

Dynamic Averaging Load Balancing on Cycles

Dan Alistarh

IST Austria, Klosterneuburg, Austria

Giorgi Nadiradze

IST Austria, Klosterneuburg, Austria

Amirmojtaba Sabour

IST Austria, Klosterneuburg, Austria

Abstract

We consider the following dynamic load-balancing process: given an underlying graph G with n nodes, in each step $t \geq 0$, one unit of load is created, and placed at a randomly chosen graph node. In the same step, the chosen node picks a random neighbor, and 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.

Variants of the above graphical balanced allocation process have been studied previously by Peres, Talwar, and Wieder [10], and by Sauerwald and Sun [12]. These authors left as open the question of characterizing the gap in the case of *cycle graphs* in the *dynamic* case, where weights are created during the algorithm’s execution. For this case, the only known upper bound is of $\mathcal{O}(n \log n)$, following from a majorization argument due to [10], which analyzes a related graphical allocation process.

In this paper, we provide an upper bound of $\mathcal{O}(\sqrt{n} \log n)$ on the expected gap of the above process for cycles of length 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 provide analytical and experimental evidence that our upper bound on the gap is tight up to a logarithmic factor.

2012 ACM Subject Classification Mathematics of computing → Combinatorics

Keywords and phrases Algorithms, Load Balancing

Digital Object Identifier 10.4230/LIPIcs.ICALP.2020.7

Category Track A: Algorithms, Complexity and Games

Related Version A full version of the paper is available at <https://arxiv.org/abs/2003.09297>.

Funding *Dan Alistarh*: European Research Council funding award PR1042ERC01

Giorgi Nadiradze: European Research Council funding award PR1042ERC01

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.

1 Introduction

We consider 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 processes has been used to model several practical allocation problems, such as load-balancing [3, 7, 11], hashing [5], or even relaxed data structures [2, 1].



© Dan Alistarh, Giorgi Nadiradze, and Amirmojtaba Sabour;
licensed under Creative Commons License CC-BY

47th International Colloquium on Automata, Languages, and Programming (ICALP 2020).

Editors: Artur Czumaj, Anuj Dawar, and Emanuela Merelli; Article No. 7; pp. 7:1–7:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



One way to put our results into context is to view them as a variation of the well-known d -choice process, in which, in each step, a new weight is generated, and is placed in the least loaded of d uniform random choices. If $d = 1$, then we have the uniform random choice scheme, whose properties are well understood. 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, Broder, Karlin, and Upfal [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 $1/\text{poly } n$. An elegant alternative proof for a slightly weaker version of this result was later provided by Talwar and Wieder [13].

More recently, Peres, Talwar, and Wieder [10] analyzed 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 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 [10], this gap is $O(\log n / \beta)$, with probability $1 - 1/\text{poly } n$.

An alternative way to frame our results is to consider *static load-balancing*, where each node starts with an arbitrary initial load, and the endpoints *average* their initial loads whenever the edge is chosen. Then, the balancing process can be mapped to a Markov chain, whose convergence is well-understood in terms of the spectral gap of the underlying graph [12]. Sauerwald and Sun [12] considered this static case in the *discrete* setting, where the fixed initial load can only be divided to *integer* tokens upon each averaging step, for which they gave strong upper bounds for a wide range of graph families.

By contrast to this previous work, in this paper we consider the less complex *continuous averaging case*, where exact averaging of the weights is possible, but in the more challenging *dynamic* scenario, where weights arrive in each step rather than being initially allocated, which is closer to the setting of the d -choice process discussed above.

One question left open by previous work concerns the evolution of the gap in the dynamic case on graphs of low expansion, such as cycles. In particular, for cycles, the only known upper bound on the expected gap in the dynamic case is of $O(n \log n)$, following from [10], whereas the only lower bound is the immediate $\Omega(\log n)$ gap lower bound for the clique. Closing this gap for cycle graphs is known to be a challenging open problem [9]. As suggested in [10], 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.

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 complement this result with a lower bound of $\Omega(n)$ on the $\mathbb{E}[Gap(t)^2]$, as well as additional experimental evidence suggesting that our upper bound is tight within a logarithmic factor. Our results extend to *weighted* input. That is, we can allow our input to come from any distribution W , such that $\mathbb{E}[W^2] \leq M^2$, for some $M > 0$.

Technical Argument. 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 node 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 on regular graphs. 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:

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

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 $\{i, i + 2^k, i + 2 \times 2^k, i + 3 \times 2^k, \dots\}$. (Since we are on a cycle, $i = i + 2^{m-k}2^k$.) Then for any $1 \leq i \leq n$ and $k > 0$, we will have

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

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 across 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, see Section 5.) We obtain that

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

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.

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 lower bound in Equation 2. 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*:

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

This second moment bound strongly suggests that our above analysis is tight within logarithmic factors. We conjecture that the bound is also tight with regards to the expected gap, and examine this claim empirically in Section 6.

Extensions and Overview. The analysis template we described above is general, and could be extended to other graph families, such as regular expanders. In particular, we note that the recursive structure of the k -hop potentials is preserved for such graphs. The main technical steps in analyzing a new graph family are to (1) identify the right upper bound on the k -hop potential (the analogue of (1)); and (2) identify the right set family for the gap covering argument, and its recursive structure (the analogue of (2)). Obtaining tight bounds for these quantities is not straightforward, since they do not seem to be immediately linked to well-studied graph properties. 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 Harary graphs [6] in a full version of the paper.

We discuss the relation between our results and bounds for the graphical power-of-two process on a cycle [10] in Section 7.

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. In the *static* case with *discrete averaging*, the problem has been considered by Sauerwald and Sun [12]. However, their techniques would not apply 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 [10]: they consider 2-choice load balancing, and one can re-do their potential analysis for (continuous) averaging load balancing, yielding the same bounds. However, as our bounds show, the resulting analysis is quite loose in the case of cycles, yielding an $\Omega(\sqrt{n})$ gap. This is a consequence of the majorization technique used, which links dynamic averaging on the cycle and a very weak form of averaging on the clique.

Our potential analysis is substantially different from that of [10], 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 notice that this is also different from the usual square potentials used for analyzing averaging load balancing, e.g. [8], 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

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, in each step, a weight $w(t)$ is generated from a same distribution W . We associate a real-valued *load* value $x_i(t)$ with each node i . In step t , an edge $(i, i + 1)$ is chosen uniformly at random, and the two endpoints 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 t balls have been thrown. 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 $\phi_1(t+1)$:

$$\begin{aligned} \mathbb{E}[\phi_1(t+1)|X(t), w(t)] &= \sum_{i=1}^n \frac{1}{n} \left(\left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i+2}(t) \right)^2 \right. \\ &\quad \left. + \left(\frac{x_i(t) + x_{i+1}(t) + w(t)}{2} - x_{i-1}(t) \right)^2 \right. \\ &\quad \left. + \sum_{j \neq i-1, i, i+1} (x_j(t) - x_{j+1}(t))^2 \right) \\ &= \frac{n-3}{n} \phi_1(t) + \frac{1}{2} + \frac{1}{2n} (\phi_1(t) + 2\phi_2(t)) \\ &= \frac{n-2}{n} \phi_1(t) + \frac{1}{2} (w(t)^2 - \frac{\phi_1(t)}{n}) + \frac{1}{n} \phi_2(t). \end{aligned}$$

Now, we proceed with calculating the expected value of $\phi_k(t+1)$, for $2 \leq k \leq \lfloor n/2 \rfloor$:

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

Note that in the above calculations for $\phi_1(t+1)$ and $\phi_k(t+1)$, for $k > 1$ the terms which contain $w(t)$ as linear multiplicative term disappear because we can assume that loads $x_1(t), x_2(t), \dots, x_n(t)$ are normalized (this will not change our potentials) and we have:

$$\sum_{i=1}^n w(t)x_i(t) = 0. \tag{6}$$

7:6 Dynamic Averaging Load Balancing on Cycles

If we remove conditioning on $w(t)$ and express these equations for $k = 1, 2, \dots, n-1$, we get:

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

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*

$$\mathbb{E}[\phi_k(t)] \leq (k(n-k) - 1)\mathbb{E}[W^2] \leq k(n-k) - 1. \quad (7)$$

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

$$\mathbb{E}[\Phi(t+1)|\Phi(t) = Y] = Y. \quad (8)$$

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

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

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 $z_i(t) = z_{n-i}(t)$):

$$\left\{ \begin{array}{l} n \times \mathbb{E}[z_1(t+1)|X(t)] = (n-2 - \frac{1}{2})z_1(t) + z_2(t). \\ n \times \mathbb{E}[z_2(t+1)|X(t)] = -z_1(t) + z_1(t) + (n-2)z_2(t) + z_3(t). \\ n \times \mathbb{E}[z_3(t+1)|X(t)] = -z_1(t) + z_2(t) + (n-2)z_3(t) + z_4(t). \\ \dots \\ n \times \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t+1)|X(t)] = -z_1(t) + z_{\lfloor \frac{n}{2} \rfloor - 1}(t) + (n-2)z_{\lfloor \frac{n}{2} \rfloor}(t) + z_{\lfloor \frac{n}{2} \rfloor + 1}(t). \end{array} \right. \quad (9)$$

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

$$0 \leq \mathbb{E}[z_1(t)] \leq \mathbb{E}[z_2(t)] \leq \dots \leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t)]. \quad (10)$$

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:

$\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor + 1}(t)] = \mathbb{E}[z_{\lceil \frac{n}{2} \rceil - 1}(t)] \geq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor - 2}(t)]$. 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 theorem. ◀

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

In this section we upper bound a gap in expectation for $n = 2^m$ case. The proof for the general case is quite technical but not necessarily more interesting, and is provided in the Section 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 $\{i, i + 2^k, i + 2 \times 2^k, i + 3 \times 2^k, \dots\}$ (Notice that $i = i + 2^{m-k}2^k$). Our proof works as follows: for each $1 \leq i \leq n$ and $0 < k \leq m$, we look at the vertices given by the sets A_k^i and $A_k^{i+2^{k-1}}$ and try to characterise the gap after we merge those sets (Note that this will give us the gap for the set $A_{k-1}^i = A_k^i \cup A_k^{i+2^{k-1}}$). Using this result, we are able to show that $\sum_{i=1}^n Gap_{A_{k-1}^i}(t)$ is upper bounded by $\sum_{i=1}^n Gap_{A_k^i}(t)$ plus n times maximum load difference between vertices at hop distance 2^{k-1} . Next, we use 2^{k-1} hop distance potential $\phi_{2^{k-1}}(t)$ to upper bound maximum load between the vertices at hop distance 2^{k-1} . Using induction on k , we are able to upper bound $\sum_{i=1}^n Gap_{A_0^i}(t)$ in terms of $\sum_{i=1}^n Gap_{A_m^i}(t)$ and $\sum_{k=1}^m \phi_{2^{k-1}}(t)$. Notice that by our definitions, for each i , $Gap_{A_m^i}(t) = 0$ (A_m^i contains only vertex i) and $Gap_{A_0^i}(t) = Gap(t)$ (A_0^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.

7:8 Dynamic Averaging Load Balancing on Cycles

► **Lemma 2.** For any $1 \leq i \leq n$ and $k > 0$, we have that

$$2\text{Gap}_{A_{k-1}^i}(t) \leq 2 \max_{j \in A_{k-1}^i} |x_j(t) - x_{j+2^{k-1}}(t)| + \text{Gap}_{A_k^{i+2^{k-1}}}(t) + \text{Gap}_{A_k^i}(t). \quad (11)$$

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

Case 1. $u \in A_k^i$ and $v \in A_k^i$. Then $\text{Gap}_{A_{k-1}^i}(t) = \text{Gap}_{A_k^i}(t)$ and we have that

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

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

$$2\text{Gap}_{A_{k-1}^i}(t) \leq 2 \max_{j \in A_{k-1}^i} |x_j(t) - x_{j+2^{k-1}}(t)| + \text{Gap}_{A_k^{i+2^{k-1}}}(t) + \text{Gap}_{A_k^i}(t). \quad (12)$$

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

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

and

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

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

$$2\text{Gap}_{A_{k-1}^i}(t) \leq 2 \max_{j \in A_{k-1}^i} |x_j(t) - x_{j+2^{k-1}}(t)| + \text{Gap}_{A_k^{i+2^{k-1}}}(t) + \text{Gap}_{A_k^i}(t). \quad (13)$$

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

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

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

► **Lemma 3.**

$$\sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+2^k}(t)| \leq \frac{n}{\sqrt{2^k}} \sqrt{\phi_{2^k}(t)}. \quad (14)$$

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

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

Where we used a fact that sets $A_k^1, A_k^2, \dots, A_k^{2^k}$ are disjoint. \blacktriangleleft

Finally, using the two Lemmas above and Theorem 1 we can upper bound the expected gap at step t :

► **Theorem 4.** *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

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

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

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

Notice that $\sum_{i=1}^n \text{Gap}_0^i(t) = n\text{Gap}(t)$ and we also have that

$$\sum_{i=1}^n \text{Gap}_{\frac{n}{2}}^i(t) = \sum_{i=1}^n |x_i(t) - x_{i+\frac{n}{2}}(t)| \leq \sqrt{n} \sqrt{\sum_{i=1}^n |x_i(t) - x_{i+\frac{n}{2}}(t)|^2} = \sqrt{n} \sqrt{\phi_{\frac{n}{2}}(t)}$$

Hence, we get that

$$\begin{aligned} n\text{Gap}(t) = \sum_{i=1}^n \text{Gap}_0^i(t) &\leq \sum_{i=1}^n \text{Gap}_{\frac{n}{2}}^i(t) + \sum_{k=1}^{m-1} \frac{n}{\sqrt{2^{k-1}}} \sqrt{\phi_{2^{k-1}}(t)} \\ &\leq \sqrt{n} \sqrt{\phi_{\frac{n}{2}}(t)} + \sum_{k=1}^{m-1} \frac{n}{\sqrt{2^{k-1}}} \sqrt{\phi_{2^{k-1}}(t)}. \end{aligned}$$

7:10 Dynamic Averaging Load Balancing on Cycles

Next, we apply Jensen and Theorem 1:

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

This gives us the proof of the theorem. ◀

4 Gap Lower Bound

Next we prove the following theorem, which provides strong evidence that our bound on the gap is tight within a logarithmic factor.

► **Theorem 5.** *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 Equations 9 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 10 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 10 means that

$$\mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + \lfloor \frac{n}{2} \rfloor + 1)] \leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + 1)] - \sum_{i=t}^{t+\lfloor \frac{n}{2} \rfloor} \frac{\mathbb{E}[z_1(i)]}{n} \quad (15)$$

$$\leq \mathbb{E}[z_{\lfloor \frac{n}{2} \rfloor}(t + 1)] - \frac{\mathbb{E}[z_1(t + \lfloor \frac{n}{2} \rfloor)]}{n} \quad (16)$$

Again by using Equations 9 and Inequalities 10, 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 15 we get that

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

Because $\left(1 - \left(\frac{1}{n}\right)^{\lfloor \frac{n}{2} \rfloor - 1}\right) < 1$ 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 on 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 the Theorem 4 for the general case, we need to redefine our sets A_i^k . 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,

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

First we prove the following lemma:

► **Lemma 6.** For any $\lfloor \log n \rfloor \geq k > 0$, we have that

(1) $\sum_{i=1}^{2^k} \delta_k^i = n$.

(2) For any $1 \leq i \leq 2^k$, $\delta_k^i \in \{\lceil \frac{n}{2^k} \rceil, \lfloor \frac{n}{2^k} \rfloor\}$ (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. ◀

7:12 Dynamic Averaging Load Balancing on Cycles

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

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

It is easy to see that for any $\lfloor \log n \rfloor \geq k > 0$ and i , we have that $|A_k^i| = 2^k$, $A_k^i = A_{k-1}^i \cup A_{k-1}^{i+\alpha_k}$ and $A_{k-1}^i \cap A_{k-1}^{i+\alpha_k} = \emptyset$. Also notice that for any $u \in A_{k-1}^i$, there exists $v \in A_{k-1}^{i+\alpha_k}$, such that $u + \alpha_k = v$ or $v + \alpha_k = u$ (For any $u \in A_{k-1}^{i+\alpha_k}$ there exists $v \in A_{k-1}^i$ with the same property).

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

► **Lemma 7.** *For any $1 \leq i \leq n$ and $\lfloor \log n \rfloor \geq k > 0$, we have that*

$$2\text{Gap}_{A_k^i}(t) \leq 2 \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| + \text{Gap}_{A_{k-1}^{i+\alpha_k}}(t) + \text{Gap}_{A_{k-1}^i}(t). \quad (17)$$

Proof. Let $u = \arg \max_{j \in A_{k-1}^i} x_j(t)$ and let $v = \arg \min_{j \in A_{k-1}^i} x_j(t)$. We consider several cases:

Case 1. $u \in A_{k-1}^i$ and $v \in A_{k-1}^i$. Notice that in this case $\text{Gap}_{A_{k-1}^i}(t) = \text{Gap}_{A_k^i}(t)$. Let $u' \in A_{k-1}^{i+\alpha_k}$ be the vertex such that $u + \alpha_k = u'$ or $u' + \alpha_k = u$ and let $v' \in A_{k-1}^{i+\alpha_k}$ be the vertex such that $v + \alpha_k = v'$ or $v' + \alpha_k = v$. We have that

$$\begin{aligned} \text{Gap}_{A_k^i}(t) &= |x_u(t) - x_v(t)| \\ &\leq |x_{u'}(t) - x_u(t)| + |x_{v'}(t) - x_v(t)| + |x_{u'}(t) - x_{v'}(t)| \\ &\leq |x_{u'}(t) - x_u(t)| + |x_{v'} - x_v(t)| + \text{Gap}_{A_{k-1}^{i+\alpha_k}}(t) \\ &\leq 2 \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| + \text{Gap}_{A_{k-1}^{i+\alpha_k}}(t). \end{aligned}$$

This gives us that

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

Case 2. $u \in A_{k-1}^i$ and $v \in A_{k-1}^{i+\alpha_k}$. Let $u' \in A_{k-1}^{i+\alpha_k}$ be the vertex such that $u + \alpha_k = u'$ or $u' + \alpha_k = u$ and let $v' \in A_{k-1}^i$ be the vertex such that $v + \alpha_k = v'$ or $v' + \alpha_k = v$. We have that:

$$\begin{aligned} \text{Gap}_{A_k^i}(t) &= |x_u(t) - x_v(t)| \leq |x_u(t) - x_{v'}(t)| + |x_{v'}(t) - x_v(t)| \\ &\leq \text{Gap}_{A_{k-1}^i}(t) + \max_{j \in A_k^i} (|x_j(t) - x_{j+\alpha_k}(t)|) \end{aligned}$$

and

$$\begin{aligned} \text{Gap}_{A_k^i}(t) &= |x_u(t) - x_v(t)| \leq |x_u(t) - x_{u'}(t)| + |x_{u'}(t) - x_v(t)| \\ &\leq \text{Gap}_{A_{k-1}^{i+\alpha_k}}(t) + \max_{j \in A_k^i} (|x_j(t) - x_{j+\alpha_k}(t)|) \end{aligned}$$

Hence, we again get that

$$2\text{Gap}_{A_k^i}(t) \leq 2 \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| + \text{Gap}_{A_{k-1}^{i+\alpha_k}}(t) + \text{Gap}_{A_{k-1}^i}(t). \quad (19)$$

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

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

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

► **Lemma 8.**

$$\sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| \leq \left\lceil \frac{n}{\lfloor \frac{n}{2^k} \rfloor} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\phi_{\alpha_k}(t)} \quad (20)$$

Proof. Notice that for any $1 \leq u \leq n$ and 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:

$$\sum_{i=1}^n \max_{j \in A_k^i} |x_j(t) - x_{j+\alpha_k}(t)| \leq \left\lceil \frac{n}{\lfloor \frac{n}{2^k} \rfloor} \right\rceil \sqrt{\lfloor \frac{n}{2^k} \rfloor} \sqrt{\phi_{\alpha_k}(t)} \quad (21)$$

◀

With the above lemmas in place, we are ready to prove Theorem 4 for general n .

From Lemma 7 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 8: 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 , $\text{Gap}_0^i(t) = 0$. Hence, we get that

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

7:14 Dynamic Averaging Load Balancing on Cycles

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

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

By combining the above two inequalities we get that

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

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

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

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 Figure 1 we ran our experiment 100 times and calculated average gap over the 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 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 of $\Theta(\sqrt{n})$. As already discussed, we also plan to extend our results to more general graph families, in particular grid graphs.

Comparison of two-choice and averaging load balancing. Finally, it is interesting to ask if it possible to extend our gap bounds in the case of the classic two-choice load balancing process. In particular, it is possible to show that the gap in the case of averaging process is



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

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 [10] can be used to upper bound the gap for averaging, 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 [10]; so what makes 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.

References

- 1 Dan Alistarh, Trevor Brown, Justin Kopinsky, Jerry Z Li, and Giorgi Nadiradze. Distributionally linearizable data structures. *arXiv preprint*, 2018. [arXiv:1804.01018](#).
- 2 Dan Alistarh, Justin Kopinsky, Jerry Li, and Giorgi Nadiradze. The power of choice in priority scheduling. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 283–292. ACM, 2017.

- 3 Yossi Azar, Andrei Z Broder, Anna R Karlin, and Eli Upfal. Balanced allocations. *SIAM journal on computing*, 29(1):180–200, 1999.
- 4 Petra Berenbrink, Artur Czumaj, Angelika Steger, and Berthold Vöcking. Balanced allocations: The heavily loaded case. In *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing*, STOC '00, pages 745–754, New York, NY, USA, 2000. ACM. doi:10.1145/335305.335411.
- 5 Alan Frieze, Páll Melsted, and Michael Mitzenmacher. An analysis of random-walk cuckoo hashing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 490–503. Springer, 2009.
- 6 Frank Harary. The maximum connectivity of a graph. *Proceedings of the National Academy of Sciences of the United States of America*, 48(7):1142, 1962.
- 7 Michael Mitzenmacher. The power of two choices in randomized load balancing. *IEEE Transactions on Parallel and Distributed Systems*, 12(10):1094–1104, 2001.
- 8 Shanmugavelayutham Muthukrishnan, Bhaskar Ghosh, and Martin H Schultz. First-and second-order diffusive methods for rapid, coarse, distributed load balancing. *Theory of computing systems*, 31(4):331–354, 1998.
- 9 Yuval Peres, 2015. Personal Communication.
- 10 Yuval Peres, Kunal Talwar, and Udi Wieder. Graphical balanced allocations and the $(1+\beta)$ -choice process. *Random Structures & Algorithms*, 47(4):760–775, 2015.
- 11 Andrea W Richa, M Mitzenmacher, and R Sitaraman. The power of two random choices: A survey of techniques and results. *Combinatorial Optimization*, 9:255–304, 2001.
- 12 Thomas Sauerwald and He Sun. Tight bounds for randomized load balancing on arbitrary network topologies. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 341–350. IEEE, 2012.
- 13 Kunal Talwar and Udi Wieder. Balanced allocations: A simple proof for the heavily loaded case. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *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*, pages 979–990. Springer, 2014. doi:10.1007/978-3-662-43948-7_81.