

(Almost) Perfect Discrete Iterative Load Balancing

Petra Berenbrink* Robert Elsässer† Tom Friedetzky‡ Hamed Hosseinpour§
Dominik Kaaser¶ Peter Kling|| Thomas Sauerwald**

Abstract. We consider discrete, iterative load balancing via matchings on arbitrary graphs. Initially each node holds a certain number of tokens, defining the load of the node, and the objective is to redistribute the tokens such that eventually each node has approximately the same number of tokens. We present results for a general class of simple local balancing schemes where the tokens are balanced via matchings. In each round the process averages the tokens of any two matched nodes. If the sum of their tokens is odd, the node to receive the one excess token is selected at random. Our class covers three popular models: in the matching model a new matching is generated randomly in each round, in the balancing circuit model a fixed sequence of matchings is applied periodically, and in the asynchronous model the load is balanced over a randomly chosen edge.

We measure the quality of a load vector by its discrepancy, defined as the difference between the maximum and minimum load across all nodes. As our main result we show that with high probability our discrete balancing scheme reaches a discrepancy of 3 in a number of rounds which asymptotically matches the spectral bound for continuous load balancing with fractional load.

This result improves and tightens a long line of previous works, by not only achieving a small constant discrepancy (instead of a non-explicit, large constant) but also holding for arbitrary instead of regular graphs. The result also demonstrates that in the general model we consider, discrete load balancing is no harder than continuous load balancing.

1 Introduction In load balancing the goal is to distribute a set of jobs to multiple computing resources. Many large-scale applications use some form of decentralized load balancing, with examples including balancing search engine queries, distributing internet traffic, updating distributed databases, or performing numerical simulations.

In a classical model for parallel and distributed computing, n processors are connected by an arbitrary, undirected graph. Initially each processor (*node*) has a given number of jobs (*tokens*). The objective is to redistribute the tokens by moving them across the edges such that eventually each node has approximately the same number of tokens, thereby minimizing the *discrepancy*¹, defined as the difference between the maximum and minimum load across all nodes. Particularly for large-scale networks, it is desirable to employ iterative protocols that are simple and local. This means each node only needs to know its current and neighboring node's loads and, based on this information, determines the number of tokens to be transferred. While load balancing on graphs is a well-studied and natural area, there are profound connections to other problems including opinion dynamics [6], gossiping [15], information aggregation [25], and clustering [4].

Two important classifications of iterative load balancing schemes are between *diffusion models* and *matching models* [19]. In the former model every node concurrently balances its load with *all* neighbors in each round. In the latter model every node balances its load with *at most one* neighbor in each round. Another, orthogonal, distinction is between *continuous* and *discrete* load. In continuous load balancing, load tokens are assumed to

*University of Hamburg, Germany (petra.berenbrink@uni-hamburg.de).

†University of Salzburg, Austria (elsa@cs.sbg.ac.at).

‡Durham University, U.K. (tom.friedetzky@durham.ac.uk).

§University of Hamburg, Germany (hamed.hosseinpour@uni-hamburg.de).

¶Technical University of Hamburg, Germany (dominik.kaaser@tuhh.de).

||Darmstadt University of Applied Sciences, Germany (peter.kling@h-da.de).

**University of Cambridge, U.K. (tms41@cam.ac.uk).

¹Note that this term is the standard definition in the field of load balancing, in contrast to its use in discrepancy theory.

be arbitrarily divisible. Continuous iterative load balancing is essentially equivalent to a Markov chain, and the spectral gap captures the convergence speed and time to reach a constant discrepancy (e.g., [32]).

While the continuous setting is well understood, the question remains whether it is a good approximation of the discrete setting, where load is composed of unit-size tokens that are not divisible [32]. The deviation between the processes is caused by the accumulation of rounding errors across different nodes and rounds, which makes the discrete process non-linear and hard to analyze. This stark contrast between the continuous setting and the discrete setting is highlighted in Rabani, Sinclair and Wanka [32], where they call the discrete setting “*true process*” and the continuous setting “*idealized process*”. In the same work from 1998, the authors also point out that “the question of a precise quantitative relationship between Markov chains and load-balancing algorithms has been posed by several authors”, notably in [23, 27, 31, 23, 35], and “seems to be of interest in its own right.” One concretization of this question, which has been the objective of many previous works in this area [32, 34, 22, 5], is as follows:

Open Question: For a given undirected, connected graph $G = (V, E)$ and an arbitrary initial load vector in $\mathbb{N}_0^n$ with initial discrepancy K , let $\tau_{\text{cont}}(K)$ denote the number of rounds needed in the continuous setting to reach discrepancy at most 1. Find a tight bound on the discrepancy after $O(\tau_{\text{cont}}(K))$ rounds in the discrete model.

In this work we nearly close the gap between the discrete and the continuous setting by bounding the discrepancy by 3, which holds for any connected graph and a large class of models. Additionally, our analytical framework is simpler and more flexible.

Our Results. We focus on the matching model, as it requires less communication per round than the diffusion model; yet it tends to perform better in theory and practice [32, 34]. We present results for a general class of matching models including the prominent *random matching model* and the *balancing circuit model*. In the former model a new matching is generated randomly in each round, while in the latter model a fixed sequence of matchings is applied periodically. We also consider an “*asynchronous*” process in which at each iteration one edge $e \in E$ is chosen uniformly at random. The very simple load balancing process works as follows. For each edge in the matching, the two incident nodes average their load as much as possible; in case their total load is odd, the so-called excess token is assigned randomly to either node. This resembles randomized rounding in optimization, which leads to the question of how well the discrete setting can approximate the continuous one with this randomized strategy. Note that the algorithm considered in this paper is natural and simple, that is, straightforward to implement. Moreover it is strictly distributed in the sense that only local information based on a node and its neighbours are used. Due to the nature of the communication patterns (matchings) there is no multi-port communication, another nod toward real-world constraints.

As is standard in the area of discrete load balancing, we use a very strict notion of smoothness of a load vector called *discrepancy*. Let $X^{(t)}$ be the load vector at the end of round t , and recall that the discrepancy, denoted by $\text{disc}(X^{(t)})$, is defined as the difference between the maximum and minimum load across all nodes. In the following theorem we focus on the three specific models described above; our general result can be found in Theorem 2.3.

THEOREM 1.1 (simplified version of Theorem 2.3). *Let G be any undirected, connected graph with n nodes and consider any initial load vector $x^{(0)}$ with $\text{disc}(x^{(0)}) \leq K$ in the discrete setting. Then, in the random matching, balancing circuit and asynchronous model there exists a time $\tau = O(\tilde{\tau}_{\text{cont}}(K))$, where $\tilde{\tau}_{\text{cont}}(K)$ is the standard spectral bound in the continuous setting (see (2.3) and (2.4)), such that for an arbitrarily small constant $c > 0$,*

$$\Pr \left[\text{disc} \left(X^{(\tau)} \right) \leq 3 \right] \geq 1 - \exp \left(-\log^{1-c}(n) \right).$$

As is commonly done we bound the convergence time in terms of the natural spectral bound $\tilde{\tau}_{\text{cont}}(K)$ on $\tau_{\text{cont}}(K)$, instead of $\tau_{\text{cont}}(K)$ directly. Furthermore, it was shown in [34, Theorem 2.10], that this upper bound is tight for the random matching model if $K \geq n^{1+\Omega(1)}$. With this slight slack in the number of rounds, our work essentially closes the lid on the long standing problem of consolidating the continuous and the discrete setting. Concretely, we obtain the following quantitative improvements. First, the previously best known result [34] used the same balancing times but only achieved a discrepancy which is a non-explicit (and large) constant. Second, the results for constant discrepancy in [34] only hold for regular graphs, whereas our theorem applies to regular and non-regular graphs alike. Third, our result holds with higher probability compared to [34].

On a very high level, the way we reduce the discrepancy from K to 3 is as follows. The most involved building block is to show a constant bound on the maximum load for instances with a *linear* number of tokens. Then

Table 1.1: Overview of related results for the balancing circuit model. The stated results from [34, 22], as well as ours, also hold for the random matching model. $\tilde{\tau}_{\text{cont}}(K)$ is the spectral bound for the continuous setting, and it is defined in (2.3) and (2.4), respectively. All runtime bounds for randomized rounding hold with probability at least $1 - o(1)$.

Reference	Rounds	Discrepancy	Rounding	Graphs
[32]	$O(\tilde{\tau}_{\text{cont}}(K))$	$O(\tilde{\tau}_{\text{cont}}(n))$	det	all
[22]	$O(\tilde{\tau}_{\text{cont}}(K))$	$O(\sqrt{\tilde{\tau}_{\text{cont}}(n)})$	rand	all
[22]	$O(\tilde{\tau}_{\text{cont}}(K) \cdot (\log \log(n))^3)$	$O(1)$	rand	expander
[29]	$2 \log_2(n)$	16	rand	hypercube
[28]	$O(\log(n))$	2	rand	hypercube
[34]	$O(\tilde{\tau}_{\text{cont}}(K))$	$O(\log^\varepsilon n)$	rand	all
[34]	$O(\tilde{\tau}_{\text{cont}}(K) \cdot \log \log(n))$	$O(\log \log(n))$	rand	all
[34]	$O(\tilde{\tau}_{\text{cont}}(K))$	$O(1)$	rand	constant degree, $\Delta = O(1)$
Thm. 2.3	$O(\tilde{\tau}_{\text{cont}}(K))$	3	rand	all

we show that for instances with an *arbitrary* number of tokens, after τ rounds the number of tokens above the average is $O(n)$. Hence we can apply the result for a linear number of tokens, which results into a discrepancy of 4 first, and using another τ rounds we finally reduce the discrepancy down to 3. See section 3 for more details and for a more detailed comparison with previous work, see Table 1.1. We note that, with slight abuse of notation, in the balancing circuit model we consider Δ to be length of the sequence of matchings, whereas in the random matching model it is the maximum degree of the given graph.

1.1 Related Work [18, 14, 13, 15] pioneer the use of Markov chains for analyzing diffusion-based load balancing schemes in the continuous model. Note that [15] considers the asynchronous process where in each round a single edge is chosen uniformly at random, whereas the other papers consider the synchronous variant. For these processes it is known that the discrepancy can be reduced from K to ℓ within $O(\tilde{\tau}_{\text{cont}}((K/\ell)))$ rounds in the diffusion model, the balancing circuit model and the random matching model (see, e.g., [32]). All these upper bounds are essentially tight, which follows from the connection between the spectral gap of the graph and mixing times of Markov chains [24].

One of the earliest rigorous analyses of the discrete setting is due to [31]. Their algorithm always computes, for each edge, the flow of load that would occur in the continuous model, and rounds that down to obtain the numbers of tokens to be sent. For the diffusion model they show that after $O(\tilde{\tau}_{\text{cont}}(K))$ rounds the discrepancy is at most $O(\Delta n / (1 - \lambda))$, where λ is the second largest eigenvalue of the diffusion matrix. [23] shows similar results for the matching model. [32] presents a more refined analysis based on Markov chains. In the same spirit as the above two papers, they assume that an excess token is always kept at the current node. Their technique works for a large class of processes including matchings and diffusion models. Among other results, they provide bounds for general graphs showing that the discrepancy is at most $O(\Delta \tilde{\tau}_{\text{cont}}(K))$ after $O(\tilde{\tau}_{\text{cont}}(K))$ rounds in the balancing circuit model as well as in the diffusion model. In [8] the authors analyzed the asynchronous process for the complete graph. They prove that after $O(n \cdot \log(Kn))$ rounds a discrepancy of two is reached. However, to the best of our knowledge, there are no results for arbitrary graphs in the discrete, asynchronous setting.

[21] introduces a quasi-random version which rounds up or down deterministically such that the accumulated rounding errors on each edge are minimized. For torus graphs they show a constant discrepancy, while for hypercubes the discrepancy is $O(\log^{3/2}(n))$. [1] analyzes a similar framework for load balancing achieving discrepancy $O(d)$ for d -regular graphs, but their algorithm requires additional memory to keep track of the decisions in the past.

Randomized rounding was first analyzed in [22] with a focus on the matching model. Their results demonstrate that for many networks randomized rounding yields roughly a square-root improvement of the discrepancy; see Table 1.1 for more details about their results for arbitrary and regular graphs. For expanders a separate and more tailored analysis leads to a constant discrepancy in $O(\tilde{\tau}_{\text{cont}}(K) \cdot (\log \log(n))^3)$ rounds in the random matching model.

Most closely related, and in many ways our direct competitor, is the paper [34]. For *arbitrary* graphs, a discrepancy of $O(\log^\varepsilon n)$ for an arbitrarily small constant $\varepsilon > 0$ in $O(\tilde{\tau}_{\text{cont}}(K))$ rounds is shown for the random matching model and the balancing circuit model. For regular graphs the authors show constant discrepancy for the random matching model and for the balancing circuit model; for the latter their result needs $\Delta = O(1)$. The number of rounds is again $O(\tilde{\tau}_{\text{cont}}(K))$. This is a non-trivial result; the intricate analysis spanning some 40 pages results in large (and non-explicit) constant factors in the discrepancy bound as well as the running time.

There are many results in related models. [29, 28] consider smoothing networks and show that dimension-exchange on hypercubes achieves a discrepancy of 16 in $2\log_2(n)$ rounds and a discrepancy of 2 in $O(\log(n))$ rounds.² Due to the special structure of the network and the matchings, there is no dependence on K .

Another interesting albeit less related class of load balancing processes are inspired by the Rotor-Router model [16, 17], where each node evenly distributes tokens to its neighbors (usually in a deterministic way). [12] considers rotor router inspired diffusion-type algorithms. Their results only hold for d -regular graphs and their best algorithm achieves a discrepancy of d .

A different direction is dynamic load balancing, where in addition to the balancing, at each round one new token is added to a random node [10, 2]. [2] considers such a process on cycles, while in [10] the authors recently presented results for arbitrary graphs and wider class of processes. However, due to the addition of new tokens, the proven discrepancy bounds all diverge in n . The same holds for so-called “selfish” load balancing models [11] where tokens act selfishly when deciding whether to migrate to a neighboring node. There are also several works on models with dynamic load generation *and* consumption [3, 9, 7], resulting in discrepancy bounds which are either not independent of n , or showing only stability.

Road Map. The remainder of this paper is organized as follows. In Section 2 we describe our process and models formally, list some preliminary results and state our main result more formally. In Section 3 we give a detailed overview of our proof techniques. We summarize our main results and point to some open problems in Section 4. The remaining sections contain the formal analysis. In section 5, we provide some negative association and concentration results for load vectors. These results enable us to reduce the task of balancing any load vector to balancing one with at most $O(n)$ tokens. Section 6 contains the backbone of our analysis, in which we analyze load vectors that have at most $O(n)$ tokens. In Section 7 we use the results from Section 6 to show results for arbitrary load vectors. Equipped with the results in sections 6 and 7, it is relatively straightforward to establish our main results, in particular Theorem 2.3, which is done in Section 8. Finally, in the appendix we present an assortment of elementary and well-known tools applied in our analysis.

2 Model and Main Results We are given an arbitrary connected and undirected graph $G = (V, E)$ with n nodes. Initially a set $\mathcal{T}$ of unit-sized *tokens* are distributed arbitrarily among the nodes. The initial load vector is denoted $x^{(0)}$, and the load vector at (the end of) round t is denoted by $X^{(t)}$. These vectors are row-vectors and the i -th entry represents the (integral) load of node i , i.e., the number of tokens on node i . Note that due to the inherent randomization in the process, $X_i^{(t)}$ for $t > 0$ is a random variable. We will use uppercase letters for random variables and matrices, but lowercase letters for fixed outcomes. We define $\bar{x} := \sum_{i \in V} x_i^{(0)} / n$ as the average load. For an n -dimensional vector X we define the discrepancy of X as $\text{disc}(X) := \max_{i \in [n]} X_i - \min_{j \in [n]} X_j$. We usually assume that the tokens on the nodes are ordered; the height of the token is its number in that order.

2.1 Process Definition In this section we first define the standard discrete balancing process which does not specify how tokens are exchanged across the matching edges (similar to [32, 34]). After that we define a so-called height-sensitive variant of the process, which also specifies the movements of individual tokens. However, both processes generate, at any point of time, exactly the same load distribution.

For both processes we are given a *sequence of matchings* $(\mathbf{M}^{(s)})_{s=1}^\infty := (\mathbf{M}^{(1)}, \mathbf{M}^{(2)}, \dots)$. The standard load balancing process updates the (discrete) load vector iteratively as follows.

```

1 for each round  $t = 1, 2, \dots$  do
2   for each edge  $\{u, v\} \in \mathbf{M}^{(t)}$  do

```

²Note that the dimension-exchange communication scheme can be regarded as one particular instance of the balancing circuit model. When we write $\log$ we mean $\log_e$ unless specified differently.

$$(X_u^{(t)}, X_v^{(t)}) \leftarrow \begin{cases} \left(\left\lfloor \frac{X_u^{(t-1)} + X_v^{(t-1)}}{2} \right\rfloor, \left\lfloor \frac{X_u^{(t-1)} + X_v^{(t-1)}}{2} \right\rfloor \right) & \text{with probability } 1/2, \\ \left(\left\lceil \frac{X_u^{(t-1)} + X_v^{(t-1)}}{2} \right\rceil, \left\lceil \frac{X_u^{(t-1)} + X_v^{(t-1)}}{2} \right\rceil \right) & \text{with probability } 1/2. \end{cases}$$

We assume the tokens in $\mathcal{T}$ are numbered from 1 to $|\mathcal{T}|$. In each round t , each token i has a *location* $W_i^{(t)} \in V$ and a *height* $H_i^{(t)} \in \{1, \dots, X_{W_i^{(t)}}^{(t)}\}$. Initially tokens are ordered arbitrarily on each node, and the initial height of a token is its position in that order. We now define the height-sensitive process which is a realization (and refinement) of the load balancing process above (an illustration can be found in Figure 2.1). The process uses the notion of siblings: two tokens with the same height on two matched nodes are called siblings.

Height-Sensitive Process (round t , matching edge $\{u, v\} \in E$ with $X_u^{(t-1)} \geq X_v^{(t-1)}$)

1. *Moving step*: Move the top $\lceil (X_u^{(t-1)} - X_v^{(t-1)})/2 \rceil$ tokens from node u to node v , preserving their relative order and adjusting their height accordingly.
2. *Shuffling step*: Swap each token on node v with its sibling on u with probability $1/2$, independently from all other tokens. In case where the topmost token at v has no sibling, this token is also moved to u with probability $1/2$.




Figure 2.1: Illustration of the height-sensitive process and the effect on the tokens, which are labelled from 1 to 9. (a): The configuration of tokens before the averaging; (b): the configuration after the moving step and (c): the configuration after the shuffling step.

In the following we will refer to the topmost token as *excess token* if it does not have a sibling. It is easy to verify that the height of a token can never increase (see Lemma 9.1, (i) in the appendix). Furthermore, *any individual* token performs a random walk with the sequence of matching matrices as transition matrices (Lemma 9.1, (ii)). Crucially, we will prove later that the movements of *different* tokens satisfy a negative association property (see Lemma 5.1).

2.2 Properties of Matchings To state the next definitions and results more formally, we adopt the notation from [32]. Assume that nodes are labelled from 1 to n . We define the balancing matrix $\mathbf{M}^{(t)} \in [0, 1]^{n \times n}$ that represents a matching of G in round t as $\mathbf{M}_{u,v}^{(t)} := 1/2$ if $u \neq v$ are matched in round t , $\mathbf{M}_{u,v}^{(t)} := 0$ if $u \neq v$ are not matched in round t , and $\mathbf{M}_{u,u}^{(t)} := 1$ if u is not matched in round t . Following [34, 32], we use the notation $[u : v] \in \mathbf{M}^{(t)}$ to indicate that nodes u and v with $u < v$ are matched in round t . Further, we define the product of all matching matrices from time t_1 to t_2 as $\mathbf{M}^{[t_1, t_2]} := \prod_{s=t_1}^{t_2} \mathbf{M}^{(s)}$ (if $t_1 > t_2$ then $\mathbf{M}^{[t_1, t_2]}$ is the identity matrix). For simplicity, we use the same symbol $\mathbf{M}^{(t)}$ for both the matching of round t and the corresponding balancing matrix, sometimes referring to both just as “matching”. Also, we will always use upper case $\mathbf{M}$ for matchings. Later we also use $\mathbf{M}_{u,D}^{[t_1+1, t_2]} := \sum_{v \in D} \mathbf{M}_{u,v}^{[t_1+1, t_2]}$. Finally, $\mathbf{M}_{u,\cdot}^{[t_1, t_2]}$ denotes the row of the matrix

$\mathbf{M}^{[t_1, t_2]}$ corresponding to node u . By $\vec{1}$ we denote the row vector of length n in which each entry is 1.

DEFINITION 2.1 ([34, Definition 2.1]). *A fixed sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^t$, $t \geq 1$ is called (K, ε) -smoothing if for any $x^{(0)} \in \mathbb{R}^n$ with $\text{disc}(x^{(0)}) \leq K$ we have $\text{disc}(x^{(0)} \cdot \mathbf{M}^{[1, t]}) \leq \varepsilon$.*

The expression $x^{(0)} \cdot \mathbf{M}^{[1, t]}$ equals the load vector of the continuous load balancing process with initial load $x^{(0)}$ applying the matchings $(\mathbf{M}^{(s)})_{s=1}^t$ (see, e.g., [32]). Hence the definition above states that the matching sequence is sufficient to balance the load up to ε in the continuous model.

In general the matching sequence can be generated deterministically or randomly, and our main result will cover both cases. The next definition states the properties of the matchings which we require for our main result.

DEFINITION 2.2. *A sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$ is called $(\tau_{\text{global}}, \tau_{\text{local}})$ -good if*

$$(2.1) \quad \inf_{t \in \mathbb{N}_0} \Pr \left[\bigcap_{u \in V} \left\| \mathbf{M}_{u, \cdot}^{[t+1, t+\tau_{\text{global}}]} - \frac{\vec{1}}{n} \right\|_2^2 \leq \frac{1}{n^7} \right] \geq 1 - \frac{1}{n^3},$$

$$(2.2) \quad \inf_{t \in \mathbb{N}_0} \min_{u \in V} \Pr \left[\left\| \mathbf{M}_{u, \cdot}^{[t+1, t+\tau_{\text{local}}]} \right\|_2^2 \leq \frac{1}{\log^{10}(n)} \right] \geq 1 - \frac{1}{\log^{11}(n)}.$$

Since each token is performing a random walk according to the matching sequence (Lemma 9.1), we can interpret these events in terms of a distribution of a token performing a time-inhomogeneous random walk and its ℓ_2 -distance to the stationary (i.e., uniform) distribution. The event in (2.1) means that the distribution of any token will be very close to uniform after τ_{global} rounds. This basically corresponds to a complete, i.e., “global” mixing of a random walk, since it holds for *any* start node u . With regards to the event in (2.2), it only requires a very coarse mixing of the distribution after τ_{local} rounds, and this condition holds only “locally”, i.e., from a specific node u . Using standard spectral techniques, for the random matching model we have $\tau_{\text{global}} = O(\log(n)/(1 - \lambda))$ and $\tau_{\text{local}} = O(\log \log(n)/(1 - \lambda))$ (assuming $p_{\min} = \Omega(1/\Delta)$).

2.3 Main Results Our main result is the following theorem; the proof can be found in section 8.

THEOREM 2.3. *Let G be any undirected, connected graph on n nodes and consider any initial load vector $x^{(0)} \in \mathbb{N}_0^n$ with $\text{disc}(x^{(0)}) \leq K$. Assume our process balances the tokens via a $(\tau_{\text{global}}, \tau_{\text{local}})$ -good sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$. Then there exists a time τ with*

$$\tau = O \left(\frac{\log(Kn)}{\log(n)} \cdot \tau_{\text{global}} + \frac{\log(n)}{\log \log(n)} \cdot \tau_{\text{local}} \right)$$

such that

$$\Pr \left[\text{disc} \left(X^{(\tau)} \right) \leq 4 \right] \geq 1 - \exp \left(-(1/200) \cdot \frac{\log(n)}{\log \log(n)} \right),$$

and for any constant $c > 0$,

$$\Pr \left[\text{disc} \left(X^{(\tau)} \right) \leq 3 \right] \geq 1 - \exp \left(-\log^{1-c}(n) \right).$$

Note that Theorem 2.3 does not make any assumption regarding the provenance of the matching sequences; all we use is the abstract property of $(\tau_{\text{global}}, \tau_{\text{local}})$ -goodness. We now provide details for three explicit ways to create such sequences, namely the *balancing circuit model*, the *random matching model* and the *asynchronous model* (a.k.a. single edge model). To state the result for these models we need some definitions. For any n by n real symmetric matrix $\mathbf{M}$, let $\lambda_1(\mathbf{M}) \geq \lambda_2(\mathbf{M}) \geq \dots \geq \lambda_n(\mathbf{M})$ be the n eigenvalues of $\mathbf{M}$. For simplicity, let $\lambda(\mathbf{M}) := \max\{|\lambda_2(\mathbf{M})|, |\lambda_n(\mathbf{M})|\}$. For a non-symmetric matrix $\mathbf{M}$, we define the symmetric matrix as $\widetilde{\mathbf{M}} = \mathbf{M} \cdot \mathbf{M}^T$ and let $\lambda(\mathbf{M}) := \max\{|\lambda_2(\widetilde{\mathbf{M}})|, |\lambda_n(\widetilde{\mathbf{M}})|\}$.

Application to Balancing Circuits. In the balancing circuit model all or a subset of the edges of G are covered using a periodic sequence of Δ fixed matchings $\mathbf{M}^{(1)}, \mathbf{M}^{(2)}, \dots, \mathbf{M}^{(\Delta)}$. The sequence $(\mathbf{M}^{(s)})_{s=1}^\infty$ is chosen deterministically and periodically such that $\mathbf{M}^{(s)} = \mathbf{M}^{((s-1) \bmod \Delta + 1)}$. Such matchings can be found, e.g., via edge-coloring, and there exist many efficient distributed algorithms that compute such a coloring [23]. In the

balancing circuit model Δ is (an upper bound on) the maximum degree of the graph induced by the union of all matchings. In this model, Theorem 2.3 provides, see Corollary 8.2,

$$(2.3) \quad \tau = O(\tilde{\tau}_{\text{cont}}(K)), \quad \text{where} \quad \tilde{\tau}_{\text{cont}}(K) := \Theta\left(\frac{\Delta \cdot \log(Kn)}{1 - \lambda(\mathbf{M}^{[1, \Delta]})}\right).$$

Application to Random Matchings. In line with previous work [23, 15, 34], we consider the following class of randomized algorithms that generate a sequence of (random) matchings $(\mathbf{M}^{(s)})_{s=1}^{\infty}$, and simply refer to this as the “random matching model”. First, we require that the matchings are mutually independent across all rounds. Second, for all edges $\{u, v\}$ and all rounds $t \geq 1$, the probability that $\{u, v\}$ is included in $\mathbf{M}^{(t)}$ is at least $p_{\min} := c/\Delta$ for some time-independent value $c > 0$ (however, c is allowed to be a function of n). Recall that we use Δ as the number of matchings for the balancing circuit model and as maximum degree of G in the random matching model. Note that the decisions whether or not to include two edges into a matching within the same round are clearly *not* independent. Some concrete distributed algorithms that satisfy both conditions are described in [23, 15]. To state the results for the random matching model we define the diffusion matrix $\mathbf{P}$ as $\mathbf{P}_{u,v} := 1/(2\Delta)$ if $(u, v) \in E$, $\mathbf{P}_{u,v} := 1 - \deg(u)/(2\Delta)$ if $u = v$, and $\mathbf{P}_{u,v} := 0$ otherwise. In this case, Theorem 2.3 provides, see Corollary 8.3,

$$(2.4) \quad \tau = O(\tilde{\tau}_{\text{cont}}(K)), \quad \text{where} \quad \tilde{\tau}_{\text{cont}}(K) := \Theta\left(\frac{\log(Kn)}{p_{\min} \cdot \Delta \cdot (1 - \lambda(\mathbf{P}))}\right)$$

Note that in case $p_{\min} = \Omega(1/\Delta)$ we have $\tilde{\tau}_{\text{cont}}(K) = \Theta(\log(Kn)/(1 - \lambda(\mathbf{P})))$.

Application to the Asynchronous (Single Edge) Model. In this model, at each round we pick a single edge e uniformly at random. This is a special case of the random matching model with $p_{\min} = 1/|E|$; and thus, the same of definition of $\tilde{\tau}_{\text{cont}}(K)$ applies here. Therefore, Theorem 2.3 provides $\tau = O\left(\frac{|E|}{\Delta} \cdot \log(Kn)/(1 - \lambda(\mathbf{P}))\right) = O\left(n \cdot \frac{d}{\Delta} \cdot \log(Kn)/(1 - \lambda(\mathbf{P}))\right)$, where $d = 2|E|/n$ is the average degree of G (see Corollary 8.3).

2.4 Comparison of our Results and Techniques with Previous Work An important novelty of our approach is that our analysis framework seamlessly covers the balancing circuit model, the random matching model and the asynchronous (single edge) model. This is in contrast to previous works [32, 31, 34, 22], which either focus only on one specific matching model or provide tailored analyses for each of them. To unify the three models covered here, our analysis is based on a coarse (and local) and a fine (and global) mixing/balancing property of these models.

In comparison with [34], the constant factors in both our running time and discrepancy are explicit and small. In addition to our analysis being tighter and simpler, we also obtain a larger success probability. We note that the tool box we use differs substantially from that of [34], not least because the pièce de résistance of our analysis, getting from a large constant bound for the discrepancy down to a very small constant is not being done in [34] at all, and is one of the hardest problems we solve. Finally, the results of [34] require regular graphs in the random matching, and a constant Δ for the balancing circuit model. Our results are covering both of these models (and additionally the asynchronous model), and hold without any of these restrictions.

While our proof method yields more general and tighter results, we also believe that our proof is more intuitive and direct. The first and more straightforward step is to carefully bound the number of tokens in each possible subset $S \subseteq V$ and then apply a union bound. This is based on a new Hoeffding-type concentration bound, which generalizes and tightens previous bounds [22, 34] in that it works for *arbitrary* convex combinations of the load vector. This leaves us with only $O(n)$ tokens to balance.

The second part of the analysis makes use of our new height-sensitive process, which constraints the movement of tokens in such a way that their heights are non-increasing. Despite this restriction, we can prove that their movements are negatively associated. This property, together with the Hoeffding-type concentration, is then used in an involved analysis to show that eventually a discrepancy 4 is reached. An extra iteration of this method finally yields a discrepancy bound of 3.

3 An Outline of the Analysis In this section we give a more detailed outline of our analysis. First, in subsection 3.1, we present a collection of the most important technical results. These technical tools are crucial for our analysis and their proofs are considerably more challenging than deriving the discrepancy bounds using

these tools. Additionally, we believe that some of these tools are of independent interest. At the end of this section, in subsection 3.2, we give a brief summary on how these technical results are combined to obtain our main result, which is a discrepancy bound of 3.

3.1 Our Technical Contributions In order to prove small discrepancy bounds, we need to keep track of the number of tokens at a specific height. Thanks to the height-sensitive process defined earlier, the height of a token is non-increasing over time, and the sequence of locations of a token $i \in \mathcal{T}$, $(W_i^{(t)})_{t \geq 0}$, form a random walk (Lemma 9.1). Crucially, we establish that these random walks are negatively associated:

LEMMA 3.1 (simplified version of Lemma 5.1). *Fix a subset of tokens $\mathcal{B} \subseteq \mathcal{T}$ at round 0. Let $t > 0$ be any round, and fix the matchings between round 0 and t . Then for any set $D \subseteq V$, the events $\{W_i^{(t)} \in D\}, i \in \mathcal{B}$ are negatively associated.*

In comparison to previous work (Lemma 4.2 from [34]), our lemma yields the same statement but here tokens move following the definition of the height-sensitive process, whereas in [34], the two nodes exchange all tokens freely, which means that the height of a token could increase. Even though in this sense our process might be slightly harder to describe and analyze, our proof is simpler and more elementary than [34], e.g., we do not need the somewhat unwieldy negative regression condition from [20]. We continue with a Hoeffding-like concentration bound.

LEMMA 3.2 (simplified version of Lemma 5.2). *Consider any load vector $x^{(0)}$ with $\text{disc}(x^{(0)}) \leq K$ and any round $t \geq 1$ such that the sequence of matchings from round 1 to t is $(K, 1/(2n))$ -smoothing. Then for any stochastic vector $(a_w)_{w \in V}$, it holds for any $\delta > 0$,*

$$\Pr \left[\left| \sum_{w \in V} a_w \cdot X_w^{(t)} - \bar{x} \right| \geq \delta \right] \leq 2 \cdot \exp \left(-\frac{(\delta - 1/(2n))^2}{4\|a\|_2^2} \right).$$

In order to appreciate this result, we first discuss a more general (but somewhat harder to apply) version, which is derived in the proof of Lemma 5.2. This states that for *any* matching sequence, and any stochastic vector $(a_w)_{w \in V}$, $\mu := \mathbf{E} \left[\sum_{w \in V} a_w \cdot X_w^{(t)} \right]$, and any $\delta > 0$ it holds

$$\Pr \left[\left| \sum_{w \in V} a_w \cdot X_w^{(t)} - \mu \right| \geq \delta \right] \leq 2 \cdot \exp \left(-\frac{\delta^2}{4\|a\|_2^2} \right).$$

This tail bound essentially matches the one from Hoeffding's inequality for $\sum_{w \in V} Y_w$ if the Y_w are all independent and $Y_w \in [-a_w, a_w]$. However, in the load balancing process the range of the X_w 's are unbounded and also the X_w 's are far from being independent; for instance, two nodes matched in round t must have a load difference of at most 1.

Our concentration bound is established by carefully aggregating all rounding errors contributing to $\sum_{w \in V} a_w \cdot X_w^{(t)} - \mu$ by means of a quadratic potential function. On a high level, our proof resembles that in [34], however, one key difference is that we employ a more general potential function which involves the coefficients $a := (a_w)_{w \in V}$.

Compared to prior work, our result generalizes [34, Lemma 3.5] to *arbitrary* stochastic vectors a . The corresponding result in [34] only works for the specific vector a with $a_w := \mathbf{M}_{w,u}^{[t+1, t_1]}$, for a fixed node $u \in V$ and round $t_1 \geq t$; in particular, if we choose $t_1 = t$ then a is a unit-vector. There are also earlier and weaker versions of this inequality, e.g., in [22, Theorem 4.6], which only match our form if the underlying graph and matchings have constant expansion. We therefore believe that our result will be the final word in the quest to find a tight and general concentration inequality of this type.

The first application of our concentration inequality (Lemma 5.2) is to bound the number of tokens above the average load:

LEMMA 3.3 (simplified version of Lemma 5.6). *Consider an initial load vector with $\text{disc}(x^{(0)}) \leq K$ and round $t \geq 0$ such that the sequence $(\mathbf{M}^{(s)})_{s=1}^t$ which is $(K, 1/(2n))$ -smoothing. Then,*

$$\Pr \left[\sum_{w \in V} \max \{X_w^{(t)} - \bar{x}, 0\} \leq 16 \cdot n \right] \geq 1 - 2 \cdot n^{-2}.$$

This lemma exploits that in Lemma 5.2 we can choose the stochastic vector $(a_w)_{w \in V}$ freely. For each possible subset $S \subseteq V$ and integer $i \geq 1$, we bound the existence of a “bad” set S of size $\Theta(n/2^i)$ in which all nodes have load at least $\bar{x} + 4 \cdot i$. This is done by choosing an appropriate vector (a_w) and threshold δ in Lemma 5.2. The proof is then concluded by a simple union bound over all possible “bad” subsets $S \subseteq V$.

The next proposition is the most involved step in our analysis. It shows that we can reduce the discrepancy to a constant if the initial load vector has at most $O(n)$ tokens.

PROPOSITION 3.4 (simplified and informal version of Proposition 6.1). *Let $t := O(\tau_{\text{global}} + \tau_{\text{local}} \cdot \log(n)/\log \log(n))$, and let $1 \leq L = O(1)$. Consider a $(\tau_{\text{global}}, \tau_{\text{local}})$ -good sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$ and an arbitrary load vector $x^{(0)}$ with at most $(L - \Omega(1/\log(n))) \cdot n$ tokens. Then,*

$$\Pr \left[\max_{w \in V} X_w^{(t)} \leq L + 1 \right] \geq 1 - \exp \left(- \Omega \left(\frac{\log(n)}{\log \log(n)} \right) \right).$$




Figure 3.1: Illustration of Phases 1 and 2 in the proof of Proposition 6.1. Phase 1 decreases $Y^{(t)}$, the number of tokens with height at least $L + 1$, from $n \cdot L$ to $\frac{n}{\log(n)}$. Then Phase 2 decreases $\hat{Y}^{(t)}$, the number of tokens with height at least $L + 2$, to 0 . Both phases first use τ_{global} rounds for a “global mixing”, and then $\ell := \frac{\log(n)}{\log \log(n)}$ short epochs of length τ_{local} .

The proof of this proposition constitutes the most challenging part of our analysis. Recall that we only have to balance $(L - \varepsilon) \cdot n$ tokens for a constant $\varepsilon > 0$ and integer $L \geq 1$. Our goal is to gradually reduce the number of tokens with height $\geq L + 1$, until no token with height $\geq L + 2$ remains. The proposition uses two phases, which are illustrated in Figure 3.1.

- **Phase 1:** Reduce the number of tokens with height $\geq L + 1$ from $(L - \varepsilon)n$ to $n/\log(n)$ (see Lemma 6.11).
- **Phase 2:** Reduce the number of tokens with height $\geq L + 2$ further to 0 (see Lemma 6.12).

Phases 1 and 2 run successively and are analyzed with the same framework, only with alternate parameters. The key difference is that after Phase 1 we only need to cope with a sublinear number, that is, $n/\log(n)$, tokens at height at least $L + 1$, and we wish to bound the number of tokens that remain at height $L + 2$. This allows us to make faster progress in Phase 2. Specifically, we prove exponential decay every τ_{local} rounds in Phase 1 and even super-exponential decay (factor $1/\log n$) in Phase 2. The analyses of both phases hinges on the key lemma stated below, which establishes a multiplicative drop on the number of tokens at height at least $L + 1$ within τ_{local} rounds:

LEMMA 3.5 (simplified and informal version of Lemma 6.8 (Key Lemma)). *Assume that a load vector $x^{(t)}$ has at most $(L - \varepsilon)n$ tokens, where $t > 0$, $0 < \varepsilon < 1$ and $1 \leq L \leq \log^7(n)$ is an integer. Let $Y^{(t)}$ be the number of tokens with height at least $L + 1$ in round t . Then,*

$$\mathbf{E} \left[Y^{(t+\tau_{\text{local}})} \mid x^{(t)} \right] \leq \left(1 - \frac{\varepsilon}{L} + \frac{2}{L \cdot \log^4(n)} \right) \cdot Y^{(t)}.$$

In the following we will sketch the central ideas needed to establish the key lemma above. We start with a token $i \in \mathcal{T}$ at height at least $L+1$, located at a node $u \in V$ in round t . Our goal is to lower bound the probability that after τ_{local} additional rounds, the token is still at this height. The location of token i at time $t + \tau_{\text{local}}$ is determined by a random walk with law $\mathbf{M}_{u,\cdot}^{[t+1, t+\tau_{\text{local}}]}$ (Lemma 9.1); we call this (random) node $v \in V$. As heights of tokens are non-decreasing in $[t, t + \tau_{\text{local}}]$, the only way for token i to remain at height at least $L+1$ is for there to be at least L many other tokens j which are also on node $v \in V$ at time $t + \tau_{\text{local}}$. Using the negative association lemma (Lemma 5.1), the expected number of tokens which collide with token i at round $t + \tau_{\text{local}}$ can be upper bounded by (see Lemma 6.4),

$$\sum_{w \in V} \underbrace{\left(\sum_{v \in V} \mathbf{M}_{u,v}^{[t+1, t+\tau_{\text{local}}]} \cdot \mathbf{M}_{w,v}^{[t+1, t+\tau_{\text{local}}]} \right)}_{=: a_w} \cdot X_w^{(t)}.$$

We proceed with upper bounding the right hand side. This is a convex combination of a load vector, which makes it amenable to our new concentration inequality for convex combinations of loads. Using that the matching sequences are of length τ_{local} , we obtain with reasonably large probability over the matching sequence that $\|a\|_2^2$ is small. Once we have established this, we apply Lemma 5.2 to the above sum; here, for fixed a the randomness is over the matching sequences and shuffling steps in $[1, t]$, where $t \geq \tau_{\text{global}}$. Taking aside some technicalities, we can then conclude that there is an expected drop in the number of tokens at height $L+1$ within τ_{local} rounds.

One significant technical challenge is to iterate this argument over consecutive epochs of length τ_{local} , as each epoch also depends on the random decisions in previous epochs (both matchings and shuffling decisions). We will overcome these dependencies by carefully defining events (corresponding to the local and global behaviour of the process; similar to τ_{local} and τ_{global}), and then integrate these events into a submartingale which shows that the number of tokens at height $L+1$ drops.

3.2 Putting the Pieces Together With the powerful Proposition 6.1 at hand, we are able to reduce the discrepancy from any K to 3 in $O(\log(Kn)/\log(n) \cdot \tau_{\text{global}} + \tau_{\text{local}} \cdot \log(n)/\log \log(n))$ rounds as follows:

1. **Discrepancy $K \rightarrow$ Linear Number of Tokens:** We first reduce the number of tokens above the average load to $16n$ using $O(\tilde{\tau}_{\text{cont}}(K))$ many rounds (see Lemma 5.6).
2. **Linear Number of Tokens $\rightarrow$ Discrepancy 38:** Applying Proposition 6.1 we then reduce the discrepancy to 38 within $O(\tau_{\text{global}} + \tau_{\text{local}} \cdot \log(n)/\log \log(n))$ many rounds (see Lemma 7.1).
3. **Discrepancy 38 $\rightarrow$ Discrepancy 4:** By an additional application of Proposition 6.1, after $O(\tau_{\text{global}} + \tau_{\text{local}} \cdot \log(n)/\log \log(n))$ further rounds the discrepancy drops to 4 (see Lemma 7.2).
4. **Discrepancy 4 $\rightarrow$ Discrepancy 3:** Now Proposition 6.1 cannot be applied directly to reduce the discrepancy any further, but we can still reduce the discrepancy to 3 by a more tailored analysis of the movement of the tokens at the two highest levels in the next $O(\tau_{\text{global}} + \tau_{\text{local}} \cdot \log(n)/\log \log(n))$ rounds, and obtain the main result of our paper.

4 Summary and Open Problems In this paper we show that, for the matching model, discrete load balancing is as efficient and effective as continuous load balancing. We showed that in the discrete setting with integer loads, a discrepancy of 3 is reached in a time that matches the standard spectral bound on the time needed by the continuous setting. In particular, this means that for expanders and a polynomial initial discrepancy, our load balancing schemes achieve a discrepancy of 3 using only $O(\log(n)) = O(\text{diam}(G))$ rounds, which is optimal not only for a distributed but also a centralized setting.

As an improvement over previous works, our general result holds for a wider class of graphs (e.g., including non-regular graphs) and models (e.g., including the asynchronous model). Also the constants in our runtime as well as in the achieved discrepancy are explicit and small, compared to large and non-explicit constants in [34].

At the heart of our analysis lie new association and concentration results, which we believe to be of independent interest. As the most involved step, showing that $O(n)$ tokens can be balanced with maximum load $O(1)$, which can be leveraged to prove a discrepancy of 3 for an arbitrary number of tokens. It should be noted that reaching a discrepancy of 1 needs $\Omega(n)$ rounds for any sequence of matchings (see [29]). Hence one may wonder whether our discrepancy bound could be improved from 3 to 2. We believe that this might be possible, but it will likely need stronger assumptions on the matchings than just being $(\tau_{\text{global}}, \tau_{\text{local}})$ -good.

5 Association and Concentration Results In this section we show two main components our analysis; first, the negative association result about token movements and secondly, the Hoeffding-like concentration inequality.

5.1 Height-Sensitive Negative Association

LEMMA 5.1 (Height-Sensitive Negative Association). *Consider any pair of rounds $0 \leq t_1 < t_2$, and let $(\mathbf{M}^{(s)})_{s=t_1}^{t_2}$ be an arbitrary but fixed sequence of matchings. Further, let $w^{(t_1)} = (w_i)_{i \in \mathcal{T}}^{(t_1)}$ be a fixed location vector in round $t_1 \geq 0$. Then for any subset of tokens $\mathcal{B} \subseteq \mathcal{T}$ and any subset of nodes $D \subseteq V$, it holds that*

$$\Pr \left[\bigcap_{i \in \mathcal{B}} \{W_i^{(t_2)} \in D\} \mid w^{(t_1)} \right] \leq \prod_{i \in \mathcal{B}} \Pr \left[W_i^{(t_2)} \in D \mid w^{(t_1)} \right] = \prod_{i \in \mathcal{B}} \mathbf{M}_{w_i^{(t_1)}, D}^{[t_1+1, t_2]}.$$

Proof. Throughout the proof we always assume that $w^{(t_1)}$ is fixed and we omit the conditioning. Note that by Lemma 9.1, we have $\prod_{i \in \mathcal{B}} \Pr \left[W_i^{(t_2)} \in D \right] = \prod_{i \in \mathcal{B}} \mathbf{M}_{w_i^{(t_1)}, D}^{[t_1+1, t_2]}$ and therefore it only remains to prove the inequality

$$\Pr \left[\bigcap_{i \in \mathcal{B}} \{W_i^{(t_2)} \in D\} \right] \leq \prod_{i \in \mathcal{B}} \mathbf{M}_{w_i^{(t_1)}, D}^{[t_1+1, t_2]}.$$

For every round $t \in [t_1, t_2]$ and a token $i \in \mathcal{B}$ we define a random variable

$$(5.1) \quad Z_i^{(t)} := \sum_{u \in V} \mathbf{1}_{W_i^{(t)}=u} \cdot \mathbf{M}_{u, D}^{[t+1, t_2]} \quad \text{and} \quad Z^{(t)} := \prod_{i \in \mathcal{B}} Z_i^{(t)}.$$

Hence, conditional on the token location $W_i^{(t)}$, $Z_i^{(t)}$ is the probability that token i starting at time t is on a node $w \in D$ at time t_2 . In the following analysis we use the simplified notation

$$\mathbf{M}_{W_i^{(t)}, D}^{[t+1, t_2]} := \sum_{u \in V} \mathbf{1}_{W_i^{(t)}=u} \cdot \mathbf{M}_{u, D}^{[t+1, t_2]}.$$

Since $\mathbf{M}^{[t_2+1, t_2]} = \mathbf{I}$,

$$Z^{(t_2)} = \prod_{i \in \mathcal{B}} \mathbf{M}_{W_i^{(t_2)}, D}^{[t_2+1, t_2]}.$$

Hence, $Z_i^{(t_2)}$ is a random variable that can take on the values zero or one. Therefore,

$$\mathbf{E} \left[Z^{(t_2)} \right] = \Pr \left[Z^{(t_2)} = 1 \right] = \Pr \left[\bigcap_{i \in \mathcal{B}} \{Z_i^{(t_2)} = 1\} \right] = \Pr \left[\bigcap_{i \in \mathcal{B}} \{W_i^{(t_2)} \in D\} \right].$$

The key idea is to prove that the sequence $Z^{(t)}$, $t \in [t_1, t_2]$ forms a supermartingale with respect to the filtration $\mathfrak{F}^{(t)}$, which reveals all random decisions of the process between rounds t_1 and t (so in particular, it reveals the location vectors $W^{(t)} = (W_1^{(t)}, \dots, W_{|\mathcal{T}|}^{(t)})$), that is,

$$(5.2) \quad \mathbf{E} \left[Z^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \leq Z^{(t-1)},$$

which would immediately imply

$$\mathbf{E} \left[Z^{(t_2)} \right] \leq Z^{(t_1)}.$$

It then follows

$$\Pr \left[\bigcap_{i \in \mathcal{B}} \{W_i^{(t_2)} \in D\} \right] = \mathbf{E} \left[Z^{(t_2)} \right] \leq Z^{(t_1)} = \prod_{i \in \mathcal{B}} \mathbf{M}_{w_i^{(t_1)}, D}^{[t_1+1, t_2]},$$

finishing the proof.

It remains to prove (5.2). Without loss of generality we may assume that each matching consists only of one edge. This is because for any sequence of matchings $(\mathbf{M}^{(s)})_{s=t_1}^{t_2}$ in which there is a round $s \in [t_1 + 1, t_2]$ whose matching includes more than one edge, we can decompose such round s into multiple “sub-rounds”, each consisting of exactly one matching. It is clear that this yields the same process, only with a larger number of rounds.

Fix round t with $t_1 \leq t \leq t_2$ and assume $\{u, v\}$ is the single edge in $\mathbf{M}^{(t)}$. First we compute $Z^{(t-1)}$. We define $T_u \subseteq \mathcal{B}$ as the tokens from $\mathcal{B}$ which are on u and $T_v \subseteq \mathcal{B}$ as the tokens from $\mathcal{B}$ which are on v at the beginning of round t . Note that it is possible that there are tokens $i \notin \mathcal{B}$ on those nodes. We get

$$\mathbf{M}_{u,D}^{[t,t_2]} = \mathbf{M}_{v,D}^{[t,t_2]} = \mathbf{M}_{u,v}^{(t)} \cdot \mathbf{M}_{v,D}^{[t+1,t_2]} + \mathbf{M}_{u,u}^{(t)} \cdot \mathbf{M}_{u,D}^{[t+1,t_2]} = \frac{\mathbf{M}_{u,D}^{[t+1,t_2]} + \mathbf{M}_{v,D}^{[t+1,t_2]}}{2}.$$

For ease of presentation we define

$$p := \mathbf{M}_{u,D}^{[t+1,t_2]} \quad \text{and} \quad q := \mathbf{M}_{v,D}^{[t+1,t_2]}.$$

From the definition (5.1) it follows that

$$(5.3) \quad Z^{(t-1)} = \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i^{(t-1)} \cdot \prod_{i \in T_u \cup T_v} Z_i^{(t-1)} = \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i^{(t-1)} \cdot \left(\frac{p+q}{2} \right)^{|T_u \cup T_v|}.$$

We will continue to analyze $\mathbf{E}[Z^{(t)} \mid \mathfrak{F}^{(t-1)}]$ and eventually upper bound it by the right hand side in (5.3), which completes the argument. For each token i located at $\{u, v\}$ let $S(i)$ be its sibling token on the other node after the moving step of round t and before the shuffling step; note that it is possible for a token to have no sibling. We partition tokens in $T_u \cup T_v$ into three sets:

$$\begin{aligned} \mathcal{B}_1 &= \{i \in T_u \cup T_v \mid S(i) \in T_u \cup T_v\}, \\ \mathcal{B}_2 &= \{i \in T_u \cup T_v \mid S(i) \in \mathcal{T} \setminus (T_u \cup T_v)\}, \\ \mathcal{B}_3 &= \{i \in T_u \cup T_v \mid S(i) \text{ does not exist.}\}. \end{aligned}$$

Note that the set $\mathcal{B}_3$ corresponds to the excess token (if there is one), and thus $|\mathcal{B}_3| \in \{0, 1\}$. We will now apply this partitioning to $\mathbf{E} \left[Z^{(t)} \mid \mathfrak{F}^t \right]$,

$$\begin{aligned} & \mathbf{E} \left[Z^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \\ &= \mathbf{E} \left[\prod_{i \in \mathcal{B}} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \\ &= \mathbf{E} \left[\prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i^{(t)} \cdot \prod_{i \in T_u \cup T_v} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \\ &\stackrel{(a)}{=} \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i^{(t-1)} \cdot \mathbf{E} \left[\prod_{i \in \mathcal{B}_1} Z_i^{(t)} \cdot \prod_{i \in \mathcal{B}_2} Z_i^{(t)} \cdot \prod_{i \in \mathcal{B}_3} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \\ (5.4) \quad &\stackrel{(b)}{=} \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i^{(t-1)} \cdot \mathbf{E} \left[\prod_{i \in \mathcal{B}_1} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \cdot \mathbf{E} \left[\prod_{i \in \mathcal{B}_2} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \cdot \mathbf{E} \left[\prod_{i \in \mathcal{B}_3} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right], \end{aligned}$$

where (a) follows from the fact that for tokens $i \in \mathcal{B} \setminus (T_u \cup T_v)$, we have $Z_i^{(t)} = Z_i^{(t-1)}$ since $W_i^{(t)}$ is not matched with any node in this round; (b) follows since tokens which are in different $\mathcal{B}_j$, $1 \leq j \leq 3$ must be at different heights and therefore move independently in the shuffling phase.

We will now analyze the three different expectations in (5.4)

Case 1: The expectation over $\mathcal{B}_1$. Using that any token $i \in \mathcal{B}_1$ only depends on its sibling token $S(i)$ we can group all tokens in $\mathcal{B}_1$ in pairs, and we denote this relation by $\sim$. Further, a token i with sibling j will take the opposite action in the shuffling step, and therefore

$$\mathbf{E} \left[\prod_{i \in \mathcal{B}_1} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] = \prod_{i < j \in \mathcal{B}_1 : i \sim j} \mathbf{E} \left[Z_i^{(t)} \cdot Z_j^{(t)} \mid \mathfrak{F}^{(t-1)} \right] = \prod_{i < j \in \mathcal{B}_1 : i \sim j} (p \cdot q).$$

Case 2: The expectation over $\mathcal{B}_2$. Since each token in $\mathcal{B}_2$ has no sibling in $\mathcal{B}$ (let alone $\mathcal{B}_2$), their movements in the shuffling step are independent. Further, each such token is on u or v with the same probability which yields,

$$\mathbf{E} \left[\prod_{i \in \mathcal{B}_2} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] = \prod_{i \in \mathcal{B}_2} \mathbf{E} \left[Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] = \prod_{i \in \mathcal{B}_2} \left(\frac{p+q}{2} \right).$$

Case 3: The expectation over $\mathcal{B}_3$. This is analogous to Case 2, as a token in $\mathcal{B}_3$ has no sibling in $\mathcal{B}$ (in fact there is not even any other token at this height). Hence,

$$\mathbf{E} \left[\prod_{i \in \mathcal{B}_3} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] = \prod_{i \in \mathcal{B}_3} \left(\frac{p+q}{2} \right).$$

Aggregating the contributions of all tokens in $\mathcal{B}$ in round t using the three cases above,

$$\begin{aligned} & \mathbf{E} \left[Z^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \\ & \stackrel{(5.4)}{=} \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i^{(t-1)} \cdot \mathbf{E} \left[\prod_{i \in \mathcal{B}_1} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \cdot \mathbf{E} \left[\prod_{i \in \mathcal{B}_2} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \cdot \mathbf{E} \left[\prod_{i \in \mathcal{B}_3} Z_i^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \\ & \stackrel{(a)}{=} \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i^{(t-1)} \cdot \prod_{i < j \in \mathcal{B}_1 : i \sim j} (p \cdot q) \cdot \prod_{i \in \mathcal{B}_2} \frac{p+q}{2} \cdot \prod_{i \in \mathcal{B}_3} \frac{p+q}{2} \\ & \stackrel{(b)}{\leq} \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i^{(t-1)} \cdot \prod_{i < j \in \mathcal{B}_1 : i \sim j} \left(\frac{p+q}{2} \right)^2 \cdot \prod_{i \in \mathcal{B}_2} \frac{p+q}{2} \cdot \prod_{i \in \mathcal{B}_3} \frac{p+q}{2} \\ & = \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i^{(t-1)} \cdot \left(\frac{p+q}{2} \right)^{|T_u \cup T_v|} \\ & \stackrel{(c)}{=} Z^{(t-1)}, \end{aligned}$$

where (a) uses Cases 1, 2 and 3 above, (b) follows from the fact that $p \cdot q \leq ((p+q)/2)^2$ and (c) follows from (5.3). Hence $\mathbf{E} \left[Z^{(t)} \mid \mathfrak{F}^{(t-1)} \right] \leq Z^{(t-1)}$, which is (5.2) and, consequently, completes the proof. $\square$

5.2 Concentration Inequality for Convex Combinations of Loads The following ‘‘Hoeffding-like’’ concentration inequality is another key tool in our analysis. We use it frequently to derive (upper) bounds on the number of tokens on a set of nodes.

LEMMA 5.2 (Hoeffding Bound for Convex Combinations of Loads). *Consider any load vector $x^{(0)}$ with $\text{disc}(x^{(0)}) \leq K$ and any round $t \geq 1$ such that $(\mathbf{M}^{(s)})_{s=1}^t$ is a (K, κ) -smoothing sequence of matchings with $K > \kappa$. Then for any stochastic vector $(a_w)_{w \in V}$, it holds for any $\delta > 0$,*

$$\Pr \left[\left| \sum_{w \in V} a_w \cdot X_w^{(t)} - \bar{x} \right| \geq \delta \mid (\mathbf{M}^{(s)})_{s=1}^t \right] \leq 2 \cdot \exp \left(- \frac{(\delta - \kappa)^2}{4 \|a\|_2^2} \right).$$

As a special case of this concentration inequality, we can take any node v , define the unit-vector $a_w = \mathbf{1}_{w=v}$ (and thus $\|a\|_2^2 = 1$), pick $\delta = \sqrt{12 \cdot \log(n)}$ and then apply a union bound over all nodes $v \in V$. This immediately yields the following bound on the discrepancy.

COROLLARY 5.3. Consider any load vector $x^{(0)}$ with $\text{disc}(x^{(0)}) \leq K$. Consider any round $t \geq 1$ and let $(\mathbf{M}^{(s)})_{s=1}^t$ be a fixed $(K, 1)$ -smoothing sequence of matchings. Then,

$$\Pr \left[\text{disc}(X^{(t)}) \leq \sqrt{48 \cdot \log(n)} + 1 \right] \geq 1 - 2 \cdot n^{-1}.$$

The exact same bound was shown in [34, Theorem 3.6].

5.3 Proof of the Concentration Inequality (Lemma 5.2) Before giving the proof of Lemma 5.2, we need to introduce some notation and auxiliary results.

We first recall a standard formula (e.g., [34, Equation 2.5]); for any node $w \in V$,

$$(5.5) \quad X_w^{(t)} - \xi_w^{(t)} = \sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} E_{u,v}^{(s)} \cdot \left(\mathbf{M}_{u,w}^{[s+1,t]} - \mathbf{M}_{v,w}^{[s+1,t]} \right),$$

where $\xi_w^{(t)} = \sum_{u \in V} x_u^{(0)} \cdot \mathbf{M}_{u,w}^{[1,t]}$ is the load after t rounds in the continuous model where load values can be fractional, and $E_{u,v}^{(t)} = \frac{1}{2} \cdot \text{Odd} \left(X_u^{(t-1)} + X_v^{(t-1)} \right) \cdot \Phi_{u,v}^{(t)}$ is the rounding error; here $\Phi_{u,v}^{(t)} \in \{-1, +1\}$ is the random orientation which has Rademacher distribution (that is, it is uniform in $\{-1, +1\}$). Therefore,

$$\begin{aligned} \sum_{w \in V} a_w \cdot \left(X_w^{(t)} - \xi_w^{(t)} \right) &= \sum_{w \in V} a_w \cdot \left(\sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} E_{u,v}^{(s)} \cdot \left(\mathbf{M}_{u,w}^{[s+1,t]} - \mathbf{M}_{v,w}^{[s+1,t]} \right) \right) \\ &= \sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} E_{u,v}^{(s)} \cdot \left[\sum_{w \in V} a_w \cdot \left(\mathbf{M}_{u,w}^{[s+1,t]} - \mathbf{M}_{v,w}^{[s+1,t]} \right) \right]. \end{aligned}$$

We will first state a Hoeffding-like concentration inequality, applied to last expression above:

LEMMA 5.4 (cf. [34, Lemma 3.4]). Consider an arbitrary but fixed sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^t$ for any rounds $t \geq 1$ and the load vector $x^{(0)}$. For any family of numbers $g_{u,v}^{(s)}, [u:v] \in \mathbf{M}^{(s)}, 1 \leq s \leq t$, define the random variable

$$Z := \sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} E_{u,v}^{(s)} \cdot g_{u,v}^{(s)}.$$

Then, $\mathbf{E}[Z] = 0$, and for any $\delta > 0$, it holds that

$$\Pr[|Z - \mathbf{E}[Z]| \geq \delta] \leq 2 \cdot \exp \left(- \frac{\delta^2}{2 \sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} \left(g_{u,v}^{(s)} \right)^2} \right).$$

The next lemma can be seen as a generalization of the first statement of Lemma 3.2 from [34], which considers the fixed vector a with $a_w := 1$ for node $w \in V$ and $a_u = 0$ for $u \neq w$. We extend it to arbitrary stochastic vectors a , which is crucial to prove Lemma 5.2 in its full generality.

LEMMA 5.5. Consider an arbitrary sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$. Let $(a_k)_{k \in V}$ be any stochastic vector. Then:

1. For any two rounds $0 \leq t_1 \leq t$ it holds,

$$\sum_{s=1}^{t_1} \sum_{[u:v] \in \mathbf{M}^{(s)}} \left(\sum_{k \in V} a_k \cdot \left(\mathbf{M}_{u,k}^{[s+1,t]} - \mathbf{M}_{v,k}^{[s+1,t]} \right) \right)^2 \leq 2 \cdot \|a\|_2^2.$$

2. For any two rounds $0 \leq t_1 \leq t$ it holds,

$$\sum_{w \in V} \left(\sum_{k \in V} a_k \cdot \mathbf{M}_{w,k}^{[t_1+1,t]} \right)^2 \leq \|a\|_2^2.$$

Proof. The proof follows closely the proof of Theorem 3.2 from [34]. We define a potential function as

$$\Psi^{(s)} := \sum_{w \in V} \left(\sum_{k \in V} a_k \mathbf{M}_{w,k}^{[s+1,t]} - \frac{1}{n} \right)^2,$$

for any round $1 \leq s \leq t$. This is a generalization of the potential function in [34], where all $a_k := 1$.

Proof of the first statement. Consider now any round $1 \leq s \leq t$, and let u, v be nodes with $[u : v] \in \mathbf{M}^{(s)}$. We have

$$\mathbf{M}_{u,k}^{[s,t]} = \mathbf{M}_{u,u}^{(s)} \cdot \mathbf{M}_{u,k}^{[s+1,t]} + \mathbf{M}_{u,v}^{(s)} \cdot \mathbf{M}_{v,k}^{[s+1,t]} = \frac{\mathbf{M}_{u,k}^{[s+1,t]} + \mathbf{M}_{v,k}^{[s+1,t]}}{2}.$$

To simplify notation, define for any two nodes $u, v \in V$,

$$y_{u,v} := \mathbf{M}_{u,v}^{[s+1,t]}.$$

With that notation, the contribution of these nodes to $\Psi^{(s)} - \Psi^{(s-1)}$ is

$$\begin{aligned} & \left(\sum_{k \in V} a_k \mathbf{M}_{u,k}^{[s+1,t]} - \frac{1}{n} \right)^2 + \left(\sum_{k \in V} a_k \mathbf{M}_{v,k}^{[s+1,t]} - \frac{1}{n} \right)^2 - \left(\sum_{k \in V} a_k \mathbf{M}_{u,k}^{[s,t]} - \frac{1}{n} \right)^2 - \left(\sum_{k \in V} a_k \mathbf{M}_{v,k}^{[s,t]} - \frac{1}{n} \right)^2 \\ &= \left(\sum_{k \in V} a_k y_{u,k} - \frac{1}{n} \right)^2 + \left(\sum_{k \in V} a_k y_{v,k} - \frac{1}{n} \right)^2 \\ & \quad - \left(\sum_{k \in V} a_k \frac{y_{u,k} + y_{v,k}}{2} - \frac{1}{n} \right)^2 - \left(\sum_{k \in V} a_k \frac{y_{u,k} + y_{v,k}}{2} - \frac{1}{n} \right)^2, \end{aligned}$$

having used that $\mathbf{M}_{u,k}^{[s,t]} = \mathbf{M}_{v,k}^{[s,t]} = \frac{y_{u,k} + y_{v,k}}{2}$. We can further rewrite the last expression as

$$\begin{aligned} &= \left(\sum_{k \in V} a_k y_{u,k} \right)^2 - \frac{2}{n} \cdot \sum_{k \in V} a_k y_{u,k} + \frac{1}{n^2} + \left(\sum_{k \in V} a_k y_{v,k} \right)^2 - \frac{2}{n} \cdot \sum_{k \in V} a_k y_{v,k} + \frac{1}{n^2} \\ & \quad - 2 \cdot \left(\sum_{k \in V} a_k \cdot \frac{y_{u,k} + y_{v,k}}{2} \right)^2 + \frac{4}{n} \cdot \sum_{k \in V} a_k \cdot \frac{y_{u,k} + y_{v,k}}{2} - \frac{2}{n^2} \\ &= \left(\sum_{k \in V} a_k y_{u,k} \right)^2 + \left(\sum_{k \in V} a_k y_{v,k} \right)^2 - \frac{1}{2} \cdot \left(\sum_{k \in V} a_k y_{u,k} + \sum_{k \in V} a_k y_{v,k} \right)^2 \\ & \quad + \frac{2}{n} \cdot \underbrace{\left(- \sum_{k \in V} a_k y_{u,k} - \sum_{k \in V} a_k y_{v,k} + \sum_{k \in V} a_k \cdot (y_{u,k} + y_{v,k}) \right)}_{=0} \\ &\stackrel{(a)}{=} \frac{1}{2} \cdot \left(\sum_{k \in V} a_k y_{u,k} - \sum_{k \in V} a_k y_{v,k} \right)^2, \end{aligned}$$

where in (a) we used the fact that $a^2 + b^2 - \frac{1}{2} \cdot (a + b)^2 = \frac{1}{2} \cdot (a - b)^2$. If a node is not matched in round s , then its contribution to $\Psi^{(s)} - \Psi^{(s-1)}$ equals zero. Accumulating the contribution of all nodes yields

$$(5.6) \quad \Psi^{(s)} - \Psi^{(s-1)} = \frac{1}{2} \cdot \sum_{[u:v] \in \mathbf{M}^{(s)}} \left(\sum_{k \in V} a_k \cdot (\mathbf{M}_{u,k}^{[s+1,t]} - \mathbf{M}_{v,k}^{[s+1,t]}) \right)^2.$$

Therefore,

$$\Psi^{(t_1)} \geq \Psi^{(t_1)} - \Psi^{(0)} = \sum_{s=1}^{t_1} (\Psi^{(s)} - \Psi^{(s-1)}) = \frac{1}{2} \cdot \sum_{s=1}^{t_1} \sum_{[u:v] \in \mathbf{M}^{(s)}} \left(\sum_{k \in V} a_k \cdot (\mathbf{M}_{u,k}^{[s+1,t]} - \mathbf{M}_{v,k}^{[s+1,t]}) \right)^2,$$

and rearranging gives,

$$\sum_{s=1}^{t_1} \sum_{[u:v] \in \mathbf{M}^{(s)}} \left(\sum_{k \in V} a_k \cdot \left(\mathbf{M}_{u,k}^{[s+1,t]} - \mathbf{M}_{v,k}^{[s+1,t]} \right) \right)^2 \leq 2 \cdot \Psi^{(t_1)} \leq 2 \cdot \Psi^{(t)},$$

where the last inequality holds since the potential $\Psi^{(s)}$ is non-decreasing over s ((5.6)). Moreover, since $\mathbf{M}^{[t+1,t]}$ is the identity matrix $\mathbf{I}$, we have $\mathbf{M}_{w,k}^{[t+1,t]} = 1$ for $k = w$ and $\mathbf{M}_{w,k}^{[t+1,t]} = 0$ for $k \neq w$. Hence

$$(5.7) \quad \Psi^{(t)} = \sum_{w \in V} \left(\sum_{k \in V} a_k \mathbf{M}_{w,k}^{[t+1,t]} - \frac{1}{n} \right)^2 = \sum_{w \in V} \left(a_w - \frac{1}{n} \right)^2 = \|a\|_2^2 - \frac{1}{n}.$$

Combining the last two statements finishes the proof of the first statement.

Proof of the second statement. Note that since $t_1 \leq t$, it follows from (5.6) that

$$\Psi^{(t_1)} \leq \Psi^{(t)}.$$

Substituting the definition of $\Psi^{(t_1)}$ and $\Psi^{(t)}$ on both sides gives

$$(5.8) \quad \sum_{w \in V} \left(\sum_{k \in V} a_k \cdot \mathbf{M}_{w,k}^{[t_1+1,t]} - \frac{1}{n} \right)^2 \leq \sum_{w \in V} \left(\sum_{k \in V} a_k \cdot \mathbf{M}_{w,k}^{[t+1,t]} - \frac{1}{n} \right)^2 \stackrel{\text{Eq. (5.7)}}{=} \|a\|_2^2 - \frac{1}{n}.$$

Finally, we have

$$\sum_{w \in V} \left(\sum_{k \in V} a_k \cdot \mathbf{M}_{w,k}^{[t_1+1,t]} \right)^2 \stackrel{(a)}{=} \sum_{w \in V} \left(\sum_{k \in V} a_k \cdot \mathbf{M}_{w,k}^{[t_1+1,t]} - \frac{1}{n} \right)^2 + \frac{1}{n} \stackrel{\text{Eq. (5.8)}}{=} \|a\|_2^2,$$

where (a) follows from Theorem 9.13 (since the matrix $\mathbf{M}^{[t_1+1,t]}$ is doubly stochastic and $\|a\|_1 = 1$). $\square$

Proof of Lemma 5.2. By (5.5),

$$\begin{aligned} \sum_{w \in V} a_w \cdot \left(X_w^{(t)} - \xi_w^{(t)} \right) &= \sum_{w \in V} a_w \cdot \left(\sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} E_{u,v}^{(s)} \cdot \left(\mathbf{M}_{u,w}^{[s+1,t]} - \mathbf{M}_{v,w}^{[s+1,t]} \right) \right) \\ &= \sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} E_{u,v}^{(s)} \cdot \left[\sum_{w \in V} a_w \cdot \left(\mathbf{M}_{u,w}^{[s+1,t]} - \mathbf{M}_{v,w}^{[s+1,t]} \right) \right]. \end{aligned}$$

Since the time interval $[1, t]$ is (K, κ) -smoothing, it holds that $-\kappa \leq \xi_w^{(t)} - \bar{x} \leq \kappa$ for all nodes $w \in V$, and therefore

$$\sum_{w \in V} a_w \cdot \xi_w^{(t)} = \bar{x} + B,$$

where $|B| \leq \kappa$. This implies

$$(5.9) \quad \sum_{w \in V} a_w \cdot X_w^{(t)} - \bar{x} = B + \sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} E_{u,v}^{(s)} \cdot \left[\sum_{w \in V} a_w \cdot \left(\mathbf{M}_{u,w}^{[s+1,t]} - \mathbf{M}_{v,w}^{[s+1,t]} \right) \right]$$

Following the notation of Lemma 5.4, let us define

$$Z := \sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} E_{u,v}^{(s)} \cdot \left[\sum_{w \in V} a_w \cdot \left(\mathbf{M}_{u,w}^{[s+1,t]} - \mathbf{M}_{v,w}^{[s+1,t]} \right) \right],$$

and

$$g_{u,v}^{(s)} := \sum_{w \in V} a_w \cdot \left(\mathbf{M}_{u,w}^{[s+1,t]} - \mathbf{M}_{v,w}^{[s+1,t]} \right).$$

Since $\mathbf{E} \left[E_{u,v}^{(s)} \right] = 0$, then $\mathbf{E} [Z] = 0$ and from Lemma 5.4 it follows that,

$$\Pr [|Z| \geq \delta] \leq 2 \cdot \exp \left(- \frac{\delta^2}{2 \sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} \left(g_{u,v}^{(s)} \right)^2} \right).$$

From the first statement of Lemma 5.5 it follows that $\sum_{s=1}^t \sum_{[u:v] \in \mathbf{M}^{(s)}} \left(g_{u,v}^{(s)} \right)^2 \leq 2 \cdot \|a\|_2^2$. Combining these two gives that $\Pr [|Z| \geq \delta] \leq 2 \cdot \exp(-\delta^2/(4 \cdot \|a\|_2^2))$. This together with (5.9) implies that,

$$\Pr \left[\left| \sum_{w \in V} a_w \cdot X_w^{(t)} - \bar{x} \right| \geq \delta + \kappa \right] \leq 2 \cdot \exp \left(- \frac{\delta^2}{4 \cdot \|a\|_2^2} \right),$$

and shifting δ by κ completes the proof. $\square$

5.4 Reducing the Number of Tokens to Linear We will now present an application of Lemma 5.2 to bound the number of tokens above the average load; this can be regarded as the number of tokens that need to be rearranged to achieve constant discrepancy. The proof relies on the flexibility of Lemma 5.2 by applying it multiple times for different coefficients $(a_w)_{w \in V}$ in order to prove that there is no large subset in which all nodes have large load.

LEMMA 5.6. *Assume our process applies a sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^{t_1}$ which is $(K, 1/(2n))$ -smoothing on an arbitrary initial load vector with $\text{disc}(x^{(0)}) \leq K$. Then it holds that*

$$\Pr \left[\sum_{w \in V} \max \left\{ X_w^{(t_1)} - \bar{x}, 0 \right\} \leq 16 \cdot n \right] \geq 1 - 2 \cdot n^{-2}.$$

Proof. For any integer $i = 1, 2, \dots, 2 \cdot \sqrt{\log(n)}$, we define the event

$$\mathcal{E}_i := \left\{ \left| \left\{ u \in V : X_u^{(t_1)} - \bar{x} \geq 4 \cdot i \right\} \right| \leq n \cdot 2^{-i} \right\} \quad \text{and} \quad \mathcal{E}_0 := \left\{ \max_{w \in V} X_w^{(t_1)} - \bar{x} \leq 6 \sqrt{\log(n)} \right\}.$$

From Lemma 5.2, second statement (with $\delta := \sqrt{36 \log(n)}$, $a_w = 1$ and $a_u = 0$ for $u \neq w$ and union bound over all nodes $w \in V$) we get

$$\Pr [\mathcal{E}_0] \geq 1 - n^{-2}.$$

It remains to analyze the probability of the other $\mathcal{E}_i$. To this end, fix any $1 \leq i \leq 2 \sqrt{\log(n)}$. By the union bound,

$$(5.10) \quad \Pr [\bar{\mathcal{E}}_i] \leq \sum_{S \subseteq V : |S|=n \cdot 2^{-i}} \Pr \left[\bigcap_{u \in S} \left\{ X_u^{(t_1)} \geq \bar{x} + 4 \cdot i \right\} \right].$$

To calculate this probability we apply Lemma 5.2 as follows. We define for a fixed subset $S \subseteq V$ a vector $(a_w)_{w \in V}$ by $a_w := 1/|S|$ for $w \in S$ and $a_w := 0$ otherwise. Then $\sum_{w \in V} a_w = 1$ and $\|a\|_2^2 = |S| \cdot (1/|S|)^2 = 1/|S|$, $\kappa = 1/n$, and we obtain that

$$\begin{aligned} \Pr \left[\bigcap_{u \in S} \left\{ X_u^{(t_1)} \geq \bar{x} + 4 \cdot i \right\} \right] &\leq \Pr \left[\sum_{u \in S} X_u^{(t_1)} \geq |S| \cdot (\bar{x} + 4 \cdot i) \right] \leq \Pr \left[\left| \sum_{w \in V} a_w \cdot X_w^{(t_1)} - \bar{x} \right| \geq 4 \cdot i \right] \\ &\stackrel{(a)}{\leq} 2 \cdot \exp \left(- \frac{(3i)^2}{4 \|a\|_2^2} \right) \leq 2 \cdot \exp \left(- (9/4) \cdot i^2 \cdot |S| \right) \end{aligned}$$

$$= 2 \cdot \exp\left(-(9/4) \cdot i^2 \cdot n \cdot 2^{-i}\right),$$

where (a) follows from Lemma 5.2. Plugging this into (5.10) yields

$$\begin{aligned} \Pr[\overline{\mathcal{E}_i}] &\leq \binom{n}{n \cdot 2^{-i}} \cdot 2 \cdot \exp\left(-(9/4) \cdot i^2 \cdot n \cdot 2^{-i}\right) \\ &\stackrel{(a)}{\leq} (e \cdot 2^i)^{n \cdot 2^{-i}} \cdot 2 \cdot \exp\left(-(9/4) \cdot i^2 \cdot n \cdot 2^{-i}\right) = 2 \cdot \left(\frac{e \cdot 2^i}{e^{(9/4) \cdot i^2}}\right)^{n \cdot 2^{-i}} \stackrel{(b)}{\leq} n^{-3}, \end{aligned}$$

where (a) used the estimate $\binom{n}{k} \leq (en/k)^k$ and (b) used that $i \in [1, 2\sqrt{\log(n)}]$. By another union bound over $i = 0, 1, \dots, 2\sqrt{\log(n)}$,

$$\Pr\left[\bigcap_{i=0}^{2\sqrt{\log(n)}} \mathcal{E}_i\right] \geq 1 - n^{-2} - 2\sqrt{\log(n)} \cdot n^{-3} \geq 1 - 2 \cdot n^{-2}.$$

We define a sequence $a_0 := -\infty$, $a_i := 4 \cdot i$ for $1 \leq i \leq 2\sqrt{\log(n)}$ and assume the event $\bigcap_{i=0}^{2\sqrt{\log(n)}} \mathcal{E}_i$ occurs and get

$$\begin{aligned} \sum_{w \in V} \max\{X_w^{(t_1)} - \bar{x}, 0\} &\leq \sum_{i=1}^{2\sqrt{\log(n)}} \left| \left\{ w \in V : X_w^{(t_1)} - \bar{x} \in (a_{i-1}, a_i] \right\} \right| \cdot a_i \\ &\leq \sum_{i=1}^{2\sqrt{\log(n)}} \left| \left\{ w \in V : X_w^{(t_1)} - \bar{x} \geq a_{i-1} \right\} \right| \cdot a_i \leq \sum_{i=1}^{\infty} (4n \cdot 2^{-i+1} \cdot i) \leq 8n \cdot \sum_{i=1}^{\infty} (i \cdot 2^{-i}) = 16n. \end{aligned}$$

This completes the proof. $\square$

6 Balancing a Linear Number of Tokens In this section we establish the most challenging step in our analysis. We consider an initial load vector $x^{(0)}$ with at most $(L - \varepsilon)n$ tokens for $\frac{4}{\log^4(n)} < \varepsilon < L$ and integer $L \geq 1$ (this is more general than requiring a linear number of tokens, which is what we will need in the final proof).

The main result of this section is given in Proposition 6.1, which states that after $O(\tau_{\text{global}} + \tau_{\text{local}} \cdot \log(n)/\log \log(n))$ rounds the maximum load is $L + 1$. An illustration of the proof method can be found in Figure 3.1.

PROPOSITION 6.1. *Let $t_2 := 2 \cdot \tau_{\text{global}} + 6 \cdot \tau_{\text{local}} \cdot \log(n)/\log \log(n)$, let L be any integer with $1 \leq L \leq \log^7(n)$ and let $\frac{4}{\log^4(n)} \leq \varepsilon < L$ (not necessarily constant). Assume our process applies a $(\tau_{\text{global}}, \tau_{\text{local}})$ -good sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^{\infty}$ to an arbitrary initial load vector $x^{(0)}$ with at most $(L - \varepsilon)n$ tokens. Then*

$$\Pr\left[\max_{w \in V} X_w^{(t_2)} \leq L + 1\right] \geq 1 - \exp\left(-\frac{\varepsilon \cdot \log(n)}{2 \cdot L \cdot \log \log(n)} + 8 \cdot \log \log(n)\right) - 2 \cdot n^{-2}.$$

The proposition makes uses of Lemma 6.11 and Lemma 6.12, which concern Phase 1 and Phase 2, respectively. Both are stated and proved at the end of this section.

Proof of Proposition 6.1. Phase 1. Let $t_1 := \tau_{\text{global}} + \tau_{\text{local}} \cdot \log(n)/\log \log(n)$ and define $Y^{(t_1)} := \sum_{u \in V} \max\{X_u^{(t_1)} - L, 0\}$ as the number of tokens at height at least $L + 1$. Define the event $\mathcal{G}_1 := \{Y^{(t_1)} \leq n/\log(n)\}$. From Lemma 6.11 (Phase 1 Lemma) it follows that

$$\Pr[\mathcal{G}_1] \geq 1 - \exp\left(-\frac{\varepsilon \cdot \log(n)}{2 \cdot L \cdot \log \log(n)} + 8 \cdot \log \log(n)\right) - n^{-2}.$$

Phase 2. For any $t \geq t_1$, we define an auxiliary load vector $\hat{X}^{(t)}$ with $\hat{X}_u^{(t)} := \max\{X_u^{(t)} - L, 0\}$. From Theorem 9.10 it follows that for any $t \geq t_1$, $X_u^{(t)} \leq \hat{X}_u^{(t)} + L$. Hence it is sufficient to show that $\max_{u \in V} \hat{X}_u^{(t_2)} \leq 1$

for some suitable round t_2 . Note that from Phase 1 we get $\sum_{u \in V} \hat{X}_u^{(t_1)} \leq n/\log(n)$, i.e. the load vector $\hat{X}^{(t_1)}$ has at most $n/\log(n)$ tokens. This time we define, for any round $t \geq t_1$, $\hat{Y}^{(t)} := \sum_{u \in V} \max\{\hat{X}_u^{(t)} - 1, 0\}$, which is equal to the number of tokens with height at least 2 in $\hat{X}^{(t)}$. Let $t_2 := t_1 + \tau_{\text{global}} + 4\tau_{\text{local}} \cdot \frac{\log(n)}{-\log(\frac{1}{\log(n)} + \frac{2}{\log^4(n)})}$. We define a second event $\mathcal{G}_2 := \{\hat{Y}^{(t_2)} = 0\}$. From Lemma 6.12 (Phase 2 Lemma) with $\varepsilon := 1 - \frac{1}{\log(n)}$ it follows that $\Pr[\mathcal{G}_2 \mid \mathcal{G}_1] \geq 1 - n^{-2}$, and by the definition of conditional probability we get

$$\begin{aligned} \Pr[\mathcal{G}_2] &\geq \Pr[\mathcal{G}_2 \mid \mathcal{G}_1] \cdot \Pr[\mathcal{G}_1] \geq 1 - n^{-2} - \Pr[\overline{\mathcal{G}_1}] \\ &\geq 1 - \exp\left(-\frac{\varepsilon \cdot \log(n)}{2 \cdot L \cdot \log \log(n)} + 8 \cdot \log \log(n)\right) - 2 \cdot n^{-2}. \end{aligned}$$

From the definition of t_2 it follows that

$$\begin{aligned} t_2 &= t_1 + \tau_{\text{global}} + \frac{4 \cdot \log(n)}{-\log\left(\frac{1}{\log(n)} + \frac{2}{\log^4(n)}\right)} \cdot \tau_{\text{local}} \leq t_1 + \tau_{\text{global}} + \frac{5 \cdot \log(n)}{\log \log(n)} \cdot \tau_{\text{local}} \\ &= 2 \cdot \tau_{\text{global}} + \frac{6 \cdot \log(n)}{\log \log(n)} \cdot \tau_{\text{local}}, \end{aligned}$$

finishing the proof. $\square$

In the rest of this section we will prove Lemma 6.11 and Lemma 6.12. To do so, we have to provide several additional definitions and statements, which ultimately leads towards the key lemma (Lemma 6.8), which proves a drop in the number of tokens with height at least $L + 1$ within τ_{local} rounds. From this, Lemma 6.11 and Lemma 6.12 can be derived relatively easily.

First of all, we define two events corresponding to τ_{global} (global mixing) and τ_{local} (local mixing).

DEFINITION 6.2. *For any round $t \geq \tau_{\text{global}}$ we define*

$$\Gamma_g^{(t)} := \left\{ \bigcap_{u \in V} \left\| \mathbf{M}_{u,\cdot}^{[1,t]} - \frac{\vec{1}}{n} \right\|_2^2 \leq \frac{1}{n^7} \right\}.$$

For any round $t \geq \tau_{\text{local}}$ and arbitrary node $u \in V$ we define

$$\Gamma_\ell^{(t)}(u) := \left\{ \left\| \mathbf{M}_{u,\cdot}^{[t-\tau_{\text{local}}+1,t]} \right\|_2^2 \leq \frac{1}{\log^{10}(n)} \right\}.$$

Note that the event $\Gamma_g^{(t)}$ depends on the matching matrices in the time interval $[1, t]$. The event implies that the balancing matrices applied during rounds $[1, t]$ are $(n \cdot \log^7(n), 1/n)$ -smoothing (see Theorem 9.9). We call that property *globally mixing*, since it implies that a token starting from any node will reach each node in round t with probability $(1 \pm o(1))/n$. In contrast, the event $\Gamma_\ell^{(t)}(u)$ adopts the viewpoint of a single node u . It implies that the matchings chosen during the last τ_{local} rounds are locally smoothing from the viewpoint u . $\Gamma_\ell^{(t)}(u)$ depends on the matchings applied in the time interval $[t - \tau_{\text{local}} + 1, t]$. We remark that for any round $t \geq 1$, the events $\Gamma_g^{(t)}$ and $\Gamma_\ell^{(t+\tau_{\text{local}})}(u)$ are statements over disjoint time-intervals, hence for the random matching model they are independent events. In the following we first calculate the probabilities that these events occur for a sequence of $(\tau_{\text{global}}, \tau_{\text{local}})$ -good matchings, the proof is a straightforward calculation.

LEMMA 6.3. *For any process generating $(\tau_{\text{global}}, \tau_{\text{local}})$ -good matchings the following holds.*

1. *For any round $t \geq \tau_{\text{global}}$, $\Pr[\Gamma_g^{(t)}] \geq 1 - 1/n^3$.*
2. *For any round $t \geq \tau_{\text{local}}$ and arbitrary node $u \in V$, $\Pr[\Gamma_\ell^{(t)}(u)] \geq 1 - 1/\log^{11} n$.*

Proof. First note that

$$(6.1) \quad \Pr[\overline{\Gamma_g^{(\tau_{\text{global}})}}] = \Pr\left[\bigcup_{u \in V} \left\{ \left\| \mathbf{M}_{u,\cdot}^{[1,\tau_{\text{global}}]} - \frac{\vec{1}}{n} \right\|_2^2 > \frac{1}{n^7} \right\}\right].$$

Applying the definition of τ_{global} (Definition 2.2), we get for $t = \tau_{\text{global}}$ that $\Pr \left[\overline{\Gamma_g^{(\tau_{\text{global}})}} \right] \leq 1/n^3$. The first statement follows from the basic fact that $\left\| \mathbf{M}_{u,\cdot}^{[1,t]} - \frac{\mathbf{1}}{n} \right\|_2^2$ is non-increasing in t (see Theorem 9.14). For the second statement, consider an arbitrary round $t \geq \tau_{\text{local}}$ and a fixed node $u \in V$. Here, we get

$$\Pr \left[\overline{\Gamma_\ell^{(t)}(u)} \right] = \Pr \left[\left\{ \left\| \mathbf{M}_{u,\cdot}^{[t-\tau_{\text{local}}+1,t]} \right\|_2^2 > \frac{1}{\log^{10}(n)} \right\} \right] \leq \frac{1}{\log^{11}(n)},$$

where the last inequality follows from the definition of τ_{local} (Definition 2.2). As above, the proof follows from Theorem 9.14. $\square$

Recall that Phase 1 starts at round τ_{global} and each of our two phases is subdivided into $\log(n)/\log \log(n)$ epochs of τ_{local} many rounds. We refer to the last round of an epoch as *milestone*. For any epoch k we denote by $e : \mathbb{N}_0 \rightarrow \mathbb{N}_0$ the function $e(k) := \tau_{\text{global}} + k \cdot \tau_{\text{local}}$ which returns the last round of k -th epoch, i.e., $e(k)$ returns the k -th milestone. We define $E(k) := [e(k-1) + 1, \dots, e(k)]$ for the interval of rounds constituting the k -th epoch.

Let the random variable $Y^{(t)}$ denote the number of tokens of height at least $L+1$ at round t , i.e.,

$$(6.2) \quad Y^{(t)} := \sum_{u \in V} \max \left\{ X_u^{(t)} - L, 0 \right\} = \sum_{j \in \mathcal{T}} \mathbf{1}_{H_j^{(t)} \geq L+1}.$$

We wish to prove that in each epoch $Y^{(t)}$ drops by a constant factor (see Lemma 6.8) such that at the end of Phase 1 the number of tokens of height at least $L+1$ is at most $n/\log(n)$ (Lemma 6.11). Since the height of the token is non-increasing this number will not increase for the rest of the process.

Let us first focus on the k -th milestone, for $k \in \mathbb{N}$. For now, let us fix an arbitrary location vector $w^{(e(k-1))}$. Consider an arbitrary token $i \in \mathcal{T}$ which has height at least $L+1$ at round $e(k-1)$. Our goal is to prove that the expected number of tokens that collide with i at round $e(k)$ is smaller than L . To this end, we define an indicator random variable $Z_{i,j}^{(e(k))}(v)$ for any token $j \in \mathcal{T}, j \neq i$ and any node $v \in V$ as

$$Z_{i,j}^{(e(k))}(v) := \mathbf{1}_{W_i^{(e(k))}=v \cap W_j^{(e(k))}=v} \quad \text{and} \quad Z_i^{(e(k))} := \sum_{v \in V} \sum_{j \in \mathcal{T}: j \neq i} Z_{ij}^{(e(k))}(v),$$

i.e., $Z_{i,j}^{(e(k))}(v) = 1$ if and only if tokens i and j are both on node v at the k -th milestone. Hence $Z_i^{(e(k))}$ counts the number of tokens colliding with token i at that time.

The next lemma bounds the expected number of collisions. Note that we assume in the lemma that the applied matrices are fixed. The randomness is due to the random decisions in the shuffling step.

LEMMA 6.4. *Let $(\mathbf{M}^{(t)})_{t=1}^{e(k)}$ be an arbitrary but fixed sequence of matchings. For any milestone $k \in \mathbb{N}$ and any token $i \in \mathcal{T}$ it holds that*

$$\mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, (\mathbf{M}^{(t)})_{t=1}^{e(k)} \right] \leq \sum_{w \in V} \left(\sum_{v \in V} \mathbf{M}_{W_i^{(e(k-1))}, v}^{[e(k-1)+1, e(k)]} \cdot \mathbf{M}_{w, v}^{[e(k-1)+1, e(k)]} \right) \cdot X_w^{(e(k-1))}.$$

Proof. First we focus on a fixed location vector $w^{(e(k-1))}$ and compute,

$$\begin{aligned} & \mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))} = w^{(e(k-1))}, (\mathbf{M}^{(t)})_{t=1}^{e(k)} \right] \\ & \stackrel{(a)}{=} \sum_{v \in V} \sum_{j \in \mathcal{T}: j \neq i} \mathbf{E} \left[Z_{ij}^{(e(k))}(v) \mid W^{(e(k-1))} = w^{(e(k-1))}, (\mathbf{M}^{(t)})_{t=1}^{e(k)} \right] \\ & \stackrel{(b)}{=} \sum_{v \in V} \sum_{j \in \mathcal{T}: j \neq i} \Pr \left[Z_{ij}^{(e(k))}(v) = 1 \mid W^{(e(k-1))} = w^{(e(k-1))}, (\mathbf{M}^{(t)})_{t=1}^{e(k)} \right] \\ (6.3) \quad & = \sum_{v \in V} \sum_{j \in \mathcal{T}: j \neq i} \Pr \left[W_i^{(e(k))} = v \cap W_j^{(e(k))} = v \mid W^{(e(k-1))} = w^{(e(k-1))}, (\mathbf{M}^{(t)})_{t=1}^{e(k)} \right], \end{aligned}$$

where (a) uses linearity of conditional expectation and (b) uses the fact that the $Z_{i,j}^{(e(k))}(v)$ are indicator random variables. Recall the definition $E(k) = [e(k-1)+1, \dots, e(k)]$. Crucially, we can now apply our negative correlation result (Lemma 5.1) to (6.3) to obtain

$$\begin{aligned}
& \mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))} = w^{(e(k-1))}, \left(\mathbf{M}^{(t)} \right)_{t=1}^{e(k)} \right] \\
& \leq \sum_{v \in V} \sum_{j \in \mathcal{T}: j \neq i} \mathbf{Pr} \left[W_i^{(e(k))} = v \mid W^{(e(k-1))} = w^{(e(k-1))}, \left(\mathbf{M}^{(t)} \right)_{t=1}^{e(k)} \right] \\
& \quad \cdot \mathbf{Pr} \left[W_j^{(e(k))} = v \mid W^{(e(k-1))} = w^{(e(k-1))}, \left(\mathbf{M}^{(t)} \right)_{t=1}^{e(k)} \right] \\
& = \sum_{v \in V} \mathbf{M}_{w_i^{(e(k-1))}, v}^{E(k)} \cdot \sum_{j \in \mathcal{T}} \mathbf{M}_{w_j^{(e(k-1))}, v}^{E(k)} \\
& = \sum_{v \in V} \mathbf{M}_{w_i^{(e(k-1))}, v}^{E(k)} \cdot \sum_{w \in V} \mathbf{M}_{w, v}^{E(k)} \cdot x_w^{(e(k-1))} \\
& = \sum_{w \in V} \left(\sum_{v \in V} \mathbf{M}_{w_i^{(e(k-1))}, v}^{E(k)} \cdot \mathbf{M}_{w, v}^{E(k)} \right) \cdot x_w^{(e(k-1))}.
\end{aligned}$$

Since the above estimate holds for all location vectors $w^{(e(k-1))}$, it follows

$$\mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \left(\mathbf{M}^{(t)} \right)_{t=1}^{e(k)} \right] \leq \sum_{w \in V} \left(\sum_{v \in V} \mathbf{M}_{w_i^{(e(k-1))}, v}^{E(k)} \cdot \mathbf{M}_{w, v}^{E(k)} \right) \cdot x_w^{(e(k-1))}. \quad \square$$

We now define two more events which we use to track the decrease of $Y^{(t)}$ from epoch to epoch.

DEFINITION 6.5. Let $\frac{4}{\log^4(n)} \leq \varepsilon < L$. Fix a node u and let i be a token located on u at milestone $e(k-1)$.

1. $\mathcal{L}^{(e(k))}(u) := \left\{ \mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] \leq L - \varepsilon + \frac{1}{\log^4(n)} \right\}$.
2. $\mathcal{E}^{(e(k))}(u) := \left\{ \Gamma_g^{(e(k-1))} \cap \left(\overline{\Gamma_\ell^{(e(k))}(u)} \cup \mathcal{L}^{(e(k))}(u) \right) \right\}$ and $\mathcal{E}^{(e(k))} := \bigcap_{u \in V} \mathcal{E}^{(e(k))}(u)$.

The idea behind the events defined above is as follows (see also Figure 6.1 for an illustration). In the definition of $\mathcal{L}^{(e(k))}(u)$ we condition on $\Gamma_\ell^{(e(k))}(u)$ which means that, from the viewpoint of node u (or token i), the randomly chosen matchings in epoch k ensure that token i mixes “locally”. If now the matchings chosen during the time interval $[1, e(k-1)]$ also suffice for a “global mixing”, then the expected number of tokens colliding with token i on node v (the location of token i at time $e(k)$) is less than L . The event $\mathcal{E}^{(e(k))}(u)$ occurs when the matchings chosen during the time interval $[1, e(k-1)]$ are globally mixing and the last epoch k was “locally mixing” for token i .




Figure 6.1: Illustration of the events $\Gamma_\ell^{(e(k))}(u)$, $\Gamma_g^{(e(k-1))}$ and $\mathcal{L}^{(e(k))}(u)$, where $L = 2$. The token i , marked in red, is located at node u in round $e(k-1)$ and has height $L+1=3$. Its location v in round $e(k)$ is random and chosen according to $\mathbf{M}_{u,\cdot}^{E(k)}$. In order to keep its height $L+1=3$, there must be at least $L=2$ other tokens on v in round $e(k)$.

Roughly speaking, the event $\mathcal{L}^{(e(k))}(u)$ corresponds to a certain upper bound on a convex combination of loads at round $e(k-1)$ holding; this event depends both on the matchings and shuffling decisions in $[0, e(k-1)]$ and the matchings in $E(k)$. The coefficients of the load vector are collision probabilities so that the **left-hand side** equals the expected number of tokens that reach the (random) location of token i in round $e(k)$.

We prove in Lemma 6.6 that if $\Gamma_\ell^{(e(k))}(u)$ and $\Gamma_g^{(e(k-1))}$ occur, then also $\mathcal{L}^{(e(k))}(u)$ occurs (with high probability). Then, conditional on this implication, we prove in our key lemma (Lemma 6.8) that the number of tokens with height $L+1=3$ drops within $E(k)$ by a suitable factor.

LEMMA 6.6. *Assume that the load vector $x^{(0)}$ has at most $(L-\varepsilon)n$ tokens for $1 \leq L \leq \log^7(n)$ being an integer and any $0 < \varepsilon < L$. Then for any milestone $k \geq 1$ we have $\Pr[\mathcal{E}^{(e(k))}] \geq 1 - \frac{2}{n^3}$.*

Proof. Consider a node $u \in V$; recall token i is located on u at round $e(k-1)$. First let us negate the two events from Definition 6.5 which gives

$$\begin{aligned} \overline{\mathcal{L}^{(e(k))}(u)} &= \left\{ \mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] > L - \varepsilon + \frac{1}{\log^4(n)} \right\} \quad \text{and} \\ \overline{\mathcal{E}^{(e(k))}(u)} &= \left\{ \overline{\Gamma_g^{(e(k-1))}} \cup \left(\Gamma_\ell^{(e(k))}(u) \cap \overline{\mathcal{L}^{(e(k))}(u)} \right) \right\}. \end{aligned}$$

For brevity, let us write $\alpha := \Pr \left[\overline{\Gamma_g^{(e(k-1))}} \right]$ and note that $\alpha \leq 1/n^3$ by Lemma 6.3. Then,

$$\begin{aligned}
& \Pr \left[\overline{\mathcal{E}^{(e(k))}}(u) \right] \\
& \leq \alpha + \Pr \left[\Gamma_\ell^{(e(k))}(u) \cap \overline{\mathcal{L}^{(e(k))}}(u) \right] \\
& \leq \alpha + \Pr \left[\overline{\mathcal{L}^{(e(k))}}(u) \right] \\
& = \alpha + \Pr \left[\mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] > L - \varepsilon + \frac{1}{\log^4(n)} \right] \\
& = \alpha + \Pr \left[\mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \Gamma_g^{(e(k-1))} \right] \cdot \Pr \left[\overline{\Gamma_g^{(e(k-1))}} \right] + \right. \\
& \quad \left. \mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \overline{\Gamma_g^{(e(k-1))}} \right] \cdot \Pr \left[\overline{\Gamma_g^{(e(k-1))}} \right] > L - \varepsilon + \frac{1}{\log^4(n)} \right] \\
& \leq \alpha + \Pr \left[\mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \Gamma_g^{(e(k-1))} \right] > L - \varepsilon + \frac{1}{\log^4(n)} - \frac{\log^7(n)}{n^2} \right],
\end{aligned}$$

where the last inequality holds since first $\Pr \left[\Gamma_g^{(e(k-1))} \right] \leq 1$, and secondly, since (deterministically) $Z_i^{(e(k))} \leq L \cdot n \leq \log^7(n) \cdot n$ we get,

$$\mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \overline{\Gamma_g^{(e(k-1))}} \right] \cdot \Pr \left[\overline{\Gamma_g^{(e(k-1))}} \right] \leq \frac{\log^7(n)}{n^2}.$$

We thus have

$$\begin{aligned}
& \Pr \left[\overline{\mathcal{E}^{(e(k))}}(u) \right] \\
(6.4) \quad & \leq \alpha + \Pr \left[\mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \Gamma_g^{(e(k-1))} \right] > L - \varepsilon + \frac{1}{\log^4(n)} - \frac{\log^7(n)}{n^2} \right].
\end{aligned}$$

To upper bound (6.4) we apply Claim 6.7 below. Since Claim 6.7 holds for any fixed sequence $(\mathbf{M}^{(t)})_{t=1}^{e(k)}$ satisfying $\Gamma_\ell^{(e(k))}(u) \cap \Gamma_g^{(e(k-1))}$, we have

$$\Pr \left[\mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \Gamma_g^{(e(k-1))} \right] > L - \varepsilon + \frac{1}{2\log^4(n)} \right] \leq \frac{1}{n^5}.$$

Applying the union bound twice implies that

$$\Pr \left[\overline{\mathcal{E}^{(e(k))}} \right] = \Pr \left[\overline{\Gamma_g^{(e(k-1))}} \cup \bigcup_{u \in V} \left(\Gamma_\ell^{(e(k))}(u) \cap \overline{\mathcal{L}^{(e(k))}}(u) \right) \right] \leq \alpha + n \cdot \frac{1}{n^5} \leq \frac{2}{n^3}. \quad \square$$

CLAIM 6.7. Fix a node u and let $(\mathbf{M}^{(t)})_{t=1}^{e(k)}$ be a fixed sequence of matchings satisfying

$$\Gamma_\ell^{(e(k))}(u) \cap \Gamma_g^{(e(k-1))}.$$

Then

$$\Pr \left[\left\{ \mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, (\mathbf{M}^{(t)})_{t=1}^{e(k)} \right] > L - \varepsilon + \frac{1}{2\log^4(n)} \right\} \right] \leq \frac{1}{n^5}.$$

Proof. As before, $E(k) = [e(k-1) + 1, \dots, e(k)]$ is the interval of rounds constituting the k -th epoch. Let $W_i^{(e(k-1))} = u$ and note

$$\Pr \left[\left\{ \mathbf{E} \left[Z_i^{(e(k))} \mid W^{(e(k-1))}, (\mathbf{M}^{(t)})_{t=1}^{e(k)} \right] > L - \varepsilon + \frac{1}{2\log^4(n)} \right\} \right]$$

$$(6.5) \quad \stackrel{(a)}{\leq} \Pr \left[\sum_{w \in V} \left(\sum_{v \in V} \mathbf{M}_{u,v}^{E(k)} \cdot \mathbf{M}_{w,v}^{E(k)} \right) \cdot X_w^{(e(k-1))} \geq L - \varepsilon + \frac{1}{2 \log^4(n)} \right],$$

where (a) follows from Lemma 6.4. To bound the probability of (6.5) we will apply our concentration inequality (Lemma 5.2). Note that, since the matchings are fixed, the only randomness remaining in Lemma 5.2 are the movements of tokens in the shuffling step. We define an n -dimensional vector $(a_w)_{w \in V}$ as $a_w := \sum_{v \in V} \mathbf{M}_{u,v}^{E(k)} \cdot \mathbf{M}_{w,v}^{E(k)}$. To apply Lemma 5.2 we need to (1) show the time interval $[1, e(k-1)]$ is $(K, 1/n)$ -smoothing for some $K \leq n \cdot L \leq n \cdot \log^7(n)$, (2) show $\|a\|_1 = 1$ and (3) compute $\|a\|_2^2$. Since the sequence of matchings $(\mathbf{M}^{(t)})_{t=1}^{e(k-1)}$ satisfies the event $\Gamma_g^{(e(k-1))}$, it follows from Theorem 9.9 that this sequence is $(n \cdot \log^7(n), 1/n)$ -smoothing. Since the matrix $\mathbf{M}^{E(k)} := \prod_{s=e(k-1)+1}^{e(k)} \mathbf{M}^{(s)}$ is doubly stochastic we get

$$\|a\|_1 = \sum_{w \in V} a_w = \sum_{w \in V} \sum_{v \in V} \mathbf{M}_{u,v}^{E(k)} \cdot \mathbf{M}_{w,v}^{E(k)} = \sum_{v \in V} \mathbf{M}_{u,v}^{E(k)} \cdot \sum_{w \in V} \mathbf{M}_{w,v}^{E(k)} = \sum_{v \in V} \mathbf{M}_{u,v}^{E(k)} = 1.$$

Furthermore,

$$\|a\|_2^2 = \sum_{w \in V} \left(\sum_{v \in V} \mathbf{M}_{u,v}^{E(k)} \cdot \mathbf{M}_{w,v}^{E(k)} \right)^2 \stackrel{(a)}{\leq} \left\| \mathbf{M}_{u,\cdot}^{E(k)} \right\|_2^2 \stackrel{(b)}{\leq} \frac{1}{\log^{10}(n)},$$

where (a) follows from the second statement of Lemma 5.5; (b) follows from the definition of the event $\Gamma_\ell^{(e(k))}(u)$. Since we have at most $(L - \varepsilon)n$ many tokens the average load $\bar{x}$ satisfies $\bar{x} \leq L - \varepsilon$. From Lemma 5.2 with $t := e(k-1)$, $\kappa := 1/n$, $\delta := 1/(2 \log^4(n))$, and the bound on $\|a\|_2^2$ from above, we have

$$(6.6) \quad \Pr \left[\sum_{w \in V} a_w \cdot X_w^{(e(k-1))} \geq L - \varepsilon + \frac{1}{2 \log^4(n)} \right] \leq 2 \cdot \exp \left(- \frac{\left(\frac{1}{2 \log^4(n)} - \frac{1}{n} \right)^2}{4 \cdot \left(\frac{1}{\log^{10}(n)} \right)} \right) \leq \frac{1}{n^5}.$$

Combining (6.5) and (6.6) finishes the proof of the claim. $\square$

Next we define two random variables which will be used in the remainder of the proof. Recall that $\mathcal{T}$ is the set of tokens, $Z_i^{(e(k))}$ counts the number of tokens colliding with token i at the k -th milestone and for any $t \geq 0$, $Y^{(t)}$ is the number of tokens with height at least $L+1$ at round t . For any milestone $k \in \mathbb{N}$ we define,

$$(6.7) \quad \tilde{Y}^{(e(k))} := Y^{(e(k))} \cdot \mathbf{1}_{\cap_{i=0}^k \mathcal{E}^{(e(i))}}.$$

Later we will show how to use $Z_i^{(e(k))}$ to bound $\tilde{Y}^{(e(k))}$, and then eventually $Y^{(e(k))}$. In the next lemma we bound the expected value of $\tilde{Y}^{(e(k))}$. We will denote by $(\mathfrak{F}^{(t)})_{t \geq 0}$ the filtration of the random process; note that in particular, $\mathfrak{F}^{(t)}$ determines not only the current load vector and previous load vectors $X^{(t)}, X^{(t-1)}, \dots, X^{(1)}$, but also all location vectors $W^{(t)}, W^{(t-1)}, \dots, W^{(1)}$. In case of randomly generated matchings $(\mathbf{M}^{(s)})_{s=1}^t$ is determined by $(\mathfrak{F}^{(t)})_{t \geq 0}$, too. In the following, to keep the notation brief, we use the convention that for any round $t \geq 0$, any random variable X and event $\mathcal{E}$,

$$\mathbf{E}^{(t)}[X] := \mathbf{E}[X \mid \mathfrak{F}^{(t)}] \quad \text{and} \quad \Pr^{(t)}[\mathcal{E}] := \Pr[\mathcal{E} \mid \mathfrak{F}^{(t)}].$$

LEMMA 6.8 (Key Lemma – Expected Drop in one Epoch). *Assume that the initial load vector $x^{(0)}$ has at most $(L - \varepsilon)n$ tokens for $1 \leq L \leq \log^7 n$ being an integer and $0 < \varepsilon < 1$. Then for any milestone $k \geq 1$,*

$$\mathbf{E}^{(e(k-1))}[\tilde{Y}^{(e(k))}] \leq \left(1 - \frac{\varepsilon}{L} + \frac{2}{L \cdot \log^4 n} \right) \cdot \tilde{Y}^{(e(k-1))}.$$

We emphasize that in the above lemma, it is possible to have an ε that depends on n .

Proof. Using the definition of $\tilde{Y}^{(e(k))}$ we get

$$\begin{aligned}
& \mathbf{E}^{(e(k-1))} \left[\tilde{Y}^{(e(k))} \right] \\
&= \mathbf{E}^{(e(k-1))} \left[Y^{(e(k))} \cdot \mathbf{1}_{\cap_{i=0}^k \mathcal{E}^{(e(i))}} \right] \\
&\stackrel{(a)}{=} \mathbf{1}_{\cap_{i=0}^{k-1} \mathcal{E}^{(e(i))}} \cdot \mathbf{E}^{(e(k-1))} \left[Y^{(e(k))} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \right] \\
&\stackrel{(b)}{=} \mathbf{1}_{\cap_{i=0}^{k-1} \mathcal{E}^{(e(i))}} \cdot \left(\mathbf{Pr}^{(e(k-1))} \left[\Gamma_g^{(e(k-1))} \right] \cdot \mathbf{E}^{(e(k-1))} \left[Y^{(e(k))} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))} \right] \right. \\
&\quad \left. + \mathbf{Pr}^{(e(k-1))} \left[\overline{\Gamma_g^{(e(k-1))}} \right] \cdot 0 \right) \\
(6.8) \quad &\leq \mathbf{1}_{\cap_{i=0}^{k-1} \mathcal{E}^{(e(i))}} \cdot \mathbf{E}^{(e(k-1))} \left[Y^{(e(k))} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))} \right],
\end{aligned}$$

where (a) is an application of the “take-out-what-is-known” rule in conditional expectations, since $\mathfrak{F}^{(e(k-1))}$ determines the random variable $\mathbf{1}_{\cap_{i=0}^{k-1} \mathcal{E}^{(e(i))}}$ (Lemma 9.3), and (b) used that $\overline{\Gamma_g^{(e(k-1))}}$ implies $\mathbf{1}_{\mathcal{E}^{(e(k))}} = 0$. Further, we have

$$(6.9) \quad Y^{(e(k))} = \sum_{j \in \mathcal{T}} \mathbf{1}_{H_j^{(e(k))} \geq L+1} \leq \sum_{j \in \mathcal{T}} \mathbf{1}_{H_j^{(e(k-1))} \geq L+1} \cdot \mathbf{1}_{Z_j^{(e(k))} \geq L},$$

where the equality follows from the definition of $Y^{(e(k))}$. To see the inequality, observe that by properties of the height-sensitive process the height of a token never increases and therefore, in order for a token j to be at height at least $L+1$ in round $e(k)$, it must have had height at least $L+1$ in round $e(k-1)$ and there must be at least L other tokens at its location in round $e(k)$.

Applying (6.9) to (6.8),

$$\begin{aligned}
& \mathbf{E}^{(e(k-1))} \left[\tilde{Y}^{(e(k))} \right] \\
&\leq \mathbf{1}_{\cap_{i=0}^{k-1} \mathcal{E}^{(e(i))}} \cdot \mathbf{E}^{(e(k-1))} \left[\sum_{j \in \mathcal{T}} \mathbf{1}_{H_j^{(e(k-1))} \geq L+1} \cdot \mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))} \right] \\
&\stackrel{(a)}{=} \mathbf{1}_{\cap_{i=0}^{k-1} \mathcal{E}^{(e(i))}} \cdot \sum_{j \in \mathcal{T}} \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{H_j^{(e(k-1))} \geq L+1} \cdot \mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))} \right] \\
(6.10) \quad &\stackrel{(b)}{=} \mathbf{1}_{\cap_{i=0}^{k-1} \mathcal{E}^{(e(i))}} \cdot \sum_{j \in \mathcal{T}} \mathbf{1}_{H_j^{(e(k-1))} \geq L+1} \cdot \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))} \right],
\end{aligned}$$

where (a) holds by linearity of conditional expectation and (b) holds since $\mathfrak{F}^{(e(k-1))}$ reveals $\mathbf{1}_{H_j^{(e(k-1))} \geq L+1}$ (“take-out-what-is-known”, Lemma 9.3). To simplify the notation we will write u instead of $w_j^{(e(k-1))}$ in what follows. Conditioning on whether $\Gamma_\ell^{(e(k))}(u)$ holds, we can bound the expectation from (6.10) by

$$\begin{aligned}
& \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))} \right] \\
&= \mathbf{Pr}^{(e(k-1))} \left[\overline{\Gamma_\ell^{(e(k))}(u)} \mid \Gamma_g^{(e(k-1))} \right] \cdot \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \overline{\Gamma_\ell^{(e(k))}(u)} \right] \\
&\quad + \mathbf{Pr}^{(e(k-1))} \left[\Gamma_\ell^{(e(k))}(u) \mid \Gamma_g^{(e(k-1))} \right] \cdot \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] \\
&\leq \mathbf{Pr}^{(e(k-1))} \left[\overline{\Gamma_\ell^{(e(k))}(u)} \mid \Gamma_g^{(e(k-1))} \right] \cdot 1 + 1 \cdot \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right].
\end{aligned}$$

We upper bound the two remaining terms in the last line separately using the following two claims, whose proofs are given after the conclusion of the proof of this lemma.

CLAIM 6.9. $\mathbf{Pr}^{(e(k-1))} \left[\overline{\Gamma_\ell^{(e(k))}}(u) \mid \Gamma_g^{(e(k-1))} \right] \leq \frac{1}{\log^{11}(n)}.$

CLAIM 6.10. $\mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] \leq 1 - \frac{\varepsilon}{L} + \frac{1}{L \cdot \log^4(n)}.$

Together Claim 6.9 and Claim 6.10 yield, using the assumption that $L \leq \log^7(n)$,

$$\mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))} \right] \leq 1 - \frac{\varepsilon}{L} + \frac{1}{L \cdot \log^4(n)} + \frac{1}{\log^{11}(n)} \leq 1 - \frac{\varepsilon}{L} + \frac{2}{L \cdot \log^4(n)}.$$

Applying this to (6.10) gives us

$$\begin{aligned} \mathbf{E}^{(e(k-1))} \left[\tilde{Y}^{(e(k))} \right] &\leq \mathbf{1}_{\cap_{i=1}^{k-1} \mathcal{E}^{(e(i))}} \cdot \sum_{j \in \mathcal{T}} \mathbf{1}_{H_j^{(e(k-1))} \geq L+1} \cdot \left(1 - \frac{\varepsilon}{L} + \frac{2}{L \cdot \log^4(n)} \right) \\ &= \mathbf{1}_{\cap_{i=1}^{k-1} \mathcal{E}^{(e(i))}} \cdot Y^{(e(k-1))} \cdot \left(1 - \frac{\varepsilon}{L} + \frac{2}{L \cdot \log^4(n)} \right) \\ &= \tilde{Y}^{(e(k-1))} \cdot \left(1 - \frac{\varepsilon}{L} + \frac{2}{L \cdot \log^4(n)} \right). \end{aligned} \quad \square$$

Proof of Claim 6.9. As before we define $E(k) = [e(k-1) + 1, \dots, e(k)]$ for the interval of rounds constituting the k -th epoch. Note that $\Gamma_\ell^{(e(k))}(u)$ depends only on the matchings in the time interval $E(k)$. Moreover, as $\Gamma_g^{(e(k-1))}$ and $\Gamma_\ell^{(e(k))}(u)$ refer to disjoint time-intervals, we have

$$\mathbf{Pr}^{(e(k-1))} \left[\overline{\Gamma_\ell^{(e(k))}}(u) \mid \Gamma_g^{(e(k-1))} \right] = \mathbf{Pr}^{(e(k-1))} \left[\overline{\Gamma_\ell^{(e(k))}}(u) \right] \stackrel{(a)}{\leq} \frac{1}{\log^{11}(n)},$$

where (a) follows Lemma 6.3 (recall that $e(k) - e(k-1) = \tau_{\text{local}}$). $\square$

Proof of Claim 6.10. Recall that $w_j^{(e(k-1))} = u$. Conditioning on whether $\mathcal{E}^{(e(k))}$ happens or not gives us

$$\begin{aligned} &\mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] \\ &= \mathbf{Pr}^{(e(k-1))} \left[\mathcal{E}^{(e(k))} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] \\ &\quad \cdot \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \mathcal{E}^{(e(k))} \right] \\ &+ \mathbf{Pr}^{(e(k-1))} \left[\overline{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] \\ &\quad \cdot \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \overline{\mathcal{E}^{(e(k))}} \right], \end{aligned}$$

and since $\overline{\mathcal{E}^{(e(k))}}$ implies that $\mathbf{1}_{\mathcal{E}^{(e(k))}} = 0$, the above is

$$\begin{aligned} &= \mathbf{Pr}^{(e(k-1))} \left[\mathcal{E}^{(e(k))} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] \\ &\quad \cdot \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \mathcal{E}^{(e(k))} \right] \\ &\leq 1 \cdot \mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \mathcal{E}^{(e(k))} \right] \\ &\stackrel{(a)}{=} \frac{\mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\Gamma_g^{(e(k-1))} \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))}} \right]}{\mathbf{Pr} \left[\Gamma_g^{(e(k-1))} \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))} \right]} \\ &= \frac{\mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L \cap \Gamma_g^{(e(k-1))} \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))}} \right]}{\mathbf{Pr} \left[\Gamma_g^{(e(k-1))} \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))} \right]} \end{aligned}$$

$$\begin{aligned}
&= \frac{\Pr^{(e(k-1))} \left[(Z_j^{(e(k))} \geq L) \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))} \right]}{\Pr \left[\Gamma_g^{(e(k-1))} \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))} \right]} \\
&= \frac{\Pr^{(e(k-1))} \left[Z_j^{(e(k))} \cdot \mathbf{1}_{\Gamma_g^{(e(k-1))} \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))}} \geq L \right]}{\Pr \left[\Gamma_g^{(e(k-1))} \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))} \right]} \\
&\stackrel{(b)}{\leq} \frac{\frac{1}{L} \cdot \mathbf{E}^{(e(k-1))} \left[Z_j^{(e(k))} \cdot \mathbf{1}_{\Gamma_g^{(e(k-1))} \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))}} \right]}{\Pr \left[\Gamma_g^{(e(k-1))} \cap \Gamma_\ell^{(e(k))}(u) \cap \mathcal{E}^{(e(k))} \right]} \\
&\stackrel{(c)}{=} \frac{1}{L} \cdot \mathbf{E} \left[Z_j^{(e(k))} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \mathcal{E}^{(e(k))} \right],
\end{aligned}$$

where (a) and (c) uses the definition of expectations conditional on events $\mathcal{E}$, $\mathbf{E} \left[Z \mid \mathcal{E} \right] = \frac{\mathbf{E}[Z \cdot \mathbf{1}_{\mathcal{E}}]}{\Pr[\mathcal{E}]}$; while (b) follows from applying the so-called “conditional Markov’s inequality” (Exercise 8.2.5 in [26]). Recall that $\mathbf{E}^{(e(k-1))} \left[Z_j^{(e(k))} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \mathcal{E}^{(e(k))} \right]$ is not a number but a random function over $\mathfrak{F}^{(e(k-1))}$ -measurable events. In the following, for any such random function f , let us write $\sup_\omega f$ for the largest value f could attain over its arguments. With this notation, we obtain the bound

$$\begin{aligned}
&\mathbf{E}^{(e(k-1))} \left[\mathbf{1}_{Z_j^{(e(k))} \geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u) \right] \\
&\leq \frac{1}{L} \cdot \sup_\omega \mathbf{E} \left[Z_j^{(e(k))} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \mathcal{E}^{(e(k))} \right] \\
&\stackrel{(a)}{=} \frac{1}{L} \cdot \sup_\omega \mathbf{E} \left[Z_j^{(e(k))} \mid \Gamma_g^{(e(k-1))}, \Gamma_\ell^{(e(k))}(u), \mathcal{E}^{(e(k))}, \mathcal{L}^{(e(k))}(u) \right] \\
&\leq \frac{1}{L} \cdot \sup_\omega \mathbf{E} \left[Z_j^{(e(k))} \mid \Gamma_\ell^{(e(k))}(u), \mathcal{L}^{(e(k))}(u) \right] \\
&\stackrel{\text{Def (6.5)}}{\leq} \frac{1}{L} \cdot \left(L - \varepsilon + \frac{1}{\log^4(n)} \right) = 1 - \frac{\varepsilon}{L} + \frac{1}{L \cdot \log^4(n)},
\end{aligned}$$

where (a) holds, since by definition, when the events $\Gamma_g^{(e(k-1))}$, $\Gamma_\ell^{(e(k))}(u)$, $\mathcal{E}^{(e(k))}$ hold, the event $\mathcal{L}^{(e(k))}(u)$ must also hold. This completes the proof of the second statement and the proof of the claim. $\square$

By repeatedly applying Lemma 6.8 over subsequent epochs, we can complete the analysis of Phase 1 and Phase 2.

LEMMA 6.11 (Phase 1). *Let $\varepsilon \geq 4/\log^4(n)$ and $1 \leq L \leq \log^7(n)$ be an integer. We assume that $x^{(0)}$ has at most $(L - \varepsilon)n$ tokens and choose $t_1 := \tau_{\text{global}} + \frac{\log(n)}{\log \log(n)} \cdot \tau_{\text{local}}$. Then,*

$$\Pr \left[Y^{(t_1)} \leq \frac{n}{\log(n)} \right] \geq 1 - \exp \left(-\frac{\varepsilon}{2L} \cdot \frac{\log(n)}{\log \log(n)} + 8 \cdot \log \log(n) \right) - 2n^{-2}.$$

Proof. Recall that $e(0) = \tau_{\text{global}}$. Hence $t_1 = e(\log(n)/\log \log(n))$. Note that $\tilde{Y}^{(e(0))} \leq Y^{(e(0))} \leq n \cdot \log^7(n)$. Applying Lemma 6.8 for $\ell := \frac{\log(n)}{\log \log(n)}$ epochs gives us

$$\begin{aligned}
\mathbf{E} \left[\tilde{Y}^{(e(\ell))} \right] &= \mathbf{E} \left[\mathbf{E} \left[\tilde{Y}^{(e(\ell))} \mid \mathfrak{F}^{(e(\ell-1))} \right] \right] \leq \left(1 - \frac{\varepsilon}{L} + \frac{2}{L \cdot \log^4(n)} \right) \cdot \mathbf{E} \left[\tilde{Y}^{(e(\ell-1))} \right] \\
&\stackrel{(a)}{\leq} \left(1 - \frac{\varepsilon}{2L} \right) \cdot \mathbf{E} \left[\tilde{Y}^{(e(\ell-1))} \right] \leq e^{-\varepsilon/(2L)} \cdot \mathbf{E} \left[\tilde{Y}^{(e(\ell-1))} \right],
\end{aligned}$$

where (a) used that $\varepsilon \geq 4/\log^4(n)$. By iterating this, it follows

$$\mathbf{E} \left[\tilde{Y}^{(e(\ell))} \right] \leq e^{-\varepsilon \cdot \ell/(2L)} \cdot \tilde{Y}^{(e(0))} \leq e^{-\varepsilon \cdot \ell/(2L)} \cdot n \cdot \log^7(n).$$

Let $\beta := \varepsilon \cdot \ell / (2L) - 8 \cdot \log \log(n)$. By Markov's inequality,

$$(6.11) \quad \Pr \left[\tilde{Y}^{(e(\ell))} \geq \exp \left(-\frac{\varepsilon \cdot \ell}{2L} + \beta \right) \cdot n \cdot \log^7(n) \right] \leq e^{-\beta}.$$

Note that

$$\begin{aligned} \exp \left(-\frac{\varepsilon \cdot \ell}{2L} + \beta \right) \cdot n \cdot \log^7(n) &= \exp \left(-\frac{\varepsilon \cdot \ell}{2L} + \frac{\varepsilon \cdot \ell}{2L} - 8 \cdot \log \log(n) \right) \cdot n \cdot \log^7(n) \\ &= \exp(-8 \cdot \log \log(n)) \cdot n \cdot \log^7(n) = \frac{n}{\log(n)}. \end{aligned}$$

By the definition of $\tilde{Y}^{(e(\ell))}$ and the union bound we get

$$\begin{aligned} \Pr \left[Y^{(e(\ell))} \geq \frac{n}{\log(n)} \right] &\leq \Pr \left[Y^{(e(\ell))} \geq \frac{n}{\log(n)} \mid \bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))} \right] + \Pr \left[\bigcup_{s=0}^{\ell} \overline{\mathcal{E}^{(e(s))}} \right] \\ &\leq \Pr \left[\tilde{Y}^{(e(\ell))} \geq \frac{n}{\log(n)} \right] + \Pr \left[\bigcup_{s=0}^{\ell} \overline{\mathcal{E}^{(e(s))}} \right] \\ &\stackrel{\text{Eq. (6.11)}}{\leq} e^{-\beta} + \sum_{s=0}^{\ell} \Pr \left[\overline{\mathcal{E}^{(e(s))}} \right] \stackrel{(a)}{\leq} e^{-\beta} + 2 \cdot (\ell + 1) \cdot n^{-3}, \end{aligned}$$

where (a) follows from Lemma 6.6. Recalling our choice of β earlier in this proof, and the choice of $\ell = \frac{\log(n)}{\log \log(n)}$, finishes the proof. $\square$

LEMMA 6.12 (Phase 2). *Assume that the load vector $x^{(0)}$ has at most $(1 - \varepsilon) \cdot n$ tokens, where $\frac{1}{2} < \varepsilon < 1$. Then it holds for $t_2 := \tau_{\text{global}} + \frac{4}{-\log(1 - \varepsilon + \frac{2}{\log^4(n)})} \cdot \log(n) \cdot \tau_{\text{local}}$,*

$$\Pr \left[\max_{w \in V} X_w^{(t_2)} \leq 1 \right] \geq 1 - n^{-2}.$$

Proof. The proof of this lemma is similar to the one of the previous lemma, but here we have the special case $L = 1$. By assumption, $Y^{(e(0))} \leq Y^{(0)} \leq n$. Furthermore, $\tilde{Y}^{(e(0))} \leq Y^{(e(0))} \leq n$. Applying Lemma 6.8 with $L = 1$ yields for any epoch $k \geq 1$,

$$\mathbf{E} \left[\tilde{Y}^{(e(k))} \right] = \mathbf{E} \left[\mathbf{E} \left[\tilde{Y}^{(e(k))} \mid \mathfrak{F}^{(e(k-1))} \right] \right] \leq \left(1 - \varepsilon + \frac{2}{\log^4(n)} \right) \cdot \mathbf{E} \left[\tilde{Y}^{(e(k-1))} \right].$$

We now consider $\ell := \frac{4}{-\log(1 - \varepsilon + \frac{2}{\log^4(n)})} \cdot \log(n)$ many epochs.

It follows that

$$\begin{aligned} \mathbf{E} \left[\tilde{Y}^{(e(\ell))} \right] &\leq \left(1 - \varepsilon + \frac{2}{\log^4(n)} \right)^{\ell} \cdot \tilde{Y}^{(e(0))} \\ &\leq \exp \left(\log \left(1 - \varepsilon + \frac{2}{\log^4(n)} \right) \cdot \frac{4}{-\log(1 - \varepsilon + \frac{2}{\log^4(n)})} \cdot \log(n) \right) \cdot n \\ &= \exp(-4 \cdot \log(n)) \cdot n = n^{-3}. \end{aligned}$$

By Markov's inequality, $\Pr \left[\tilde{Y}^{(e(\ell))} \geq 1 \right] \leq n^{-3}$. By the definition of $\tilde{Y}^{(e(\ell))}$ and the union bound,

$$\Pr \left[Y^{(e(\ell))} \geq 1 \right] \leq \Pr \left[Y^{(e(\ell))} \geq 1 \mid \bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))} \right] + \Pr \left[\bigcup_{s=0}^{\ell} \overline{\mathcal{E}^{(e(s))}} \right]$$

$$\begin{aligned}
&\leq \Pr \left[\tilde{Y}^{(e(\ell))} \geq 1 \right] + \sum_{s=0}^{\ell} \Pr \left[\overline{\mathcal{E}^{(e(s))}} \right] \\
&\leq n^{-3} + \sum_{s=0}^{\ell} \Pr \left[\overline{\mathcal{E}^{(e(s))}} \right] \stackrel{(a)}{\leq} n^{-3} + 2 \cdot (\ell + 1) \cdot n^{-3} \leq n^{-2},
\end{aligned}$$

where the last line used that $\ell = O(\log(n))$ and (a) follows from Lemma 6.6,. $\square$

7 From Linear Number of Tokens to Constant Discrepancy

7.1 Reducing Discrepancy to 38 Here we show that once the load vector consists of only $O(n)$ tokens, then after additional $O(\tau_{\text{global}} + \frac{\log(n)}{\log \log(n)} \cdot \tau_{\text{local}})$ rounds the discrepancy is reduced to 38 (see Lemma 7.1). The proof of the result relies on Lemma 5.6 showing that the number of tokens with height at least $\bar{x} + 1$ is at most $16n$.

In this section we derive discrepancy bounds for load vectors where the number of tokens with height above the average load is $O(n)$. In the following lemma we show that the discrepancy can be reduced to 38.

LEMMA 7.1. *Let $t^* := 2 \cdot \tau_{\text{global}} + \frac{6 \log(n)}{\log \log(n)} \cdot \tau_{\text{local}}$. Assume our process applies a $(\tau_{\text{global}}, \tau_{\text{local}})$ -good sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^{\infty}$ to a load vector $x^{(0)}$, which satisfies*

$$\sum_{w \in V} \max\{x_w^{(0)} - \bar{x}, 0\} \leq 16 \cdot n.$$

Then

$$\Pr \left[\text{disc}(X^{(t^*)}) \leq 38 \right] \geq 1 - \exp(-(1/80)) \cdot \log(n) / \log \log(n).$$

Proof. For any $t \geq 0$, define an auxiliary load vector $\tilde{X}_u^{(t)} := \max\{X_u^{(t)} - \lceil \bar{x} \rceil, 0\}$, $u \in V$, that is, we subtract $\lceil \bar{x} \rceil$ tokens from any node (as long as its load is large enough). By assumption, we have $\sum_{w \in V} \tilde{x}_w^{(0)} \leq 16 \cdot n$. Here we can apply Proposition 6.1 with $L = 17$, $\varepsilon := 1/2$ and $t^* := 2 \cdot \tau_{\text{global}} + \frac{6 \log(n)}{\log \log(n)} \cdot \tau_{\text{local}}$ and obtain for $c := 1/70$,

$$\Pr \left[\max_{w \in V} \tilde{X}_w^{(t^*)} \leq 18 \right] \geq 1 - \exp \left(-c \cdot \frac{\log(n)}{\log \log(n)} \right).$$

Applying Theorem 9.10 we get that the same holds for $X^{(t^*)}$, resulting in

$$\Pr \left[\max_{w \in V} X_w^{(t^*)} \leq \lceil \bar{x} \rceil + 18 \right] \geq 1 - \exp \left(-c \cdot \frac{\log(n)}{\log \log(n)} \right).$$

Using a simple symmetry argument (see Theorem 9.7),

$$\Pr \left[\min_{w \in V} X_w^{(t^*)} \geq \lfloor \bar{x} \rfloor - 19 \right] \geq 1 - \exp \left(-c \cdot \frac{\log(n)}{\log \log(n)} \right).$$

A final union bound gives

$$\Pr \left[\text{disc}(X^{(t^*)}) \leq 38 \right] \geq 1 - 2 \exp \left(-c \cdot \frac{\log(n)}{\log \log(n)} \right) \geq 1 - \exp \left(-(1/80) \cdot \frac{\log(n)}{\log \log(n)} \right),$$

which yields the statement. $\square$

7.2 Reducing Discrepancy from 38 to 4

LEMMA 7.2. *Let $x^{(0)}$ be any load vector with $\text{disc}(x^{(0)}) \leq 38$. Then for $t^* := 2 \cdot \tau_{\text{global}} + 6\tau_{\text{local}} \cdot \log(n) / \log \log(n)$ we have,*

$$\Pr \left[\text{disc}(x^{(t^*)}) \leq 4 \right] \geq 1 - \exp(-(1/160) \cdot \log(n) / \log \log(n)).$$

Proof. Assume without loss of generality that the load values are $\{0, 1, 2, \dots, 38\}$. Let us pick $L := \lceil \bar{x} + \frac{1}{2} \rceil$ and $\varepsilon := \frac{1}{2}$; clearly $L \leq 39$ since $\bar{x} \leq 38$. Using Proposition 6.1, we obtain that at round $t^* := 2 \cdot \tau_{\text{global}} + 6\tau_{\text{local}} \cdot \log(n)/\log \log(n)$ the maximum load is at most $L + 1 = \lceil \bar{x} + \frac{1}{2} \rceil + 1$ with probability at least

$$1 - \exp\left(-\frac{\log(n)}{156 \cdot \log \log(n)} + 8 \cdot \log \log(n)\right) - 2 \cdot n^{-2}.$$

Let us now consider the load vector $\tilde{x}^{(0)} := 38 - x^{(0)}$. By Theorem 9.11, $\tilde{x}^{(t)} = 38 - x^{(t)}$ for all $t \geq 1$. Also, $\overline{(\tilde{x})} = 38 - \bar{x}$. Repeating the above argument, but now applied to $\tilde{x}$, yields for any $u \in V$,

$$\tilde{x}_u^{(t^*)} \leq L + 1 = \left\lceil \overline{(\tilde{x})} + \frac{1}{2} \right\rceil + 1,$$

which implies

$$x_u^{(t^*)} \geq 38 - \left\lceil \overline{(\tilde{x})} + \frac{1}{2} \right\rceil - 1 = 37 - \left\lceil 38 - \bar{x} + \frac{1}{2} \right\rceil \stackrel{(a)}{=} 37 - \left(38 + \left\lceil -\bar{x} + \frac{1}{2} \right\rceil\right) \stackrel{(b)}{=} \left\lfloor \bar{x} - \frac{1}{2} \right\rfloor - 1,$$

where (a) used the fact that for any integer k and real z , $\lceil k + z \rceil = k + \lceil z \rceil$ and (b) used the fact that $\lceil -z \rceil = -\lfloor z \rfloor$ for any real z . Hence we can also conclude that the minimum load in $x^{(t^*)}$ is at least $\lfloor \bar{x} - \frac{1}{2} \rfloor - 1$. Since all load values at round t^* are integers in the interval $[\lfloor \bar{x} - \frac{1}{2} \rfloor - 1, \lceil \bar{x} + \frac{1}{2} \rceil + 1]$, which are at most 5 values, the discrepancy is at most 4. Hence a final union bound gives,

$$\Pr \left[\text{disc}(X^{(t^*)}) \leq 4 \right] \geq 1 - 2 \exp\left(-\frac{\log(n)}{156 \cdot \log \log(n)}\right) - 4 \cdot n^{-2} \geq 1 - \exp\left(-\frac{\log(n)}{160 \cdot \log \log(n)}\right). \quad \square$$

7.3 Reducing Discrepancy from 4 to 3 To obtain discrepancy 3, we may not be able to apply Proposition 6.1 directly. For example, if $\bar{x}$ is an integer (or close to an integer), the highest load is $\bar{x} + 2$, and the minimum load is $\bar{x} - 2$, then there is no L and ε such that the application of Proposition 6.1 would result in a reduced discrepancy. To overcome this problem, we first remove a small number of tokens at the two highest levels (they will be called “secondary”), and focus on the movement of the remaining tokens (called “primary”). We can then show, by Proposition 6.1, that most of the primary tokens reduce their height to at most $\bar{x}$. Then we can show that the amount of these primary tokens together with the secondary tokens is smaller than $(L - \varepsilon) \cdot n$ (for $L = 2$ and proper ε), and by applying Lemma 6.12 we obtain that all tokens at height $\bar{x} + 2$ will reduce their height.

LEMMA 7.3. *Consider any initial load vector with $\text{disc}(x^{(0)}) \leq 4$. Then for any $0 < \delta < 1/2$ and $t^* := 2\tau_{\text{global}} + \frac{5}{\delta \log \log(n)} \cdot \log(n) \cdot \tau_{\text{local}}$,*

$$\Pr \left[\text{disc}(X^{(t^*)}) \leq 3 \right] \geq 1 - 2 \cdot \exp\left(-\log^{1-2\delta}(n)\right).$$

Proof. For simplicity, we assume that the load values at time 0 are $\{0, 1, 2, 3, 4\}$ (this can be achieved by reducing the load by the same value at each node).

To show the lemma, we consider two different cases (1 and 2). In Case 1, we assume that $\bar{x} \leq 2$. Ideally, we would like to apply Proposition 6.1 for $L = 2$, but this requires that the total number of tokens is at most $(L - \varepsilon) \cdot n$ for $\varepsilon > \frac{4}{\log^4 n}$, but this may not be satisfied if $\bar{x} \in (2 - \frac{4}{\log^4 n}, 2]$. To overcome this problem, we will first apply the first phase in the proof of Proposition 6.1 to a load vector with slightly fewer tokens. To this end, we first mark $\frac{n}{\log^\delta n}$ tokens at height 3 and 4, and call these tokens “secondary” tokens (all other tokens are called “primary”). We now apply a $(\tau_{\text{global}}, \tau_{\text{local}})$ -good sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$, and consider the load balancing process with a fixed sequence of orientations on the graph w.r.t. the primary tokens only (in which the secondary tokens are removed) versus the process with the same sequence of orientations w.r.t. both types of tokens. The load vector at some time t , which results from the load balancing process w.r.t. the primary tokens only, is denoted by $p^{(t)}$.

Clearly, the number of tokens in $p^{(0)}$ is at most $2n - \frac{n}{\log^\delta n} = (L - \varepsilon) \cdot n$, where $L := 2$ and $\varepsilon := \frac{1}{\log^\delta n}$. Then we can apply Lemma 6.11 for $t_1 := \tau_{\text{global}} + \log(n)/\log \log(n) \cdot \tau_{\text{local}}$, (note that $\varepsilon \geq 4/\log^4(n)$ is satisfied) and conclude that the number of nodes with load at least 3 in $p^{(t_1)}$ satisfies

$$\begin{aligned} \Pr \left[\left| \left\{ u \in V : p_u^{(t_1)} \geq 3 \right\} \right| \leq \frac{n}{\log(n)} \right] &\geq 1 - \exp \left(-\frac{(1/\log^\delta n)}{4} \cdot \frac{\log(n)}{\log \log(n)} + 8 \cdot \log \log(n) \right) - 2n^{-2} \\ &\geq 1 - \exp \left(-\log^{1-2\delta}(n) \right). \end{aligned}$$

According to the second statement of Theorem 9.8, we have for all rounds $t \geq 0$ and $u \in V$,

$$x_u^{(t)} \geq p_u^{(t)}.$$

As there are $n/\log^\delta(n)$ secondary tokens, the number of tokens at height at least 3 in $x^{(t_1)}$ is at most

$$2n/\log(n) + n/\log^\delta(n) \leq 2n/\log^\delta(n) = (1 - \varepsilon) \cdot n,$$

for $\varepsilon := 1 - \frac{2}{\log^\delta(n)}$. Then we consider an additional phase, applied to $x^{(t)}$, $t \geq t_1$, where we consider only these $(1 - \varepsilon) \cdot n$ tokens at height 3 and 4. By Lemma 6.12, by round $t^* := t_1 + \tau_{\text{global}} + \frac{4}{-\log(1 - \varepsilon + \frac{2}{\log^4(n)})} \cdot \log(n) \cdot \tau_{\text{local}}$, all tokens at height 4 are eliminated with probability $1 - n^{-2}$. Hence the total number of rounds is

$$t^* = 2\tau_{\text{global}} + \frac{\log(n)}{\log \log(n)} \cdot \tau_{\text{local}} + \frac{4}{\delta \log \log(n) - 2} \cdot \log(n) \cdot \tau_{\text{local}} \leq 2\tau_{\text{global}} + \frac{5}{\delta \log \log(n)} \cdot \log(n) \cdot \tau_{\text{local}},$$

and by the union bound the success probability for both phases is

$$1 - \exp \left(-\log^{1-2\delta}(n) \right) + n^{-2} \geq 1 - 2 \cdot \exp \left(-\log^{1-2\delta}(n) \right).$$

In Case 2, where $\bar{x} \geq 2$, we consider the flipped load vector with the entries $y_i^{(0)} = 4 - x_i^{(0)}$, and apply the analysis of the first case to this vector. $\square$

8 Proofs of our Main Results

Here we provide a proof of our main Theorem 2.3.

Before proving the theorem, we present a simple lemma which relates $(\tau_{\text{global}}, \tau_{\text{local}})$ -good sequences of matchings to the property of (K, ε) -smoothing.

LEMMA 8.1. *Consider a sequence of $(\tau_{\text{global}}, \tau_{\text{local}})$ -good matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$. Let $K \geq 2n^2$ and $0 < \varepsilon \leq 1$. Then for $t^* := (3 \cdot \log(\frac{K}{\varepsilon})/\log(n)) \cdot \tau_{\text{global}}$,*

$$\Pr \left[\left(\mathbf{M}^{(s)} \right)_{s=1}^{t^*} \text{ is } (K, \varepsilon)\text{-smoothing} \right] \geq 1 - n^{-2}.$$

Proof. We first define $x := 3 \lceil \log_{2n^2}(\frac{K}{\varepsilon}) \rceil$. We will consider x subsequent and disjoint subsequences of matchings, $(\mathbf{M}^{(s)})_{s=(i-1) \cdot \tau_{\text{global}}+1}^{i \cdot \tau_{\text{global}}}$, where $i \in [1, x]$. For any $i \in [1, x]$ we define a random variables Z_i to be zero if the sequence of matchings $(\mathbf{M}^{(s)})_{s=(i-1) \cdot \tau_{\text{global}}+1}^{i \cdot \tau_{\text{global}}}$ is $(1, 1/(2n^2))$ -smoothing and one otherwise. From Definition 2.2 together with Theorem 9.9 we get

$$\Pr[Z_i = 1] \leq n^{-3}.$$

We define $Z := \sum_{i=1}^x Z_i$ and an event $\Psi := \{Z \leq x/2\}$. By linearity of expectation we get that $\mathbf{E}[Z] \leq x/n^3$. From Markov's inequality it follows that $\Pr[Z \geq \frac{x}{2}] \leq \Pr[Z \geq \frac{x}{n}] \leq n^{-2}$ implying that $\Pr[\Psi] \geq 1 - n^{-2}$.

In the remainder of the proof we assume that the event Ψ occurs. This implies that at least $x/2$ matching subsequences of length τ_{global} are $(1, 1/(2n^2))$ -smoothing. Note that the discrepancy is non-increasing over the time. Hence conditioning on the event Ψ , we get that after $x \cdot \tau_{\text{global}}$ rounds the discrepancy is at most

$$K \cdot \left(\frac{1}{2n^2} \right)^{x/2} \leq K \cdot \left(\frac{1}{2n^2} \right)^{\log_{2n^2}(K/\varepsilon)} = \varepsilon.$$

Finally, we have that

$$3 \cdot \left\lceil \log_{2n^2} \left(\frac{K}{\varepsilon} \right) \right\rceil = 3 \cdot \left\lceil \frac{\log(\frac{K}{\varepsilon})}{\log(2n^2)} \right\rceil \stackrel{(\star)}{\leq} 3 \cdot \frac{\log(\frac{K}{\varepsilon})}{\log(n)} = t^\star \cdot \frac{1}{\tau_{\text{global}}},$$

where $(\star)$ used that $K \geq 2n^2$ and $\varepsilon \leq 1$, which implies that the argument of $\lceil \cdot \rceil$ is at least 1. This completes the proof. $\square$

Proof of Theorem 2.3. To ease the notation, we define $\ell := \frac{\log(n)}{\log \log(n)}$. Let $t_0 := 3\tau_{\text{global}} \cdot \log(2Kn)/\log(n)$. First, we define two events

$$\mathcal{G}^\star := \left\{ \left(\mathbf{M}^{(s)} \right)_{s=1}^{t_0} \text{ is } (K, 1/(2n))\text{-smoothing} \right\},$$

and

$$\mathcal{G}_0 := \left\{ \sum_{w \in V} \max\{X_w^{(t_0)} - \bar{x}, 0\} \leq 16 \cdot n \right\}.$$

From Lemma 8.1 it follows that $\Pr[\mathcal{G}^\star] \geq 1 - n^{-2}$ and from Lemma 5.6 it follows that

$$\Pr[\mathcal{G}_0 \mid \mathcal{G}^\star] \geq 1 - 2 \cdot n^{-2}.$$

Let $t_1 := t_0 + 2 \cdot \tau_{\text{global}} + 6\ell \cdot \tau_{\text{local}}$. Here we define an event $\mathcal{G}_1 := \{\text{disc}(X^{(t_1)}) \leq 38\}$ and from Lemma 7.1 it follows that

$$\Pr[\mathcal{G}_1 \mid \mathcal{G}_0] \geq 1 - \exp(-(1/80) \cdot \ell).$$

Let $t_2 := t_1 + 2 \cdot \tau_{\text{global}} + 6\ell \cdot \tau_{\text{local}}$. We define another event $\mathcal{G}_2 := \{\text{disc}(X^{(t_2)}) \leq 4\}$, and from Lemma 7.2 it follows that

$$\Pr[\mathcal{G}_2 \mid \mathcal{G}_1] \geq 1 - \exp(-(1/160) \cdot \ell).$$

Finally, by the definition of conditional probability we obtain

$$\begin{aligned} \Pr[\mathcal{G}_2] &\geq \Pr[\mathcal{G}_2 \mid \mathcal{G}_1] \cdot \Pr[\mathcal{G}_1] \geq \Pr[\mathcal{G}_2 \mid \mathcal{G}_1] \cdot \Pr[\mathcal{G}_1 \mid \mathcal{G}_0] \cdot \Pr[\mathcal{G}_0 \mid \mathcal{G}^\star] \cdot \Pr[\mathcal{G}^\star] \\ &\geq 1 - \exp(-(1/160) \cdot \ell) - \exp(-(1/80) \cdot \ell) - 2 \cdot n^{-2} - n^{-2} \geq 1 - \exp(-(1/200) \cdot \ell). \end{aligned}$$

Note that

$$\begin{aligned} t_2 &:= t_1 + 2 \cdot \tau_{\text{global}} + 6\ell \cdot \tau_{\text{local}} = t_0 + 4 \cdot \tau_{\text{global}} + 12\ell \cdot \tau_{\text{local}} \\ &= \left(\frac{3 \log(2Kn)}{\log(n)} + 4 \right) \cdot \tau_{\text{global}} + 12\ell \cdot \tau_{\text{local}}, \end{aligned}$$

finishing the proof of the first statement.

Let $c > 0$ be any constant and $\tau := t_2 + 2 \cdot \tau_{\text{global}} + \lceil 10/c \rceil \cdot \ell \cdot \tau_{\text{local}}$. Here we define an event $\mathcal{G}_3 := \{\text{disc}(X^{(\tau)}) \leq 3\}$. From Lemma 7.3 with $\delta := c/2$ it follows that

$$\Pr[\mathcal{G}_3 \mid \mathcal{G}_2] \geq 1 - 2 \exp(-\log^{1-c}(n)).$$

By the definition of conditional probability we get,

$$\begin{aligned} \Pr[\mathcal{G}_3] &\geq \Pr[\mathcal{G}_3 \mid \mathcal{G}_2] \cdot \Pr[\mathcal{G}_2] \geq 1 - 2 \exp(-\log^{1-c}(n)) - \Pr[\overline{\mathcal{G}_2}] \\ &\geq 1 - 2 \exp(-\log^{1-c}(n)) - \exp(-(1/200) \cdot \ell) \geq 1 - \exp(-\log^{1-2c}(n)). \end{aligned}$$

Finally,

$$\begin{aligned} \tau &= t_2 + 2 \cdot \tau_{\text{global}} + \frac{\lceil 10/c \rceil \log(n)}{\log \log(n)} \cdot \tau_{\text{local}} \\ &= \left(\frac{3 \log(2Kn)}{\log(n)} + 6 \right) \cdot \tau_{\text{global}} + \frac{(\lceil 10/c \rceil + 12) \log(n)}{\log \log(n)} \cdot \tau_{\text{local}}. \end{aligned}$$

Since c is a constant, this finishes the proof. $\square$

We now provide formal statements for the remarks in Subsection 2.3 regarding application of the main theorem in various concrete settings: balancing circuit and random matchings.

COROLLARY 8.2. *Assuming the premise of Theorem 2.3, for the balancing circuit model we have $\tau = O(\Delta \cdot \log(Kn)/(1 - \lambda(\mathbf{M}^{[1, \Delta]})))$.*

Proof. To simplify notation, we define $\lambda := \lambda(\mathbf{M}^{[1, \Delta]})$. Using standard spectral arguments (e.g., [32, Theorem 1]) yields that the sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^t$ for $t := \frac{4\Delta}{1-\lambda} \cdot \log\left(\frac{Kn}{\varepsilon}\right)$ is (K, ε) -smoothing. Note that this corresponds to the multiplication of t/Δ round matrices.

In the following analysis we will show that the balancing circuit model is also $(\tau_{\text{global}}, \tau_{\text{local}})$ -good with $\tau_{\text{global}} := \frac{7\Delta \cdot \log(n)}{1-\lambda}$ and $\tau_{\text{local}} := \frac{11\Delta \cdot \log \log(n)}{1-\lambda}$. Fix an arbitrary $t \geq 1$. From [34, Lemma 2.4] (restated as Lemma 9.4) it follows for any $u \in V$ and for the choice of τ_{global} as above that we have

$$\left\| \mathbf{M}_{u, \cdot}^{[t, t+\tau_{\text{global}}]} - \frac{\vec{1}}{n} \right\|_2^2 \leq (1-\lambda)^{\log(n^7)/(1-\lambda)} \leq e^{-\log(n^7)} = \frac{1}{n^7}.$$

Now we consider τ_{local} . We fix a node $u \in V$ and an arbitrary $t \geq 1$. Then we have

$$\left\| \mathbf{M}_{u, \cdot}^{[t, t+\tau_{\text{local}}]} \right\|_2^2 - \frac{1}{n} \stackrel{\text{Obs. 9.12}}{=} \left\| \mathbf{M}_{u, \cdot}^{[1, \tau_{\text{local}}]} - \frac{\vec{1}}{n} \right\|_2^2 \stackrel{\text{Lem. 9.4}}{\leq} \frac{1}{\log^{11} n}.$$

From this we get

$$\left\| \mathbf{M}_{u, \cdot}^{[t, t+\tau_{\text{local}}]} \right\|_2^2 \leq \frac{1}{\log^{11} n} + \frac{1}{n} \leq \frac{1}{\log^{10} n}. \quad \square$$

COROLLARY 8.3. *Assuming the premise of Theorem 2.3, for the random matching model we have $\tau = O(\log(Kn)/(p_{\min} \cdot \Delta \cdot (1 - \lambda(\mathbf{P}))))$.*

Proof. It suffices to show that the random matching model is $(\tau_{\text{global}}, \tau_{\text{local}})$ -good with $\tau_{\text{global}} := \frac{14 \cdot \log(n)}{p_{\min} \cdot \Delta \cdot (1 - \lambda(\mathbf{P}))}$ and $\tau_{\text{local}} := \frac{22 \cdot \log \log(n)}{p_{\min} \cdot \Delta \cdot (1 - \lambda(\mathbf{P}))}$.

Fix an arbitrary $t \geq 1$. From [34, Corollary 2.7] (see Lemma 9.5) it follows for any $u \in V$ and the choice of τ_{global} as above (and sufficiently large n) that we have

$$\Pr \left[\left\| \mathbf{M}_{u, \cdot}^{[t, t+\tau_{\text{global}}]} - \frac{\vec{1}}{n} \right\|_2^2 \leq \frac{1}{n^7} \right] \geq 1 - \frac{1}{n^7}.$$

Applying the union bound over all nodes $u \in V$ implies that

$$\Pr \left[\bigcap_{u \in V} \left\| \mathbf{M}_{u, \cdot}^{[t, t+\tau_{\text{global}}]} - \frac{\vec{1}}{n} \right\|_2^2 \leq \frac{1}{n^7} \right] \geq 1 - \frac{1}{n^6}.$$

Now we consider τ_{local} . We fix a node $u \in V$ and an arbitrary $t \geq 1$. Then we have

$$\begin{aligned} \Pr \left[\left\| \mathbf{M}_{u, \cdot}^{[t, t+\tau_{\text{local}}]} \right\|_2^2 \leq \frac{1}{\log^{10} n} \right] &\geq \Pr \left[\left\| \mathbf{M}_{u, \cdot}^{[t, t+\tau_{\text{local}}]} \right\|_2^2 \leq \frac{1}{\log^{11} n} + \frac{1}{n} \right] \\ &\stackrel{\text{Obs. 9.12}}{=} \Pr \left[\left\| \mathbf{M}_{u, \cdot}^{[1, \tau_{\text{local}}]} - \frac{\vec{1}}{n} \right\|_2^2 \leq \frac{1}{\log^{11} n} \right] \stackrel{\text{Lem. 9.5}}{\geq} 1 - \frac{1}{\log^{11} n}. \end{aligned} \quad \square$$

References

- [1] H. AKBARI, P. BERENBRINK, AND T. SAUERWALD, *A simple approach for adapting continuous load balancing processes to discrete settings*, Distributed Comput., 29 (2016), pp. 143–161.
- [2] D. ALISTARH, G. NADIRADZE, AND A. SABOUR, *Dynamic averaging load balancing on cycles*, Algorithmica, 84 (2022), pp. 1007–1029, <https://doi.org/10.1007/S00453-021-00905-9>, <https://doi.org/10.1007/s00453-021-00905-9>.
- [3] A. ANAGNOSTOPOULOS, A. KIRSCH, AND E. UPFAL, *Load balancing in arbitrary network topologies with stochastic adversarial input*, SIAM Journal on Computing, 34 (2005), pp. 616–639, <https://doi.org/10.1137/S0097539703437831>, <https://doi.org/10.1137/S0097539703437831>.
- [4] L. BECCHETTI, A. CLEMENTI, E. NATALE, F. PASQUALE, AND L. TREVISAN, *Find your place: Simple distributed algorithms for community detection*, in Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, P. N. Klein, ed., SIAM, 2017, pp. 940–959, <https://doi.org/10.1137/1.9781611974782.59>, <https://doi.org/10.1137/1.9781611974782.59>.
- [5] P. BERENBRINK, C. COOPER, T. FRIEDETZKY, T. FRIEDRICH, AND T. SAUERWALD, *Randomized diffusion for indivisible loads*, J. Comput. Syst. Sci., 81 (2015), pp. 159–185, <https://doi.org/10.1016/J.JCSS.2014.04.027>, <https://doi.org/10.1016/j.jcss.2014.04.027>.
- [6] P. BERENBRINK, C. COOPER, C. GAVA, D. K. MARZAGÃO, F. MALLMANN-TRENN, T. RADZIK, AND N. RIVERA, *Distributed averaging in opinion dynamics*, in Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing, PODC, R. Oshman, A. Nolin, M. M. Halldórsson, and A. Balliu, eds., ACM, 2023, pp. 211–221, <https://doi.org/10.1145/3583668.3594593>, <https://doi.org/10.1145/3583668.3594593>.
- [7] P. BERENBRINK, T. FRIEDETZKY, AND L. A. GOLDBERG, *The natural work-stealing algorithm is stable*, SIAM J. Comput., 32 (2003), pp. 1260–1279.
- [8] P. BERENBRINK, T. FRIEDETZKY, D. KAASER, AND P. KLING, *Tight & simple load balancing*, in 2019 IEEE International Parallel and Distributed Processing Symposium, IPDPS, IEEE, 2019, pp. 718–726, <https://doi.org/10.1109/IPDPS.2019.00080>, <https://doi.org/10.1109/IPDPS.2019.00080>.
- [9] P. BERENBRINK, T. FRIEDETZKY, AND R. A. MARTIN, *Dynamic diffusion load balancing*, in Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings, 2005.
- [10] P. BERENBRINK, L. HINTZE, H. HOSSEINPOUR, D. KAASER, AND M. RAU, *Dynamic averaging load balancing on arbitrary graphs*, in 50th International Colloquium on Automata, Languages, and Programming, ICALP, K. Etessami, U. Feige, and G. Puppis, eds., vol. 261 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, pp. 18:1–18:18, <https://doi.org/10.4230/LIPICS.ICALP.2023.18>, <https://doi.org/10.4230/LIPICS.ICALP.2023.18>.
- [11] P. BERENBRINK, M. HOEFER, AND T. SAUERWALD, *Distributed selfish load balancing on networks*, ACM Trans. Algorithms, 11 (2014), pp. 2:1–2:29, <https://doi.org/10.1145/2629671>, <https://doi.org/10.1145/2629671>.
- [12] P. BERENBRINK, R. KLASING, A. KOSOWSKI, F. MALLMANN-TRENN, AND P. UZNANSKI, *Improved analysis of deterministic load-balancing schemes*, ACM Trans. Algorithms, 15 (2019), pp. 10:1–10:22, <https://doi.org/10.1145/3282435>, <https://doi.org/10.1145/3282435>.
- [13] D. P. BERTSEKAS AND J. N. TSITSIKLIS, *Parallel and Distributed Computation: Numerical Methods*, Prentice Hall, 1989.
- [14] J. E. BOILLAT, *Load balancing and poisson equation in a graph*, Concurr. Pract. Exp., 2 (1990), pp. 289–314, <https://doi.org/10.1002/CPE.4330020403>, <https://doi.org/10.1002/cpe.4330020403>.

- [15] S. P. BOYD, A. GHOSH, B. PRABHAKAR, AND D. SHAH, *Randomized gossip algorithms*, IEEE Trans. Inf. Theory, 52 (2006), pp. 2508–2530, <https://doi.org/10.1109/TIT.2006.874516>, <https://doi.org/10.1109/TIT.2006.874516>.
- [16] J. N. COOPER, B. DOERR, T. FRIEDRICH, AND J. SPENCER, *Deterministic random walks on regular trees*, Random Struct. Algorithms, 37 (2010), pp. 353–366, <https://doi.org/10.1002/RSA.20314>, <https://doi.org/10.1002/rsa.20314>.
- [17] J. N. COOPER AND J. SPENCER, *Simulating a random walk with constant error*, Comb. Probab. Comput., 15 (2006), pp. 815–822, <https://doi.org/10.1017/S0963548306007565>, <https://doi.org/10.1017/S0963548306007565>.
- [18] G. CYBENKO, *Dynamic load balancing for distributed memory multiprocessors*, J. Parallel Distributed Comput., 7 (1989), pp. 279–301, [https://doi.org/10.1016/0743-7315\(89\)90021-X](https://doi.org/10.1016/0743-7315(89)90021-X), [https://doi.org/10.1016/0743-7315\(89\)90021-X](https://doi.org/10.1016/0743-7315(89)90021-X).
- [19] R. DIEKMANN, A. FROMMER, AND B. MONIEN, *Efficient schemes for nearest neighbor load balancing*, Parallel Comput., 25 (1999), pp. 789–812, [https://doi.org/10.1016/S0167-8191\(99\)00018-6](https://doi.org/10.1016/S0167-8191(99)00018-6), [https://doi.org/10.1016/S0167-8191\(99\)00018-6](https://doi.org/10.1016/S0167-8191(99)00018-6).
- [20] D. P. DUBHASHI AND D. RANJAN, *Balls and bins: A study in negative dependence*, Random Struct. Algorithms, 13 (1998), pp. 99–124.
- [21] T. FRIEDRICH, M. GAIRING, AND T. SAUERWALD, *Quasirandom load balancing*, SIAM J. Comput., 41 (2012), pp. 747–771, <https://doi.org/10.1137/100799216>, <https://doi.org/10.1137/100799216>.
- [22] T. FRIEDRICH AND T. SAUERWALD, *Near-perfect load balancing by randomized rounding*, in Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC, M. Mitzenmacher, ed., ACM, 2009, pp. 121–130, <https://doi.org/10.1145/1536414.1536433>, <https://doi.org/10.1145/1536414.1536433>.
- [23] B. GHOSH AND S. MUTHUKRISHNAN, *Dynamic load balancing by random matchings*, J. Comput. Syst. Sci., 53 (1996), pp. 357–370, <https://doi.org/10.1006/JCSS.1996.0075>, <https://doi.org/10.1006/jcss.1996.0075>.
- [24] V. GURUSWAMI, *Rapidly mixing markov chains: A comparison of techniques (a survey)*, 2016, <https://arxiv.org/abs/1603.01512>, <https://arxiv.org/abs/1603.01512>.
- [25] D. KEMPE, A. DOBRA, AND J. GEHRKE, *Gossip-based computation of aggregate information*, in 44th Symposium on Foundations of Computer Science, FOCS, IEEE Computer Society, 2003, pp. 482–491, <https://doi.org/10.1109/SFCS.2003.1238221>, <https://doi.org/10.1109/SFCS.2003.1238221>.
- [26] A. KLENKE, *Probability theory - a comprehensive course*, Universitext, Springer, 2008.
- [27] L. LOVÁSZ AND P. WINKLER, *Mixing of random walks and other diffusions on a graph*, in Surveys in Combinatorics, P. Rowlinson, ed., vol. 218 of London Mathematical Society Lecture Note Series, Cambridge University Press, 1995, pp. 119–154, <https://doi.org/10.1017/CBO9780511662096.007>.
- [28] M. MAVRONICOLAS AND T. SAUERWALD, *A randomized, $o(\log w)$ -depth 2 smoothing network*, in SPAA 2009: Proceedings of the 21st Annual ACM Symposium on Parallelism in Algorithms and Architectures, F. M. auf der Heide and M. A. Bender, eds., ACM, 2009, pp. 178–187, <https://doi.org/10.1145/1583991.1584043>, <https://doi.org/10.1145/1583991.1584043>.
- [29] M. MAVRONICOLAS AND T. SAUERWALD, *The impact of randomization in smoothing networks*, Distributed Comput., 22 (2010), pp. 381–411, <https://doi.org/10.1007/S00446-009-0087-3>, <https://doi.org/10.1007/s00446-009-0087-3>.
- [30] R. MOTWANI AND P. RAGHAVAN, *Randomized Algorithms*, Cambridge University Press, 1995, <https://doi.org/10.1017/CBO9780511814075>.

- [31] S. MUTHUKRISHNAN, B. GHOSH, AND M. H. SCHULTZ, *First- and second-order diffusive methods for rapid, coarse, distributed load balancing*, Theory Comput. Syst., 31 (1998), pp. 331–354, <https://doi.org/10.1007/S002240000092>, <https://doi.org/10.1007/s002240000092>.
- [32] Y. RABANI, A. SINCLAIR, AND R. WANKA, *Local divergence of markov chains and the analysis of iterative load balancing schemes*, in 39th Annual Symposium on Foundations of Computer Science, FOCS, IEEE Computer Society, 1998, pp. 694–705, <https://doi.org/10.1109/SFCS.1998.743520>, <https://doi.org/10.1109/SFCS.1998.743520>.
- [33] J. S. ROSENTHAL, *A First Look at Rigorous Probability Theory*, World Scientific, second edition ed., 2006, <https://doi.org/10.1142/6300>.
- [34] T. SAUERWALD AND H. SUN, *Tight bounds for randomized load balancing on arbitrary network topologies*, in 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS, IEEE Computer Society, 2012, pp. 341–350, <https://doi.org/10.1109/FOCS.2012.86>, <https://doi.org/10.1109/FOCS.2012.86>. Full version at <https://arxiv.org/abs/1201.2715>.
- [35] R. SUBRAMANIAN AND I. D. SCHERSON, *An analysis of diffusive load-balancing*, in Proceedings of the 6th Annual ACM Symposium on Parallel Algorithms and Architectures, SPAA, L. Snyder and C. E. Leiserson, eds., ACM, 1994, pp. 220–225, <https://doi.org/10.1145/181014.181361>, <https://doi.org/10.1145/181014.181361>.

Appendix

9 Assorted Tools

9.1 Basic Properties of the Height-Sensitive Process

LEMMA 9.1. *Consider any pair of rounds $t_1 < t_2$ and let $(\mathbf{M}^{(s)})_{s=t_1+1}^{t_2}$ be an arbitrary but fixed sequence of matchings. For any token $i \in \mathcal{T}$ it holds that (1) $H_i^{(t_1)} \geq H_i^{(t_2)}$, and (2) $\Pr \left[W_i^{(t_2)} = v \mid w_i^{(t_1)} = u \right] = \mathbf{M}_{u,v}^{[t_1+1, t_2]}$.*

Proof. The height of a token can only change in the moving step of a round t . This only happens if the token is on one endpoint of an edge $[u : v] \in \mathbf{M}^{(t)}$ with $x_u^{(t-1)} \geq x_v^{(t-1)}$. Then tokens on node u at height $x_v^{(t-1)} + \lceil \frac{x_u^{(t-1)} - x_v^{(t-1)}}{2} \rceil + i$ for some $i > 0$ will decrease their height by $\lceil \frac{x_u^{(t-1)} - x_v^{(t-1)}}{2} \rceil$ and move to node v . No other token on u or v will change its height during the moving step, and hence the first statement follows.

Similar to the proof of [34, Lemma 4.1], we prove the second statement by induction over $t \in [t_1 + 1, t_2]$, that is, for all $u, v \in V$, we have $\Pr \left[W_i^{(t)} = v \mid w_i^{(t_1)} = u \right] = \mathbf{M}_{u,v}^{[t_1+1, t]}$. For $t = t_1 + 1$, $\mathbf{M}^{[t_1+1, t]}$ is the identity matrix, which means that the induction base holds. For the induction hypothesis, consider $\Pr \left[W_i^{(t)} = v \mid w_i^{(t_1)} = u \right]$. If v is not part of a matching in round t , then

$$\Pr \left[W_i^{(t)} = v \mid w_i^{(t_1)} = u \right] = \Pr \left[W_i^{(t-1)} = v \mid w_i^{(t_1)} = u \right] \stackrel{(\star)}{=} \mathbf{M}_{u,v}^{[t_1+1, t-1]} = \mathbf{M}_{u,v}^{[t_1, t]},$$

where $(\star)$ used the induction hypothesis. If v is matched with a node w in round t , then by definition of the height-sensitive process the token i has a probability of $\frac{1}{2}$ of reaching v from either v or w , and hence

$$\begin{aligned} \Pr \left[W_i^{(t)} = v \mid w_i^{(t_1)} = u \right] &= \Pr \left[W_i^{(t-1)} = v \mid w_i^{(t_1)} = u \right] \cdot \frac{1}{2} + \Pr \left[W_i^{(t-1)} = w \mid w_i^{(t_1)} = u \right] \cdot \frac{1}{2} \\ &\stackrel{(\star)}{=} \mathbf{M}_{u,v}^{[t_1, t]} \cdot \frac{1}{2} + \mathbf{M}_{u,w}^{[t_1, t]} \cdot \frac{1}{2} \\ &= \mathbf{M}_{u,v}^{[t_1, t]} \cdot \mathbf{M}_{v,v}^{(t+1)} + \mathbf{M}_{u,w}^{[t_1, t]} \cdot \mathbf{M}_{w,v}^{(t+1)} \\ &= \mathbf{M}_{u,v}^{[t_1, t+1]}, \end{aligned}$$

where $(\star)$ used the induction hypothesis. □

9.2 Tail Bounds For completeness we state the following tail bound that we use in our proofs.

THEOREM 9.2 (cf. [30, page 92, Theorem 4.16, Azuma's Inequality]). *Let $X_0, X_1, \dots$ be a martingale sequence such that for each k , $|X_k - X_{k-1}| \leq c_k$, where c_k may depend on k . Then, for all $t \geq 0$ and any $\lambda > 0$,*

$$\Pr[|X_t - X_0| \geq \lambda] \leq 2 \cdot \exp\left(-\frac{\lambda^2}{2 \sum_{k=1}^t c_k^2}\right).$$

LEMMA 9.3 (cf. [33, Proposition 13.2.6]). *Let X and Y be random variables, and let $\mathfrak{F}$ be a sub- σ -algebra. Suppose that $\mathbf{E}[X]$ and $\mathbf{E}[XY]$ are finite, and furthermore that X is $\mathfrak{F}$ -measurable. Then with probability 1,*

$$\mathbf{E}[X \cdot Y \mid \mathfrak{F}] = X \cdot \mathbf{E}[Y \mid \mathfrak{F}].$$

9.3 Basic Results from [34] In this appendix we list basic results from [34] that we use in our analysis.

LEMMA 9.4 (cf. [34, Lemma 2.4]). *Consider the balancing circuit model with sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$. Let $\lambda := \lambda(\mathbf{M}^{[1,\Delta]})$. Then for any node $u \in V$ it holds*

$$\left\| \mathbf{M}_{u,\cdot}^{[1,t,\Delta]} - \frac{\vec{1}}{n} \right\|_2^2 \leq \lambda^t.$$

LEMMA 9.5 (cf. [34, Corollary 2.7]). *Consider the random matching model with sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$. Then for any node $u \in V$ and any round $t \geq 1$ we have*

$$\Pr \left[\left\| \mathbf{M}_{u,\cdot}^{[1,t]} - \frac{\vec{1}}{n} \right\|_2^2 \leq e^{-\frac{p_{\min} \cdot \Delta}{2} \cdot (1 - \lambda(\mathbf{P})) \cdot t} \right] \geq 1 - e^{-\frac{p_{\min} \cdot \Delta}{2} \cdot (1 - \lambda(\mathbf{P})) \cdot t},$$

We remark that in contrast to [34, Corollary 2.7], we have dropped the constraint that $p_{\min} = \Omega(\frac{1}{\Delta})$ as it is easy to see that the proof works without this constraint.

THEOREM 9.6 (cf. [34, Theorem 2.9]). *Let G be any graph with maximum degree Δ and consider the random matching model. Then with probability at least $1 - n^{-1}$ the sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^t$ is $(K, 1/(2n))$ -smoothing for*

$$t := \frac{8}{\Delta \cdot p_{\min}} \cdot \frac{1}{1 - \lambda(\mathbf{P})} \cdot \log(4Kn^2).$$

Next observation bounds the tail of the lower gap by that of upper gap.

OBSERVATION 9.7 (cf. [34, Observation 2.12]). *Assume $\text{disc}(x^{(0)}) = K$. Fix a sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$. Then for arbitrary positive integers α and t it holds that*

$$\max_{\substack{y \in \mathbb{Z}^n: \\ \text{disc}(y) \leq K}} \left\{ \Pr \left[X_{\min}^{(t)} \leq \lfloor \bar{x} \rfloor - \alpha \mid x^{(0)} = y \right] \right\} \leq \max_{\substack{y \in \mathbb{Z}^n: \\ \text{disc}(y) \leq K}} \left\{ \Pr \left[X_{\max}^{(t)} \geq \lfloor \bar{x} \rfloor + \alpha \mid x^{(0)} = y \right] \right\}.$$

Occasionally, we may add/subtract the same number of tokens to/from each node.

OBSERVATION 9.8 (cf. [34, Observation 2.11]). *Fix a sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$. Consider two executions of the discrete load balancing protocol with the same matchings and the same random choices for the excess tokens but with different initial load vectors $x^{(0)}$ and $\tilde{x}^{(0)}$. Then the following two statements hold:*

1. *If $\tilde{x}^{(0)} = x^{(0)} + \alpha \cdot \mathbf{1}$ for some $\alpha \in \mathbb{Z}$, then $\tilde{X}^{(t)} = X^{(t)} + \alpha \cdot \mathbf{1}$ for all $t \geq 1$.*
2. *If $x_u^{(0)} \leq \tilde{x}_u^{(0)}$ for all $u \in V$, then $X_u^{(t)} \leq \tilde{X}_u^{(t)}$ for all $u \in V$ and $t \geq 1$.*

The following observation shows an important relationship between Definitions 2.1 and 2.2

OBSERVATION 9.9 (cf. [34, Lemma 2.2]). *Assume the sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^t$ satisfies for all $u \in V$ that $\left\| \mathbf{M}_{u,\cdot}^{[1,t]} - \frac{\vec{1}}{n} \right\|_2^2 \leq \left(\frac{\varepsilon}{2K \cdot n} \right)^2$. Then this sequence is (K, ε) -smoothing.*

We will omit the proof, since the result follows immediately by noting that $\max_{v \in V} \left\| \mathbf{M}_{u,v}^{[1,t]} - \frac{1}{n} \right\| \leq \left\| \mathbf{M}_{u,\cdot}^{[1,t]} - \frac{\vec{1}}{n} \right\|_2$ and then applying the third statement of [34, Lemma 2.2].

9.4 Other Tools In this section we collect a few other tools and results that are frequently used in our analysis.

OBSERVATION 9.10. *Fix a sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$. Consider two executions of the discrete load balancing protocol with the same matchings and the same random choices for the excess tokens' movements, but with different initial load vectors $x^{(0)}$ and $\tilde{x}^{(0)}$ such that for some $\alpha \in \mathbb{N}$ and for all $u \in V$ it holds*

$$\tilde{x}_u^{(0)} := \max \left\{ x_u^{(0)} - \alpha, 0 \right\}.$$

Then for any round $t \geq 0$ it holds

$$X_u^{(t)} \leq \tilde{X}_u^{(t)} + \alpha.$$

It then directly follows that the height of any token in load vector $X^{(t)}$ is at most the maximum height of any token in load vector $\tilde{X}^{(t)}$ plus α .

Proof. We consider two auxiliary executions of discrete load balancing with initial load vectors $\hat{x}^{(0)}$ and $\tilde{x}^{(0)}$ in which for each node $u \in V$ we have

$$\hat{x}_u^{(0)} := \max \{ x_u^{(0)}, \alpha \} \quad \text{and} \quad \tilde{x}_u^{(0)} := \hat{x}_u^{(0)} - \alpha.$$

We run the discrete load balancing for these executions with the same matching $\mathbf{M}^{(t)}$ and the same random choices for the excess tokens in each round $t \in \mathbb{N}$ but with different initial load vectors constructed as above. From the second statement of Theorem 9.8, it follows that for any round $t \geq 0$ and each node $u \in V$ we have $X_u^{(t)} \leq \hat{X}_u^{(t)}$. Similarly, from the first statement of Theorem 9.8 it follows that $\tilde{X}_u^{(t)} = \hat{X}_u^{(t)} - \alpha$. A combination of these two implies that

$$X_u^{(t)} \leq \tilde{X}_u^{(t)} + \alpha,$$

for any node $u \in V$ and any round $t \geq 0$. □

The next observation uses the concept of orientation, which is defined in (5.5).

OBSERVATION 9.11. *Fix a sequence of matchings $(\mathbf{M}^{(s)})_{s=1}^\infty$. Consider two executions of the discrete load balancing protocol with the same matchings but with different initial load vectors $x^{(0)}$ and $\tilde{x}^{(0)}$, where $\tilde{x}^{(0)} := K \cdot \vec{1} - x^{(0)}$ for some $K \in \mathbb{N}_0$, and flipped orientations, i.e., $\tilde{\Phi}_{u,v}^{(t)} = -\Phi_{u,v}^{(t)}$. Then for any round $t \geq 1$,*

$$\tilde{X}^{(t)} = K \cdot \vec{1} - X^{(t)}.$$

Proof. The proof is by induction over $t \in \mathbb{N}_0$. The base case $t = 0$ holds by assumption. Now assuming $\tilde{x}^{(t-1)} = K \cdot \vec{1} - X^{(t-1)}$ for some $t-1 \in \mathbb{N}_0$, we will prove that the same equation holds for round t . If a node u is not matched in round t then

$$\tilde{X}_u^{(t)} = \tilde{X}_u^{(t-1)} = K - X_u^{(t-1)} = K - X_u^{(t)}.$$

Now assume the edge $[u : v]$ is in the matching $\mathbf{M}^{(t)}$. Then we have

$$X_u^{(t)} = \frac{X_u^{(t-1)} + X_v^{(t-1)}}{2} + \frac{1}{2} \cdot \text{Odd} \left(X_u^{(t-1)} + X_v^{(t-1)} \right) \cdot \Phi_{u,v}^{(t)}$$

and

$$X_v^{(t)} = \frac{X_u^{(t-1)} + X_v^{(t-1)}}{2} + \frac{1}{2} \cdot \text{Odd} \left(X_u^{(t-1)} + X_v^{(t-1)} \right) \cdot \Phi_{v,u}^{(t)}.$$

Furthermore, we have

$$\begin{aligned} \tilde{X}_u^{(t)} &= \frac{\tilde{X}_u^{(t-1)} + \tilde{X}_v^{(t-1)}}{2} + \frac{1}{2} \cdot \text{Odd} \left(\tilde{X}_u^{(t-1)} + \tilde{X}_v^{(t-1)} \right) \cdot \tilde{\Phi}_{u,v}^{(t)} \\ &= \frac{2K - \left(X_u^{(t-1)} + X_v^{(t-1)} \right)}{2} + \frac{1}{2} \cdot \text{Odd} \left(2K - \left(X_u^{(t-1)} + X_v^{(t-1)} \right) \right) \cdot \tilde{\Phi}_{u,v}^{(t)} \end{aligned}$$

$$\begin{aligned}
&= \frac{2K - (X_u^{(t-1)} + X_v^{(t-1)})}{2} + \frac{1}{2} \cdot \text{Odd}(X_u^{(t-1)} + X_v^{(t-1)}) \cdot \tilde{\Phi}_{u,v}^{(t)} \\
&= K - \frac{(X_u^{(t-1)} + X_v^{(t-1)})}{2} - \frac{1}{2} \cdot \text{Odd}(X_u^{(t-1)} + X_v^{(t-1)}) \cdot \Phi_{u,v}^{(t)} \\
&= K - X_u^{(t)}.
\end{aligned}$$

The case for v is analogous to the case for u . Hence for each node $u \in V$, we get $\tilde{X}_u^{(t)} = K - X_u^{(t)}$ and consequently, $\tilde{X}^{(t)} = K \cdot \vec{1} - X^{(t)}$. This finishes the induction step and completes the proof. $\square$

OBSERVATION 9.12. Let $\mathbf{M}$ be an $n \times n$ doubly stochastic matrix. Then for any node $u \in V$,

$$\left\| \mathbf{M}_{u,\cdot} - \frac{\vec{1}}{n} \right\|_2^2 = \|\mathbf{M}_{u,\cdot}\|_2^2 - \frac{1}{n}.$$

Proof. We calculate

$$\begin{aligned}
\left\| \mathbf{M}_{u,\cdot} - \frac{\vec{1}}{n} \right\|_2^2 &= \sum_{v \in V} \left(\mathbf{M}_{u,v} - \frac{1}{n} \right)^2 = \sum_{v \in V} \left((\mathbf{M}_{u,v})^2 - \frac{2}{n} \cdot \mathbf{M}_{u,v} + \frac{1}{n^2} \right) \\
&= \sum_{v \in V} (\mathbf{M}_{u,v})^2 - \frac{2}{n} \cdot \sum_{v \in V} \mathbf{M}_{u,v} + \sum_{v \in V} \frac{1}{n^2} = \|\mathbf{M}_{u,\cdot}\|_2^2 - \frac{2}{n} + \frac{1}{n}.
\end{aligned}$$

$\square$

OBSERVATION 9.13. Let $\mathbf{M}$ be any doubly stochastic matrix and $(a_k)_{k \in V}$ be any stochastic vector. Then

$$\sum_{w \in V} \left(\sum_{k \in V} a_k \cdot \mathbf{M}_{w,k} \right)^2 = \sum_{w \in V} \left(\sum_{k \in V} a_k \cdot \mathbf{M}_{w,k} - \frac{1}{n} \right)^2 + \frac{1}{n}.$$

Proof.

$$\begin{aligned}
\sum_{w \in V} \left(\sum_{k \in V} a_k \cdot \mathbf{M}_{w,k} \right)^2 &= \sum_{w \in V} \left(\sum_{k \in V} \left(a_k - \frac{1}{n} \right) \cdot \mathbf{M}_{w,k} + \frac{1}{n} \right)^2 \\
&= \sum_{w \in V} \left(\sum_{k \in V} \left(a_k - \frac{1}{n} \right) \cdot \mathbf{M}_{w,k} \right)^2 + \frac{2}{n} \cdot \sum_{w \in V} \sum_{k \in V} \left(a_k - \frac{1}{n} \right) \cdot \mathbf{M}_{w,k} + \sum_{w \in V} \frac{1}{n^2} \\
&= \sum_{w \in V} \left(\sum_{k \in V} \left(a_k - \frac{1}{n} \right) \cdot \mathbf{M}_{w,k} \right)^2 + \frac{2}{n} \cdot \sum_{k \in V} \left(a_k - \frac{1}{n} \right) \cdot \sum_{w \in V} \mathbf{M}_{w,k} + \frac{1}{n} \\
&\stackrel{(a)}{=} \sum_{w \in V} \left(\sum_{k \in V} a_k \cdot \mathbf{M}_{w,k} - \frac{1}{n} \right)^2 + \frac{1}{n},
\end{aligned}$$

where (a) holds since $\mathbf{M}_{w,\cdot}$ is stochastic, from $\sum_{k \in V} a_k = 1$ and $|V| = n$ we get $\sum_{k \in V} (a_k - \frac{1}{n}) = 0$. $\square$

The next observation is a well-known statement and in fact, it is implied by Equation (5.6).

OBSERVATION 9.14. For each node $w \in V$ the expression $\left\| \mathbf{M}_{w,\cdot}^{[1,t]} - \frac{\vec{1}}{n} \right\|_2^2$ is non-increasing over t .