the cycle case, there is a need for a new approach, which takes the structure of the load balancing graph into account.

1.1 Contribution

In this paper, we address this question for the case where averaging is performed on a cycle graph. Let $Gap(t)$ be a difference between highest and lowest loads of the nodes at time step t . We provide the upper bound on the gap in the dynamic, heavily-loaded case, via a new potential argument. More formally, for any $t > 0$, we show that for a cycle graph with n vertices:

$$\mathbb{E}[Gap(t)] = O(\sqrt{n} \log(n)). \tag{1}$$

We show that our technique can be used to upper bound the gap for the 4-connected Harary Graph. We complement this result with a lower bound of $\Omega(n)$ on $\mathbb{E}[(Gap(t))^2]$. Further, we provide experimental evidence that the gap is of order $\Theta(\sqrt{n})$, making our upper bound accurate up to a $\log n$ factor. Our results extend to the case where the load generated at each node is *weighted* according to some distribution whose second moment is bounded. Formally, we allow our input to come from any distribution W , such that $\mathbb{E}[W^2] \leq M^2$, for some $M > 0$.

1.2 Technical Overview

Our upper bound result is based on two main ideas. The first introduces a new parametrized hop-potential function, which measures the squared difference in load between any k -hop neighbors on the graph, where $k \geq 1$ is a fixed hop parameter. Let $G = (V, E)$ be our input graph, where $V = \{1, 2, \dots, n\}$. Throughout the paper, for any $1 \leq i \leq n$ we assume that the nodes $i + n$ and $i - n$ are the same as the node i . Let $x_i(t)$ be the load of node i at step t . Then, we define the k -hop potential as:

$$\phi_k(t) = \sum_{i=1}^n (x_i(t) - x_{i+k}(t))^2.$$

The first technical step in the proof is to understand the expected (“steady-state”) value of the k -hop potential. We show that, in expectation, the k -hop potential has a recursive structure. While the expected values of k -hop potentials cannot be computed precisely, we can isolate upper and lower bounds on their values for cycles. In particular, for the k -hop potential on an n -cycle, we prove the following bound:

$$\mathbb{E}[\phi_k(t)] \leq k(n - k) - 1, \quad \forall k \geq 1. \tag{2}$$

In the second technical step, we shift gears, aiming to bound the *maximum possible value* of the gap between any two nodes, leveraging the fact that we understand the hop potential for any $k \geq 1$. We achieve this via a “gap covering” technique, which characterizes the maximum value of the gap across all possible subsets of a certain type.

More precisely, in the case of a cycle of length $n = 2^m$, for each node i and hop count k , we define the set family A_k^i to be formed of nodes $\{i, i + 2^{m-k}, i + 2 \times 2^{m-k}, i + 3 \times 2^{m-k}, \dots\}$ (since we are on a cycle, $i = i + 2^{m-k}2^k$). Then for any $1 \leq k \leq m$, we will have

$$\sum_{i=1}^n \text{Gap}_{A_k^i}(t) \leq \sum_{i=1}^n \text{Gap}_{A_{k-1}^i}(t) + \frac{n}{\sqrt{2^{m-k}}} \sqrt{\phi_{2^{m-k}}(t)}, \tag{3}$$

where $\text{Gap}_X(t)$ is the maximal gap inside the set X at time t . Intuitively, this result allows us recursively characterize the gap value at various “resolutions” across the graph.

Finally, we notice that we can “cover” the gap between *any* two nodes by carefully unwinding the recursion in the above inequality, considering all possible subsets of a well-chosen structure, and recursively bounding the gaps within each subset (this step is particularly delicate in the case where n is not a power of two, please see Sect. 5). We obtain that

$$\mathbb{E}[\text{Gap}(t)] = O(\sqrt{n} \log(n)), \tag{4}$$

as claimed. The logarithmic slack is caused by the second term on the right-hand-side of (2). We note that this technique extends to the case where inserted items are *weighted*, where the weights are coming from some distribution of bounded second moment.

1.3 Lower Bound

It is interesting to ask whether this upper bound is tight. To examine this question, we revisit the recursive structure of the k -hop potential, which we used to obtain the upper bound in Eq. (3). We can leverage this structure to obtain a *lower bound* on the expected k -hop potential as well. Starting from this lower bound, we can turn the upper bound argument “inside out,” to obtain a linear lower bound on the *expected squared gap*:

$$\mathbb{E}[\text{Gap}(t)^2] = \Omega(n). \tag{5}$$

We conjecture that both upper and lower bounds on the expected gap are of order $O(\sqrt{n})$ (given that $E[W^2]$ is constant), and examine this claim empirically in Sect. 6.

1.4 Extensions and Discussion

We believe that the analysis template we described above is general, and can be extended to other graph families, such as regular expanders. Here, we focus on obtaining tight bounds on the gap for cycles, which is technically non-trivial, and leave the extensions for other graph families as future work. To substantiate our generality claim, we exhibit an application of our analysis technique to the specific instance of Harary graphs [13] in Sect. 7. More precisely, we provide the upper bound on a gap for a graph on n vertices, where each vertex i is connected with edges to vertices $i - 2, i - 1, i + 1$ and $i + 2$.

We discuss the relation between our results and bounds for the graphical power-of-two process on a cycle [17] in Sect. 8.

1.5 Related Work

As we have already discussed broad background, we will now mainly focus on the technical differences from previous work. As stated, we are the first to specifically consider the *dynamic* case for *continuous averaging* on cycles. The *static* case has been studied both with *continuous averaging* [5–7,10–12,19,22] and *discrete averaging* [21]. However, their techniques would not apply (or would result in a worse bound) in our case, since we consider that weights would be introduced *dynamically*, during the processes' execution.

To our knowledge, the only non-trivial upper bound on the gap of the process we consider which would follow from previous work is of $\mathcal{O}(n \log n)$, by the potential analysis of [17]: they consider 2-choice load balancing, and one can re-do their potential analysis for (continuous) averaging load balancing, yielding the same bounds. (Specifically, one can use the same definition of the potential $\Gamma(t)$ as in this reference, and will obtain the same upper bounds, since the pairwise load balancing we consider has slightly stronger guarantees.) However, as our bounds show, the resulting analysis is quite loose in the case of cycle graphs, yielding an $\Omega(\sqrt{n})$ gap between the bounds yielded by these techniques. Technically, this is a consequence of the majorization technique used in [17], which links dynamic averaging on the cycle with a very weak form of averaging on the clique. Reference [8] studies the performance differences between various load-balancing techniques; specifically, it shows that continuous averaging cannot improve the bound on the gap relative to d -choice on the clique. Unfortunately, this result does not seem to extend to arbitrary graphs.

Our potential analysis is substantially different from that of [17], as they track a sum of exponential potentials across the entire graph. By contrast, our analysis tracks the squared load differences between k -hop neighbors, establishing recurrences between these potentials. We note that this is also different from the usual square potentials used for analyzing averaging load balancing, e.g. [15], which usually compare against the *global mean*, as opposed to pairwise potential differences. Our approach is also different from the classic analyses of e.g. [3], which perform probabilistic induction on the number of bins at a given load, assuming a clique.

Generally, our technique can be seen as performing the induction needed to bound the gap not on the bin loads, as is common in previous work, e.g. [3], but *over the topology of the graph*. This approach is natural, since we wish to obtain tight, topology-specific bounds, but we believe we are the first to propose and analyze it successfully.

2 Averaging on the Cycle: Upper Bounding the Gap

2.1 Preliminaries

We consider a cycle graph $G = (V, E)$ where $V = \{1, 2, \dots, n\}$, such that each node i is connected to its left and right neighbors, $i - 1$ and $i + 1$ (recall that for any $1 \leq i \leq n$ the nodes $i + n$ and $i - n$ are the same as the node i).

We consider a stochastic process following real time $t \geq 0$, in which, at each step $t + 1$, a ball of weight $w(t) \geq 0$ is generated from a same distribution W . We associate a real-valued *load* $x_i(t)$ with each node i ($x_i(t)$ is the value after t steps). Initially, we have that $x_i(0) = 0$ for every node i . At step $t + 1$, an edge $(i, i + 1)$ is chosen uniformly at random, and the two endpoint nodes update their weights as follows:

$$x_i(t + 1) = x_{i+1}(t + 1) = \frac{x_i(t) + x_{i+1}(t) + w(t)}{2}.$$

We will assume that the second moment of the distribution W is bounded. That is: $E[W^2] \leq M^2$, for some $M > 0$. For simplicity, we will assume that weights are normalized by M . This gives us that $\mathbb{E}[W^2] \leq 1$.

Let $X(t) = (x_1(t), x_2(t), \dots, x_n(t))$ be the vector of the bin weights after step t . First, we define the following potential functions:

$$\forall k \in \{1, 2, \dots, n - 1\} : \phi_k(t) := \sum_{i=1}^n (x_i(t) - x_{i+k}(t))^2.$$

Notice that for every $1 \leq i \leq n$, we have that $\phi_i(t) = \phi_{n-i}(t)$. We want to analyze what is the value of these functions in expectation after an additional ball is thrown, for a given load vector $X(t)$.

We start with $\phi_1(t + 1)$:

$$\begin{aligned} \mathbb{E}[\phi_1(t + 1)|X(t), w(t)] &= \sum_{i=1}^n \frac{1}{n} \left(\left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i+2}(t) \right)^2 \right. \\ &\quad + \left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i-1}(t) \right)^2 \\ &\quad \left. + \sum_{j \neq i-1, i, i+1} (x_j(t) - x_{j+1}(t))^2 \right). \end{aligned} \tag{6}$$

Notice that:

$$\sum_{i=1}^n \sum_{j \neq i-1, i, i+1} (x_j(t) - x_{j+1}(t))^2 = (n - 3)\phi_1(t). \tag{7}$$

Hence we need to bound the remaining terms:

$$\begin{aligned} & \sum_{i=1}^n \left(\left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i+2}(t) \right)^2 + \left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i-1}(t) \right)^2 \right) \\ &= \sum_{i=1}^n \left(\frac{2x_i(t)^2 + 2x_{i+1}(t)^2 + 2w(t)^2 + 4x_i(t)x_{i+1}(t)}{4} \right. \\ & \quad \left. + x_{i+2}(t)^2 + x_{i-1}(t)^2 - (x_i(t) + x_{i+1}(t))(x_{i-1}(t) + x_{i+2}(t)) \right). \end{aligned}$$

where we used the fact that terms, which are linear in $w(t)$, cancel out.

The right side of the above equation can be rewritten as:

$$\begin{aligned} & \sum_{i=1}^n \frac{x_i(t)^2}{2} + \sum_{i=1}^n \frac{x_{i+1}(t)^2}{2} + \sum_{i=1}^n x_{i+2}(t)^2 + \sum_{i=1}^n x_{i-1}(t)^2 \\ & + \sum_{i=1}^n x_i(t)x_{i+1}(t) - \sum_{i=1}^n x_i(t)x_{i-1}(t) - \sum_{i=1}^n x_{i+1}(t)x_{i+2}(t) \\ & - \sum_{i=1}^n x_i(t)x_{i+2}(t) - \sum_{i=1}^n x_{i+1}(t)x_{i-1}(t) + \frac{nw(t)^2}{2} \\ &= 3 \sum_{i=1}^n x_i(t)^2 - \sum_{i=1}^n x_i(t)x_{i+1}(t) - 2 \sum_{i=1}^n x_i(t)x_{i+2}(t) + \frac{nw(t)^2}{2} \\ &= \sum_{i=1}^n \frac{(x_i(t) - x_{i+1}(t))^2}{2} + \sum_{i=1}^n (x_i(t) - x_{i+2}(t))^2 + \frac{nw(t)^2}{2} \\ &= \frac{\phi_1(t)}{2} + \phi_2(t) + \frac{nw(t)^2}{2}. \end{aligned}$$

By using the above equation and Eq. (7) in Eq. (6) we get that

$$\mathbb{E}[\phi_1(t + 1) | X(t), w(t)] = \frac{n - 2}{n} \phi_1(t) + \frac{1}{2} \left(w(t)^2 - \frac{\phi_1(t)}{n} \right) + \frac{\phi_2(t)}{n}.$$

Now, we proceed with calculating the expected value of $\phi_k(t + 1)$,

for $2 \leq k \leq \lfloor n/2 \rfloor$:

$$\begin{aligned} \mathbb{E}[\phi_k(t+1)|X(t), w(t)] &= \sum_{i=1}^n \frac{1}{n} \left(\left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i-k}(t) \right)^2 \right. \\ &\quad + \left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i+1-k}(t) \right)^2 \\ &\quad + \left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i+k}(t) \right)^2 \\ &\quad + \left. \left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i+1+k}(t) \right)^2 \right) \\ &\quad + \sum_{j \neq i-k, i+1, i, i+1} (x_j(t) - x_{j+k}(t))^2. \end{aligned}$$

Notice that:

$$\frac{1}{n} \sum_{i=1}^n \sum_{j \neq i-k, i+1-k, i+k, i+1+k} (x_j(t) - x_{j+k}(t))^2 = \frac{n-4}{n} \phi_k(t).$$

In the similar way as for $\phi_1(t)$ the remaining terms can be rewritten as:

$$\begin{aligned} &\sum_{i=1}^n \frac{1}{n} \left(x_i(t)^2 + x_{i+1}(t)^2 + w(t)^2 + 2x_i(t)x_{i+1}(t) + 4x_i^2(t) \right. \\ &\quad \left. - (x_i(t) + x_{i+1}(t)) (x_{i+k}(t) + x_{i+k+1}(t) + x_{i-k}(t) + x_{i-k+1}(t)) \right) \\ &= \frac{1}{n} \left(nw(t)^2 + 6 \sum_{i=1}^n x_i(t)^2 + 2 \sum_{i=1}^n x_i(t)x_{i+1}(t) - 4 \sum_{i=1}^n x_i(t)x_{i+k}(t) \right. \\ &\quad \left. - 2 \sum_{i=1}^n x_i(t)x_{i+k+1}(t) - 2 \sum_{i=1}^n x_i(t)x_{i+k-1}(t) \right) \\ &= \frac{2}{n} \phi_k(t) + (w(t)^2 - \frac{\phi_1(t)}{n}) + \frac{\phi_{k+1}(t)}{n} + \frac{\phi_{k-1}(t)}{n}. \end{aligned}$$

Hence, we get that:

$$\mathbb{E}[\phi_k(t+1)|X(t), w(t)] = \frac{n-2}{n} \phi_k(t) + (w(t)^2 - \frac{\phi_1(t)}{n}) + \frac{\phi_{k+1}(t)}{n} + \frac{\phi_{k-1}(t)}{n}.$$

If we remove conditioning on $w(t)$ (note that $\mathbb{E}[w(t)^2] = \mathbb{E}[W^2]$) and express these equations for

$\phi_1(t + 1), \phi_2(t + 1), \dots, \phi_{n-1}(t + 1)$ (recall that $\phi_k(t) = \phi_{n-k}(t)$), we get:

$$\left\{ \begin{array}{l} \mathbb{E}[\phi_1(t + 1)|X(t)] = \left(\frac{n-2}{n}\right)\phi_1(t) + \frac{1}{2}(\mathbb{E}[W^2] - \frac{\phi_1(t)}{n}) + \frac{\phi_2(t)}{n}. \\ \mathbb{E}[\phi_2(t + 1)|X(t)] = \left(\frac{n-2}{n}\right)\phi_2(t) + (\mathbb{E}[W^2] - \frac{\phi_1(t)}{n}) + \frac{\phi_1(t)}{n} + \frac{\phi_3(t)}{n}. \\ \dots \\ \mathbb{E}[\phi_{\lfloor \frac{n}{2} \rfloor}(t + 1)|X(t)] = \left(\frac{n-2}{n}\right)\phi_{\lfloor \frac{n}{2} \rfloor}(t) + (\mathbb{E}[W^2] - \frac{\phi_1(t)}{n}) \\ \quad + \frac{\phi_{\lfloor \frac{n}{2} \rfloor - 1}(t)}{n} + \frac{\phi_{\lfloor \frac{n}{2} \rfloor + 1}(t)}{n}. \\ \dots \\ \mathbb{E}[\phi_{n-2}(t + 1)|X(t)] = \left(\frac{n-2}{n}\right)\phi_{n-2}(t) \\ \quad + (\mathbb{E}[W^2] - \frac{\phi_1(t)}{n}) + \frac{\phi_{n-3}(t)}{n} + \frac{\phi_{n-1}(t)}{n}. \\ \mathbb{E}[\phi_{n-1}(t + 1)|X(t)] = \left(\frac{n-2}{n}\right)\phi_{n-1}(t) + \frac{1}{2}(\mathbb{E}[W^2] - \frac{\phi_1(t)}{n}) + \frac{\phi_{n-2}(t)}{n}. \end{array} \right. \tag{8}$$

Using the above equations we can prove the following:

Lemma 1 For every $t \geq 0$ and $1 \leq k \leq n - 1$, we have that

$$\mathbb{E}[\phi_k(t)] \leq (k(n - k) - 1)\mathbb{E}[W^2] \leq k(n - k) - 1. \tag{9}$$

Proof Let $\Phi(t) = (\phi_1(t), \phi_2(t), \dots, \phi_{n-1}(t))$ be the vector of values of our potentials at time step t and let $Y = (y_1, y_2, \dots, y_{n-1})$, be the vector containing our desired upper bounds for each potential. That is: for each $1 \leq i \leq n - 1$, we have that $y_i = (i(n - i) - 1)\mathbb{E}[W^2]$.

An interesting and easily checkable thing about the vector Y is that

$$\mathbb{E}[\Phi(t + 1)|\Phi(t) = Y] = Y. \tag{10}$$

Next, consider the vector $Z(t) = (z_1(t), z_2(t), \dots, z_{n-1}(t)) = Y - \Phi(t)$. Our goal is to show that for every step t and coordinate i , $\mathbb{E}[z_i(t)] \geq 0$.

We have that

$$\begin{aligned} \mathbb{E}[z_1(t + 1)|X(t)] &= y_1 - \mathbb{E}[\phi_1(t + 1)|X(t)] \\ &= \left(\frac{n-2}{n}\right)y_1 + \frac{1}{2}(\mathbb{E}[W^2] - \frac{y_1}{n}) + \frac{y_2}{n} \\ &\quad - \left(\left(\frac{n-2}{n}\right)\phi_1(t) + \frac{1}{2}(\mathbb{E}[W^2] - \frac{\phi_1(t)}{n}) + \frac{\phi_2(t)}{n}\right) \\ &= \left(\frac{n-2}{n}\right)z_1(t) - \frac{z_1(t)}{2n} + \frac{z_2(t)}{n}. \end{aligned}$$

and for $2 \leq i \leq \lfloor \frac{n}{2} \rfloor$, we have that

$$\mathbb{E}[z_i(t + 1)|X(t)] = \left(\frac{n-2}{n}\right)z_i(t) - \frac{z_1(t)}{n} + \frac{z_{i+1}(t)}{n} + \frac{z_{i-1}(t)}{n}.$$

Hence we get the following equations (recall that $z_i(t) = z_{n-i}(t)$):

$$\begin{cases} n \times \mathbb{E}[z_1(t+1)|X(t)] = (n-2-\frac{1}{2})z_1(t) + z_2(t). \\ n \times \mathbb{E}[z_2(t+1)|X(t)] = -z_1(t) + z_1(t) + (n-2)z_2(t) + z_3(t). \\ n \times \mathbb{E}[z_3(t+1)|X(t)] = -z_1(t) + z_2(t) + (n-2)z_3(t) + z_4(t). \\ \dots \\ n \times \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t+1)|X(t)] = -z_1(t) + z_{\lfloor \frac{n}{2} \rfloor - 1}(t) + (n-2)z_{\lfloor \frac{n}{2} \rfloor}(t) + z_{\lfloor \frac{n}{2} \rfloor + 1}(t). \end{cases} \tag{11}$$

Next, using induction on t , we show that for every $t \geq 0$

$$0 \leq \mathbb{E}[z_1(t)] \leq \mathbb{E}[z_2(t)] \leq \dots \leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t)]. \tag{12}$$

The base case holds trivially since $Z(0) = Y$. For the induction step, assume that $0 \leq \mathbb{E}[z_1(t)] \leq \mathbb{E}[z_2(t)] \leq \dots \leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t)]$. First, we have that

$$n\mathbb{E}[z_1(t+1)] = n\mathbb{E}_{X(t)}[\mathbb{E}[z_1(t+1)|X(t)]] = (n-2-\frac{1}{2})\mathbb{E}[z_1(t)] + \mathbb{E}[z_2(t)] \geq 0.$$

Additionally, we have that:

$$\begin{aligned} n\mathbb{E}[z_1(t+1)] &= (n-2-\frac{1}{2})\mathbb{E}[z_1(t)] + \mathbb{E}[z_2(t)] \leq (n-2)\mathbb{E}[z_1(t)] + \mathbb{E}[z_2(t)] \\ &\leq (n-2)\mathbb{E}[z_2(t)] + \mathbb{E}[z_3(t)] = n\mathbb{E}[z_2(t+1)]. \end{aligned}$$

For $2 \leq i \leq \lfloor \frac{n}{2} \rfloor - 2$, we have that

$$\begin{aligned} n\mathbb{E}[z_i(t+1)] &= -\mathbb{E}[z_1(t)] + \mathbb{E}[z_{i-1}(t)] + (n-2)\mathbb{E}[z_i(t)] + \mathbb{E}[z_{i+1}(t)] \\ &\leq -\mathbb{E}[z_1(t)] + \mathbb{E}[z_i(t)] + (n-2)\mathbb{E}[z_{i+1}(t)] + \mathbb{E}[z_{i+2}(t)] \\ &= n\mathbb{E}[z_{i+1}(t+1)]. \end{aligned}$$

Next, observe that by our assumption:

$\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor + 1}(t)] = \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor - 1}(t)] \geq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor - 2}(t)]$. Finally, by using this observation we get that

$$\begin{aligned} n\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor - 1}(t+1)] &= -\mathbb{E}[z_1(t)] + \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor - 2}(t)] + (n-2)\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor - 1}(t)] + \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t)] \\ &\leq -\mathbb{E}[z_1(t)] + \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor + 1}(t)] + \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor - 1}(t)] + (n-3)\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor - 1}(t)] + \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t)] \\ &\leq -\mathbb{E}[z_1(t)] + \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor + 1}(t)] + \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor - 1}(t)] + (n-2)\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t)] \\ &= n\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t+1)]. \end{aligned}$$

This completes the proof of the lemma. □

3 Upper Bound on the Gap for $n = 2^m$

In this section we upper bound a gap in expectation for the case when $n = 2^m$. is quite technical but not necessarily more interesting, and is provided in the Sect. 5.

We begin with some definitions. For a set $A \subseteq \{1, 2, \dots, n\}$, let

$$Gap_A(t) = \max_{i \in A} x_i(t) - \min_{i \in A} x_i(t).$$

Also, let A_k^i be $\{i, i + 2^{m-k}, i + 2 \times 2^{m-k}, i + 3 \times 2^{m-k}, \dots\}$ (Notice that $i = i + 2^{m-k}2^k$). Our proof works as follows: for each $1 \leq i \leq n$ and $1 \leq k \leq m$, we look at the vertices given by the sets A_{k-1}^i and $A_{k-1}^{i+2^{m-k}}$ and try to characterise the gap after we merge those sets (note that this will give us the gap for the set $A_k^i = A_{k-1}^i \cup A_{k-1}^{i+2^{m-k}}$). Using this result, we are able to show that $\sum_{i=1}^n Gap_{A_k^i}(t)$ is upper bounded by $\sum_{i=1}^n Gap_{A_{k-1}^i}(t)$ plus n times maximum load difference between vertices at hop distance 2^{m-k} . Next, we use 2^{m-k} hop distance potential $\phi_{2^{m-k}}(t)$ to upper bound maximum load difference between the vertices at hop distance 2^{m-k} . By summing up the derived inequality for $k = 1$ to m , we are able to upper bound $\sum_{i=1}^n Gap_{A_m^i}(t)$ in terms of $\sum_{i=1}^n Gap_{A_0^i}(t)$ and $\sum_{k=1}^m \phi_{2^{m-k}}(t)$. Notice that by our definitions, for each i , $Gap_{A_0^i}(t) = 0$ (A_0^i contains only vertex i) and $Gap_{A_m^i}(t) = Gap(t)$ (A_m^i contains all vertices). Hence, what is left is to use the upper bounds for the hop distance potentials, which we derived in the previous section.

We start by proving the following useful lemma.

Lemma 2 For any $1 \leq i \leq n$ and $1 \leq k \leq m$, we have that

$$2Gap_{A_k^i}(t) \leq 2 \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)| + Gap_{A_{k-1}^{i+2^{m-k}}}(t) + Gap_{A_{k-1}^i}(t). \tag{13}$$

Proof Fix vertex i . Recall that $A_k^i = A_{k-1}^i \cup A_{k-1}^{i+2^{m-k}}$. Let $u = \arg \max_{j \in A_k^i} x_j(t)$ and let $v = \arg \min_{j \in A_k^i} x_j(t)$. We consider several cases on the membership of nodes u and v , and bound the gap in each one:

Case 1 $u \in A_{k-1}^i$ and $v \in A_{k-1}^{i+2^{m-k}}$. Then $Gap_{A_k^i}(t) = Gap_{A_{k-1}^i}(t)$ and we have that

$$\begin{aligned} Gap_{A_k^i}(t) &= |x_u(t) - x_v(t)| \\ &\leq |x_{u+2^{m-k}}(t) - x_u(t)| + |x_{v+2^{m-k}}(t) - x_v(t)| \\ &\quad + |x_{u+2^{m-k}}(t) - x_{v+2^{m-k}}(t)| \\ &\leq |x_{u+2^{m-k}}(t) - x_u(t)| + |x_{v+2^{m-k}}(t) - x_v(t)| \\ &\quad + Gap_{A_{k-1}^{i+2^{m-k}}}(t) \\ &\leq 2 \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)| + Gap_{A_{k-1}^{i+2^{m-k}}}(t), \end{aligned}$$

where we used the fact that both $u + 2^{m-k}$ and $v + 2^{m-k}$ belong to $A_{k-1}^{i+2^{m-k}}$. This gives us that

$$2\text{Gap}_{A_k^i}(t) \leq 2 \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)| + \text{Gap}_{A_{k-1}^{i+2^{m-k}}}(t) + \text{Gap}_{A_{k-1}^i}(t). \tag{14}$$

Case 2 $u \in A_{k-1}^i$ and $v \in A_{k-1}^{i+2^{m-k}}$. Then we have that:

$$\begin{aligned} \text{Gap}_{A_k^i}(t) &= |x_u(t) - x_v(t)| \leq |x_u(t) - x_{v+2^{m-k}}(t)| + |x_{v+2^{m-k}}(t) - x_v(t)| \\ &\leq \text{Gap}_{A_{k-1}^i}(t) + \max_{j \in A_k^i} (|x_j(t) - x_{j+2^{m-k}}(t)|) \end{aligned}$$

and

$$\begin{aligned} \text{Gap}_{A_k^i}(t) &= |x_u(t) - x_v(t)| \leq |x_u(t) - x_{u+2^{m-k}}(t)| + |x_{u+2^{m-k}}(t) - x_v(t)| \\ &\leq \text{Gap}_{A_{k-1}^{i+2^{m-k}}}(t) + \max_{j \in A_k^i} (|x_j(t) - x_{j+2^{m-k}}(t)|), \end{aligned}$$

where we used $v + 2^{m-k} \in A_{k-1}^i$ and $u + 2^{m-k} \in A_{k-1}^{i+2^{m-k}}$. Hence, we again get that

$$2\text{Gap}_{A_k^i}(t) \leq 2 \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)| + \text{Gap}_{A_{k-1}^{i+2^{m-k}}}(t) + \text{Gap}_{A_{k-1}^i}(t). \tag{15}$$

Case 3 $u \in A_{k-1}^{i+2^{m-k}}$ and $v \in A_{k-1}^{i+2^{m-k}}$, is similar to Case 1.

Case 4 $v \in A_{k-1}^i$ and $u \in A_{k-1}^{i+2^{m-k}}$, is similar to Case 2.

□

Next, we upper bound the quantity $\sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)|$.

Lemma 3 For any $1 \leq k \leq m$:

$$\sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)| \leq \frac{n}{\sqrt{2^{m-k}}} \sqrt{\phi_{2^{m-k}}(t)}. \tag{16}$$

Proof Notice that for any i and $i' \in A_k^i$, we have that $A_k^i = A_k^{i'}$, hence $\max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)| = \max_{j \in A_k^{i'}} |x_j(t) - x_{j+2^{m-k}}(t)|$ and this means that

$$\begin{aligned} \sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)| &= \frac{n}{2^{m-k}} \sum_{i=1}^{2^{m-k}} \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)| \\ &\stackrel{\text{Cauchy-Schwarz}}{\leq} \frac{n}{2^{m-k}} \sqrt{2^{m-k}} \sqrt{\sum_{i=1}^{2^{m-k}} \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)|^2} \end{aligned}$$

$$\begin{aligned} &\leq \frac{n}{2^{m-k}} \sqrt{2^{m-k}} \sqrt{\sum_{j=1}^n |x_j(t) - x_{j+2^{m-k}}(t)|^2} \\ &= \frac{n}{\sqrt{2^{m-k}}} \sqrt{\phi_{2^{m-k}}(t)}, \end{aligned}$$

where in the last inequality we used a fact that sets $A_k^1, A_k^2, \dots, A_k^{2^{m-k}}$ are disjoint. \square

Finally, using the above lemmas we can upper bound the expected gap at step t :

Theorem 1 *For every $t \geq 0$, we have that*

$$\mathbb{E}[\text{Gap}(t)] = O(\sqrt{n} \log(n)).$$

Proof From Lemma 2 we have that

$$\begin{aligned} \sum_{i=1}^n 2\text{Gap}_{A_k^i}(t) &\leq \sum_{i=1}^n \text{Gap}_{A_{k-1}^i}(t) + \sum_{i=1}^n \text{Gap}_{A_{k-1}^{i+2^{m-k}}}(t) \\ &\quad + \sum_{i=1}^n 2 \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)| \\ &= 2 \sum_{i=1}^n \text{Gap}_{A_{k-1}^i}(t) + 2 \sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)|. \end{aligned}$$

After dividing the above inequality by 2 and applying Lemma 3 we get that:

$$\sum_{i=1}^n \text{Gap}_{A_k^i}(t) \leq \sum_{i=1}^n \text{Gap}_{A_{k-1}^i}(t) + \frac{n}{\sqrt{2^{m-k}}} \sqrt{\phi_{2^{m-k}}(t)}.$$

Next, we sum up the above inequality for $k = 1$ to m :

$$\sum_{k=1}^m \sum_{i=1}^n \text{Gap}_{A_k^i}(t) \leq \sum_{k=1}^m \sum_{i=1}^n \text{Gap}_{A_{k-1}^i}(t) + \sum_{k=1}^m \frac{n}{\sqrt{2^{m-k}}} \sqrt{\phi_{2^{m-k}}(t)}.$$

Recall that $\sum_{i=1}^n \text{Gap}_{A_m^i}(t) = n\text{Gap}(t)$ and $\sum_{i=1}^n \text{Gap}_{A_0^i}(t) = 0$. Hence, we get that

$$n\text{Gap}(t) \leq \sum_{k=1}^m \frac{n}{\sqrt{2^{m-k}}} \sqrt{\phi_{2^{m-k}}(t)}.$$

Next, we apply Jensen’s inequality and Lemma 1:

$$n\mathbb{E}[\text{Gap}(t)] \leq \sum_{k=1}^m \frac{n}{\sqrt{2^{m-k}}} \mathbb{E} \sqrt{\phi_{2^{m-k}}(t)}$$

$$\begin{aligned}
 &\leq \sum_{k=1}^m \frac{n}{\sqrt{2^{m-k}}} \sqrt{\mathbb{E}[\phi_{2^{m-k}}(t)]} \\
 &\leq \sum_{k=1}^m \frac{n}{\sqrt{2^{m-k}}} \sqrt{2^{m-k}(n - 2^{m-k})} \\
 &\leq mn\sqrt{n} = n(\log n)\sqrt{n}.
 \end{aligned}$$

This gives us the proof of the theorem. □

4 Gap Lower Bound

Next we prove the following theorem, which lower bounds the second moment of the gap in expectation.

Theorem 2 *The following limit holds:*

$$\lim_{t \rightarrow \infty} \mathbb{E}[Gap(t)^2] = \Omega(n\mathbb{E}[W^2]).$$

Proof In this case we want to prove that not only does vector $Z(t)$ have positive coordinates in expectation, but also $\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}]$ converges to 0. This will give us that $\phi_{\lfloor \frac{n}{2} \rfloor}$ approaches its upper bound $(\lfloor \frac{n}{2} \rfloor \lceil \frac{n}{2} \rceil - 1)\mathbb{E}[W^2]$ in expectation. Then, we can show that there exist two nodes (at distance $\lfloor \frac{n}{2} \rfloor$) such that the expected square of difference between their loads is $\Omega(n\mathbb{E}[w^2])$.

Recall from Eq. (11) that

$$n\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t+1)] = -\mathbb{E}[z_1(t)] + \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor+1}(t)] + \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor-1}(t)] + (n - 2)\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t)].$$

We also know that Inequalities (12) hold for every t , hence we get that

$$\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t+1)] \leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t)] - \frac{\mathbb{E}[z_1(t)]}{n}.$$

The above inequality in combination with Inequalities (12) means that

$$\begin{aligned}
 \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + \lfloor \frac{n}{2} \rfloor + 1)] &\leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + 1)] - \sum_{i=t}^{t+\lfloor \frac{n}{2} \rfloor} \frac{\mathbb{E}[z_1(i)]}{n} \\
 &\leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + 1)] - \frac{\mathbb{E}[z_1(t + \lfloor \frac{n}{2} \rfloor)]}{n} \tag{17}
 \end{aligned}$$

Again by using Eq. (11) and Inequalities (12), we can show that for every $1 \leq i \leq \lfloor \frac{n}{2} \rfloor - 1$:

$$\mathbb{E}[z_i(t + 1)] \geq \frac{\mathbb{E}[z_{i+1}(t)]}{n}.$$

This gives us that:

$$\begin{aligned} \mathbb{E}[z_1(t + \lfloor \frac{n}{2} \rfloor)] &\geq \left(\frac{1}{n}\right) \mathbb{E}[z_2(t + \lfloor \frac{n}{2} \rfloor - 1)] \geq \left(\frac{1}{n}\right)^2 \mathbb{E}[z_3(t + \lfloor \frac{n}{2} \rfloor - 2)] \\ &\geq \dots \\ &\geq \left(\frac{1}{n}\right)^{\lfloor \frac{n}{2} \rfloor - 1} \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + \lfloor \frac{n}{2} \rfloor - (\lfloor \frac{n}{2} \rfloor - 1))] \\ &= \left(\frac{1}{n}\right)^{\lfloor \frac{n}{2} \rfloor - 1} \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + 1)]. \end{aligned}$$

By plugging the above inequality in Inequality (17), we get that

$$\begin{aligned} \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + \lfloor \frac{n}{2} \rfloor + 1)] &\leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + 1)] - \frac{\mathbb{E}[z_1(t + \lfloor \frac{n}{2} \rfloor)]}{n} \\ &\leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + 1)] - \left(\frac{1}{n}\right)^{\lfloor \frac{n}{2} \rfloor} \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + 1)] \\ &= \left(1 - \left(\frac{1}{n}\right)^{\lfloor \frac{n}{2} \rfloor}\right) \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + 1)]. \end{aligned}$$

Because $\left(1 - \left(\frac{1}{n}\right)^{\lfloor \frac{n}{2} \rfloor}\right) < 1$ and does not depend on t , we get that

$$\lim_{t \rightarrow \infty} \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t)] = 0.$$

This means that $\lim_{t \rightarrow \infty} \mathbb{E}[\phi_{\lfloor \frac{n}{2} \rfloor}(t)] = \Omega(n^2 \mathbb{E}[W^2])$.

Let $Gap_{\lfloor \frac{n}{2} \rfloor}(t) = \max_{1 \leq i \leq n} |x_i(t) - x_{i+\lfloor \frac{n}{2} \rfloor}(t)|$. Note that:

$$Gap(t)^2 \geq Gap_{\lfloor \frac{n}{2} \rfloor}(t)^2 \geq \frac{\phi_{\lfloor \frac{n}{2} \rfloor}(t)}{n}.$$

Hence

$$\lim_{t \rightarrow \infty} \mathbb{E}[Gap(t)^2] = \Omega(n \mathbb{E}[W^2]).$$

Unfortunately we are not able to obtain the lower bound on the gap, since our approach uses the fact that the upper bounds on k -hop potentials are 'tight'. Since our potentials are quadratic, we are not able to derive any kind of lower bound for the gap itself. Intuitively, this will be an issue with any argument which uses convex potential. □

5 Upper Bound on the Gap, General Case

To prove Theorem 1 for the general case, we need to redefine our sets A_k^i . In order to do this, for each k we define 2^k dimensional vector $\Delta_k = (\delta_k^1, \delta_k^2, \dots, \delta_k^{2^k})$. For $k = 0$,

we have that $\Delta_k = (n)$. For $\lfloor \log n \rfloor \geq k > 0$ we set $\Delta_k = (\alpha_k, \delta_{k-1}^1 - \alpha_k, \alpha_k, \delta_{k-1}^2 - \alpha_k, \dots, \alpha_k, \delta_{k-1}^{2^{k-1}} - \alpha_k)$. where

$$\alpha_k = \begin{cases} \lfloor \frac{\lfloor \frac{n}{2^{k-1}} \rfloor}{2} \rfloor, & \text{if } \lfloor \frac{n}{2^{k-1}} \rfloor \text{ is even.} \\ \lceil \frac{\lceil \frac{n}{2^{k-1}} \rceil}{2} \rceil, & \text{otherwise.} \end{cases}$$

First we prove the following Lemma:

Lemma 4 For any $\lfloor \log n \rfloor \geq k \geq 0$, we have that

1. $\sum_{i=1}^{2^k} \delta_k^i = n$.
2. For any $1 \leq i \leq 2^k$, $\delta_k^i \in \{ \lfloor \frac{n}{2^k} \rfloor, \lceil \frac{n}{2^k} \rceil \}$ (notice that this means $\alpha_k = \lfloor \frac{n}{2^k} \rfloor$ or $\alpha_k = \lceil \frac{n}{2^k} \rceil$).

Proof We prove the lemma using induction on k . Base case $k = 0$ holds trivially. For the induction step, assume that Properties 1 and 2 hold for $k - 1$, we aim to prove that they hold for k as well. We have that $\sum_{i=1}^{2^k} \delta_k^i = \sum_{i=1}^{2^{k-1}} (\alpha_k + \delta_{k-1}^i - \alpha_k) = \sum_{i=1}^{2^{k-1}} \delta_{k-1}^i = n$. To prove Property 2 we consider several cases:

Case 1 $\frac{n}{2^{k-1}} = 2q$, for some integer q .

We have that $\alpha_k = q$, and hence for any $1 \leq i \leq 2^{k-1}$, $\delta_{k-1}^i - \alpha_k = q$. Since $\lfloor \frac{n}{2^k} \rfloor = q$, Property 2 holds.

Case 2 $\frac{n}{2^{k-1}} = 2q + 1$, for some integer q .

We have that $\alpha_k = q$, and hence for any $1 \leq i \leq 2^{k-1}$, $\delta_{k-1}^i - \alpha_k = q + 1$. Since $\lfloor \frac{n}{2^k} \rfloor = q$ and $\lceil \frac{n}{2^k} \rceil = q + 1$, Property 2 holds.

Case 3 $\frac{n}{2^{k-1}} = 2q + \epsilon$, for some integer q and $0 < \epsilon < 1$.

We have that $\lfloor \frac{n}{2^{k-1}} \rfloor = 2q$ and $\lceil \frac{n}{2^{k-1}} \rceil = 2q + 1$. Additionally, $\alpha_k = q$, and hence for any $1 \leq i \leq 2^{k-1}$, $(\delta_{k-1}^i - \alpha_k) \in \{q, q + 1\}$. Since $\lfloor \frac{n}{2^k} \rfloor = q$ and $\lceil \frac{n}{2^k} \rceil = q + 1$, Property 2 holds.

Case 4 $\frac{n}{2^{k-1}} = 2q + 1 + \epsilon$, for some integer q and $0 < \epsilon < 1$.

We have that $\lfloor \frac{n}{2^{k-1}} \rfloor = 2q + 1$ and $\lceil \frac{n}{2^{k-1}} \rceil = 2q + 2$. Additionally, $\alpha_k = q + 1$, and hence for any $1 \leq i \leq 2^{k-1}$, $(\delta_{k-1}^i - \alpha_k) \in \{q, q + 1\}$. Since $\lfloor \frac{n}{2^k} \rfloor = q$ and $\lceil \frac{n}{2^k} \rceil = q + 1$, Property 2 holds. □

Next, for $\lfloor \log n \rfloor \geq k > 0$ we set

$$A_k^i = \{i, i + \delta_k^1, i + \delta_k^1 + \delta_k^2, \dots, i + \sum_{j=1}^{2^k-1} \delta_k^j\}.$$

It is easy to see that for any $\lfloor \log n \rfloor \geq k > 0$ and i , we have that $|A_k^i| = 2^k$, $A_k^i = A_{k-1}^{i+\alpha_k} \cup A_{k-1}^i$ and $A_{k-1}^i \cap A_{k-1}^{i+\alpha_k} = \emptyset$. Also notice that for any $u \in A_{k-1}^i$, $u + \alpha_k \in A_{k-1}^{i+\alpha_k}$ and for any $u \in A_{k-1}^{i+\alpha_k}$, $u - \alpha_k \in A_{k-1}^i$.

Next we prove the lemma which is similar to the Lemma 2 for $n = 2^m$ case:

Lemma 5 For any $1 \leq i \leq n$ and $\lfloor \log n \rfloor \geq k > 0$, we have that

$$2\text{Gap}_{A_k^i}(t) \leq 2 \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| + \text{Gap}_{A_{k-1}^{i+\alpha_k}}(t) + \text{Gap}_{A_{k-1}^i}(t). \quad (18)$$

Proof Notice that the statement we need to prove is identical to that of Lemma 2, with the exception that in the letter we use 2^{m-k} instead of α^k . The proofs are also almost identical (2^{m-k} can be simply replaced with α_k). The only difference is that the proof of Lemma 2 uses the property that for any $u \in A_{k-1}^{i+2^{m-k}}$, $|x_u(t) - x_{u+2^{m-k}}(t)| \leq \max_{j \in A_k^i} |x_j(t) - x_{j+2^{m-k}}(t)|$ and $u + 2^{m-k} \in A_{k-1}^i$. Instead, we will use the property that for any $u \in A_{k-1}^{i+\alpha_k}$, $|x_u(t) - x_{u-\alpha_k}(t)| \leq \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)|$ and $u - \alpha_k \in A_{k-1}^i$. \square

Next, we upper bound $\sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)|$, by proving the following lemma, which is the analogue of Lemma 3.

Lemma 6

$$\sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| \leq \left\lceil \frac{n}{\lfloor \frac{n}{2^k} \rfloor} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\phi_{\alpha_k}(t)} \quad (19)$$

Proof Notice that for any $1 \leq u \leq n$ and the sets $A_k^u, A_k^{u+1}, \dots, A_k^{u+\lfloor \frac{n}{2^k} \rfloor - 1}$ are disjoint, because for any $1 \leq j \leq 2^k$, $\delta_k^j \geq \lfloor \frac{n}{2^k} \rfloor$ (this means that for any $1 \leq i \leq n$, distances between consecutive vertices in A_k^i are at least $\lfloor \frac{n}{2^k} \rfloor$). Using this fact and Cauchy–Schwarz inequality we get that

$$\begin{aligned} & \sum_{i=u}^{u+\lfloor \frac{n}{2^k} \rfloor - 1} \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| \\ & \leq \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\sum_{i=u}^{u+\lfloor \frac{n}{2^k} \rfloor - 1} \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)|^2} \\ & \leq \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\sum_{j=1}^n |x_j(t) - x_{j+\alpha_k}(t)|^2} = \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\phi_{\alpha_k}(t)} \end{aligned}$$

Since the above inequality holds for any u we can write that:

$$\sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| \leq \left\lceil \frac{n}{\lfloor \frac{n}{2^k} \rfloor} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\phi_{\alpha_k}(t)}. \quad \square$$

\square

With the above three lemmas in place, we are ready to prove Theorem 1 for general n .

From Lemma 5 we have that

$$\begin{aligned} \sum_{i=1}^n 2\text{Gap}_{A_k^i}(t) &\leq \sum_{i=1}^n \text{Gap}_{A_{k-1}^i}(t) + \sum_{i=1}^n \text{Gap}_{A_{k-1}^{i+\alpha_k}}(t) \\ &\quad + \sum_{i=1}^n 2 \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| \\ &= 2 \sum_{i=1}^n \text{Gap}_{A_{k-1}^i}(t) + 2 \sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)|. \end{aligned}$$

After dividing the above inequality by 2 and applying Lemma 6, we get that:

$$\sum_{i=1}^n \text{Gap}_{A_k^i}(t) \leq \sum_{i=1}^n \text{Gap}_{A_{k-1}^i}(t) + \left\lceil \frac{n}{\lfloor \frac{n}{2^k} \rfloor} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\phi_{\alpha_k}(t)}.$$

Notice that for any i , $\text{Gap}_{A_0^i}(t) = 0$. Hence, after summing up the above inequality for $k = 1$ to $\lfloor \log n \rfloor$ we get that

$$\sum_{i=1}^n \text{Gap}_{A_{\lfloor \log n \rfloor}^i}(t) \leq \sum_{k=1}^{\lfloor \log n \rfloor} \left\lceil \frac{n}{\lfloor \frac{n}{2^k} \rfloor} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\phi_{\alpha_k}(t)}.$$

Let $i' = \arg \min_i \text{Gap}_{A_{\lfloor \log n \rfloor}^i}(t)$. Notice that consecutive vertices in $A_{\lfloor \log n \rfloor}^{i'}$ are 1 or 2 edges apart, hence for any $1 \leq i \leq n$, either $i \in A_{\lfloor \log n \rfloor}^{i'}$ or $i + 1 \in A_{\lfloor \log n \rfloor}^{i'}$. This gives us that

$$\begin{aligned} \text{Gap}(t) &\leq \text{Gap}_{A_{\lfloor \log n \rfloor}^{i'}}(t) + 2 \max_i |x_i(t) - x_{i+1}(t)| \\ &= \text{Gap}_{A_{\lfloor \log n \rfloor}^{i'}}(t) + 2 \sqrt{\max_i |x_i(t) - x_{i+1}(t)|^2} \leq \text{Gap}_{A_{\lfloor \log n \rfloor}^{i'}}(t) + 2\sqrt{\phi_1(t)}. \end{aligned}$$

By combining the above two inequalities we get that

$$\begin{aligned} n\text{Gap}(t) &\leq n\text{Gap}_{A_{\lfloor \log n \rfloor}^{i'}}(t) + 2n\sqrt{\phi_1(t)} \leq \sum_{i=1}^n \text{Gap}_{A_{\lfloor \log n \rfloor}^i}(t) + 2n\sqrt{\phi_1(t)} \\ &\leq \sum_{k=1}^{\lfloor \log n \rfloor} \left\lceil \frac{n}{\lfloor \frac{n}{2^k} \rfloor} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\phi_{\alpha_k}(t)} + n\sqrt{\phi_1(t)}. \end{aligned}$$




Fig. 1 The evolution of average gap divided by square root of n , where n is the number of bins

Next, we apply Jensen’s inequality and Lemma 1 (we are going to use a looser upper bound: $\mathbb{E}[\phi_i(t)] \leq i(n - i) - 1 \leq in$)

$$\begin{aligned}
 n\mathbb{E}[Gap(t)] &\leq 2n\mathbb{E}[\sqrt{\phi_1(t)}] + \sum_{k=1}^{\lfloor \log n \rfloor} \left\lceil \frac{n}{2^k} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \mathbb{E}\sqrt{\phi_{\alpha_k}(t)} \\
 &\leq 2n\sqrt{\mathbb{E}[\phi_1(t)]} + \sum_{k=1}^{\lfloor \log n \rfloor} \left\lceil \frac{n}{2^k} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\mathbb{E}[\phi_{\alpha_k}(t)]} \\
 &\leq 2n\sqrt{n} + \sum_{k=1}^{\lfloor \log n \rfloor} \left\lceil \frac{n}{2^k} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\alpha_k n} = O(n\sqrt{n} \log n).
 \end{aligned}$$

This completes the proof.

6 Experimental Validation

On the practical side, we implemented our load balancing algorithm with unit weight increments on a cycle. The results confirm our hypothesis that the gap is of order $\Theta(\sqrt{n})$. In our experiment we observe the evolution of gap as we perform up to 10^9 increment operations. In Fig. 1 we ran our experiment 100 times and calculated average gap over all runs. x -axis shows number of balls thrown (which is the same as the number of increments) and y -axis is current average gap divided by \sqrt{n} . The experiment shows that once the number of thrown balls is large enough, the gap stays between \sqrt{n} and $1.4\sqrt{n}$.

7 Harary Graph, Upper Bound on the Gap

Recall that the Harary graph $H_{k,n}$ is a k -connected graph with n vertices, which has the smallest possible number of edges. In this section we show that our approach can be extended to the Harary graph $H_{4,n}$ (unless specified we will assume $H_{4,n}$ to be the Harary graph): each vertex i is connected with edges to vertices $i - 1, i + 1$ (called cycle edges), $i - 2$ and $i + 2$ (called extra edges). As before, the operation consists of picking an edge u.a.r and doing increment and averaging (for the simplicity we assume that increments have unit weights, the result can be extended to the random weights, in the similar fashion to the cycle case). After careful calculations which mimic the calculations for the cycle case, but by taking extra edges of the Harary graph into the account, we can derive the following equations for the hop potentials (hop distance is counted over the cycle edges only):

$$\begin{cases} \mathbb{E}[\phi_1(t + 1)] = \frac{n-2}{n} \mathbb{E}[\phi_1(t)] + \frac{3}{4} + \frac{\mathbb{E}[\phi_1(t)]}{4n} + \frac{\mathbb{E}[\phi_3(t)]}{2n} \\ \mathbb{E}[\phi_2(t + 1)] = \frac{n-2}{n} \mathbb{E}[\phi_2(t)] + \frac{3}{4} - \frac{\mathbb{E}[\phi_2(t)]}{4n} + \frac{\mathbb{E}[\phi_3(t)]}{2n} + \frac{\mathbb{E}[\phi_4(t)]}{2n} \\ \dots \\ \mathbb{E}[\phi_k(t + 1)] = \frac{n-2}{n} \mathbb{E}[\phi_k(t)] + 1 - \frac{\mathbb{E}[\phi_1(t)]}{2n} - \frac{\mathbb{E}[\phi_2(t)]}{2n} + \frac{\mathbb{E}[\phi_{k-2}(t)]}{2n} + \frac{\mathbb{E}[\phi_{k-1}(t)]}{2n} \\ \quad + \frac{\mathbb{E}[\phi_{k+1}(t)]}{2n} + \frac{\mathbb{E}[\phi_{k+2}(t)]}{2n} \end{cases}$$

Recall that for the cycle potential $\phi_k(t + 1)$ depends on the potentials $\phi_k(t), \phi_{k+1}(t), \phi_{k-1}(t)$ and $\phi_1(t)$ (please see Eq. (8)). In the case of Harary graph $\phi_k(t + 1)$ depends on the potentials $\phi_{k+2}(t), \phi_{k+1}(t), \phi_k(t), \phi_{k-1}(t), \phi_{k-2}(t), \phi_1(t)$ and $\phi_2(t)$. The reason is that we are able to perform load balancing operation on two hop neighbours. Similarly, if we have a graph where each vertex is connected with all vertices which are at hop distance at most ℓ (this is the Harary graph $H_{2\ell,n}$), then $\phi_k(t + 1)$ will depend on $\phi_{k+\ell}(t), \dots, \phi_{k+1}(t), \phi_k(t), \phi_{k-1}(t), \dots, \phi_{k-\ell}(t)$, and $\phi_1(t), \phi_2(t), \dots, \phi_\ell(t)$. Next step is to find the stationary points for the hop potentials. That is: the values which stay the same after we apply step given by the above equations. As before (please see Lemma 1), these values will be used as the upper bounds for the expected values of potentials. In the case of $H_{4,n}$, we get that for every t and k :

$$\mathbb{E}[\phi_k(t)] \leq \frac{2}{5}k(n - k) + \alpha. \tag{20}$$

Here, extra term α has a closed form, which we omit and instead concentrate on the property that it is upper bounded by $2n$. Observe that since the Harary graph contains cycle and we defined our hop potentials based on the hop counts of that cycle, we can upper bound the gap by using:

$$\mathbb{E}[Gap(t)] \leq \sum_{k=1}^m \frac{1}{\sqrt{2^{m-k}}} \mathbb{E}\sqrt{\phi_{2^{m-k}}(t)} \leq O(\sqrt{n} \log n). \tag{21}$$

Notice that since $\alpha \leq 2n$, for large enough k and n , the upper bound for $\mathbb{E}[\phi_k(t)]$ can be two times smaller than the upper bound for the cycle case, which was shown in Lemma 1. Hence, we can use this to slightly improve the constant hidden by big O notation in the upper bound. In general, we conjecture that for any pair of parameters $1 \leq l_1 < l_2 \leq n$, a gap for the Harary graph $H_{2l_2, n}$ is smaller than a gap for the Harary graph $H_{2l_1, n}$. Unfortunately, we are not able to provide the exact dependence of a gap on the Harary graph parameter.

8 Discussion and Future Work

We have shown that in the case of dynamic averaging on a cycle the gap between highest and lowest loaded bins is upper bounded by $O(\sqrt{n} \log n)$ in expectation. Additionally we showed that the expected square of the gap is lower bounded by $\Omega(n)$. In the future, it would be interesting to further tighten our results, matching our experimental analysis. We conjecture that the “correct” bound on the expected gap is $\Theta(\sqrt{n})$. As already discussed, we also plan to extend our results to more general graph families, in particular grid graphs.

8.1 Comparison of Two-Choice and Averaging Load Balancing

Finally, it is interesting to ask if it is possible to extend our gap bounds to the case of the classic two-choice load balancing process. In particular, is it possible to show that the gap in the case of averaging process is always smaller in expectation than the gap in the case of two choice process? Intuitively this should be the case, since the load balancing operation in the case of averaging can be viewed as picking up a random edge, incrementing the load of the less loaded endpoint, and then averaging the values. The extra averaging step should not make the gap larger. Indeed, the exponential potential used to analyse the gap in [17] can be used to upper bound the gap for the averaging process, since the exponential function is convex and averaging values does not increase it (by Jensen’s inequality).

Unfortunately, it is not clear if averaging helps to actually *decrease* the exponential potential. Additionally, this argument shows that averaging does not make the gap worse if applied to the particular technique of upper bounding the gap, and it is not clear if the gap itself is actually smaller, if we use averaging on top of the two-choice process. We conjecture that there exists a majorization argument which is based on *how often* the process performs the averaging step. More precisely, we consider the setting where after the increment step (using two choice), we perform averaging with probability β . The gap should decrease in expectation as we increase β . Note that the only result which lower bounds the gap for the two-choice process on the cycle is the straightforward $\Omega(\log n)$ lower bound which can be shown for the clique [17]. Where The lower bound comes from the observation that if $\Omega(n \log n)$ balls of weight one are placed into n bins according to the two choice process, then the average load is $\Omega(\log n)$ and with constant probability there exists a bin which is empty. What would make the existence of the majorization argument interesting is that it would allow us to

show that the lower bound we derived on the second moment of the gap while always performing averaging step on the cycle ($\beta = 1$) can be automatically used as the lower bound on the gap for two choice on the cycle ($\beta = 0$). We plan to investigate this connection in future work.

Acknowledgements The authors sincerely thank Thomas Sauerwald and George Giakkoupis for insightful discussions, and Mohsen Ghaffari, Yuval Peres, and Udi Wieder for feedback on earlier versions of this draft. We also thank the ICALP anonymous reviewers for their very useful comments.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alistarh, D., Brown, T., Kopinsky, J., Li, J.Z., Nadiradze, G.: Distributionally linearizable data structures. arXiv preprint [arXiv:1804.01018](https://arxiv.org/abs/1804.01018) (2018)
2. Alistarh, D., Kopinsky, J., Li, J., Nadiradze, G.: The power of choice in priority scheduling. In: Proceedings of the ACM Symposium on Principles of Distributed Computing, pp. 283–292. ACM (2017)
3. Azar, Y., Broder, A.Z., Karlin, A.R., Upfal, E.: Balanced allocations. *SIAM J. Comput.* **29**(1), 180–200 (1999)
4. Berenbrink, P., Czumaj, A., Steger, A., Vöcking, B.: Balanced allocations: the heavily loaded case. In: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, STOC'00, New York, NY, USA. ACM, pp. 745–754 (2000)
5. Berenbrink, P., Friedetzky, T., Hu, Z.: A new analytical method for parallel, diffusion-type load balancing. vol. 2006 (2006)
6. Boillat, J.E.: Load balancing and Poisson equation in a graph. *Concurr. Pract. Exp.* **2**, 289–314 (1990)
7. Boyd, S., Ghosh, A., Prabhakar, B., Shah, D.: Randomized gossip algorithms. *IEEE/ACM Trans. Netw.* **14**(SI), 2508–2530 (2006)
8. Czumaj, A., Stemann, V.: Randomized allocation processes. In: Proceedings 38th Annual Symposium on Foundations of Computer Science. IEEE, pp. 194–203 (1997)
9. Frieze, A., Melsted, P., Mitzenmacher, M.: An analysis of random-walk cuckoo hashing. In: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques. Springer, pp. 490–503 (2009)
10. Ghosh, B., Leighton, F., Maggs, B., Muthukrishnan, S., Plaxton, C., Rajaraman, R., Richa, A., Tarjan, R., Zuckerman, D.: Tight analyses of two local load balancing algorithms, vol. 29, pp. 548–558 (1995)
11. Ghosh, B., Muthukrishnan, S.: Dynamic load balancing by random matchings. *J. Comput. Syst. Sci.* **53**(3), 357–370 (1996)
12. Ghosh, B., Muthukrishnan, S., Schultz, M.H.: First and second order diffusive methods for rapid, coarse, distributed load balancing (extended abstract). In: Proceedings of the Eighth Annual ACM Symposium on Parallel Algorithms and Architectures, SPAA'96, New York, NY, USA. Association for Computing Machinery, pp. 72–81 (1996)
13. Harary, F.: The maximum connectivity of a graph. *Proc. Natl. Acad. Sci. USA* **48**(7), 1142 (1962)
14. Mitzenmacher, M.: The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distrib. Syst.* **12**(10), 1094–1104 (2001)
15. Muthukrishnan, S., Ghosh, B., Schultz, M.H.: First-and second-order diffusive methods for rapid, coarse, distributed load balancing. *Theory Comput. Syst.* **31**(4), 331–354 (1998)
16. Peres, Y.: Personal Communication (2015)

17. Peres, Y., Talwar, K., Wieder, U.: Graphical balanced allocations and the $(1 + \beta)$ -choice process. *Random Struct. Algorithms* **47**(4), 760–775 (2015)
18. Raab, M., Steger, A.: “balls into bins”—a simple and tight analysis. In: *International Workshop on Randomization and Approximation Techniques in Computer Science*. Springer, pp. 159–170 (1998)
19. Rabani, Y., Sinclair, A., Wanka, R.: Local divergence of Markov chains and the analysis of iterative load-balancing schemes. In: *Proceedings 39th Annual Symposium on Foundations of Computer Science* (Cat. No. 98CB36280), pp. 694–703 (1998)
20. Richa, A.W., Mitzenmacher, M., Sitaraman, R.: The power of two random choices: a survey of techniques and results. *Comb. Optim.* **9**, 255–304 (2001)
21. Sauerwald, T., Sun, H.: Tight bounds for randomized load balancing on arbitrary network topologies. In: *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*. IEEE, pp. 341–350 (2012)
22. Sinclair, A., Jerrum, M.: Approximate counting, uniform generation and rapidly mixing Markov chains. *Inf. Comput.* **82**(1), 93–133 (1989)
23. Talwar, K., Wieder, U.: Balanced allocations: a simple proof for the heavily loaded case. In: Esparza, J., Fraigniaud, P., Husfeldt, T., Koutsoupias, E. (eds.) *Automata, Languages, and Programming—41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8–11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*. Springer, pp. 979–990 (2014)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Dan Alistarh¹ · Giorgi Nadiradze¹  · Amirmojtaba Sabour¹

✉ Giorgi Nadiradze
giorgi.nadiradze@ist.ac.at

Dan Alistarh
dan.alistarh@ist.ac.at

Amirmojtaba Sabour
amirmojtaba.sabour@ist.ac.at

¹ IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria