Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

Benefit of Random-Local Updates in Networks

(Load Balancing)

Dissertation zur Erlangung des akademischen Grades
Dr. rer. nat.

Betreut von: Prof. Dr. Petra Berenbrink

eingereicht von

Hamed Hosseinpour

Hamburg 2026

An der Fakultät für Mathematik, Informatik und Naturwissenschaften
Fachbereich Informatik
Universität Hamburg

Gutachterin 1: Prof. Dr. Petra Berenberink,
Gutachter 2 : Prof. Dr. Soeren Laue

Datum der Disputation : 23.01.2026

Hamed Hosseinpour, 7335036
Vogt-Kölln-Straße 30, 22527 Hamburg
hamed.hosseinpour@uni-hamburg.de

Abstract

This thesis studies the discrete load balancing problem on graphs in both *static* and *dynamic* settings. In discrete load balancing, a graph $G = (V, E)$ is given, and each node initially holds an integer number of load items (tokens). At each step, nodes apply the same balancing rule, and load items cannot be divided. In the static setting, no new tokens are added, whereas in the dynamic setting, a fixed number of tokens are generated and distributed uniformly at random to the nodes at the start of each round. The goal is to make the loads across all nodes approximately equal. We focus on two main models: the *matching* model, where each node interacts with at most one neighbor per round, and the *diffusion* model, where each node interacts with all its neighbors in each round.

We propose and analyze simple *local* balancing rules, which use *randomization*. In the matching model, two matched nodes take the average of their loads. If the sum of their loads is odd, the node that receives the one excess token is selected at random. In the diffusion model, each node spreads its loads as evenly as possible among its neighbors and itself. Any extra tokens are distributed randomly and without replacement between itself and its neighbors.

In part two of this thesis, we investigate dynamic load balancing on graphs using matchings, focusing on three popular models: the *balancing circuit*, *random matching*, and *(asynchronous) single random edge* models. We provide upper bounds on the *discrepancy*, defined as the maximum difference in load between any two nodes. Furthermore, we establish a lower bound on the discrepancy in the balancing circuit model, showing that the upper bound is tight up to an additional $\mathrm{O}(\sqrt{\log(n)})$ factor.

In part three, we study static load balancing using matchings. We provide a general framework that includes the three models mentioned above. We show that the discrepancy reaches 3 after a sufficiently large number of rounds, which depends on the underlying graph. We develop a novel technique here, which may be of independent interest.

In part four, we first look at static load balancing on d-regular graphs in the diffusion model. We provide upper bounds on the discrepancy, which are $\mathrm{O}(\sqrt{d \cdot \log(n)} + \log(n))$, after sufficient time. We develop a simple method that helps us prove a tail concentration bound for the discrepancy, which could be of independent interest. Using our methods from part two, we extend the result to the dynamic setting. This is the first work considering the dynamic setting of discrete diffusion.

In part five, we tackle the *Token Distribution problem* on matchings. If the load difference between two matched nodes is at least one, one load item is moved from the higher-loaded node to the other. We prove that, after sufficiently many rounds, the discrepancy falls below the graph's diameter for any connected graph. Moreover, we derive a lower bound on the discrepancy.

Zusammenfassung

Diese Arbeit untersucht das diskrete Lastenausgleichsproblem auf Graphen sowohl in *statischer* als auch in *dynamischer* Umgebung. Beim diskreten Lastenausgleich ist ein Graph $G = (V, E)$ gegeben, und jeder Knoten hält anfangs eine ganzzahlige Anzahl von Lastobjekten (Token). In jedem Schritt wenden die Knoten dieselbe Ausgleichsregel an, und Lastobjekte können nicht geteilt werden. In der statischen Umgebung werden keine neuen Token hinzugefügt, während in der dynamischen Umgebung eine feste Anzahl von Token erzeugt und zu Beginn jeder Runde gleichmäßig zufällig auf die Knoten verteilt wird. Ziel ist es, die Lasten über alle Knoten hinweg annähernd gleich zu machen. Wir konzentrieren uns auf zwei Hauptmodelle: das *Matching*-Modell, bei dem jeder Knoten pro Runde mit höchstens einem Nachbarn interagiert, und das *Diffusions*-Modell, bei dem jeder Knoten in jeder Runde mit allen seinen Nachbarn interagiert. Wir schlagen einfache *lokale* Ausgleichsregeln vor und analysieren diese, die *Randomisierung* nutzen. Im Matching-Modell nehmen zwei gematchte Knoten den Durchschnitt ihrer Lasten. Ist die Summe ihrer Lasten ungerade, wird der Knoten, der das eine überschüssige Token erhält, zufällig ausgewählt. Im Diffusions-Modell verteilt jeder Knoten seine Lasten so gleichmäßig wie möglich auf seine Nachbarn und sich selbst. Etwaige überschüssige Token werden zufällig und ohne Zurücklegen zwischen ihm und seinen Nachbarn verteilt.

Im zweiten Teil dieser Arbeit untersuchen wir den dynamischen Lastenausgleich auf Graphen unter Verwendung von Matchings, wobei wir uns auf drei populäre Modelle konzentrieren: das *Balancing Circuit*-Modell, das *Random Matching*-Modell und das *(asynchrone) Single Random Edge*-Modell. Wir geben obere Schranken für die *Diskrepanz* an, definiert als die maximale Differenz der Last zwischen zwei beliebigen Knoten. Darüber hinaus zeigen wir eine untere Schranke für die Diskrepanz im Balancing-Circuit-Modell, was beweist, dass die obere Schranke bis auf einen zusätzlichen Faktor von $\mathrm{O}(\sqrt{\log(n)})$ scharf ist.

Im dritten Teil untersuchen wir den statischen Lastenausgleich mittels Matchings. Wir stellen einen allgemeinen Rahmen vor, der die oben genannten drei Modelle umfasst. Wir zeigen, dass die Diskrepanz nach einer hinreichend großen Anzahl von Runden, abhängig vom zugrunde liegenden Graphen, den Wert 3 erreicht. Dabei entwickeln wir eine neue Technik, die von eigenem Interesse sein könnte.

Im vierten Teil betrachten wir zunächst den statischen Lastenausgleich auf d-regulären Graphen im Diffusionsmodell. Wir geben obere Schranken für die Diskrepanz an, die nach ausreichend vielen Runden $\mathrm{O}(\sqrt{d \cdot \log(n)} + \log(n))$ betragen. Wir entwickeln eine einfache Methode, die uns hilft, eine Tail-Konzentrationsschranke für die Diskrepanz zu beweisen, die von eigenem Interesse sein könnte. Unter Verwendung der Methoden aus Teil zwei erweitern wir dieses Ergebnis auf die dynamische Umgebung. Dies ist die erste Arbeit, die die dynamische Umgebung der diskreten Diffusion betrachtet.

Im fünften Teil behandeln wir das *Token-Distribution-Problem* auf Matchings. Wenn der Lastunterschied zwischen zwei gematchten Knoten mindestens eins beträgt, wird ein Lastobjekt vom höherbelasteten Knoten auf den anderen übertragen. Wir beweisen, dass nach ausreichend vielen Runden die Diskrepanz für jeden zusammenhängenden Graphen unter den Durchmesser des Graphen fällt. Zusätzlich geben wir eine dazu untere Schranke für die Diskrepanz an.

Acknowledgments

I thank my advisor, Prof. Dr. Petra Berenbrink, for her guidance and support. I am grateful to my co-authors for fruitful collaborations. I thank my colleagues and friends at the university for a stimulating environment. I also thank the wider research community for many enjoyable discussions. I thank my family, especially my parents, for their constant encouragement. Finally, I thank Germany and its people for providing a welcoming and supportive environment for my studies.

Contents

List of Figures

List of Tables

1 Introduction

Efficiently distributing indivisible tasks (tokens) across the processors of a networked system is a central problem in distributed computing with immediate applications to datacenter scheduling, cloud platforms [76], large-scale computations [87], numerical simulations [73], finite element analysis [79], and GPU farms used for training modern machine learning models. The discrete nature of jobs, the locality of communication, and the potential for continuous arrivals create a rich design space of local load-balancing rules. The practical performance and theoretical guarantees of these rules depend sensitively on two orthogonal design choices: (1) the interaction model, which specifies which neighbors exchange load at each step, and (2) the operating setting, which distinguishes between a static scenario (a one-time redistribution) and a dynamic one (with a continuous flow of new work). This thesis focuses on two canonical families of local rules - matching-style exchanges, where each node interacts with at most one neighbor per round, and diffusion-style exchanges, where each node interacts with all neighbors every round and analyzes their behavior in both static and dynamic regimes. By treating these axes together we aim both to isolate the fundamental trade-offs between convergence speed and communication locality, and to provide modular techniques that yield tight bounds for discrete and dynamic settings.

We formalize the setting on an undirected graph $G = (V, E)$ with $n = |V|$ nodes. Each node $v \in V$ holds an integer number of tokens, and we denote the load vector at (logical) time t by $X(t) = (X_v(t))_{v \in V}$. The principal measure of imbalance that we study throughout is the *discrepancy* which captures the worst gap between any two nodes, i.e., $\mathrm{disc}(X(t)) := \max_{v \in V} X_v(t) - \min_{v \in V} X_v(t)$.

Two operating settings are considered. In the static setting a fixed multiset of tokens is present at time $t = 0$ and no new tokens arrive; here the questions concern how quickly and with what migration cost local iterative rules can reduce discrepancy starting from an arbitrary initial configuration. In the dynamic setting, a fixed number of new tokens is generated in each round and distributed to nodes (in our main model uniformly at random); here the objective becomes maintaining a small instantaneous discrepancy in steady state while controlling the long-run migration and communication volume induced by continual arrivals.

The two interaction models we analyze capture different practical constraints and trade-offs. In matching-style models the system proceeds by repeatedly selecting a collection of disjoint edges (a matching) and letting each matched pair perform a local balancing operation; this limits the number of simultaneous transfers incident on any node and is appropriate when pairwise channels or connection limits constrain per-round communication. In the concrete matching rule we use, two matched nodes replace their loads by the integer average of the pair, and when the pair sum is odd a single extra token is assigned uniformly at random to one endpoint. In diffusion-style models every node exchanges with all neighbors in each round and attempts to split its tokens as evenly as possible among itself and its neighbors; leftover tokens due to indivisibility are allocated uniformly at random without replacement across the local neighborhood. Diffusion generally promotes faster global smoothing because each node acts on all incident edges in parallel, but it can induce larger per-round migration and requires careful rounding control when loads are discrete.

On the matching side we analyze three natural scheduling variants that arise in practice and in prior theoretical work. The balancing-circuit model applies a fixed, deterministic sequence of matchings periodically; the random-matching model samples a fresh uniform random matching each round; and the asynchronous single-random-edge model activates exactly one uniformly random edge in each logical round and lets its endpoints balance, a model that abstracts asynchronous or extremely sparse-communication environments. On the diffusion side we concentrate our strongest results on d-regular graphs, both because regularity simplifies several probabilistic symmetries and because many network topologies of interest admit locally approximately-regular neighborhoods; where possible we also explain how the techniques extend beyond regular graphs and how parameter choices must be adapted in nonregular settings.

Our analytic goals are to bound discrepancy (both after a finite number of rounds in the static case and in steady state for the dynamic case), to quantify convergence time from worst-case initial states, and to measure migration and communication costs induced by balancing. Intuitively, matching-based rules limit per-round migration at the expense of slower mixing, whereas diffusion-style rules mix faster but create more simultaneous transfers; a recurring theme of the thesis is making this intuition precise, showing how graph structure (spectral gap, conductance, diameter) and the chosen schedule affect the attainable guarantees in discrete and dynamic environments.

From a technical standpoint, discrete and dynamic processes bring several complications that do not appear in the continuous (divisible-load) idealization. Continuous diffusion admits a clean linear-algebraic treatment: the load vector evolves under a linear operator closely related to the graph Laplacian, and spectral gap bounds the rate at which the system approaches the uniform distribution. When loads are indivisible, however, rounding errors accumulate and may create persistent local imbalances; random tie-breaking reduces bias but requires careful potential-based accounting to bound long-term effects. Asynchrony and concurrent activations (especially in single-edge or partially overlapping matchings) complicate coupling arguments and demand analyses that tolerate overlapping moves and delayed effects. In the dynamic setting continuous injection of new tokens produces a steady-state that balances smoothing against fresh imbalance; establishing nontrivial steady-state discrepancy bounds therefore requires drift arguments that quantify how quickly the protocol removes injected imbalance relative to the arrival rate. Finally, for matching-based processes, a precise combinatorial tracking of how discrete averaging propagates through the sequence of matchings is necessary to prove sharp static bounds; developing such a tracking invariant is one of the technical contributions of this work.

In the remainder of this section, we first provide an overview of the problems and results in Subsection 1.1, followed by a discussion of publications in Subsection 1.2. Subsection 1.3 reviews related work relevant to our problems. Subsection 1.4 introduces some basic stochastic elements. Finally, Subsection 1.5 presents the notation used across all models.

1.1 Overview of Problems and Results

Part Two. In this part, we study dynamic load balancing on general graphs. We consider infinite-time and dynamic processes, where in each step new load items are assigned to randomly chosen nodes. A matching is selected, and the load is averaged over the edges of that matching. We analyze the discrete case, where load items are indivisible. Moreover, our results also carry over to the continuous case, where load items can be split arbitrarily. Regarding the choice of matchings, we consider three different models: random matchings of linear size, random matchings consisting of individual edges, and deterministic sequences of matchings that cover the entire graph. We bound the *discrepancy*, defined as the difference between the maximum and minimum load.

Our results cover a broad range of graph classes and, to the best of our knowledge, represent the first analysis of discrete and dynamic averaging load balancing processes. Table 1 shows our simplified upper bounds for some specific graphs, and Table 2 provides the corresponding lower bounds. In fact, we show that in the deterministic sequence model, the gap between the upper and lower bounds is $\mathrm{O}(\sqrt{\log(n)})$, which is nearly tight.

Part Three. Here, we investigate discrete static load balancing via matchings on arbitrary graphs. Initially, each node holds a certain number of tokens. The objective is to redistribute the tokens so that eventually each node has approximately the same number of tokens. We present results for a general class of simple local balancing schemes covering the three matching models mentioned above. In each round, the process averages the tokens of any two matched nodes. If the sum of their tokens is odd, the node that receives the excess token is selected at random.

As our main result, we show that, with high probability, our discrete balancing scheme reaches a discrepancy of 3 in a number of rounds that matches the spectral bound for continuous load balancing with fractional load. The result improves and tightens a long line of previous work, both by achieving a small constant discrepancy (rather than a non-explicit, large constant) and by applying to arbitrary, rather than only regular graphs. It also demonstrates that, in the general model we consider, discrete load balancing is no harder than continuous load balancing. A summary of our results, along with comparisons to related work, can be found in Table 3.

Part Four. In this part, we investigate the *vertex-based diffusion process*. A d-regular graph $G(V, E)$ is given, and each node initially holds some integer load items. In each round, every node distributes its load as evenly as possible among its neighbors and itself. If it is not possible to do so without splitting some tokens, the node distributes its excess tokens randomly without replacement among all its neighbors and itself.

We first prove a new bound of $\mathrm{O}(d \log(n))$ and then improve it to $\mathrm{O}(\sqrt{d \log(n)} + \log(n))$ for the discrepancy in d-regular graphs. Our bound improves upon existing results for the regime $d = \omega(\sqrt[4]{\log(n)})$. Moreover, we propose the first bound on the discrepancy in the dynamic setting of the discrete vertex-based diffusion process. A summary of our results is provided in Table 4.

Part Five. In this part, we study the token distribution problem via matchings on arbitrary graphs. Initially, each node holds a certain number of load items. In each step, a matching is given, and for two matched nodes, if the load difference is at least one, then only one token is moved from the more loaded node to the other. Previous work has shown that the discrepancy converges to the diameter of a given regular graph, but no concrete bound on the runtime has been provided ([49] and [3]).

We show that the discrepancy drops below the diameter (a) while providing an explicit bound on the number of rounds, and (b) in a way that applies to arbitrary connected graphs. Our analysis is simple and relies solely on a well-known quadratic potential function. A summary of our results, together with a comparison to the most closely related work, is presented in Table 5. Furthermore, we establish a lower bound on the discrepancy for regular graphs, which matches the bound given in [49].

General Outline. The focus is on Parts Two to Five. Part Six collects important and useful related results that serve as the foundation for our calculations. Each main part follows a consistent structure:

- **Introduction:** briefly presents the problem, main results, and an outline.
- **Model and Definitions:** introduces notations, the model, and key definitions.
- **Detailed Analysis:** states the main theorems, supporting propositions, and technical lemmas.
- **Bounds:** provides results for specific graph classes, models, or rounds, depending on the problem.
- **Conclusion:** summarizes the results and highlights open problems.

1.2 Publications

Most of our work is based on peer-reviewed papers, and the results presented here advance the state of the art. The following list provides an overview of each paper, including all co-authors, as well as the conferences and journals where the work was published or submitted. In addition, we include two of our previously published papers, which are not discussed in this thesis; interested readers are referred to the original publications for further information.

	Part Two: ***Dynamic Averaging Load Balancing on Arbitrary Graphs***
Authors:	Petra Berenbrink, Lukas Hintze, Hamed Hosseinpour, Dominik Kaaser, Malin Rau
Conference:	ICALP 2023 - The EATCS International Colloquium on Automata, Languages and Programming, Germany, July 10-14 2023, [24]
Journal:	ACM Transactions on Algorithms (TALG), **Submitted**
Full Version:	arXiv CoRR, [23]

	Part Three: ***(Almost) Perfect Discrete Iterative Load Balancing***
Authors:	Petra Berenbrink, Robert Elsaesser, Tom Friedetzky, Hamed Hosseinpour, Dominik Kaaser, Peter Kling, Thomas Sauerwald
Conference:	SODA 2026 - The ACM-SIAM Symposium on Discrete Algorithms, Vancouver, Canada, January 11-14 2026 [17]
Full Version:	arXiv CoRR, [16]

	Fast Consensus via Unconstrained Undecided State Dynamics
Authors:	Gregor Bankhamer, Petra Berenbrink, Felix Biermeier, Robert Elsaesser, Hamed Hosseinpour, Dominik Kaaser, Peter Kling
Conference:	SODA 2022 - The ACM-SIAM Symposium on Discrete Algorithms, Virtual / Alexandria, VA, January 9-12 2022, [10]
Full Version:	arXiv CoRR, [9]

	Population Protocols for Exact Plurality Consensus
Authors:	Gregor Bankhamer, Petra Berenbrink, Felix Biermeier, Robert Elsaesser, Hamed Hosseinpour, Dominik Kaaser, Peter Kling
Conference:	PODC 2022 - The ACM Symposium on Principles of Distributed Computing, Italy, July 25-29 2022, [12]
Full Version:	arXiv CoRR, [11]

The contribution of the author of this thesis to the results of each part is described at the end (**??**).

1.3 Related Works

We provide an overview of existing results related to the problems considered in the main four parts of this thesis. There is a vast body of literature on iterative load balancing schemes on graphs, where nodes balance (or average) their load only with neighbors. A common distinction is between *diffusion* load balancing, where nodes balance with all neighbors simultaneously, and the *matching* (or *dimension-exchange*) model, where edges used for balancing form a matching. In the latter, each node participates in at most one balancing action per step, which simplifies the analysis.

This overview focuses on theoretical results for discrete load balancing; for continuous load balancing, see, e.g., [39, 60]. Related literature also studies selfish load balancing, where tokens act as independent agents; see [45] for a survey and [2, 27, 57] for recent results. Recall that the discrepancy (at round t) is the maximum difference between the loads of any pair of nodes at round t. Let $\widetilde{\tau}_{\mathrm{S}}(K) := \log(Kn)/(1-\lambda)$, where λ is the second-largest eigenvalue of the diffusion matrix and K is the initial discrepancy. Moreover, Δ is the maximum degree (of the graph) in the random matching model and the sequence length in the balancing circuit model.

Continuous Model. Bertsekas and Tsitsiklis [29], Boillat [31], Boyd et al. [32], and Cybenko [38] pioneered the use of Markov chains for analyzing diffusion-based load balancing schemes in the continuous model (where the load can be arbitrarily divided). Note that [32] considers the asynchronous process, in which a single edge is chosen uniformly at random in each round, whereas the other papers study the synchronous variant. For these processes, it is known that the discrepancy can be reduced from K to ℓ within $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K/\ell))$ rounds in the diffusion model, the balancing circuit model, and the random matching model (see, e.g., [82]). All these upper bounds are essentially tight, which follows from the connection between the spectral gap of the graph and the mixing times of Markov chains [54].

Discrete Models. One of the earliest rigorous analyses of the discrete setting is due to Muthukrishnan et al. [78]. Their algorithm computes, for each edge, the flow of load that would occur in the continuous model and rounds that value down to obtain the number of tokens to be sent. For the diffusion model, they show that after $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds the discrepancy is at most $\mathrm{O}(\Delta n/(1-\lambda))$. Ghosh and Muthukrishnan [51] show similar results for the matching model. Rabani et al. [82] present a more refined analysis based on Markov chains. They introduce the so-called *local divergence*, which aggregates the sum of load differences over all edges in all rounds. The authors prove that the local divergence provides an upper bound on the maximum deviation between the continuous and discrete versions of a protocol. In the same spirit as the previous papers, they assume that an excess token is always kept at the current node. Their technique applies to a large class of processes, including matching and diffusion models. Among other results, they provide bounds for general graphs, showing that the discrepancy is at most $\mathrm{O}(\Delta\widetilde{\tau}_{\mathrm{S}}(K))$ after $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds in both the balancing circuit model and the diffusion model.

Berenbrink et al. [20] analyzed the asynchronous process for the complete graph. In each round, a pair of nodes is selected uniformly at random and completely balances their loads up to a rounding error of ± 1. They prove that after $\mathrm{O}(n\log(Kn))$ 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.

Friedrich et al. [46] introduce a quasi-random version that rounds up or down deterministically so that the accumulated rounding errors on each edge are minimized. For torus graphs, they show a constant discrepancy, while for hypercubes the discrepancy is $\mathrm{O}(\log^{3/2}(n))$. Akbari et al. [4] analyze a similar framework for load balancing, achieving a discrepancy of $\mathrm{O}(d)$ for d-regular graphs, but their algorithm requires additional memory to track previous decisions.

In [47], the authors present several results for a randomized protocol with rounding in the matching model. For complete graphs, their results show a discrepancy of $\mathrm{O}(n\sqrt{\log(n)})$ after $\Theta(\log(Kn))$ steps. For expanders, a more tailored analysis yields a constant discrepancy in $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K)(\log\log(n))^3)$ rounds in the random matching model. Later, [14] extended some of these results to the diffusion model. The bounds they provide depend on the spectral gap and a refined measure of local divergence. For constant-degree expanders and torus graphs, they obtain exponential improvements in the discrepancy bounds, and for hypercubes, polynomial improvements. Subsequently, Sun and Sauerwald [85] extended these results and showed that the discrepancy is bounded by $\mathrm{O}(d^2\sqrt{\log(n)})$ after $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds for the vertex-based discrete diffusion model.

Sauerwald and Sun [85] also studied discrete load balancing on matchings for *arbitrary* graphs and proved that a discrepancy of $\mathrm{O}(\log^{\epsilon} n)$ for an arbitrarily small constant $\epsilon > 0$ can be achieved in $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds in both the random matching and balancing circuit models. For regular graphs, the authors show constant discrepancy for both models, though for the latter this requires $\Delta = \mathrm{O}(1)$. They show that the number of rounds needed to reach constant discrepancy is, with high probability, $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$.

Berenbrink et al. [19] propose a very simple potential-function technique to analyze discrete and continuous diffusion load balancing.

The authors of [33] study load balancing via matchings assuming random placement of load items. The initial load distribution is sampled from exponentially concentrated distributions (including uniform, binomial, geometric, and Poisson). They show that in this setting, convergence time is smaller than in the worst case. Regardless of the graph's topology, the discrepancy decreases by a factor of $\sqrt[4]{t}$ within t synchronous rounds. Their approach of using concentration inequalities to bound the discrepancy (in terms of the squared 2-norm of the columns of the matrices underlying the mixing process) strongly influenced our approach in part Two.

Dynamic Models. There are far fewer results for the dynamic setting, where new load enters the system over time. In [6], the authors study a model similar to our asynchronous model. In each step, one load item is allocated to a chosen node. The chosen node then selects a random neighbor, and the two nodes balance their loads by averaging them (continuous model). The authors show that the expected discrepancy is bounded by $\mathrm{O}(\sqrt{n}\log(n))$, as well as a lower bound on the square of the discrepancy of $\Omega(n)$. Anagnostopoulos et al. [7] consider load balancing via matchings in a dynamic model where the load is distributed by an adversary. They show that the system is stable for sufficiently limited adversaries and provide upper bounds on the maximum load for more restricted adversaries. Berenbrink et al. [22] consider discrete dynamic diffusion load balancing on arbitrary graphs. In each step, up to n load items are generated on arbitrary nodes (the allocation is determined by an adversary). Then, the nodes balance their loads with each neighbor, and finally, one load item is deleted from every non-empty node. They show that the system is stable, meaning that the total load remains bounded over time (as a function of n alone and independently of t).

In the *graphical balanced allocations* setting, the initial allocation of a load item is constrained to a randomly chosen edge of a graph, and load items cannot be moved after allocation (in contrast to our setting). For d-regular graphs, Peres et al. [81] show that for the greedy algorithm, which allocates a load item to the less-loaded node of each edge (with the edge distribution being uniform), the discrepancy is $\mathrm{O}(\log(n)d/\alpha)$ with high probability, where α is the edge expansion of the graph. They further generalize this result to distributions over arbitrary subsets of nodes. Bansal and Feldheim [13] present a non-greedy algorithm using limited non-local information that achieves a discrepancy of $\mathrm{O}((d/k)\log^4(n)\log(\log(n)))$ for k-edge-connected d-regular graphs, as well as a lower bound for the graphical balanced allocation setting, showing that the discrepancy is $\Omega(d/k + \log(n))$ with constant probability at any given time for any allocation strategy.

Balls into Bins. In the *Balls-into-Bins* setting, at each step m balls are distributed among n nodes. In the *One-choice* process, where each ball is placed on a uniformly random node, the discrepancy is $\mathrm{O}(\log(n)/\log\log(n))$ for $m = n$ and $\mathrm{O}(\sqrt{(m/n)\log(n)})$ for $m > n$ [75]. In the *d-Choices* process, each ball samples d nodes uniformly at random, and the least-loaded node receives the ball. In [58], the authors show that the maximum load deviation from the average (the upper gap) is $\log_d \log(n) + \mathrm{O}(1)$ for $m = n$. Later, Berenbrink et al. [15] show that for the 2-Choice protocol and $m \gg n$, the upper gap is $\log_2 \log(n) + \mathrm{O}(1)$. Mitzenmacher et al. [74] introduce the *Memory Process*, which maintains a one-slot memory storing a bin. In each step, a bin is sampled; if its load is smaller than that of the bin in memory, it replaces the stored bin. The ball is then placed into the bin currently in memory. They show that the upper gap is $\mathrm{O}(\log\log(n))$ for $m = n$. Los et al. [66] show the same bound for $m \gg n$ and extend the results to other classes of protocols using memory.

Token Distribution A related line of research focuses on *token distribution*, where nodes transfer individual tokens to neighboring nodes with smaller loads rather than fully balancing their loads. Early works address both static settings [50, 55, 80] and dynamic ones [8]. In [49], the process is analyzed under two models: the *Single-Port* model, where each node sends one token per round, and the *Multiple-Port* model, where a node may send one token to each neighbor if their load difference exceeds $2(d + 1)$. In the single-port model and for any d-regular graph G, the discrepancy becomes $\mathrm{O}(d\log(n)/\alpha)$ after $\mathrm{O}(dK/\alpha + d^2\log(n)/\alpha^2)$ steps with high probability, where α denotes the *edge expansion* and K the initial imbalance. These bounds are shown to be asymptotically tight.

Further studies extend these results to specific topologies and dynamic settings. Gehrke et al. [48] show that on a ring, nodes achieve perfect balance within $4\mathrm{OPT}(b) + n$ steps in which $\mathrm{OPT}(b)$ is the time taken by the optimal centralized algorithm to balance b completely. Aiello et al. [3] prove that in networks whose active edges form a μ-expander, balancing to an additive $\mathrm{O}(K\log(n)/\mu)$ is possible in $\mathrm{O}(K\log(nK)/\mu)$ steps, with a matching lower bound of $\Omega(K/\mu)$.

There are many results in related models. Mavronicolas and Sauerwald [70, 71] 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 $\mathrm{O}(\log(n))$ rounds.[1] 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 is inspired by the Rotor-Router model [36, 37], where each node evenly distributes tokens to its neighbors, usually in a deterministic manner. Berenbrink, Klasing, Kosowski, Mallmann-Trenn, and Uznanski [26] consider rotor-router-inspired diffusion-type algorithms. Their results hold only for d-regular graphs, and their best algorithm achieves a discrepancy of d. There are also so-called "selfish" load balancing models [25], where tokens act selfishly when deciding whether to migrate to a neighboring node. Several works also study models with dynamic load generation and consumption [7, 18, 21], resulting in discrepancy bounds that are either not independent of n or that establish only stability.

[1] Note that the dimension-exchange communication scheme can be regarded as a particular instance of the balancing circuit model. When we write log, we mean $\log_e$ unless specified otherwise.

1.4 Stochastic Ingredients

In this section, we briefly introduce several well-known stochastic elements that form the foundation of our analysis. We begin with a general setup. The processes we study are *stochastic processes.*

Definition 1.1 (Stochastic Process (Section 7.1 in [75]))**.** *A stochastic process* $\mathbf{X} = \{X(t) : t \in \mathbb{N}\}$ *is a collection of random variables indexed by* t*, which typically represents time. The process* $\mathbf{X}$ *models how the value of a random variable evolves over time.*

We refer to $X(t)$ as the *state* of the process at time t. In what follows, we use X_t interchangeably with $X(t)$. If $X(t)$ takes values from a countably infinite set, then $\mathbf{X}$ is a discrete-space process. If it takes values from a finite set, the process is finite. When the index set t is countably infinite, $\mathbf{X}$ is a discrete-time process.

We focus on a special class of discrete-time, discrete-space stochastic processes $X_0, X_1, \ldots$ in which the value of X_t depends only on X_{t-1} and not on the sequence of preceding states.

Definition 1.2 (Markov Chain (Definition 7.1 in [75]))**.** *A discrete-time stochastic process* $X_0, X_1, \ldots$ *is a* Markov chain *if*

$$\mathbf{Pr}[X_t = a_t \mid X_{t-1} = a_{t-1}, X_{t-2} = a_{t-2}, \ldots, X_0 = a_0] = \mathbf{Pr}[X_t = a_t \mid X_{t-1} = a_{t-1}].$$

This property states that the state X_t depends only on the immediately preceding state X_{t-1}, regardless of how the process arrived there. This is known as the *Markov property* or the *memoryless property.* Note that X_t is generally not independent of all previous steps; any dependence on the past is entirely captured by X_{t-1}.

We define the *transition probability* as

$$P_{i,j}^{(t)} := \mathbf{Pr}[X_t = j \mid X_{t-1} = i]. \tag{1}$$

If $P^{(t)}$ is identical for all $t \in \mathbb{N}$, the chain is called *time-homogeneous*; otherwise, it is *time-inhomogeneous.* A time-homogeneous Markov chain is completely described by its *transition matrix* $P = (P_{i,j})_{i,j \in [1,\infty]}$.

Markov chains naturally arise as models of random movement, for example in *random walks* on graphs. Let $G = (V, E)$ be a finite, undirected, and connected graph.

Definition 1.3 (Random Walk (Section 7.1 in [75]))**.** *A random walk on* G *is a Markov chain in which a single particle moves between the nodes of* G*. If the particle is at node* i *at step* $t-1$*, the probability that it moves to node* j *is defined as in Equation (1).*

In our analysis, the runtime of some algorithms depends on the underlying graph and its transition matrix. For any $n \times n$ real symmetric matrix P, let $\lambda_1(P) \geq \lambda_2(P) \geq \ldots \geq \lambda_n(P)$ denote its eigenvalues, and define

$$\lambda(P) := \max\{|\lambda_2(P)|, |\lambda_n(P)|\}.$$

For non-symmetric matrices, let $\widehat{P} = P \cdot P^T$ and define $\lambda(P) = \max\{|\lambda_2(\widehat{P})|, |\lambda_n(\widehat{P})|\}$. Consider a particle at an arbitrary node of G at round 0, using P as its transition matrix. After $O(\log(n)/(1-\lambda(P)))$ steps, the particle is at any node with probability $\Theta(1/n)$.

The term *with high probability (w.h.p.)* denotes probability at least $1 - n^{-\Omega(1)}$. Since our processes are stochastic, we frequently need to bound error probabilities. The following lemma provides a fundamental tool: the probability that at least one of several events occurs is at most the sum of their individual probabilities.

Lemma 1.1 (Union Bound (Lemma 1.2 in [75]))**.** *For any finite or countably infinite sequence of events* $E_1, E_2, \ldots,$

$$\mathbf{Pr}\left[\bigcup_{i\geq 1} E_i\right] \leq \sum_{i\geq 1} \mathbf{Pr}[E_i].$$

We often require tail concentration results for random variables or their convex combinations to show that they are tightly concentrated around their expectation. A common tool is the *Chernoff bound*, which is closely related to *McDiarmid's inequality*.

Lemma 1.2 (Chernoff Bounds (Theorems 4.4 and 4.5 in [75]))**.** *Let* $X_1, \ldots, X_n$ *be independent Bernoulli trials with* $\mathbf{Pr}[X_i = 1] = p_i$*, and let* $X = \sum_{i=1}^{n} X_i$ *with* $\mu = \mathbf{E}[X]$*. Then:*

1. *For any* $\delta > 0$,
$$\mathbf{Pr}[X \geq (1+\delta)\mu] \leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu}.$$

2. *For any* $0 < \delta < 1$,
$$\mathbf{Pr}[X \leq (1-\delta)\mu] \leq \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{\mu}.$$

In many cases, the random variables in our processes are dependent, so Chernoff bounds cannot always be applied. When this occurs, we use the *Azuma–Hoeffding inequality* or derive custom bounds.

Definition 1.4 (Martingale (Definition 13.1 in [75]))**.** *A sequence of random variables* $Z_0, Z_1, \ldots$ *is a* martingale *with respect to a sequence* $X_0, X_1, \ldots$ *if, for all* $n \geq 0$*:*

1. Z_n *is a function of* $X_0, \ldots, X_n$*;*

2. $\mathbf{E}[|Z_n|] < \infty$*;*

3. $\mathbf{E}[Z_{n+1} \mid X_0, \ldots, X_n] = Z_n$*.*

Lemma 1.3 (Azuma–Hoeffding Inequality (Theorem 13.4 in [75]))**.** *Let* $X_0, \ldots, X_n$ *be a martingale such that* $|X_k - X_{k-1}| \leq c_k$*. Then for all* $t \geq 1$ *and any* $\lambda > 0$,

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

We may also exploit the property of *negative association* to obtain Chernoff-like tail bounds.

Definition 1.5 (Negative Association (Definition 1 in [42]))**.** *A vector* $(X_i)_{i\in[n]}$ *of random variables is* negatively associated *if, for any disjoint index subsets* $\Gamma, R \subseteq [n]$ *and any non-decreasing function* f,

$$\mathbf{E}[\,f(X_\ell, \ell \in \Gamma) \mid X_r = x_r, r \in R\,]$$

is non-increasing in each x_r*.*

To establish negative association in the diffusion protocol, we may use the following result.

Lemma 1.4 (Zero–One Lemma (Lemma 8 in [42]))**.** *If* $X_1, \ldots, X_n$ *are zero–one random variables such that* $\sum_i X_i = 1$*, then* $X_1, \ldots, X_n$ *satisfy the negative association property.*

1.5 Notation and Preliminaries

In this subsection, we introduce the basic notation, a formula for the load vector and the models containing random matching, balancing circuit and diffusion, along with several important preliminaries such as the *smoothing time*, which serves as a benchmark for analyzing the convergence time of the discrete process. For each specific problem, we recall and define the further notions in the corresponding section.

We consider an arbitrary graph $G = (V, E)$ with n nodes. Each process is modeled as a Markov chain $(X(t))_{t\in\mathbb{N}_0}$, where the *load vector* $X(t) = (X_i(t))_{i\in[n]} \in \mathbb{R}^n$ represents the *state* of the process at the end of step t, and $X_i(t)$ denotes the load of node i at time t. We measure the imbalance of a load vector using the *discrepancy*:

$$\operatorname{disc}(X(t)) := \max_{i\in[n]} X_i(t) - \min_{j\in[n]} X_j(t),$$

and define the average load as

$$\overline{x} := \frac{1}{n}\sum_{i\in V} x_i(0).$$

Uppercase letters, such as $X_i(t)$ and $\mathbf{M}^{(t)}$, denote random variables and random matrices, while lowercase letters (e.g., $x_i(t)$, $\mathbf{m}^{(t)}$) denote fixed realizations.

For matching-based problems, the idealized (continuous) balancing step in round t can be represented by multiplication with a matrix $\mathbf{M}^{(t)} \in \mathbb{R}^{n\times n}$ defined as

$$\mathbf{M}^{(t)}_{i,j} := \begin{cases} 1, & \text{if } i = j \text{ and } i \text{ is not matched at time } t,\\ 1/2, & \text{if } i = j \text{ and } i \text{ is matched at time } t,\\ 1/2, & \text{if } i \text{ and } j \text{ are matched at time } t,\\ 0, & \text{otherwise.} \end{cases}$$

With a slight abuse of notation, we use the same symbol $\mathbf{M}^{(t)}$ to refer both to the matching and to the associated balancing matrix, and we refer to both simply as *matchings*. Following [82, 85], we write $[u : v] \in \mathbf{M}^{(t)}$ to indicate that nodes u and v, with $u < v$, are matched in round t.

For the product of all matching matrices from time t_1 to t_2, we write

$$\mathbf{M}^{[t_1,t_2]} := \prod_{s=t_1}^{t_2} \mathbf{M}^{(s)},$$

where for $t_1 > t_2$ this product is defined as the identity matrix. We generally refer to these matrices as *mixing matrices*. Moreover, we write $\mathbf{M}^{[t]}$ for the sequence of matching matrices $(\mathbf{M}^{(\tau)})_{\tau\in[t]}$ and analogously $\mathbf{m}^{[t]}$ for a fixed sequence $(\mathbf{m}^{(\tau)})_{\tau\in[t]}$. For a matrix $\mathbf{M}$, we denote by $\mathbf{M}_{k,\cdot}$ its kth row (which we often treat as a column vector when convenient).

In line with previous work [32, 51, 85], we consider the class of randomized algorithms that generate a sequence of (random) matchings $\mathbf{M}^{[\infty]}$, referred to as the *random matching model*. The matchings are mutually independent across all rounds. Furthermore, 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 constant $c > 0$ (possibly depending on n). Here, Δ denotes the maximum degree of G. Note that within the same round, the inclusion decisions for different edges are generally *not* independent. Concrete distributed algorithms satisfying these conditions are described in [32, 51].

In the *balancing circuit model*, the graph G is covered by Δ fixed matchings $\mathbf{m}^{(1)}, \ldots, \mathbf{m}^{(\Delta)}$. In step

$t \in \mathbb{N}$, the selected matching is

$$\mathbf{m}^{(t)} := \mathbf{m}^{((t-1) \bmod \Delta+1)}.$$

The *round matrix* is then defined as

$$\mathbf{R} := \mathbf{M}^{[1,\Delta]},$$

which corresponds to the product of the matching matrices forming one complete period of the circuit. Here, Δ denotes both the number of matchings in the circuit model and the maximum degree of G in the random matching model. We write $\lambda(\mathbf{R})$ for the spectral gap of $\mathbf{R}$, i.e., the difference between its two largest eigenvalues.

Let $\varepsilon(t) \in \mathbb{R}^n$ denote the vector of additive rounding errors in round t, where $\varepsilon_k(t)$ is the difference between the load at node k after step t and the load in the idealized scheme with arbitrarily divisible load. Moreover, let $A(t) \in \mathbb{N}^n$ denote the allocated load vector in round t, with $A_k(t)$ items added to node k. In this work, all vectors are considered as column vectors.

With this notation, the load vector at the end of step t can be expressed as

$$X(t) = \mathbf{M}^{(t)} \cdot (X(t-1) + A(t)) + \varepsilon(t). \tag{2}$$

For problems in the diffusion model with diffusion matrix $\mathbf{P}$, we have

$$X(t) = \mathbf{P} \cdot (X(t-1) + A(t)) + \varepsilon(t). \tag{3}$$

In the diffusion process on d-regular graphs, the diffusion matrix $\mathbf{P}$ follows the standard random walk on a d-regular graph G, i.e.,

$$\mathbf{P}_{u,v} := \begin{cases} \frac{1}{d+1}, & \text{if } (u,v) \in E \text{ or } u = v, \\ 0, & \text{otherwise.} \end{cases}$$

Idealized Continuous Process. In the idealized process, the load is assumed to be arbitrarily divisible. Hence, both the allocation vector $A(t)$ and the rounding error vector $\varepsilon(t)$ are zero. This process serves as a useful analytical benchmark: when the continuous load discrepancy becomes negligible, the discrete load is determined primarily by the newly added load items and rounding errors.

In the matching process, we have $X(t) = \mathbf{M}^{(t)]} \cdot X(t-1)$, that is, for $[u,v] \in \mathbf{M}^{(t)}$,

$$X_u(t) = X_v(t) = \frac{X_u(t-1) + X_v(t-1)}{2}.$$

In the diffusion process, we have $X(t) = \mathbf{P} \cdot X(t-1)$, i.e., for each node $u \in V$,

$$X_u(t) = \sum_{v \in N(u) \cup \{u\}} \frac{X_v(t-1).}{d+1}.$$

Throughout the proofs, we use the notion of (K,κ)-*smoothing* time (also known as the *continuous balancing* time). Specifically, it denotes the time required by the continuous process to reduce the discrepancy from K to $\kappa < K$.

Definition 1.6 ([85, Definition 2.1])**.** *A fixed sequence of adjacency matrices* $\mathbf{M}^{[t]}$*,* $t \geq 1$*, is called*

(K, ϵ)-smoothing *if for any* $x(0) \in \mathbb{R}^n$ *with* $\mathrm{disc}(x(0)) = K$*, we have*

$$\mathrm{disc}(\mathbf{M}^{[1,t]} \cdot x(0)) \le \epsilon.$$

For the random matching, balancing circuit, and asynchronous models, there exists a round $t^\star = \mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K/\kappa))$ for which the sequence of matchings $\mathbf{M}^{[t^\star]}$ is (K, κ)-smoothing, w.h.p., [82, 85].

In the diffusion process, the sequence of matrices $\mathbf{M}^{[t]}$ is replaced by $\mathbf{P}^{[t]}$, the sequence of t transition matrices $\mathbf{P}$. In this case, we define the *smoothing time* of matrix $\mathbf{P}$ as follows.

Definition 1.7 (Smoothing Time)**.** *For a diffusion matrix* $\mathbf{P}$ *and an arbitrary initial load vector with* $\mathrm{disc}(X(0)) = K$ *and* $\kappa < K$*, the* (K, κ)*-smoothing time is defined as*

$$\tau_S(\mathbf{P}, K, \kappa) := \min\bigl\{t \bigm| \mathrm{disc}(\mathbf{P}^t \cdot x(0)) \le \kappa\bigr\}.$$

It is the time required by the continuous process to decrease the initial discrepancy K to κ, while applying $\mathbf{P}$ in each round. Note that Theorem 1 in [82] states that

$$\tau_{\mathrm{S}}(\mathbf{P}, K, \kappa) = \mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K/\kappa)) = \mathrm{O}\biggl(\frac{\log(Kn/\kappa)}{1 - \lambda(\mathbf{P})}\biggr).$$

Part Two:

2 Discrete Dynamic Load Balancing on Matchings

P. Berenbrink et al.: *Dynamic Averaging Load Balancing on Arbitrary Graphs*. In *50th International Colloquium on Automata, Languages, and Programming, ICALP*. LIPIcs, volume 261, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, pages 18:1–18:18. DOI: 10.4230/LIPICS.ICALP.2023.18

2.1 Introduction

In this section we study dynamic averaging load balancing on general graphs. We consider infinite time and dynamic processes, where in every step new load items are assigned to randomly chosen nodes. A matching is chosen, and the load is averaged over the edges of that matching. We analyze the discrete case where load items are indivisible, moreover our results also carry over to the continuous case where load items can be split arbitrarily. For the choice of the matchings we consider three different models, random matchings of linear size, random matchings containing only single edges, and deterministic sequences of matchings covering the whole graph. We bound the discrepancy, which is defined as the difference between the maximum and the minimum load. Our results cover a broad range of graph classes and, to the best of our knowledge, our analysis is the first result for discrete and dynamic averaging load balancing processes.

Results in a Nutshell. We present, for the three models introduced above, bounds on the expected discrepancy and bounds that hold with high probability. Our bounds for the synchronous model with balancing circuits hold for arbitrary graphs G, the bounds for the asynchronous model and the synchronous model with random matchings hold for regular graphs G only. For the asynchronous model and the model with random matchings our bounds on the discrepancy are expressed in terms of hitting times of a standard random walk on G, as well as in terms of the spectral gap of the Laplacian of G. For the synchronous model with balancing circuits we express our bounds in terms of the *global divergence*. This can be thought of as a measure of the convergence speed of the Markov chains modeling a random walk on G. However, it does not directly measure the speed of convergence of the chain. It accounts for the time period in which the chain keeps a given distance from the stationary (and uniform) distribution. In physics terminology, it is a measure of total *absement*, which is the time-integral of displacement.

For all three infinite processes our bounds on the discrepancy hold at an arbitrary point of time as long as the system is initially empty. Otherwise, the bounds hold after an initial time period, its length is a function of the initial discrepancy. In the following we give some exemplary results assuming that the system is initially empty $m = n$. For the synchronous model with random matchings and the asynchronous model we can bound the discrepancy by $\mathrm{O}(\sqrt{n}\log(n))$ for *any* regular graph G. Our results show a polylogarithmic bound on the discrepancy for all regular graphs with a hitting time at most $\mathrm{O}(n\mathrm{poly}\log(n))$ (e.g., the two-dimensional torus or the hypercube). In all models we can bound the discrepancy by $\mathrm{O}(\sqrt{n\log(n)})$ for arbitrary constant-degree regular graphs. Moreover, in the balancing circuit model we provide a lower bound on the discrepancy showing that the upper bound is tight up to a multiplicative factor $\sqrt{\log(n)}$. We give a detailed overview on the results on specific graph classes in table 1. All bounds presented in this paper also hold for the corresponding continuous processes without rounding.

Comparisons to last works. The authors of [6] consider the asynchronous process on cycles in the continuous setting where the load items can be divided into arbitrary small pieces. They bound the expected discrepancy and show that $\mathbf{E}[\mathrm{disc}(G)] = \mathrm{O}(\sqrt{n}\log(n))$ for a cycle G with n nodes. In contrast, we improve that bound for the cycle to $\mathrm{disc}(G) = \mathrm{O}(\sqrt{n\log(n)})$. Note that our result not only bounds the expected discrepancy but it also holds with high probability. Moreover, our results work for arbitrary graphs in balancing circuit model and regular graphs in random matching and single edge model.

Outline. The remainder of this section is structured as follows. Section 2.2 introduces the notation and provides the formal definitions of the synchronous and asynchronous models. In Section 2.3, we analyze the balancing circuit model, establishing upper bound on the discrepancy and presenting the intermediate

Table 1: Asymptotic upper bounds on the discrepancy in specific graph classes.

Graph	$\textsc{SBal}(\mathcal{D}_{\mathrm{BC}}(G), 1, m)$	$\textsc{SBal}(\mathcal{D}_{\mathrm{RM}}(G), 1, m)$	$\textsc{ABal}(\mathcal{D}_1(G), 1)$
d-regular graph (const. d)	$\log(n) + \sqrt{m \cdot \log(n)}$	$\log(n) + \sqrt{m \cdot \log(n)}$	$\sqrt{n \cdot \log(n)}$
cycle C_n	$\log(n) + \sqrt{m \cdot \log(n)}$	$\log(n) + \sqrt{m \cdot \log(n)}$	$\sqrt{n \cdot \log(n)}$
2-D torus	$(1 + \sqrt{m/n}) \cdot \log(n)$	$\log(n) + \sqrt{m/n} \cdot \log^{3/2}(n)$	$\log^{3/2}(n)$
r-D torus (const. $r \geq 3$)	$\log(n) + \sqrt{m/n \cdot \log(n)}$	$(1 + \sqrt{m/n}) \cdot \log(n)$	$\log(n)$
hypercube	$(1 + \sqrt{m/n}) \cdot \log(n)$	$(1 + \sqrt{m/n}) \cdot \log(n)$	$\log(n)$
expander	$\log(n) + \sqrt{\Delta/\lambda(\mathbf{R})} \cdot \sqrt{m/n \cdot \log(n)}$	$\log(n) + \sqrt{m/n \cdot \log(n)}$	$\log(n)$

results that support this bound, followed by corresponding lower bound. Section 2.4 examines the random matching model and derives an upper bound on the discrepancy in terms of structural properties of the underlying graph. Section 2.5 extends the analysis to the asynchronous setting. Section 2.6 compiles the technical lemmas that underpin our results, encompassing both fundamental and intermediate findings. Finally, Section 2.7 discusses results for specific graph families, and Section 2.8 concludes with a summary of our contributions and a discussion of open problems.

2.2 Model and Definitions

We begin by introducing notations and the balancing process. Subsection 2.2.1 introduces the synchronous process while Subsection 2.2.2 defines the asynchronous process and recall useful definitions.

Each process is modeled by a Markov chain $(X(t))_{t\in\mathbb{N}_0}$, where the *load vector* $X(t) = (X_i(t))_{i\in[n]} \in \mathbb{R}^n$ is the *state* of the process at the end of step t, and $X_i(t)$ is the load of node i at time t. We measure a load vector's imbalance by the discrepancy $\mathrm{disc}(\vec{x})$, which is the difference between the maximum load and the minimum load $\mathrm{disc}(\vec{x}) := \max_{i\in[n]} x_i - \min_{j\in[n]} x_j$.

We investigate two balancing processes, the synchronous process SBal and the asynchronous process ABal. In the former, a matching of linear size is given and the matched nodes balance their load items. In the latter, a matching of size one is given and the two matched nodes balance their load items.

For balancing action, we introduce a *balancing parameter* β determining the balancing speed which measures what fraction of load items should move. For two matched nodes i, j, the balancing with speed β, $\textsc{Bal}(i, j, \beta)$, is defined as follows.

$\textsc{Bal}(i, j, \beta)$:

1. Assume w.l.o.g. that $X_i(t) \geq X_j(t)$.
2. Let $p = \frac{\beta\cdot(X_i(t)-X_j(t))}{2} - \left\lfloor \frac{\beta\cdot(X_i(t)-X_j(t))}{2} \right\rfloor$.
3. Then, node i sends $L_{i,j}$ load items to node j where

$$L_{i,j} := \begin{cases} \left\lceil \frac{\beta\cdot(X_i(t)-X_j(t))}{2} \right\rceil, & \text{with probability } p, \\ \left\lfloor \frac{\beta\cdot(X_i(t)-X_j(t))}{2} \right\rfloor, & \text{with probability } 1-p. \end{cases}$$

To capture the matching distribution, we introduce $\mathcal{D}(G)$ for SBAL which is a distribution over linear-sized matchings of G. For ABAL, $\mathcal{D}_1(G)$ is a distribution over edges of G. SBAL is additionally parameterized by the number of load items $m \in \mathbb{N}^+$ allocated in each round. ABAL allocates only one new load item per step.
In the idealized setting, where the load is continuously divisible, a load of $\beta(X_i(t) - X_j(t))/2$ is sent from node i to node j.

2.2.1 Synchronous Processes

The synchronous process SBAL$(\mathcal{D}(G), \beta, m)$ works as follows. The process first allocates m items to randomly chosen nodes. Then it uses the matching distribution $\mathcal{D}(G)$ to determine the matching which is applied. Finally it balances the load over the edges of the matching. The parameter $\beta \in (0, 1]$ controls the fraction of the load difference that is sent over an edge in a step.

For the synchronous process SBAL we consider two families of matching distributions, balancing circuits $(\mathcal{D}_{\mathrm{BC}}(G))$ and random matchings $(\mathcal{D}_{\mathrm{RM}}(G))$. In the balancing circuit model we assume G is covered by Δ fixed matchings $\mathbf{m}^{(1)}, \ldots, \mathbf{m}^{(\Delta)}$. $\mathcal{D}_{\mathrm{RM}}(G)$ is generated according to the following method described in [51]. First an edge set S is formed by including each edge with probability $1/(4d) - 1/(16d^2) = \Theta(1/d)$, independently from all other edges. Then a linear-sized matching $\mathbf{M}^{(t)} \subseteq S$ is computed locally.

SBAL$(\mathcal{D}(G), \beta, m)$: In each round $t \in \mathbb{N}^+$:

1. Allocate m discrete, unit-sized load items to the nodes uniformly and independently at random. Define $A_i(t)$ as the number of tokens assigned to node $i \in [n]$ and set $X_i(t) := X_i(t-1) + A_i(t)$.
2. Sample $\mathbf{M}^{(t)}$ according to $\mathcal{D}(G)$.
3. For each edge $\{i, j\} \in \mathbf{M}^{(t)}$ do the balancing action with BAL(i, j, β).

2.2.2 Asynchronous Process

The asynchronous process ABAL$(\mathcal{D}_1(G), \beta)$ works as follows. The process first uses $\mathcal{D}_1(G)$ to generate a matching, this time containing one edge only. The distribution we consider, $\mathcal{D}_1(G)$, first chooses a node i uniformly at random and then it chooses one of the nodes' edges (i, j) uniformly at random. Finally one new token is assigned to either node i or j and then the edge (i, j) is used for balancing. Note that for ABAL$(\mathcal{D}_1(G), \beta)$ the load allocation heavily depends on the edges which are used for balancing. This makes the analysis for this model quite challenging. In contrast, in SBAL$(\mathcal{D}_{\mathrm{RM}}(G), \beta, m)$ the load allocation and the balancing are independent. Note that in the case of d-regular graphs $\mathcal{D}_1(G)$ is equivalent to the uniform distribution over all edges or to choosing a random matching of size one.

ABAL$(\mathcal{D}_1(G), \beta)$: In each round $t \in \mathbb{N}^+$:

1. Select an edge $\{i, j\}$ according to $\mathcal{D}_1(G)$.
2. Allocate a single unit-size load item to either node i or j with a probability of $1/2$.

   I.e., with prob. $1/2$ set $A_i(t) = 1$ and $A_k(t) = 0$ for all $k \neq i$, otherwise set $A_j(t) = 1$ and $A_k(t) = 0$ for all $k \neq j$ and set $X_i(t) := X_i(t-1) + A_i(t)$ for $i \in [n]$.
3. For the edge $\{i, j\}$ do the balancing with BAL(i, j, β).

We are given an arbitrary graph $G = (V, E)$ with n nodes. We mainly assume that G is regular and write d for the node degree. We model the idealized balancing step in round t by multiplication with a matrix $\mathbf{M}^{\beta}(t) \in \mathbb{R}^{n\times n}$ given by

$$\mathbf{M}^{\beta}_{i,j}(t) := \begin{cases} 1, & \text{if } i = j \text{ and } i \text{ is not matched at time } t,\\ 1-\beta/2, & \text{if } i=j \text{ and } i \text{ is matched at time } t,\\ \beta/2, & \text{if } i \text{ and } j \text{ are matched at time } t,\\ 0, & \text{otherwise.}\end{cases}$$

Recall that from Equation (2) we have

$$X(t) = \mathbf{M}^{\beta}(t)\cdot(X(t-1)+A(t)) + \varepsilon(t), \tag{4}$$

in which $X(t), \mathbf{M}^{\beta}(t), A(t)$ and $\varepsilon(t)$ are the load vector, the balancing matrix, the newly added load items and rounding error vectors in round t. We will omit the parameter β if it is clear from context.

We write $\mathrm{t_{hit}(G)}$ for the *hitting time* of G, which is the maximum expected time it takes for a standard random walk on G (i.e., the walk moves to a neighbor chosen uniformly at random in each step) to reach a given node i from a given node j, with the maximum taken over all such pairs of nodes. We write $\mathrm{t^{*}_{hit}(G)}$ for the *edge hitting time* of G, which is defined like the hitting time, except that the maximum is taken over adjacent nodes only. We write $\mathbf{L}(G)$ for the normalized Laplacian matrix of a graph G. For regular graphs it may be defined as $\mathbf{L}(G) := \mathbf{I} - \mathbf{A}(G)/d$, where $\mathbf{A}(G)$ is the adjacency matrix of G. Writing $\lambda_0 \le \lambda_1 \le \ldots \le \lambda_{n-1}$ for the real eigenvalues of $\mathbf{L}(G)$, we let $\lambda(\mathbf{L}(G)) := \lambda_1 - \lambda_0$ be the spectral gap of the Laplacian of G. The convergence time of idealized load balancing, is a function of spectral gap of the Laplacian graph, i.e., $1-\mathbf{L}(G)$ and logarithm of the initial discrepancy.

2.3 Balancing Circuit Model

In this section, we analyze the synchronous process for the balancing circuit model. The main results are Theorem 2.1 and Theorem 2.5, which provide an upper and lower bound on the discrepancy. Note that these results hold for arbitrary graphs.

Recall that in the balancing circuit model, the graph G is covered by Δ fixed matchings $\mathbf{m}^{(1)}, \ldots, \mathbf{m}^{(\Delta)}$. The matching distribution $\mathcal{D}_{\mathrm{BC}}(G)$ deterministically selects the matching $\mathbf{m}^{(t)} = \mathbf{m}^{((t-1) \bmod \Delta+1)}$ at step t. The round matrix is defined as $\mathbf{R} = \mathbf{M}^{[1,\Delta]}$. For a sequence of matchings $\mathbf{m}^{[t]} := (\mathbf{m}^{(s)})_{s=1}^{t}$, we define the *global divergence* as

$$\Upsilon(\mathbf{m}^{[t]}) := \max_{k\in[n]} \sqrt{\sum_{\tau=1}^{t} \left\| \mathbf{M}^{[\tau,t]}_{k,\cdot} - \frac{\vec{1}}{n} \right\|_2^2}.$$

The global divergence can be interpreted as a measure of the convergence speed of a random walk that uses the matching matrices as transition probabilities. In [14, 47, 85], the authors introduce a related notion called the *local p-divergence*, also defined on a sequence of matchings. The difference is that the global divergence measures deviations from the global average across all nodes, whereas the local divergence measures differences between neighboring nodes.

Theorem 2.1. *Let G be an arbitrary graph and $X(t)$ be the state of process* $\textsc{SBal}(\mathcal{D}_{\mathrm{BC}}(G), 1, m)$ *at time t with* $\mathrm{disc}(X(0)) =: K$. *For all $t \in \mathbb{N}$ with $t \ge 2\cdot(\Delta/\lambda(\mathbf{R}))\cdot \ln(K\cdot n)$ it holds w.h.p. and in expectation*

$$\mathrm{disc}(X(t)) = \mathrm{O}\left(\log(n) + \sqrt{\frac{m}{n}}\cdot \Upsilon(\mathbf{m}^{[t]})\cdot\sqrt{\log(n)}\right).$$

Proof. Note that $\beta = 1$. We define $\mathbf{m}^{(t)} := \mathbf{M}^{\beta}(t)$. We first expand the recurrence of Equation (4) (cf. [82]). After one step we get

$$\begin{aligned} X(t) &= \mathbf{m}^{(t)} \cdot (X(t-1) + A(t)) + \varepsilon(t) \\ &= \mathbf{m}^{(t)} \cdot \Big(\underbrace{\Big(\mathbf{m}^{(t-1)} \cdot (X(t-2) + A(t-1)) + \varepsilon(t-1) \Big)}_{X(t-1)} + A(t) \Big) + \varepsilon(t) \\ &= \mathbf{m}^{[t-1,t]} \cdot X(t-2) + \sum_{s=t-1}^{t} \mathbf{m}^{[s,t]} \cdot A(s) + \sum_{s=t-1}^{t} \mathbf{m}^{[s+1,t]} \cdot \varepsilon(s) \end{aligned}$$

We repeatedly expand this form up to the beginning of the process and get

$$X(t) = \underbrace{\mathbf{m}^{[1,t]} \cdot X(0)}_{\substack{I(t) \\ \text{initial load contribution}}} + \underbrace{\sum_{s=1}^{t} \mathbf{m}^{[s,t]} \cdot A(s)}_{\substack{D(t) \\ \text{dynamically allocated Load contribution}}} + \underbrace{\sum_{s=1}^{t} \mathbf{m}^{[s+1,t]} \cdot \varepsilon(s)}_{\substack{R(t) \\ \text{rounding error contribution}}} . \quad (5)$$

We write $I(t)$, $D(t)$, and $R(t)$ for the three terms as indicated. Note that, in general, these terms are vectors of real numbers. The sum $I(t) + D(t)$ can be regarded as the contribution of the (dynamic) continuous process, where $I(t)$ is the contribution of the initial load and $D(t)$ is the contribution of the dynamically allocated load. Thus, $R(t)$ is the deviation between the idealized process without rounding and the discrete process.

To bound the discrepancy $\operatorname{disc}(X(t))$ of the load vector $X(t)$ at time t we use the fact that the discrepancy is sub-additive such that $\operatorname{disc}(\vec{x} + \vec{y}) \le \operatorname{disc}(\vec{x}) + \operatorname{disc}(\vec{y})$ (see Observation 2.13). Hence, to bound $\operatorname{disc}(X(t))$ we individually bound the discrepancies of the three terms in Equation (5) and get

$$\operatorname{disc}(X(t)) \le \operatorname{disc}(I(t)) + \operatorname{disc}(D(t)) + \operatorname{disc}(R(t)). \quad (6)$$

If the system is initially empty, then $\operatorname{disc}(I(t)) = 0$. Moreover, in the idealized setting without rounding $\operatorname{disc}(R(t)) = 0$. In the rest of this Subsection we bound these three terms. Bounds on $\operatorname{disc}(I(t))$ and $\operatorname{disc}(R(t))$ is quite well-known and simple and are done in Subsections 2.3.1 and 2.3.2, respectively. The interesting part is bounding $\operatorname{disc}(D(t))$ which is done in Subsection 2.3.3.

From Proposition 2.2 for the specified t in the statement we have,

$$\operatorname{disc}(I(t)) = 1. \quad (7)$$

From Proposition 2.3 we get

$$\mathbf{Pr}\Big[\operatorname{disc}(R(t)) = \mathrm{O}\Big(\sqrt{\log(n)}\Big)\Big] \ge 1 - 2 \cdot n^{-2}. \quad (8)$$

And from Lemma 2.4 (with $\gamma := 3$) it follows that

$$\mathbf{Pr}\left[\operatorname{disc}(D(t)) = \mathrm{O}\left(\log(n) + \sqrt{\log(n) \cdot \frac{m}{n}} \cdot \Upsilon(\mathbf{m}^{[t]})\right)\right] \ge 1 - 2 \cdot n^{-2}. \quad (9)$$

The statement of the theorem then follows from a union bound over the Equation (7), Equation (8) and Equation (9). The bound on expectation follows from the linearity of expectation and the bounds on the expected discrepancies in the aforementioned propositions. ☐

In the rest of this section, we first prove the three propositions used in the theorem above. Then, we provide a lower bound on the discrepancy in this model.

2.3.1 Bounding the Contribution of the Initial Load, $\operatorname{disc}(I(t))$

The following proposition shows that, in this model, for sufficiently large rounds t, the contribution of the initial load to the overall discrepancy becomes negligible. In particular, we determine the value of t for which the sequence $\mathbf{m}^{[t]}$ is $(K, 1)$-smoothing in this model (see Definition 1.6).

Proposition 2.2 (Smoothing Property)**.** *Consider the balancing circuit model with round matrix* $\mathbf{R}$*. For all* $t \in \mathbb{N}$*,* $K \geq 1$ *if* $t \geq 2 \cdot (\Delta / \lambda(\mathbf{R})) \cdot \ln(K \cdot n)$*, then* $\operatorname{disc}(I(t)) \leq 1$*.*

Proof. Recall that the round matrix is given by $\mathbf{R} = \prod_{s=1}^{\Delta} \mathbf{m}^{(s)}$. The *quadratic node potential* is defined as

$$\Phi(\vec{x}) \coloneqq \sum_{i \in [n]} (x_i - \overline{x})^2, \quad \text{where} \quad \overline{x} \coloneqq \frac{1}{n} \sum_{j \in [n]} x_j.$$

Since $\Phi(X(0)) \leq K^2 n$, it follows from Lemma 2 in [53] (restated as Lemma 6.22) that

$$\begin{aligned} \Phi\Big(\mathbf{m}^{[1,t]} \cdot X(0)\Big) &\leq (1 - \lambda(\mathbf{R}))^{2\lfloor t \rfloor / \Delta} \cdot \Phi(X(0)) \\ &\leq (1 - \lambda(\mathbf{R}))^{2\lfloor t \rfloor / \Delta} \cdot K^2 n \\ &\leq e^{-2\lfloor t \rfloor\, \lambda(\mathbf{R})/\Delta + 2\ln(Kn)}. \end{aligned}$$

Setting $t \geq 2\Delta \cdot \ln(Kn) / \lambda(\mathbf{R})$ implies that $\Phi\big(\mathbf{m}^{[1,t]} \cdot x(0)\big) \leq n^{-2}$. Consequently, for all $t \in \mathbb{N}$ satisfying this bound, we have $\operatorname{disc}(I(t)) \leq 1$. □

2.3.2 Bounding the Contribution of the Rounding Errors, $\operatorname{disc}(R(t))$

In this subsection, we analyze the effect of cumulative rounding errors on the discrepancy. The proof follows similarly to the proof of Theorem 3.6 in [85] which is based on work in [14].

Proposition 2.3 (Insignificance of Rounding Errors)**.** *Let* G *be an arbitrary graph. Then for all* $t \in \mathbb{N}$*, and* $k \in [n]$ *we get with probability at least* $1 - 2n^{-2}$ *and in expectation* $\operatorname{disc}(R(t)) = \mathrm{O}\Big(\sqrt{\log(n)}\Big)$.

Proof. The proof directly follows from the proof of [85, Theorem 3.6]. In the next part, using a different technique, we prove a generalized version of this theorem. In fact, we show that with slightly smaller probability the discrepancy of rounding errors is at most 3 (after sufficiently many steps). However, for our analysis, the bound of $\mathrm{O}(\sqrt{\log(n)})$ is sufficient. To establish the expectation bound on $\mathbf{E}[\operatorname{disc}(R(t)]$, we apply Lemma 6.2 with $X = \operatorname{disc}(R(t))$, $c = 2$ and $C = 2\sqrt{\log(n)}$ to see that

$$\mathbf{E}[\operatorname{disc}(R(t))] \leq 2\sqrt{\log(n)} \cdot \left(1 + \frac{2}{\log(n)}\right) = \mathrm{O}\Big(\sqrt{\log(n)}\Big). \qquad \square$$

2.3.3 Bounding the Contribution of Dynamically Allocated Loads by Global Divergence

In the following lemma, we bound the discrepancy of dynamically allocated loads in terms of the global divergence. In this setting, the matchings are assumed to be fixed, and the only source of randomness comes from the random placement of loads. By fixing the sequence of matchings, the result can later be applied directly to the random matching model $\mathcal{D}_{\mathrm{RM}}$. The proof relies on two high-level ideas. First, the placement of newly added items is independent of the sequence of matching matrices. Second, the matching matrices themselves are doubly stochastic. As we will see in a later section, this lemma can also be extended to the diffusion model (see Lemma 4.14).

Lemma 2.4 (Load concentration via global divergence)**.** *Let* $\mathbf{m}^{[t]}$ *be an arbitrary sequence of matchings. Then for all* $\gamma > 0$ *and* $t \in \mathbb{N}$ *we get with probability at most* $2 \cdot n^{-\gamma+1}$

$$\mathrm{disc}(D(t)) \geq \frac{8}{3} \cdot \gamma \log(n) + \sqrt{32\gamma \log(n) \cdot \frac{m}{n}} \cdot \Upsilon(\mathbf{m}^{[t]}).$$

Proof. Fix a node $k \in [n]$. First we establish a concentration inequality on $D_k(t)$ in terms of $\Upsilon_k(\mathbf{m}^{[t]})$.

Our goal is to decompose $D_k(t)$ into a sum of independent random variables. Recall that we assume that the matching matrices are fixed and all randomness is due to the random choices of the load items. This will enable us to apply a concentration inequality to this sum. For the decomposition observe that $D(t) = \sum_{\tau=1}^{t} \mathbf{m}^{[\tau,t]} \cdot A(\tau)$, where $A(\tau)$ is the random load vector corresponding to the m load items allocated at time τ. So the kth coordinate of $D(t)$ is $D_k(t) = \sum_{\tau=1}^{t} \sum_{w\in[n]} \mathbf{m}_{k,w}^{[\tau,t]} \cdot A_w(\tau)$. We define the indicator random variable $B(\tau, j, w)$ for $\tau \in [t], j \in [m]$ and $w \in [n]$ as , starting from their token locations, $\Pr[B(\tau, j, w) = 1] = 1/n$ and $\mathbf{E}[B(\tau, j, w)] = 1/n$. Observe that $A_w(\tau)$, the load allocated to node w at step τ, can be expressed as $\sum_{j\in[m]} B(\tau, j, w)$. Merging this with the value of $D_k(t)$ gives

$$D_k(t) = \sum_{\tau=1}^{t} \sum_{w\in[n]} \mathbf{m}_{k,w}^{[\tau,t]} \cdot \left(\sum_{j\in[m]} B(\tau, j, w) \right) = \sum_{\tau=1}^{t} \sum_{j\in[m]} \underbrace{\left(\sum_{w\in[n]} \left(\mathbf{m}_{k,w}^{[\tau,t]} \cdot B(\tau, j, w) \right) \right)}_{=: C_k(\tau,j)}.$$

For a fixed $\tau \in [t]$ and $j \in [m]$ we define $C_k(\tau, j) := \sum_{w\in[n]} \mathbf{m}_{k,w}^{[\tau,t]} \cdot B(\tau, j, w)$. This random variable measures the contribution of j-th load item of round τ to $D_k(t)$. Note that the load items are allocated independently from each other. Since $\mathbf{m}^{[\tau,t]}$ are fixed matrices, then $C_k(\tau, j)$ and $C_k(\tau', j')$ are independent for all τ and τ' and $j \neq j'$. To apply the concentration inequality from Theorem 6.9 we need to show that $C_k(\tau, j) \leq 1$ and compute an upper bound on $\mathrm{Var}[C_k(\tau, j)]$. Showing the first condition is easy since exactly one of the indicator random variables $B(\tau, j, w)$ is one and $\mathbf{m}_{k,w}^{[\tau,t]}$ has a value between zero and one.

It remains to consider the variance of $C_k(\tau, j)$. First note that by linearity of expectation

$$\mathbf{E}[C_k(\tau, j)] = \mathbf{E}\left[\sum_{w\in[n]} \left(\mathbf{m}_{k,w}^{[\tau,t]} \cdot B(\tau, j, w) \right) \right] = \sum_{w\in[n]} \mathbf{m}_{k,w}^{[\tau,t]} \cdot \mathbf{E}[B(\tau, j, w)] = \sum_{w\in[n]} \mathbf{m}_{k,w}^{[\tau,t]} \cdot \frac{1}{n} = \frac{1}{n},$$

where the last equality follows form the fact that $\mathbf{m}^{[\tau,k]}$ is doubly stochastic. Now we get

$$\mathrm{Var}[C_k(\tau, j)] = \mathbf{E}\Big[(C_k(\tau, j) - \mathbf{E}[C_k(\tau, j)])^2 \Big] = \mathbf{E}\left[\left(\Big(\sum_{w\in[n]} \mathbf{m}_{k,w}^{[\tau,t]} \cdot B(\tau, j, w) \Big) - \frac{1}{n} \right)^2 \right]$$

$$= \sum_{w'\in[n]} \frac{1}{n} \cdot \left(\mathbf{m}_{k,w'}^{[\tau,t]} - \frac{1}{n} \right)^2 = \frac{1}{n} \cdot \left\| \mathbf{m}_{k,\cdot}^{[\tau,t]} - \frac{\vec{1}}{n} \right\|_2^2,$$

where we used that for each τ and each j exactly one of the $B(\tau, j, w)$ is one and all others are zero, and each of the n possible cases has uniform probability.

Recall that $C_k(\tau, j)$ and $C_k(\tau', j')$ are independent for all τ, τ' and $j \neq j'$. Hence we get

$$\begin{aligned}\mathrm{Var}\left[\sum_{\tau=1}^{t}\sum_{j\in[m]} C_k(\tau,j)\right] &= \sum_{\tau=1}^{t}\sum_{j\in[m]} \mathrm{Var}[C_k(\tau,j)] = \frac{1}{n}\cdot\sum_{\tau=1}^{t}\sum_{j\in[m]}\left\|\mathbf{m}_{k,\cdot}^{[\tau,t]} - \frac{\vec{1}}{n}\right\|_2^2 \\ &= \frac{m}{n}\cdot\left(\Upsilon_k(\mathbf{m}^{[t]})\right)^2,\end{aligned}$$

where the final equality uses the definition of the global divergence $\Upsilon_k(\mathbf{m}^{[t]})$. Applying Theorem 6.9 with $M = 1$ and $X = D_k(t) = \sum_{\tau=1}^{t}\sum_{j\in[m]} C_k(\tau, j)$ with $\lambda = 2\gamma\log(n)/3 + \Upsilon_k(\mathbf{m}^{[t]})\cdot\sqrt{2\gamma m/n}$ results in

$$\Pr\left[D_k(t) - t\cdot\frac{m}{n} \geq \frac{2}{3}\cdot\gamma\log(n) + \sqrt{2\gamma\log(n)\cdot\frac{m}{n}}\cdot\Upsilon_k(\mathbf{m}^{[t]})\right] \leq n^{-\gamma}.$$

The lower bound can be established using Theorem 6.10 (with $a_i = 0$ and $M = 1$) instead of Theorem 6.9. Via a union bound we get

$$\mathbf{Pr}\left[\left|D_k(t) - t\cdot\frac{m}{n}\right| \geq \frac{4}{3}\cdot\gamma\log(n) + \sqrt{8\gamma\log(n)\cdot\frac{m}{n}}\cdot\Upsilon_k(\mathbf{m}^{[t]})\right] \leq 2\cdot n^{-\gamma}. \qquad (10)$$

Applying the union bound over all nodes $k \in [n]$ together with observation 2.14 (showing that $\mathrm{disc}(D(t)) \leq 2|D_k(t) - t\cdot m/n|$) finishes the proof. □

The next theorem provides a lower bound on the discrepancy in the balancing circuit model with $\beta = 1$, $\mathrm{SBAL}(\mathcal{D}_{\mathrm{BC}}(G), 1, m)$. The proof idea is to apply Berry-Essen Theorem (Theorem 6.8) in a way similar to Theorem 1.2 in [33]. For this, we rely on the following facts: (1) the placement of newly added items is independent of the sequence of matrices, (2) the matching matrices are doubly stochastic, and (3) the variable $D_k(t)$ can be decomposed into independent random variables with finite and bounded second and third moments.

Theorem 2.5. *Let G be an arbitrary graph and $X(t)$ be the state of process $\mathrm{SBAL}(\mathcal{D}_{\mathrm{BC}}(G), 1, m)$ at time t. Then for all $t \in \mathbb{N}$ and $m \geq 4n\cdot\log(n)/\Upsilon(\mathbf{m}^{[t]})$ it holds with constant probability*

$$\mathrm{disc}(X(t)) = \Omega\left(\sqrt{\frac{m}{n}}\cdot\Upsilon(\mathbf{m}^{[t]})\right).$$

Proof. First we show a lower bound on $D_k(t)$. The idea is to decompose $D_k(t)$ into sum of independent random variables (Y_ℓ) which have expected value zero and also show that $\sum_\ell \mathbf{E}\left[|Y_\ell^3|\right]$ is properly bounded. It allows us to apply a concentration inequality to the sum. To do so, we define several intermediate random variables similar to the proof of Lemma 2.4.

Fix round t and consider node $k \in [n]$ such that $\Upsilon_k(\mathbf{m}^{[t]}) = \Upsilon(\mathbf{m}^{[t]})$. Recall that,

$$D_k(t) = \sum_{\tau=1}^{t}\sum_{w\in[n]} \mathbf{m}_{k,w}^{[\tau,t]}\cdot A_w(\tau).$$

We define indicator random variables $B(\tau, j, w)$ for $\tau \in [t]$, $j \in [m]$ and $w \in [n]$ as follows.

$$B(\tau,j,w) := \begin{cases} 1, & \text{if } j\text{-th load item of step } \tau \text{ goes to node } w,\\ 0, & \text{otherwise.}\end{cases}$$

Note that for fixed j and τ, $\sum_{w\in[n]} B(\tau, j, w) = 1$ and $\Pr[B(\tau, j, w) = 1] = 1/n$. Recall that $A_w(\tau)$ can be expressed as $\sum_{j\in[m]} B(\tau,j,w)$. It then follows that

$$D_k(t) = \sum_{\tau=1}^{t}\sum_{j\in[m]}\sum_{w\in[n]}\Big(\mathbf{m}_{k,w}^{[\tau,t]}\cdot B(\tau,j,w)\Big).$$

We define the derivative from the average for $D_k(t)$ as

$$\widetilde{D}_k(t) \coloneqq \sum_{\tau=1}^{t}\sum_{k\in[m]}\underbrace{\sum_{w\in[n]}\left(\mathbf{m}_{k,w}^{[\tau,t]}\cdot B(\tau,j,w)-\frac{1}{n^2}\right)}_{C_k(\tau,j)}.$$

It immediately follows that $\widetilde{D}_k(t) = D_k(t) - t\cdot m/n$. We call

$$C_k(\tau,j) \coloneqq \sum_{w\in[n]}\left(\mathbf{m}_{k,w}^{[\tau,t]}\cdot B(\tau,j,w)-\frac{1}{n^2}\right)$$

the contribution of the j-th load item (of step τ) to $\widetilde{D}_k(t)$. For a fixed τ and j, from the linearity of expectation, it follows that

$$\mathbf{E}[C_k(\tau,j)] = \sum_{w\in[n]}\mathbf{E}\left[\mathbf{m}_{k,w}^{[\tau,t]}\cdot B(\tau,j,w)-\frac{1}{n^2}\right] = \left(\sum_{w\in[n]}\mathbf{m}_{k,w}^{[\tau,t]}\cdot\frac{1}{n}\right)-\frac{1}{n}=0,$$

where the last inequality follows since $\mathbf{m}^{[\tau,t]}$ is a doubly stochastic matrix.

Here for $\ell = (\tau-1)\cdot m + j$ such that $\tau\in[t]$ and $j\in[m]$ we define $Y_\ell \coloneqq C_k(\tau,j)$ and it follows $\widetilde{D}_k(t) = \sum_{\ell=1}^{t\cdot m} Y_\ell$. Note that Y_ℓ's are independent. We want to apply the Berry-Esseen Theorem [28, 43] (see Theorem 6.8). To do so, we need to compute $\mathrm{Var}[Y_\ell]$ and $\mathbf{E}\big[|Y_\ell|^3\big]$. Then we get

$$\begin{aligned}\mathrm{Var}[Y_\ell] &= \mathbf{E}\left[\Big(C_k(\tau,j)-\underbrace{\mathbf{E}[C_k(\tau,j)]}_{=0}\Big)^2\right] = \mathbf{E}\left[\left(\sum_{w\in[n]}\left(\mathbf{m}_{k,w}^{[\tau,t]}\cdot B(\tau,j,w)-\frac{1}{n^2}\right)\right)^2\right]\\ &= \mathbf{E}\left[\left(\left(\sum_{w\in[n]}\mathbf{m}_{k,w}^{[\tau,t]}\cdot B(\tau,j,w)\right)-\frac{1}{n}\right)^2\right] = \frac{1}{n}\sum_{w'\in[n]}\left(\mathbf{m}_{k,w'}^{[\tau,t]}-\frac{1}{n}\right)^2 = \frac{1}{n}\cdot\left\|\mathbf{m}_{k,\cdot}^{[\tau,t]}-\frac{\vec{1}}{n}\right\|_2^2,\end{aligned}$$

where in the second last equality we used the fact that for each τ and each j exactly one of the $B(\tau,j,w)$ is one and all others are zero, and that each of the n possible cases has uniform probability. Similarly,

$$\begin{aligned}\mathbf{E}\big[|Y_\ell|^3\big] = \mathbf{E}\Big[|C_k(\tau,j)|^3\Big] &= \mathbf{E}\left[\left|\sum_{w\in[n]}\left(\mathbf{m}_{k,w}^{[\tau,t]}\cdot B(\tau,j,w)-\frac{1}{n^2}\right)\right|^3\right]\\ &\overset{(a)}{=}\sum_{w'\in[n]}\mathbf{E}\left[\left|\sum_{w\in[n]}\left(\mathbf{m}_{k,w}^{[\tau,t]}\cdot B(\tau,j,w)-\frac{1}{n^2}\right)\right|^3 \,\middle|\, B(\tau,j,w')=1\right]\cdot\Pr[B(\tau,j,w')=1],\end{aligned}$$

where (a) follows form the law to total expectation. Using the fact that for any $w'\in[n]$, $|\mathbf{m}_{k,w'}^{[\tau,t]}-1/n|<1$

in (b), from above, we get

$$\mathbf{E}\left[|Y_\ell|^3\right] = \frac{1}{n}\cdot\sum_{w'\in[n]}\left|\mathbf{m}^{[\tau,t]}_{k,w'} - \frac{1}{n}\right|^3 \overset{(b)}{\leq} \frac{1}{n}\cdot\sum_{w'\in[n]}\left(\mathbf{m}^{[\tau,t]}_{k,w'} - \frac{1}{n}\right)^2 \leq \frac{1}{n}\cdot\left\|\mathbf{m}^{[\tau,t]}_{k,\cdot} - \frac{\vec{1}}{n}\right\|_2^2,$$

Recall that $\left\|\mathbf{m}^{[\tau,t]}_{k,\cdot} - \frac{\vec{1}}{n}\right\|_2^2 = \Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)$. By defining $F_{t\cdot m}(x)$ as the distribution of $\frac{\widetilde{D}_k(t)}{\sqrt{\sum_{\ell=1}^{t\cdot m}\mathrm{Var}[Y_\ell]}}$, from Theorem 6.8 it follows that,

$$|F_{t\cdot m}(x) - \Phi_N(x)| \leq C_0\cdot\frac{\sum_{\ell=1}^{t\cdot m}\mathbf{E}\left[|Y_\ell|^3\right]}{\left(\sum_{\ell=1}^{t\cdot m}\mathrm{Var}[Y_\ell]\right)^{3/2}} \leq C_0\cdot\frac{\frac{m}{n}\cdot\sum_{\tau=1}^{t}\Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)}{\left(\frac{m}{n}\cdot\sum_{\tau=1}^{t}\Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)\right)^{3/2}} = o(1),$$

in which the last inequality follows from the assumption, $m \geq 4n\log(n)/\sum_{\tau=1}^{t}\Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)$, and C_0 is some constant. Note that $\Phi_N(x)$ is the standard normal distribution. Therefore it holds that,

$$F_{t\cdot m}(x) \geq \Phi_N(x) - o(1) \geq \frac{1}{\sqrt{\pi}(x+\sqrt{x^2+2})e^{x^2}} - o(1),$$

where the last inequality follows from [[1], Formula 7.1.13] which states

$$\frac{1}{\sqrt{\pi}(x+\sqrt{x^2+2})e^{x^2}} \leq \Phi_N(x) \leq \frac{1}{\sqrt{\pi}(x+\sqrt{x^2+4/\pi})e^{x^2}}.$$

Hence with $x = 1$ we have

$$F_{t\cdot m}(1) \geq \frac{1}{\sqrt{\pi}(1+\sqrt{3})e} - o(1) \geq \frac{1}{16}.$$

Therefore by replacing the definition of $F_{t\cdot m}(1)$ we get that

$$\Pr\left[\frac{\widetilde{D}_k(t)}{\sqrt{\frac{m}{n}\cdot\sum_{\tau=1}^{t}\Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)}} \geq 1\right] = \Pr\left[\widetilde{D}_k(t) \geq \sqrt{\frac{m}{n}\sum_{\tau=1}^{t}\Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)}\right] \geq \frac{1}{16}.$$

Recall that $\widetilde{D}_k(t) = D_k(t) - \mathbf{E}[D_k(t)]$, then it follows that

$$\Pr\left[D_k(t) \geq \mathbf{E}[D_k(t)] + \sqrt{\frac{m}{n}\cdot\sum_{\tau=1}^{t}\Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)}\right] \geq \frac{1}{16}.$$

Moreover, when node k receives more than expectation from the allocated load items, there is (at least) one node w receiving less than expectation. Hence,

$$\Pr\left[D_k(t) - D_w(t) \geq \sqrt{\frac{m}{n}\cdot\sum_{\tau=1}^{t}\Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)}\right] \geq \frac{1}{16}\cdot 1.$$

Since $X(0) = \vec{0}$, then $I_k(t) = I_w(t) = 0$. From Proposition 2.3 it follows that $|R_k(t) - R_w(t)| = \mathrm{O}(\sqrt{\log(n)})$ with probability $1 - o(1)$. Since $m \geq 4n\cdot\log(n)/\sum_{\tau=1}^{t}\Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)$ and $X_k(t) = I_k(t) + D_k(t) + R_k(t)$, then it follows

$$\Pr\left[X_k(t) - X_w(t) \geq \frac{1}{2}\cdot\sqrt{\frac{m}{n}\cdot\sum_{\tau=1}^{t}\Phi\left(\mathbf{m}^{[\tau,t]}_{k,\cdot}\right)}\right] \geq \frac{1}{16}\cdot(1-o(1)) \geq \frac{1}{17},$$

completing the proof. □

2.4 Random Matching Model

Building on the upper bound established for the balancing circuit model, we now extend our results to the process $\mathrm{SBAL}(\mathcal{D}_{\mathrm{RM}}(G), \beta, m)$ on d-regular graphs G. The matching distribution $\mathcal{D}_{\mathrm{RM}}(G)$ is generated according to the algorithm described in [51]. As with the previous model, the upper bound holds at any point in time t, assuming the system is initially empty. Furthermore, the same results remain valid in the idealized setting where load items can be divided into arbitrarily small pieces (see [6]). To achieve this extension, we adapt the proof of Theorem 2.1 so that it (a) holds for an arbitrary β and (b) applies to the matching distribution $\mathcal{D}_{\mathrm{RM}}(G)$

Theorem 2.6. *Let G be a d-regular graph and define $T(G) := \min\left\{\frac{\mathrm{t_{hit}(G)}}{n} \cdot \log(n), \sqrt{\frac{d}{\lambda(\mathbf{L}(G))}}, \frac{1}{\lambda(\mathbf{L}(G))}\right\}$. Let $X(t)$ be the state of process $\mathrm{SBAL}(\mathcal{D}_{\mathrm{RM}}(G), \beta, m)$ at time t with $\mathrm{disc}(X(0)) =: K \geq 1$. There exists a constant $c > 0$ such that for all $t \geq c \cdot \log(K \cdot n)/(\lambda(\mathbf{L}(G)) \cdot \beta)$ it holds w.h.p.and in expectation*

$$\mathrm{disc}(X(t)) = \mathrm{O}\left(\log(n) \cdot \left(1 + \sqrt{\frac{m}{n} \cdot \frac{\mathrm{t^*_{hit}(G)}}{n}}\right) + \sqrt{\frac{\log(n)}{\beta} \cdot \frac{m}{n} \cdot T(G)}\right).$$

Proof. The proof follows along the lines of the proof of Theorem 2.1. For a fixed β throughout this proof we let $\mathbf{M}^{(t)} := \mathbf{M}^{\beta}(t)$. Similar to Equation (5), we get

$$X(t) = \underbrace{\mathbf{M}^{[1,t]} \cdot X(0)}_{\substack{I(t) \\ \text{initial load contribution}}} + \underbrace{\sum_{s=1}^{t} \mathbf{M}^{[s,t]} \cdot A(s)}_{\substack{D(t) \\ \text{dynamically allocated Load contribution}}} + \underbrace{\sum_{s=1}^{t} \mathbf{M}^{[s+1,t]} \cdot \varepsilon(s)}_{\substack{R(t) \\ \text{rounding error contribution}}}, \quad (11)$$

and

$$\mathrm{disc}(X(t)) \leq \mathrm{disc}(I(t)) + \mathrm{disc}(D(t)) + \mathrm{disc}(R(t)).$$

In the rest of this section, we bound each term in the right hand side of the inequality above. We first consider $\mathrm{disc}(I(t))$ (in subsection 2.4.1) and $\mathrm{disc}(R(t))$ (in subsection 2.4.2) since the techniques to bound these terms are well-established. The hardest part of our analysis is that of $\mathrm{disc}(D(t))$; hence, we consider this part at the end of the section (in subsection 2.4.3).

With these bounds we are ready to finish the proof. Let now $\gamma > 1$. First, it follows from Proposition 2.7 that for all $t \geq c\log(Kn)/(\lambda(\mathbf{L}(G))\beta)$ we have $\mathrm{disc}(I(t)) \leq 1$ with probability at least $1 - n^{-\gamma}$. Second, it follows from Proposition 2.8 that $\mathrm{disc}(R(t)) \leq 2\sqrt{\gamma\log(n)/\beta}$ with probability at least $1 - 3n^{-\gamma+1}$. Third, it follows from Proposition 2.9 that

$$\mathrm{disc}(D(t)) = \mathrm{O}\left(\gamma\log(n) \cdot \left(1 + \sqrt{\frac{m}{n} \cdot \frac{\mathrm{t^*_{hit}(G)}}{n}}\right) + \sqrt{\frac{\gamma\log(n)}{\beta} \cdot \frac{m}{n} \cdot T(G)}\right)$$

with probability at least $1 - 2 \cdot n^{-\gamma+1}$. The statement of the theorem therefore follows from a union bound over the statements of Proposition 2.7, Proposition 2.8, and Proposition 2.9. The bound on expectation follows analogously from the linearity of expectation and the bounds on the expected discrepancies in the aforementioned propositions. □

2.4.1 Bounding the Contribution of the Initial Load, $\mathrm{disc}(I(t))$

Here, we show that the contribution of the initial load to the discrepancy becomes negligible when t is sufficiently large. We generalize the analysis of Theorem 1 in [82] to establish a bound on the discrepancy of the initial load as a function of β. In particular, we demonstrate that a sequence of matchings of sufficient length t is $(K, 1)$-smoothing.

Proposition 2.7 (Smoothing Property)**.** *Let G be a d-regular graph and $K = \mathrm{disc}(X(0))$. There exists a constant $c > 0$ such that for all $\gamma > 0$ and $t \in \mathbb{N}$ with $t \geq t_0(\gamma) := c \cdot \max\{\gamma \log(n), \log(K \cdot n)\}/(\lambda(\mathbf{L}(G)) \cdot \beta)$ we get with probability at least $1 - n^{-\gamma}$ and in expectation $\mathrm{disc}(I(t)) \leq 1$.*

Proof. Note that $\max_{i \in [n]} |x_i - \overline{x}| \leq \sqrt{\Phi(\vec{x})}$ by definition of Φ. Hence, $\mathrm{disc}(\vec{x}) \leq 2\sqrt{\Phi(\vec{x})}$. By Lemma 2.22 (presented in technical lemma section), if $t \in \mathbb{N}$ with $t \geq t_0(\gamma) := c \cdot \max\{\gamma \log(n), \log(K \cdot n)\}/(\lambda(\mathbf{L}(G)) \cdot \beta)$, then $\Phi(I(t)) \leq 1/4$ with probability at least $1 - n^{-\gamma}$, and hence $\mathrm{disc}(I(t)) \leq 2\sqrt{\Phi(I(t))} \leq 2\sqrt{1/4} = 1$. Moreover, by the lemma and the Jensen's inequality we get for $t \geq t_0(0)$ that,

$$\mathbf{E}[\mathrm{disc}(I(t))] \leq \mathbf{E}\Big[2\sqrt{\Phi(I(t))}\Big] \leq 2\sqrt{\mathbf{E}[\Phi(I(t))]} \leq 2\sqrt{\frac{1}{4}} = 1,$$

finishing the proof. □

2.4.2 Bounding the Contribution of the Rounding Errors, $\mathrm{disc}(R(t))$

Here we bound the contribution of cumulative rounding errors to the discrepancy. The following proposition is not restricted to the random matching model but applies to all three models considered in this work.

Proposition 2.8 (Insignificance of Rounding Errors)**.** *Let G be an arbitrary graph. Then for all $\gamma > 1$, $t \in \mathbb{N}$, and $k \in [n]$ we get with probability at least $1 - 2n^{-\gamma+1}$ and in expectation* $\mathrm{disc}(R(t)) = \mathrm{O}\Big(\sqrt{\gamma \log(n)/\beta}\Big)$.

Proof. The proof is similar to the proof of [85, Theorem 3.6] for $\beta = 1$. In Lemma 2.23 (presented in technical lemmas) we generalized it to arbitrary β. To show the bound on $\mathbf{E}[\mathrm{disc}(R(t)]$, we apply Lemma 6.2 with $X = \mathrm{disc}(R(t))$, $c = 2$ and $C = 2\sqrt{\log(n)/\beta}$ to see that

$$\mathbf{E}[\mathrm{disc}(R(t))] \leq 2\sqrt{\frac{\log(n)}{\beta}} \cdot \left(1 + \frac{2}{\log(n)}\right) = \mathrm{O}\left(\sqrt{\frac{\log(n)}{\beta}}\right).$$ □

2.4.3 Bounding the Contribution of the Dynamically Allocated Load, $\mathrm{disc}(D(t))$

In this subsection, we establish a tail bound on the discrepancy contribution of dynamically allocated load (Proposition 2.9). This, in fact, constitutes the most challenging proof in this section.

Proposition 2.9 (Contribution of Dynamically Allocated Load)**.** *Let G be a d-regular graph. Define*

$$T(G) := \min\Big\{\mathrm{t_{hit}(G)} \cdot \log(n)/n, \sqrt{d/\lambda(\mathbf{L}(G))}, 1/\lambda(\mathbf{L}(G))\Big\}.$$

Then for all $\gamma > 1$ and $t \in \mathbb{N}$ we get with probability at least $1 - 3n^{-\gamma+1}$ and in expectation

$$\mathrm{disc}(D(t)) = \mathrm{O}\left(\gamma \log(n) \cdot \left(1 + \sqrt{\frac{m}{n} \cdot \frac{\mathrm{t^*_{hit}(G)}}{n}}\right) + \sqrt{\frac{\gamma \log(n)}{\beta} \cdot \frac{m}{n} \cdot T(G)}\right).$$

Proof. The proof is partitioned into two parts. First we bound the discrepancy in terms of global divergence. Then we bound the global divergence by the quantity $T(G)$ defined in the statement.

Let $\mathbf{m}^{[t]}$ be an arbitrary but fixed sequence of matchings generated by $\mathcal{D}_{\mathrm{RM}}(G)$ and let $\gamma > 0$. From Lemma 2.4 it follows, with a probability at least $1 - 2n^{-\gamma+1}$, that

$$\mathrm{disc}(D(t)) = \mathrm{O}\left(\gamma \log(n) + \sqrt{\gamma \log(n) \cdot \frac{m}{n}} \cdot \Upsilon(\mathbf{m}^{[t]})\right). \tag{12}$$

Since the statement holds for any arbitrary and fixed matching sequence $\mathbf{m}^{[t]}$ generated by $\mathcal{D}_{\mathrm{RM}}(G)$, then it holds for any sequence of random matchings $\mathbf{M}^{[t]}$ generated by $\mathcal{D}_{\mathrm{RM}}(G)$ as well.

From Lemma 2.24 (shown in technical lemmas) it follows, with probability at least $1 - n^{-\gamma}$, that

$$\Upsilon(\mathbf{M}^{[t]}) = \mathrm{O}\left(\sqrt{\gamma \log(n) \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}} + \sqrt{\frac{T(G)}{\beta}}\right). \tag{13}$$

Applying union bound over Equation (12) and Equation (13), gives us that with probability at least $1 - 3n^{-\gamma+1}$

$$\begin{aligned}\mathrm{disc}(D(t)) &= \mathrm{O}\left(\gamma \log(n) + \sqrt{\gamma \log(n) \cdot \frac{m}{n}} \cdot \left(\sqrt{\gamma \log(n) \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}} + \sqrt{\frac{T(G)}{\beta}}\right)\right) \\ &= \mathrm{O}\left(\gamma \log(n) \cdot \left(1 + \sqrt{\frac{m}{n} \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}}\right) + \sqrt{\frac{\gamma \log(n)}{\beta} \cdot \frac{m}{n} \cdot T(G)}\right).\end{aligned}$$

The corresponding bound on $\mathbf{E}[\mathrm{disc}(D(t)]$ follows by applying the bound presented in Lemma 6.2. □

2.5 Asynchronous Model

In this section, we analyze the asynchronous process. The following theorem presents our main result for the asynchronous model. The bounds provided by Theorem 2.10 for the asynchronous model differ from those in Theorem 2.6 for the random matching model in two respects. First, the lower bound on the smoothing time is larger by a factor of n, reflecting the fact that the asynchronous model balances across only one edge per round, in contrast to $\Theta(n)$ edges in the random matching model. Second, the upper bound on $\mathrm{disc}(X(t))$ is considerably simpler. Note, however, that setting $m = n$ in Theorem 2.6 and further simplifying the result using $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})/n = \Omega(1)$ (see also Claim 2.26 used in the proof of Proposition 2.9) yields the same asymptotic bound as in Theorem 2.10.

Theorem 2.10. *Let G be a d-regular graph and define $T(G) := \min\left\{\frac{\mathrm{t}_{\mathrm{hit}}(\mathrm{G})}{n} \cdot \log(n), \sqrt{\frac{d}{\lambda(\mathbf{L}(G))}}, \frac{1}{\lambda(\mathbf{L}(G))}\right\}$. Let $X(t)$ be the state of process $\mathrm{ABAL}(\mathcal{D}_1(G), \beta)$ at time t with $\mathrm{disc}(X(0)) =: K \geq 1$. There exists a constant $c > 0$ such that for all $t \geq c \cdot n \cdot \log(K \cdot n)/(\lambda(\mathbf{L}(G)) \cdot \beta)$ it holds w.h.p. and in expectation*

$$\mathrm{disc}(X(t)) = \mathrm{O}\left(\log(n)\sqrt{\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}} + \sqrt{\frac{\log(n)}{\beta} \cdot T(G)}\right).$$

Proof. The proof follows the approach of Theorem 2.1, using the fact that

$$\mathrm{disc}(X(t)) \leq \mathrm{disc}(I(t)) + \mathrm{disc}(D(t)) + \mathrm{disc}(R(t)),$$

where $I(t)$, $D(t)$, and $R(t)$ denote the contributions of the initial load, dynamically allocated load, and rounding errors, respectively. Proposition 2.8, which bounds $\mathrm{disc}(R(t))$, also applies in the asynchronous model, as it only requires that the subgraph used for balancing is a matching. The contribution of $\mathrm{disc}(I(t))$ is bounded in Proposition 2.11 (replacing Proposition 2.7, see below). However, the proof of Lemma 2.4 (which bounds $\mathrm{disc}(D(t))$ in terms of the global divergence) and, most importantly, the proof of Proposition 2.9 (which provides a concentration bound on $\mathrm{disc}(D(t))$) cannot be directly applied to ABal. We therefore replace Proposition 2.9 with Proposition 2.12. A final union bound over Proposition 2.8, Proposition 2.11, and Proposition 2.12 completes the proof. The result concerning the expectation follows similarly from Lemma 6.2. □

2.5.1 Bounding the Contribution of the Initial Load $I(t)$

The following lemma is a version of Proposition 2.7 adapted to this model. Its proof requires recalculating the functions g_G and σ_G to demonstrate that the matching distribution used by $\mathcal{D}_1(G)$ is (g_G, σ_G^2)-good (Lemma 2.20). For the precise definition of goodness, see Definition 2.1 in the technical lemmas section.

Proposition 2.11 (Memorylessness Property)**.** *Let G be a d-regular graph. Let $K = \mathrm{disc}(X(0))$. Then there exists a constant $c > 0$ such that for all $\gamma > 0$ and $t \in \mathbb{N}$ with $t \geq t_0(\gamma) := cn \cdot \max\{\gamma \log(n), \log(K \cdot n)\}/(\lambda(\mathbf{L}(G)) \cdot \beta)$ we get with probability at least $1 - n^{-\gamma}$ and in expectation $\mathrm{disc}(I(t)) \leq 1$.*

Proof. The proof follows the same approach as that of Proposition 2.7, with the exception that in Equation (25) and Equation (26), we replace Lemma 2.17 with Lemma 2.21, which is presented in Section 2.6. In essence, to apply Theorem 6.1 (the Drift Theorem) and derive a bound analogous to that in Proposition 2.7, we must show that the sequence of matchings in this model is (g_G, σ_G^2)-good for specific functions g_G and σ_G^2. This is established in Lemma 2.21. □

2.5.2 Bounding the Contribution of the Dynamically Allocated Load $D(t)$

Here we bound the contribution of dynamically allocated load to the discrepancy.

Proposition 2.12 (Contribution of Dynamically Allocated Load)**.** *Let G be a d-regular graph. Define*

$$T(G) := \min\Big\{\mathrm{t_{hit}(G)} \cdot \log(n)/n, \sqrt{d/\lambda(\mathbf{L}(G))}, 1/\lambda(\mathbf{L}(G))\Big\}.$$

Then for all $\gamma > 1$ and $t \in \mathbb{N}$ we get with probability at least $1 - 9n^{-\gamma+1}$ and in expectation

$$\mathrm{disc}(D(t)) = \mathrm{O}\left(\log(n)\sqrt{\frac{\mathrm{t^*_{hit}(G)}}{n}} + \sqrt{\frac{\log(n)}{\beta} \cdot T(G)}\right).$$

Proof. In Lemma 2.27 (presented in the technical lemmas), we bound the discrepancy in terms of the global divergence. In Lemma 2.29 (also presented in the technical lemmas), we bound the global divergence in terms of $T(G)$. Applying a union bound over these two results completes the proof. □

2.6 Technical Lemmas

This section is divided into two subsections. In the first subsection (2.6.1), we present all the basic lemmas. In the second subsection (2.6.2), we use these basic lemmas to prove the developed tools that are subsequently employed to establish our main propositions.

2.6.1 Basic Tools

We start with a well-known fact that the discrepancy is sub-additive.

Observation 2.13. *For two vectors $\vec{x}, \vec{y} \in \mathbb{R}^n$,*

$$\mathrm{disc}(\vec{x} + \vec{y}) \leq \mathrm{disc}(\vec{x}) + \mathrm{disc}(\vec{y}).$$

Proof. For any $\vec{a}, \vec{b} \in \mathbb{R}^n$,

$$\max_{i \in [n]}(a_i + b_i) \leq \max_{i \in [n]} a_i + \max_{i \in [n]} b_i,$$

and thus

$$\begin{aligned}
\mathrm{disc}(\vec{x} + \vec{y}) &= \max_{i \in [n]}(x_i + y_i) - \min_{i \in [n]}(x_i + y_i) = \max_{i \in [n]}(x_i + y_i) + \max_{i \in [n]}((-x_i) + (-y_i)) \\
&\leq \max_{i \in [n]} x_i + \max_{i \in [n]} y_i + \max_{i \in [n]}(-x_i) + \max_{i \in [n]}(-y_i) \\
&= \left(\max_{i \in [n]} x_i - \min_{i \in [n]} x_i\right) + \left(\max_{i \in [n]} y_i - \min_{i \in [n]} y_i\right) \\
&= \mathrm{disc}(\vec{x}) + \mathrm{disc}(\vec{y}),
\end{aligned}$$

as claimed. □

Next observation is an immediate result of it.

Observation 2.14. *It holds that* $\mathrm{disc}(D(t)) \leq 2 \cdot \max_{k \in [n]} |D_k(t) - t \cdot m/n|$.

To bound the global divergence of the matching sequence used by the process we use two potential functions. Recall that the quadratic node potential $\Phi(\vec{x})$ is given by

$$\Phi(\vec{x}) = \sum_{i \in [n]} (x_i - \overline{x})^2, \quad \text{where} \quad \overline{x} = \frac{1}{n} \cdot \sum_{j \in [n]} x_j.$$

For a set of edges S on the nodes $[n]$ and a vector $\vec{x} \in \mathbb{R}^n$, the *quadratic edge potential* is

$$\Psi_S(\vec{x}) \coloneqq \sum_{\{i,j\} \in S} (x_i - x_j)^2.$$

We may also write $\Psi_G \coloneqq \Psi_{E(G)}$ whenever G is a graph, and $\Psi_{\mathbf{M}} \coloneqq \Psi_{E(\mathbf{M})}$ whenever $\mathbf{M}$ is a matching matrix. The following observation relates the drop of node potential to the edge potential in terms of β.

Observation 2.15. *Let $\mathbf{M}^\beta$ be a matching matrix with parameter $\beta \in (0, 1]$. Then for any $\vec{x} \in \mathbb{R}^n$ we have $\Phi(\vec{x}) - \Phi(\mathbf{M}^\beta \cdot \vec{x}) = (1 - (1 - \beta)^2)/2 \cdot \Psi_{E(\mathbf{M}^\beta)}(\vec{x})$.*

Proof. We assume w.l.o.g. that the entries of $\vec{x}$ sum to 0, meaning that $\overline{x} = 0$, so that $\Phi(\vec{x}) = \sum_{i \in [n]} x_i^2$. As loads only change at matched nodes, let us investigate the potential change at two matched nodes i and j, where w.l.o.g. $x_i \geq x_j$. The amount of load transferred from i to j under idealized balancing (without rounding) is $(x_i - x_j) \cdot \beta/2$. So with

$$a \coloneqq \frac{x_i + x_j}{2}, \quad b \coloneqq \frac{x_i - x_j}{2}, \quad c \coloneqq (1 - \beta) \cdot \frac{x_i - x_j}{2},$$

the loads before balancing are $x_i = a + b$ and $x_j = a - b$, and the loads after idealized balancing are $x_i' = a + c$ and $x_j' = a - c$. So the change of the potential contributions at i and v is

$$(a+b)^2 + (a-b)^2 - ((a+c)^2 + (a-c)^2) = 2(a^2 + b^2) - 2(a^2 + c^2) = 2(b^2 - c^2),$$

where we used $(x+y)^2 + (x-y)^2 = (x^2 + 2xy + y^2) + (x^2 - 2xy + y^2) = 2x^2 + 2y^2$. Now,

$$2(b^2 - c^2) = 2(1^2 - (1-\beta)^2)\left(\frac{x_i - x_j}{2}\right)^2 = \frac{1 - (1-\beta)^2}{2}(x_i - x_j)^2.$$

Summing this over all edges in the matching gives, as claimed,

$$\Phi(\vec{x}) - \Phi(\mathbf{M}^\beta \cdot \vec{x}) = \frac{1 - (1-\beta)^2}{2} \sum_{[i,j] \in \mathbf{M}^\beta} (x_i - x_j)^2 = \frac{1 - (1-\beta)^2}{2} \cdot \Psi_{E(\mathbf{M}^\beta)}(\vec{x}). \qquad \square$$

We now define a notion of a matching distribution being *good*. In Lemma 2.19 below we show that the notion is sufficient for showing that matching sequences generated from such distributions have bounded global divergence. Note that the "goodness" of a distribution does not depend on β but on graph properties and the random choices with which the matchings are chosen. Hence, we assume $\beta = 1$.

Definition 2.1. *Assume G is an arbitrary d-regular graph. Let $g\colon \mathbb{R}_0^+ \to \mathbb{R}^+$ be an increasing function and let $\sigma^2 > 1$. Then a matching distribution $\mathcal{D}_{\mathrm{RM}}(G)$ is* (g, σ^2)-good *if the following conditions hold for $\mathbf{M}^1 \sim \mathcal{D}_{\mathrm{RM}}(G)$ and all stochastic vectors $\vec{x} \in \mathbb{R}^n$.*

1. $\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] \ge g(\Phi(\vec{x}))$.
2. *Var*$\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] \le (\sigma^2 - 1) \cdot \big(\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big]\big)^2$.

In the next lemma, we calculate a function g_G and the values of σ_G for which the matching distribution $\mathcal{D}_{\mathrm{RM}}(G)$ is (g_G, σ_G^2)-good. This lemma together with Lemma 2.19 allows us to bound the global divergence of random matching model (see Lemma 2.24).

Lemma 2.16. *Assume G is an arbitrary d-regular graph. Let*

$$g_G(x) := \frac{1}{16d} \cdot \max\left\{d \cdot \lambda(\mathbf{L}(G)) \cdot x, \frac{x^2}{\mathrm{Res}(G)}, \frac{4}{27} \cdot x^3\right\} \text{ and } \sigma_G^2 = 32 \cdot (\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})/n) + 5.$$

Then $\mathcal{D}_{\mathrm{RM}}(G)$ is (g_G, σ_G^2)-good.

Proof. First, note that the function $g_G(x)$ is increasing in x. Applying the first part of Lemma 2.17 (see below) we get that for any vector $\vec{x} \in \mathbb{R}^n$ it holds that

$$\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] \ge \frac{1}{16d} \cdot \Psi_G(\vec{x}).$$

From the first two statements of Lemma 2.18 (stated behind Lemma 2.18) we see that for $\mathbf{M}^1 \sim \mathcal{D}_{\mathrm{RM}}(G)$ and all stochastic vectors $\vec{x} \in \mathbb{R}^n$

$$\Psi_G(\vec{x}) \ge \max\left\{d \cdot \lambda(\mathbf{L}(G)) \cdot \Phi(\vec{x}), \frac{\Phi(\vec{x})^2}{\mathrm{Res}(G)}, \frac{4}{27} \cdot \Phi(\vec{x})^3\right\}.$$

Hence,

$$\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] \ge \frac{1}{16d} \cdot \max\left\{d \cdot \lambda(\mathbf{L}(G)) \cdot \Phi(\vec{x}), \frac{\Phi(\vec{x})^2}{\mathrm{Res}(G)}, \frac{4}{27} \cdot \Phi(\vec{x})^3\right\},$$

and as a consequence, $\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] \geq g_G(\Phi(\vec{x}))$ by the definition of g_G.

It remains to check the second condition of Definition 2.1 with our claimed value σ_G^2. Inserting its value as stated in the lemma, the condition requires that

$$\mathrm{Var}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] \leq (32(\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})\,/n) + 5 - 1) \cdot \big(\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big]\big)^2,$$

which is given in the second part of Lemma 2.17 (see below). □

In Lemma 2.17 we first relate the drop of Φ to the quadratic edge potential Ψ. In the second part we bound the variance of the potential drop as a function of the edge hitting time.

Lemma 2.17. *Let G be a d-regular graph, let $\mathbf{M}^1 \sim \mathcal{D}_{\mathrm{RM}}(G)$, and let $\vec{x} \in \mathbb{R}^n$, then*

1. $\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] \geq \frac{1}{16d} \cdot \Psi_G(\vec{x})$.
2. $Var\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] \leq (32 \cdot (\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}) + 4) \cdot \big(\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big]\big)^2$.

Proof. By observation 2.15, we have

$$\Phi(\vec{x}) - \Phi(\mathbf{M}^1 \cdot \vec{x}) = \frac{1-(1-1)^2}{2} \cdot \Psi_{E(\mathbf{M}^1)}(\vec{x}).$$

Rearranging this lower bound into $\Phi(\mathbf{M}^1 \cdot \vec{x}) = \Phi(\vec{x}) - \frac{1}{2} \cdot \Psi_{E(\mathbf{M}^1)}$, and expanding the definition of $\Psi_{E(\mathbf{M}^1)}$ we have by linearity of expectation

$$\begin{aligned}
\mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] &= \mathbf{E}\left[\Phi(\vec{x}) - \frac{1}{2} \cdot \sum_{[i,j] \in \mathbf{M}^1} (x_i - x_j)^2\right] \\
&= \Phi(\vec{x}) - \frac{1}{2} \cdot \sum_{\{i,j\} \in E(G)} \mathbf{E}\big[\mathbf{1}_{\{i,j\} \in E(\mathbf{M}^1)} \cdot (x_i - x_j)^2\big] \\
&= \Phi(\vec{x}) - \frac{1}{2} \cdot \sum_{\{i,j\} \in E(G)} \mathbf{Pr}\big[\{i,j\} \in E(\mathbf{M}^1)\big] \cdot (x_i - x_j)^{\cdot}
\end{aligned} \tag{14}$$

From [51, Lemma 2], it follows that for $\mathbf{M}^1 \sim \mathcal{D}_{\mathrm{RM}}(G)$ and all edges $e \in E(G)$, it holds that $\mathbf{Pr}\big[e \in E(\mathbf{M}^1)\big] \geq 1/(8d)$. Applying to Equation (14) gives us

$$\mathbf{E}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] \leq \Phi(\vec{x}) - \frac{1}{2} \cdot \sum_{\{i,j\} \in E(G)} \frac{1}{8d} \cdot (x_i - x_j)^2 = \Phi(\vec{x}) - \frac{1}{16d} \cdot \Psi_G(\vec{x}),$$

finishing the proof of the first statement.

For the second statement observe that by observation 2.15 we have

$$\Phi(\vec{x}) - \Phi(\mathbf{M}^1 \cdot \vec{x}) = \frac{1}{2} \cdot \Psi_{E(\mathbf{M}^1)}(\vec{x})$$

Then, as $\Phi(\vec{x})$ is constant for a given $\vec{x}$,

$$\mathrm{Var}\big[\Phi(\mathbf{M}^1 \cdot \vec{x})\big] = \mathrm{Var}\big[\Phi(\vec{x}) - \Phi(\mathbf{M}^1 \cdot \vec{x})\big] = \mathrm{Var}\left[\frac{1}{2} \cdot \Psi_{\mathbf{M}^1}(\vec{x})\right] = \frac{1}{4}\mathrm{Var}[\Psi_{\mathbf{M}^1}(\vec{x})]. \tag{15}$$

Recall that the matching distribution $\mathcal{D}_{\mathrm{RM}}(G)$ is obtained as follows. First, generate a random edge set S as follows. For each $e \in E(G)$, $e \in S$ with probability $p_{\max} := \mathbf{Pr}[e \in S] = 1/(4d) - 1/(64d^2) \leq 1/(4d)$, independently of all other edges. Then, some edges of S are deleted to create a proper matching, resulting in $E(\mathbf{M}^1) \subseteq S$. Hence

$$0 \leq \Psi_{E(\mathbf{M}^1)}(\vec{x}) = \sum_{[i,j] \in \mathbf{M}^1} (x_i - x_j)^2 \leq \sum_{\{i,j\} \in S} (x_i - x_j)^2 = \Psi_S(\vec{x}),$$

and

$$\operatorname{Var}\left[\Psi_{E(\mathbf{M}^1)}(\vec{x})\right] \leq \mathbf{E}\left[(\Psi_{E(\mathbf{M}^1)}(\vec{x}))^2\right] \leq \mathbf{E}\left[(\Psi_S(\vec{x}))^2\right] = \operatorname{Var}[\Psi_S(\vec{x})] + (\mathbf{E}[\Psi_S(\vec{x})])^2. \tag{16}$$

Observe that $\Psi_S(\vec{x})$ can be expressed as $\Psi_S(\vec{x}) = \sum_{\{i,j\} \in E(G)} (x_i - x_j)^2 \, \mathbf{1}_{\{i,j\} \in S}$ with $\mathbf{Pr}\left[\mathbf{1}_{\{i,j\} \in S} = 1\right] = p_{\max}$. Thus,

$$\begin{aligned}
\mathbf{E}[\Psi_S(\vec{x})] &= \sum_{\{i,j\} \in E} (x_i - x_j)^2 \cdot \mathbf{E}\left[\mathbf{1}_{\{i,j\} \in S}\right] = p_{\max} \cdot \sum_{\{i,j\} \in E} (x_i - x_j)^2 = p_{\max} \cdot \Psi_G(\vec{x}); \\
\operatorname{Var}[\Psi_S(\vec{x})] &= \sum_{\{i,j\} \in E} (x_i - x_j)^4 \cdot \operatorname{Var}\left[\mathbf{1}_{\{i,j\} \in S}\right] = \sum_{\{i,j\} \in E} (x_i - x_j)^4 \cdot p_{\max}(1 - p_{\max}) \\
&= p_{\max}(1 - p_{\max}) \cdot \sum_{\{i,j\} \in E} (x_i - x_j)^4 \leq p_{\max} \cdot \sum_{\{i,j\} \in E} (x_i - x_j)^2 \cdot \max_{\{k,l\} \in E} (x_k - x_l)^2 \\
&\leq p_{\max} \cdot \Psi_G(\vec{x}) \cdot \max_{\{k,l\} \in E} (x_k - x_l)^2.
\end{aligned}$$

By using Lemma 2.18(3) and the first statement of Lemma 6.21 we get that $\max_{\{k,l\} \in E} (x_i - x_j)^2 \leq \operatorname{Res}^*(G) \cdot \Psi_G(\vec{x}) \leq \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{|E|} \cdot \Psi_G(\vec{x})$. Hence,

$$\begin{aligned}
\operatorname{Var}\left[\Psi_{E(\mathbf{M}^1)}(\vec{x})\right] &\overset{\text{Eq. (16)}}{\leq} \operatorname{Var}[\Psi_S(\vec{x})] + (\mathbf{E}[\Psi_S(\vec{x})])^2 \\
&\leq p_{\max} \cdot \Psi_G(\vec{x}) \cdot \max_{\{k,l\} \in E} (x_k - x_l)^2 + (p_{\max} \cdot \Psi_G(\vec{x}))^2 \\
&\leq p_{\max} \cdot \Psi_G(\vec{x}) \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{|E|} \cdot \Psi_G(\vec{x}) + p_{\max}^2 \cdot (\Psi_G(\vec{x}))^2 \\
&\leq \frac{1}{4d} \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{dn/2} \cdot (\Psi_G(\vec{x}))^2 + \frac{1}{16d^2} \cdot (\Psi_G(\vec{x}))^2 = \frac{1}{2d^2} \cdot \left(\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n} + \frac{1}{8} \right) \cdot \Psi_G(\vec{x})^2.
\end{aligned} \tag{17}$$

Applying the first statement of this lemma we get

$$\Psi_G(\vec{x}) \leq 16d \cdot \left(\Phi(\vec{x}) - \mathbf{E}\left[\Phi(\mathbf{M}^1 \cdot \vec{x})\right]\right). \tag{18}$$

Putting everything together the second statement follows from

$$\begin{aligned}
\operatorname{Var}\left[\Phi(\mathbf{M}^1 \cdot \vec{x})\right] &\overset{\text{Eq. (15)}}{\leq} \frac{1}{4} \cdot \operatorname{Var}[\Psi_{\mathbf{M}^1}(\vec{x})] \overset{\text{Eq. (17)}}{\leq} \frac{1}{4} \cdot \frac{1}{2d^2} \cdot \left(\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{s} + \frac{1}{8} \right) \cdot (\Psi_G(\vec{x}))^2 \\
&\overset{\text{Eq. (18)}}{\leq} \frac{1}{8d^2} \cdot \left(\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n} + \frac{1}{8} \right) \cdot \left(16d \cdot \left(\Phi(\vec{x}) - \mathbf{E}\left[\Phi(\mathbf{M}^1 \cdot \vec{x})\right]\right)\right)^2 \\
&= 32 \cdot \left(\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n} + \frac{1}{8} \right) \cdot \left(\Phi(\vec{x}) - \mathbf{E}\left[\Phi(\mathbf{M}^1 \cdot \vec{x})\right]\right)^2 \\
&\leq \left(32 \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n} + 4 \right) \cdot \left(\Phi(\vec{x}) - \mathbf{E}\left[\Phi(\mathbf{M}^1 \cdot \vec{x})\right]\right)^2 \qquad \square
\end{aligned}$$

In Lemma 2.18 we relate the size of the quadratic edge potential Ψ_G to the second-largest eigenvalue of $\mathbf{L}(G)$, the effective resistance of G and node potential. To state it, we need some additional definitions.

For any two nodes i and j of the graph G $\mathrm{Res}(i,j)$ is the *effective resistance* (or *resistive distance*) between i and j in G (for a detailed definition see Section 6.3). Furthermore, we write $\mathrm{Res}(G)$ for the *resistive diameter* of G, i.e., the largest resistive distance between any pair of nodes in G, and write $\mathrm{Res}^*(G)$ for the maximum effective resistance between any pair of nodes adjacent in G. I.e., $\mathrm{Res}(G) \coloneqq \max_{i,j\in[n]} \mathrm{Res}(i,j)$ and $\mathrm{Res}^*(G) \coloneqq \max_{\{i,j\}\in E(G)} \mathrm{Res}(i,j)$. The first part of the following lemma was previously shown in [51, 85].

Lemma 2.18. *Let $\vec{x} \in \mathbb{R}^n$, and let G be a connected d-regular graph.*

1. $\Psi_G(\vec{x}) \geq d \cdot \lambda(\mathbf{L}(G)) \cdot \Phi(\vec{x})$.
2. *If $\vec{x}$ is stochastic, then* $\Psi_G(\vec{x}) \geq \max\left\{\frac{1}{\mathrm{Res}(G)} \cdot \Phi(\vec{x})^2, \frac{4}{27} \cdot \Phi(\vec{x})^3\right\}$
3. $\max_{\{i,j\}\in E(G)}(x_i - x_j)^2 \leq \mathrm{Res}^*(G) \cdot \Psi_G(\vec{x})$.

Proof. First note that for all $\vec{x} \in \mathbb{R}^n$, $a, b \in \mathbb{R}$, and $S \subseteq E(G)$,

$$\begin{aligned}\Psi_S(a\cdot\vec{x}+b) &= \sum_{\{i,j\}\in S}((a\cdot x_i+b)-(a\cdot x_j+b))^2 = \sum_{\{i,j\}\in S}(a\cdot x_i+b-a\cdot x_j-b)^2 \\ &= \sum_{\{i,j\}\in S} a^2(x_i-x_j)^2 = a^2\,\Psi(\vec{x}).\end{aligned} \tag{19}$$

The proof of the first part is similar to that of Theorem 2.6 in [85]. First, see that

$$\begin{aligned}\Psi_G(\vec{x}) &= \sum_{\{i,j\}\in E(G)}(x_i-x_j)^2 = \sum_{\{i,j\}\in E(G)}(x_i^2-2x_ix_j+x_j^2) \\ &= \sum_{i\in[n]} d\cdot x_i^2 - \sum_{i,j\in[n]}\mathbf{A}_{i,j}\,x_ix_j = d\cdot\langle\vec{x},\vec{x}\rangle - \sum_{i\in[n]}x_i\left(\sum_{j\in[n]}\mathbf{A}_{i,j}\,x_j\right) \\ &= d\cdot\langle\vec{x},\mathbf{I}\,\vec{x}\rangle - \langle\vec{x},\mathbf{A}\,\vec{x}\rangle = d\cdot\langle\vec{x},(\mathbf{I}-\mathbf{A}\,/d)\vec{x}\rangle = d\cdot\langle\vec{x},\mathbf{L}(G)\vec{x}\rangle.\end{aligned}$$

As $\Psi_G(\vec{x}-b) = \Psi_G(\vec{x})$ by observation 19, we may assume w.l.o.g. that $\langle\vec{x},\vec{1}\rangle = 0$ by subtracting $b \coloneqq \langle\vec{x},\vec{1}\rangle/n$ from every coordinate of $\vec{x}$. For such a vector we have $\Phi(\vec{x}) = ||x||_2^2 = \langle\vec{x},\vec{x}\rangle$, and

$$\begin{aligned}\Psi_G(\vec{x}) &= d\langle\vec{x},\mathbf{L}(G)\vec{x}\rangle = d\cdot\frac{\langle\vec{x},\mathbf{L}(G)\vec{x}\rangle}{\langle\vec{x},\vec{x}\rangle}\cdot\Phi(\vec{x}) \geq d\cdot\Phi(\vec{x})\cdot\min_{\substack{\vec{a}\in\mathbb{R}^n\setminus\{\vec{0}\}\\ \langle\vec{a},\vec{1}\rangle=0}}\frac{\langle\vec{a},\mathbf{L}(G)\vec{a}\rangle}{\langle\vec{a},\vec{a}\rangle} \\ &= d\cdot\lambda(\mathbf{L}(G))\cdot\Phi(\vec{x}),\end{aligned}$$

where the final equality is due to the min-max theorem and the fact that the smallest eigenvalue of $\mathbf{L}(G)$ is 0, with its associated eigenvector being $\vec{1}$.

For the second part, let $i, j \in [n]$ be two distinct nodes of the graph with $x_i \neq x_j$. Then

$$\Psi_G(\vec{x}) = (x_i-x_j)^2\cdot\Psi_G\left(\frac{\vec{x}-x_j}{x_i-x_j}\right) \geq (x_i-x_j)^2\cdot\min_{\substack{\vec{a}\in\mathbb{R}^n\\ a_i=1\\ a_j=0}}\Psi_G(\vec{a}) = \frac{(x_i-x_j)^2}{\mathrm{Res}(i,j)}, \tag{20}$$

where the first equality uses observation 19, the central inequality holds because the argument of Ψ_G is a vector $\vec{a} \in \mathbb{R}^n$ with $a_i = 1$ and $a_j = 0$, and the final equality is by Dirichlet's principle (Theorem 6.19). Note that the bound also holds when $x_i = x_j$.

Given Equation (20), we now show that $\Psi_G(\vec{x})$ is larger than the first, resp. second, term inside the maximum of the second part's statement. For the first term, we choose i and j such that $x_i - x_j = \operatorname{disc}(\vec{x})$, and recall that $\operatorname{Res}(i,j) \le \operatorname{Res}(G)$ for all $i,j \in [n]$. Then, Equation (20) states that $\Psi_G(\vec{x}) \ge \operatorname{disc}(\vec{x})^2/\operatorname{Res}(G)$, and it remains to bound $\operatorname{disc}(\vec{x})$ from below by $\Phi(\vec{x})$. To that end, as the vector $\vec{x}$ is stochastic by assumption, the sum over all its entries is 1, and there is at least one $k \in [n]$ with $x_k \le 1/n$. Hence, $\operatorname{disc}(\vec{x}) \ge \max_{k\in[n]}(x_k - 1/n)$, and so

$$\operatorname{disc}(\vec{x}) \ge \operatorname{disc}(\vec{x}) \cdot \underbrace{\sum_{k\in[n]} x_k}_{=1} \ge \sum_{k\in[n]} \underbrace{\left(x_k - \frac{1}{n}\right)}_{\le \operatorname{disc}(\vec{x})} x_k - \frac{1}{n} \cdot \underbrace{\sum_{k\in[n]} \left(x_k - \frac{1}{n}\right)}_{=\frac{1}{n}\cdot 0 = 0} = \sum_{k\in[n]} \left(x_k - \frac{1}{n}\right)^2 = \Phi(\vec{x}),$$

as needed to complete the bound for the first term.

For the second term, we choose i and j such that $x_i = \max_{k\in[n]} x_k$, $x_j \le x_i - 2/3 \cdot \operatorname{disc}(\vec{x})$ with the distance D between i and j being minimal. As $x_i \ge \operatorname{disc}(\vec{x})$, each of the entries of $\vec{x}$ for the $D-1$ non-terminal nodes on a shortest path between i and j is at least $\operatorname{disc}(\vec{x})/3$. As $\vec{x}$ is stochastic by assumption, the sum of all loads is at most 1, and we have

$$\operatorname{disc}(\vec{x}) + (D-1) \cdot \frac{\operatorname{disc}(\vec{x})}{3} = \frac{D+2}{3} \cdot \operatorname{disc}(\vec{x}) \le 1,$$

which implies $D \le 3/\operatorname{disc}(\vec{x})$.

Since $\operatorname{Res}(i,j)$ is bounded by the standard distance between i and j (see Lemma 6.16), and $x_i - x_j \ge 2/3 \cdot \operatorname{disc}(\vec{x})$, we thus have, by Equation (20),

$$\Psi_G(\vec{x}) \ge \frac{(x_j - x_i)^2}{\operatorname{Res}(i,j)} \ge \frac{(2/3 \cdot \operatorname{disc}(\vec{x}))^2}{3/\operatorname{disc}(\vec{x})} = \frac{4 \cdot \operatorname{disc}(\vec{x})^3}{27} \ge \frac{4 \cdot \Phi(\vec{x})}{27},$$

where the final inequality uses $\operatorname{disc}(\vec{x}) \ge \Phi(\vec{x})$ as shown above.

For the third statement we first rearrange Equation (20) to see that, for all $i \ne j$,

$$(x_i - x_j)^2 \le \Psi_G(\vec{x}) \cdot \operatorname{Res}(i,j).$$

Taking the maximum over all $\{i,j\} \in E(G)$ on both sides gives us

$$\max_{\{i,j\}\in E(G)} (x_j - x_i)^2 \le \Psi_G(\vec{x}) \cdot \max_{\{i,j\}\in E(G)} \operatorname{Res}(i,j) = \Psi_G(\vec{x}) \cdot \operatorname{Res}^*(G),$$

as claimed, where the final equality is by definition of $\operatorname{Res}^*(G)$. □

In the next lemma we show assuming a matching distribution is (g,σ^2)-good, the global divergence of a matching sequence generated by that distribution can be bounded in terms of g and σ.

Lemma 2.19 (Bounding Global Divergence with Goodness). *Assume G is an arbitrary graph. Let $g\colon \mathbb{R}_0^+ \to \mathbb{R}^+$ be an increasing function, $\sigma^2 > 1$, and $\beta \in (0,1]$. Let $\mathbf{M}^{[t]} = (\mathbf{M}^{\beta}(\tau))_{\tau=1}^{t}$ be an i.i.d. sequence of matching matrices generated by $\mathcal{D}_{\mathrm{RM}}(G)$ and assume $\mathcal{D}_{\mathrm{RM}}(G)$ is a (g,σ^2)-good matching distribution. Then for all $\gamma > 0$ and $k \in [n]$ we get with probability at least $1 - n^{-\gamma}$*

$$\left(\Upsilon_k(\mathbf{M}^{[t]})\right)^2 \le 8\sigma^2(\gamma \log(n) + \log(8\sigma^2)) + \frac{2}{\beta} \cdot \int_0^1 \frac{x}{g(x)}\, \mathrm{d}x.$$

Proof. First recall that

$$\left(\Upsilon_k(\mathbf{M}^{[t]})\right)^2 = \sum_{\tau=1}^{t} \left\| \mathbf{M}_{k,\cdot}^{[\tau,t]} - \frac{\vec{1}}{n} \right\|_2^2 .$$

As the mixing matrices are doubly stochastic, each row is a stochastic vector $\vec{x}$. By definition of the node potential Φ we know

$$\left\| \mathbf{M}_{k,\cdot}^{[\tau,t]} - \frac{\vec{1}}{n} \right\|_2^2 = \sum_{w=1}^{n} \left(\mathbf{M}_{k,w}^{[\tau,t]} - \frac{1}{n} \right)^2 = \Phi\left(\mathbf{M}_{k,\cdot}^{[\tau,t]} \right)$$

and hence

$$\left(\Upsilon_k(\mathbf{M}^{[t]})\right)^2 = \sum_{\tau=1}^{t} \Phi(\mathbf{M}_{k,\cdot}^{[\tau,t]}).$$

To bound this sum we will apply the second statement of Theorem 6.1 to the sequence of values $\Phi(\mathbf{M}^{[\tau,t]})$ for $\tau = t, \dots, 1$. Since the matching matrices $\mathbf{M}^{\beta}(1) \dots, \mathbf{M}^{\beta}(t)$ are symmetric we get

$$\Phi(\mathbf{M}_{k,\cdot}^{[\tau,t]}) = \Phi\left(\mathbf{M}_{k,\cdot}^{[\tau+1,t]} \cdot \mathbf{M}^{\beta}(\tau) \right) = \Phi\left(\mathbf{M}^{\beta}(\tau) \cdot \mathbf{M}_{k,\cdot}^{[\tau+1,t]} \right).$$

By observation 2.15 with $S = E(\mathbf{M}^{\beta}(\tau))$ defined as the edges of $\mathbf{M}^{\beta}(\tau)$ we get

$$\Phi(\mathbf{M}_{k,\cdot}^{[\tau+1,t]}) - \Phi(\mathbf{M}_{k,\cdot}^{[\tau,t]}) = \frac{1-(1-\beta)^2}{2} \cdot \Psi_S(\mathbf{M}_{k,\cdot}^{[\tau+1,t]}) \geq 0. \tag{21}$$

This shows that $\Phi(\mathbf{M}_{k,\cdot}^{[\tau,t]}) \leq \Phi(\mathbf{M}_{k,\cdot}^{[\tau+1,t]})$ for all τ. Expressing Equation (21) with balancing parameter 1 and, for the ease of presentation, setting $\vec{V} := \mathbf{M}_{k,\cdot}^{[\tau+1,t]}$ gives us

$$\Phi(\mathbf{M}_{k,\cdot}^{[\tau+1,t]}) - \Phi(\mathbf{M}_{k,\cdot}^{[\tau,t]}) = \Phi(\vec{V}) - \Phi(\mathbf{M}^{\beta}(\tau) \cdot \vec{V}) = (1-(1-\beta)^2) \cdot \left(\Phi(\vec{V}) - \Phi(\mathbf{M}^{1}(\tau) \cdot \vec{V}) \right).$$

Since $\beta \leq 1-(1-\beta)^2 \leq 2\beta$ for $\beta \in (0,1]$ we get

$$\mathbf{E}\left[\Phi(\vec{V}) - \Phi(\mathbf{M}^{\beta}(\tau) \cdot \vec{V}) \right] \geq \beta \cdot \mathbf{E}\left[\Phi(\vec{V}) - \Phi(\mathbf{M}^{1}(\tau) \cdot \vec{V}) \right], \tag{22}$$

$$\mathrm{Var}\left[\Phi(\vec{V}) - \Phi(\mathbf{M}^{\beta}(\tau) \cdot \vec{V}) \right] \leq 4\beta^2 \cdot \mathrm{Var}\left[\Phi(\vec{V}) - \Phi(\mathbf{M}^{1}(\tau) \cdot \vec{V}) \right]. \tag{23}$$

As $\mathcal{D}_{\mathrm{RM}}(G)$ is (g,σ^2)-good, for any stochastic vector $\vec{v} \in \mathbb{R}^n$ we have $\mathbf{E}\left[\Phi(\vec{v}) - \Phi(\mathbf{M}^1(\tau) \cdot \vec{v})\right] \geq g(\Phi(\vec{v}))$. Combining this with Equation (22) gives

$$\mathbf{E}\left[\Phi(\vec{v}) - \Phi(\mathbf{M}^{\beta}(\tau) \cdot \vec{v})\right] \geq \beta \cdot g(\Phi(\vec{v})).$$

And thus,

$$\mathbf{E}\left[\Phi(\mathbf{M}_{k,\cdot}^{[\tau,t]}) \,\middle|\, \Phi(\vec{V}) = \varphi \right] = \mathbf{E}\left[\Phi(\mathbf{M}^{\beta}(\tau) \cdot \vec{V}) \,\middle|\, \Phi(\vec{V}) = \varphi \right] \leq \varphi - \beta \cdot g(\varphi).$$

Similarly, as $\mathcal{D}_{\mathrm{RM}}(G)$ is (g,σ^2)-good, for any stochastic vector $\vec{v} \in \mathbb{R}^n$ we have

$$\mathrm{Var}\left[\Phi(\mathbf{M}^1 \cdot \vec{v})\right] \leq (\sigma^2 - 1) \cdot \left(\Phi(\vec{v}) - \mathbf{E}\left[\Phi(\mathbf{M}^1 \cdot \vec{v})\right]\right)^2.$$

Combining this with Equation (23) gives us

$$\mathrm{Var}\left[\Phi(\mathbf{M}^{\beta} \cdot \vec{v})\right] \leq 4\beta^2(\sigma^2 - 1)\left(\Phi(\vec{v}) - \mathbf{E}\left[\Phi(\mathbf{M}^1 \cdot \vec{v})\right]\right)^2,$$

and thus

$$\begin{aligned}\mathrm{Var}\Big[\Phi(\mathbf{M}^{[\tau,t]}_{k,\cdot}) \,\Big|\, \Phi(\vec{V}) = \varphi\Big] &= \mathrm{Var}\Big[\Phi(\mathbf{M}^{\beta}(\tau)\cdot\vec{V}) \,\Big|\, \Phi(\vec{V}) = \varphi\Big] \\ &\le 4\beta^2\cdot(\sigma^2-1)\cdot\Big(\varphi - \mathbf{E}\Big[\Phi\Big(\mathbf{M}^1(\tau)\cdot\vec{V}\Big)\,\Big|\,\Phi(\vec{V})=\varphi\Big]\Big)^2 \\ &= 4(\sigma^2-1)\cdot\Big(\beta\cdot\mathbf{E}\Big[\varphi - \Phi\Big(\mathbf{M}^1(\tau)\cdot\vec{V}\Big)\,\Big|\,\Phi(\vec{V})=\varphi\Big]\Big)^2 \\ &\overset{\text{Eq. (22)}}{\le} 4(\sigma^2-1)\cdot\Big(\mathbf{E}\Big[\varphi - \Phi\Big(\mathbf{M}^{\beta}(\tau)\cdot\vec{V}\Big)\,\Big|\,\Phi(\vec{V})=\varphi\Big]\Big)^2 \\ &= 4(\sigma^2-1)\cdot\Big(\mathbf{E}\Big[\Phi\Big(\mathbf{M}^{\beta}(\tau)\cdot\vec{V}\Big)\,\Big|\,\Phi(\vec{V})=\varphi\Big] - \varphi\Big)^2.\end{aligned}$$

We apply the second statement of Theorem 6.1 with $p := n^{-\gamma}$, $\delta := 0.5$, and $h(x) := \beta\cdot g(x)$, which is an increasing function as g is increasing by the definition of (g,σ^2)-good, and get

$$\mathbf{Pr}\left[\sum_{\tau=1}^{t-t_0}\Phi(\mathbf{M}^{[\tau,t]}_{k,\cdot}) \le 2\cdot\int_0^1 \frac{x}{\beta\cdot g(x)}\,\mathrm{d}x\right] \ge 1-n^{-\gamma},$$

where $t_0 = 8\sigma^2(\gamma\log(n)+\log(8\sigma^2))$. From this follows that with probability at least $1-n^{-\gamma}$

$$\Big(\Upsilon_k(\mathbf{M}^{[t]})\Big)^2 = \sum_{\tau=1}^{t-t_0}\Phi(\mathbf{M}^{[\tau,t]}_{k,\cdot}) + \sum_{\tau=t-t_0+1}^{t}\Phi(\mathbf{M}^{[\tau,t]}_{k,\cdot}) \overset{(a)}{\le} 2\cdot\int_0^1\frac{x}{\beta\cdot g(x)}\,\mathrm{d}x + t_0,$$

where (a) follows from the fact that $\Phi(\mathbf{M}_{k,\cdot}) < 1$ for k-th row of any stochastic matrix $\mathbf{M}$. The lemma follows applying the definition of t_0. □

The next result is analogous to Lemma 2.16:

Lemma 2.20. *Assume G is an arbitrary d-regular graph. Then $\mathcal{D}_1(G)$ is (g_G,σ_G^2)-good, where*

$$g_G(x) := \frac{1}{dn}\cdot\max\Big\{d\cdot\lambda(\mathbf{L}(G))\cdot x,\ \frac{1}{\mathrm{Res}(G)}\cdot x^2,\ \frac{4}{27}\cdot x^3\Big\};\quad \sigma^2 = 2\cdot \mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})\,.$$

The proof of Lemma 2.20 is analogous to that of Lemma 2.16, except that we use Lemma 2.21 stated below instead of Lemma 2.17.

Lemma 2.21. *Let G be a d-regular graph, let $\mathbf{M}^1 \sim \mathcal{D}_1(G)$, and let $\vec{x}\in\mathbb{R}^n$, Then*

1. $\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1\cdot\vec{x})\big] = \frac{1}{dn}\cdot\Psi_G(\vec{x})$.
2. *Var*$\big[\Phi(\mathbf{M}^1\cdot\vec{x})\big] \le (2\cdot\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})-1)\cdot\big(\Phi(\vec{x})-\mathbf{E}\big[\Phi(\mathbf{M}^1\cdot\vec{x})\big]\big)^2$.

Proof. For the first statement, we use observation 2.15 as well as the fact that $\mathcal{D}_1(G)$ is the uniform distribution over the edges of G to see that, as claimed.

$$\begin{aligned}\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1\cdot\vec{x})\big] &= \mathbf{E}\big[\Phi(\vec{x})-\Phi(\mathbf{M}^1\cdot\vec{x})\big] = \mathbf{E}\Big[\frac{1}{2}\cdot\Psi_{\mathbf{M}^1}(\vec{x})\Big] \\ &= \frac{1}{2}\cdot\sum_{\{i,j\}\in E(G)}\frac{1}{|E|}\cdot(x_i-x_j)^2 = \frac{1}{2}\cdot\frac{1}{dn/2}\cdot\Psi_G(\vec{x}) = \frac{1}{dn}\cdot\Psi_G(\vec{x}).\end{aligned}$$

For the second statement we first observe that $\Phi(\vec{x})$ is constant and by observation 2.15 we have

$$\mathrm{Var}\big[\Phi(\mathbf{M}^1\cdot\vec{x})\big] = \mathrm{Var}\big[\Phi(\vec{x})-\Phi(\mathbf{M}^1\cdot\vec{x})\big] = \mathrm{Var}\Big[\frac{1}{2}\cdot\Psi_{\mathbf{M}^1}(\vec{x})\Big].$$

We bound this variance using the Bhatia-Davis inequality (Theorem 6.3). It states that, for a random variable X taking values in $[m, M]$, and with $\mu := \mathbf{E}[X]$, it is the case that $\mathrm{Var}[X] \le (M-\mu)(\mu-m)$. Now from the definition of Ψ, it is immediate that $\Psi_{\mathbf{M}^1}(\vec{x}) \ge 0$. For the upper bound on $\Psi_{\mathbf{M}^1}(\vec{x})$, recall that the matchings $\mathbf{M}^1 \sim \mathcal{D}_1(G)$ consist of just one edge, and so $\Psi_{\mathbf{M}^1} \le \max_{\{i,j\}\in E(G)}(x_i - x_j)^2$. The latter is bounded from above by the third statement of Lemma 2.18, yielding

$$\Psi_{\mathbf{M}^1}(\vec{x}) \le \max_{\{i,j\}\in E(G)} (x_i - x_j)^2 \le \mathrm{Res}^*(G) \cdot \Psi_G(\vec{x}).$$

And so, by the Bhatia-Davis inequality (Theorem 6.3),

$$\begin{aligned}
\mathrm{Var}\left[\frac{1}{2}\cdot\Psi_{\mathbf{M}^1}(\vec{x})\right] &\le \left(\mathrm{Res}^*(G)\cdot\Psi_G(\vec{x}) - \frac{1}{dn}\cdot\Psi_G(\vec{x})\right)\cdot\frac{1}{dn}\cdot\Psi_G(\vec{x}),\\
&= (\mathrm{Res}^*(G)\cdot dn - 1)\cdot\left(\frac{1}{dn}\cdot\Psi_G(\vec{x})\right)^2\\
&\le 2\cdot \mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})\cdot\big(\Phi(\vec{x}) - \mathbf{E}\big[\Phi(\mathbf{M}^1\cdot\vec{x}]\big)\big)^2,
\end{aligned}$$

where the last inequality used the fact that $\mathrm{Res}^*(G)\cdot dn = 2\cdot\mathrm{Res}^*(G)\cdot|E| \le 2\cdot\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})$ by Lemma 6.21. □

2.6.2 Developed Tools

We now use the basic technical lemmas introduced above to prove the key tools that serve as the main ingredients in our propositions.

The following lemma shows that, after a sufficiently long time (which depends on the spectral gap and the logarithm of the initial discrepancy) the contribution of the initial load becomes negligible with high probability. This lemma is used in Proposition 2.7, which applies to both the random matching and asynchronous models.

Lemma 2.22. *For $t \in \mathbb{N}$ with $t = t_0(\gamma) := c\cdot\max\{\gamma\log(n), \log(K\cdot n)\}/(\lambda(\mathbf{L}(G))\cdot\beta)$ then $\mathbf{E}[\Phi(I(t))] \le 1/4$, and if $t \ge t_0(\gamma)$, then $\mathbf{Pr}\big[\Phi(I(t)) \le \frac{1}{4}\big] \ge 1 - n^{-\gamma}$.*

Proof. We aim to use the first statement of Theorem 6.1 on $\Phi(I(t))$ and therefore need to check its preconditions. By the definition of $I(t)$, for all $t \ge 1$,

$$I(t) = \mathbf{M}^{[1,t]}\cdot X(0) = \mathbf{M}^{\beta}(t)\cdot\mathbf{M}^{[1,t-1]}\cdot X(0) = \mathbf{M}^{\beta}(t)\cdot I(t-1).$$

Entirely analogous to the calculations in the proof of Lemma 2.19 (Equations (22) and (23)), we have, writing $\vec{V} = I(t-1)$ (so that $I(t) = \mathbf{M}^{\beta}\cdot\vec{V}$),

$$\mathbf{E}\Big[\Phi(\vec{V}) - \Phi(\mathbf{M}^{\beta}(t)\cdot\vec{V})\Big] \ge \beta\cdot\mathbf{E}\Big[\Phi(\vec{V}) - \Phi(\mathbf{M}^{1}(t)\cdot\vec{V})\Big], \tag{24}$$

and

$$\mathrm{Var}\Big[\Phi(\vec{V}) - \Phi(\mathbf{M}^{\beta}(t)\cdot\vec{V})\Big] \le 4\beta^2\cdot\mathrm{Var}\Big[\Phi(\vec{V}) - \Phi(\mathbf{M}^{1}(t)\cdot\vec{V})\Big],$$

and from the latter it immediately follows that for all φ

$$\begin{aligned}\mathrm{Var}\Big[\Phi(I(t)) \mid \Phi(\vec{V})=\varphi\Big] &= \mathrm{Var}\Big[\Phi(\mathbf{M}^\beta(t)\cdot\vec{V}) \mid \Phi(\vec{V})=\varphi\Big]\\ &= \mathrm{Var}\Big[\varphi-\Phi(\mathbf{M}^\beta(t)\cdot\vec{V}) \mid \Phi(\vec{V})=\varphi\Big]\\ &\le 4\beta^2\cdot\mathrm{Var}\Big[\varphi-\Phi(\mathbf{M}^1(t)\cdot\vec{V}) \mid \Phi(\vec{V})=\varphi\Big],\\ &= 4\beta^2\cdot\mathrm{Var}\Big[\Phi(\mathbf{M}^1(t)\cdot\vec{V}) \mid \Phi(\vec{V})=\varphi\Big].\end{aligned}$$

Combining the first statement of Lemma 2.18 and the first statement of Lemma 2.17 gives us, for all $\vec{x}\in\mathbb{R}^n$,

$$\Phi(\vec{x})-\mathbf{E}\big[\Phi(\mathbf{M}^1(t)\cdot\vec{x})\big] \ge \frac{\lambda(\mathbf{L}(G))}{16}\cdot\Phi(\vec{x}), \tag{25}$$

so that, for all φ,

$$\mathbf{E}\Big[\Phi(I(t)) \mid \Phi(\vec{V})=\varphi\Big] = \mathbf{E}\Big[\Phi(\mathbf{M}^\beta(t)\cdot\vec{V}) \mid \Phi(\vec{V})=\varphi\Big] \le \varphi-\beta\cdot\frac{\lambda(\mathbf{L}(G))}{16}\cdot\varphi.$$

By the second statement of Lemma 2.17, for all $\vec{x}\in\mathbb{R}^n$:

$$\mathrm{Var}\big[\Phi(\mathbf{M}^1(t)\cdot\vec{x})\big] \le (32\cdot(\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})/n)+4)\cdot\big(\Phi(\vec{x})-\mathbf{E}\big[\Phi(\mathbf{M}^1(t)\cdot\vec{x})\big]\big)^2. \tag{26}$$

And so,

$$\begin{aligned}\mathrm{Var}\Big[\Phi(I(t)) \mid \Phi(\vec{V})=\varphi\Big] &\le 4\beta^2\cdot\mathrm{Var}\Big[\Phi(\mathbf{M}^1(t)\cdot\vec{V}) \mid \Phi(\vec{V})=\varphi\Big]\\ &\le 4\beta^2\cdot\left(32\cdot\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}+4\right)\cdot\Big(\varphi-\mathbf{E}\Big[\Phi(\mathbf{M}^1(t)\cdot\vec{V}) \mid \Phi(\vec{V})=\varphi\Big]\Big)^2\\ &= \left(128\cdot\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}+16\right)\cdot\Big(\beta\cdot\mathbf{E}\Big[\varphi-\Phi\Big(\mathbf{M}^1(t)\cdot\vec{V}\Big) \mid \Phi(\vec{V})=\varphi\Big]\Big)^2\\ &\overset{\text{Eq. (24)}}{\le} \left(128\cdot\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}+16\right)\cdot\Big(\mathbf{E}\Big[\varphi-\Phi\Big(\mathbf{M}^\beta(t)\cdot\vec{V}\Big) \mid \Phi(\vec{V})=\varphi\Big]\Big)^2\\ &= \left(128\cdot\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}+16\right)\cdot\Big(\mathbf{E}\Big[\Phi\Big(\mathbf{M}^\beta(t)\cdot\vec{V}\Big) \mid \Phi(\vec{V})=\varphi\Big]-\varphi\Big)^2.\end{aligned}$$

So we can now apply Theorem 6.1 with $h(x):=\beta\cdot\lambda(\mathbf{L}(G))\cdot x/16$ and $\sigma:=128\cdot\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})/n+16$. With these values and $\delta:=1/2$, the first statement of Theorem 6.1 gives us

$$\mathbf{Pr}\left[\int_{\Phi(I(t))}^{\Phi(I(0))}\frac{1}{h(\varphi)}\,\mathrm{d}\varphi\le t/2\right]\le\exp\left(-\frac{t}{8(\sigma+1)}\right).$$

The integral evaluates to

$$\int_{\Phi(I(t))}^{\Phi(I(0))}\frac{1}{h(\varphi)}\,\mathrm{d}\varphi=\frac{16}{\beta\,\lambda(\mathbf{L}(G))}\cdot\int_{\Phi(I(t))}^{\Phi(I(0))}\frac{1}{\varphi}\,\mathrm{d}\varphi=\log\left(\frac{\Phi(I(0))}{\Phi(I(t))}\right)\cdot\frac{16}{\beta\cdot\lambda(\mathbf{L}(G))}.$$

This is at least $t/2$ if and only if

$$\Phi(I(t))\le\Phi(I(0))\cdot\exp\left(-\frac{\beta\cdot\lambda(\mathbf{L}(G))}{32}\cdot t\right),$$

which follows after rearranging the initial inequality and exponentiation. So

$$\mathbf{Pr}\left[\Phi(I(t)) \leq \Phi(I(0)) \cdot \exp\left(-\frac{\beta \cdot \lambda(\mathbf{L}(G))}{32} \cdot t\right)\right] \geq 1 - \exp\left(-\frac{t}{8(\sigma+1)}\right). \tag{27}$$

Let $K \coloneqq \mathrm{disc}(I(0)) = \mathrm{disc}(X(0))$ which implies that $\Phi(I(0)) \leq n \cdot K^2$, so that $\log(\Phi(I(0))) \leq 2\log(K \cdot n)$. Furthermore, it is the case that $0.5 \leq \mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})/n \leq 1/\lambda(\mathbf{L}(G))$ (by Theorem 6.20) and that $\beta \in (0, 1]$. Therefore, there is a sufficiently large constant $c > 0$ such that if $t \geq t_0(\gamma) = c \cdot \max\{\gamma \log(n), \log(K \cdot n)\}/(\beta \cdot \lambda(\mathbf{L}(G)))$, then

$$t \geq \frac{\beta \cdot \lambda(\mathbf{L}(G))}{32} \cdot \log(8 \cdot \Phi(I(0))),$$

as well as

$$t \geq \max\{\gamma \log(n), \log(\Phi(I(0)))\} \cdot 8 \cdot \left(128 \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n} + 33\right) = \max\{\gamma \log(n), \log(\Phi(I(0)))\} \cdot 8(\sigma+1).$$

From $t \geq (\beta \cdot \lambda(\mathbf{L}(G))/32) \cdot \log(8 \cdot \Phi(I(0)))$, it follows that

$$\Phi(I(0)) \cdot \exp\left(-\frac{\beta \cdot \lambda(\mathbf{L}(G))}{32} \cdot t\right) \leq \frac{1}{8}.$$

From $t \geq \max\{\gamma \log(n), \log(\Phi(I(0)))\} \cdot 8(\sigma+1)$, it follows that

$$\exp\left(-\frac{t}{8(\sigma+1)}\right) \leq \min\left\{n^{-\gamma}, \frac{1}{8 \cdot \Phi(I(t))}\right\}.$$

And so, for $t \geq t_0(\gamma)$, Equation (27) entails

$$\mathbf{Pr}\left[\Phi(I(t)) \leq \frac{1}{8}\right] \geq 1 - n^{-\gamma},$$

which is the remaining claim for the high-probability statement.

For the statement concerning the expectation, note that for $t \geq t_0(0)$, the calculations above and Equation (27) entail that

$$\mathbf{Pr}\left[\Phi(I(t)) \leq \frac{1}{8}\right] \geq 1 - \frac{1}{8 \cdot \Phi(I(0))}.$$

Hence, as $\Phi(I(\tau)) \leq \Phi(I(0))$ for all $\tau \in \mathbb{N}$, we have, for all $t \geq t_0(0)$,

$$\begin{aligned}\mathbf{E}[\Phi(I(t))] &\leq \frac{1}{8} \cdot \mathbf{Pr}\left[\Phi(I(t)) \leq \frac{1}{8}\right] + \Phi(I(0)) \cdot \mathbf{Pr}\left[\Phi(I(t)) > \frac{1}{8}\right] \leq \frac{1}{8} + \Phi(I(0)) \cdot \frac{1}{8 \cdot \Phi(I(0))} \\ &= \frac{1}{8} + \frac{1}{8} = \frac{1}{4},\end{aligned}$$

as claimed. □

The following lemma shows that the rounding errors in all three matching models are small. The proof relies on the following facts: (1) the expected rounding error at each node is zero, (2) the matching matrices are doubly stochastic, and (3) the variable $R(t)$ can be decomposed into components with bounded second moments.

Lemma 2.23 (Rounding Errors)**.** *In all three models on the matchings and any $\gamma > 1$ and $t \in \mathbb{N}$ it holds* *that*

$$\mathbf{Pr}\Big[\mathrm{disc}(R(t)) = \mathrm{O}\Big(\sqrt{\gamma\log(n)/\beta}\Big) \geq 1 - 2n^{-\gamma+1}\Big]$$

Proof. We show the concentration bound on $\mathrm{disc}(R(t))$ by proving concentration bounds on the absolute values $|R_k(t)|$ for each $k \in [n]$ and then applying a union bound over all k. To show the concentration bound on $R_k(t)$ holds for any fixed sequence of matchings $\mathbf{m}^{[t]} = (\mathbf{m}^{\beta}(\tau))_{\tau=1}^{t}$; this implies a concentration bound on a random sequence of matchings by the law of total probability.

So we fix $\mathbf{m}^{[t]}$. Recall that

$$R(t) = \sum_{\tau=1}^{t} \mathbf{m}_{k,\cdot}^{[\tau+1,t]} \cdot \varepsilon(\tau),$$

where $\varepsilon(\tau) = (\varepsilon_k(t))_{k\in[n]}$ is the vector of additive rounding errors incurred in round τ: it is the difference between the load vector step t, and what the load vector *would* be after step t if the balancing in this step were idealized. This additive rounding error stems from the constraint that only whole items can be transferred across the edges $[i,j]$ of the matching at time τ. From the description of the protocol, it is immediate that the rounding errors at matched nodes sum to 0, so that $\varepsilon_i(\tau) := \varepsilon_{i,j}(\tau) = -\varepsilon_{j,i}(\tau) := -\varepsilon_j(\tau)$ for all edges $[i,j] \in \mathbf{m}^{(\tau)}$ matched in round τ. Thus,

$$\begin{aligned} R_k(t) &= \sum_{\tau=1}^{t} \mathbf{m}_{k,\cdot}^{[\tau+1,t]} \cdot \varepsilon(\tau) = \sum_{\tau=1}^{t} \mathbf{m}_{k,\cdot}^{[\tau+1,t]} \cdot \sum_{[i,j]\in\mathbf{m}^{(\tau)}} (\varepsilon_i(\tau) + \varepsilon_j(\tau)) \\ &= \sum_{\tau=1}^{t} \sum_{[i,j]\in\mathbf{m}^{(\tau)}} \Big(\mathbf{m}_{k,i}^{[\tau+1,t]} - \mathbf{m}_{k,j}^{[\tau+1,t]}\Big) \cdot \varepsilon_i(\tau). \end{aligned}$$

We will derive the claimed tail bound on $R_k(t)$ by applying the Azuma-Hoeffding inequality (Theorem 6.6) to a sequence of partial sums as follows. We sequence the rounding actions with τ increasing and arbitrarily within rounds. If i is the representative node of the kth edge in round τ (with $k \in [\lfloor n/2 \rfloor]$ and $\tau \in [t]$), for $l = (\tau-1)\cdot\lfloor n/2\rfloor + k$ let us write

$$Y_l = \Big(\mathbf{m}_{k,i}^{[\tau+1,t]} - \mathbf{m}_{k,j}^{[\tau+1,t]}\Big) \cdot \varepsilon_i(\tau),$$

and let $Y_l = 0$ if there are fewer than k edges are in the matching in round τ. Se sequence of partial sums is then $S_l := \sum_{a\in[l]} Y_l$, which we consider with respect to the filtration $(\mathcal{F}(l))_{l=0}^{t\cdot\lfloor n/2\rfloor}$ in which $\mathcal{F}(l-1)$ completely determines the state right before the rounding action corresponding to the term Y_l. Note that $S_{t\cdot\lfloor n/2\rfloor} = R_k(t)$. To apply Theorem 6.6, it is enough to show that the conditional expectation of the difference between successive terms is zero, and that we can bound the differences between terms.

To check these preconditions, let us write F_l for the fractional value of the load at node i before the rounding action (i.e., the fractional value of the load i if balancing were idealized and no rounding was necessary). Then the load will be rounded up with probability F_l, resulting in a positive rounding error of $\varepsilon_i(\tau) = 1 - F_l$, or rounded down with probability $1 - F_l$, resulting in a negative rounding error of $\varepsilon_i(\tau) = -F_l$. Hence,

$$\mathbf{E}[\varepsilon_i(\tau) \mid \mathcal{F}(l-1)] = F_l \cdot (1 - F_l) + (1 - F_l)\cdot(-F_l) = 0,$$

so that, as required,

$$\mathbf{E}[Y_l \mid \mathcal{F}(l-1)] = \mathbf{E}\Big[\Big(\mathbf{m}_{k,i}^{[\tau+1,t]} - \mathbf{m}_{k,j}^{[\tau+1,t]}\Big)\cdot \varepsilon_i(\tau) \,\Big|\, \mathcal{F}(l-1)\Big] = 0.$$

From this description, it is also clear that writing $g_{i,j}^{(\tau)} := \mathbf{m}_{k,i}^{[\tau+1,t]} - \mathbf{m}_{k,j}^{[\tau+1,t]}$, the term Y_l is bounded from above by $a_l := g_{i,j}^{(\tau)} \cdot (1 - F_i(\tau))$, and from below by $-b_l := -g_{i,j}^{(\tau)} \cdot F_i(\tau)$, so that $a_l + b_l = g_{i,j}^{(\tau)}$.

So we may apply Theorem 6.6; to use it we require (an upper bound on) the value of the sum $\sum_{l=1}^{\tau\cdot\lfloor n/2\rfloor}(a_l + b_l)^2$, which we bound by applying observation 2.15 and collapsing the ensuing telescoping sum (analogously to the proof of Theorem 3.2 in [85]):

$$\begin{aligned}\sum_{l=1}^{\tau\cdot\lfloor n/2\rfloor}(a_l + b_l)^2 &= \sum_{\tau=1}^{t}\underbrace{\sum_{[i,j]\in\mathbf{m}^{(\tau)}}\left(\mathbf{m}_{k,i}^{[\tau+1,t]} - \mathbf{m}_{k,j}^{[\tau+1,t]}\right)^2}_{=\Psi_{E(\mathbf{m}^{(\tau)})}}\\ &= \sum_{\tau=1}^{t}\frac{2}{1-(1-\beta)^2}\left(\Phi\left(\mathbf{m}_{k,\cdot}^{[\tau+1,t]}\right) - \Phi\left(\mathbf{m}_{k,\cdot}^{[\tau+1,t]}\cdot\mathbf{m}^{(\tau)}\right)\right)\\ &\overset{(a)}{\leq} \sum_{\tau=1}^{t}\frac{2}{\beta}\left(\Phi\left(\mathbf{m}_{k,\cdot}^{[\tau+1,t]}\right) - \Phi\left(\mathbf{m}_{k,\cdot}^{[\tau+1,t]}\cdot\mathbf{m}^{(\tau)}\right)\right)\end{aligned}$$

where (a) follows from the fact that $\beta \in (0,1]$ and therefore, $1-(1-\beta)^2 \geq \beta$. Hence,

$$\begin{aligned}\sum_{l=1}^{\tau\cdot\lfloor n/2\rfloor}(a_l + b_l)^2 &= \frac{2}{\beta}\cdot\sum_{\tau=1}^{t}\left(\Phi\left(\mathbf{m}_{k,\cdot}^{[\tau+1,t]}\right) - \Phi\left(\mathbf{m}_{k,\cdot}^{[\tau,t]}\right)\right) = \frac{2}{\beta}\cdot\left(\Phi\left(\mathbf{m}_{k,\cdot}^{[t+1,t]}\right) - \Phi\left(\mathbf{m}_{k,\cdot}^{[1,t]}\right)\right)\\ &= \frac{2}{\beta}\cdot\left(\Phi(\mathbf{I}_{k,\cdot}) - \Phi\left(\mathbf{m}_{k,\cdot}^{[1,t]}\right)\right) \leq \frac{2}{\beta}\cdot(1-0) = \frac{2}{\beta},\end{aligned}$$

So by Theorem 6.6 (with $\varepsilon := \sqrt{\gamma\log(n)/\beta}$ and $\mathbf{E}[R_k(t)] = 0$) we have

$$\mathbf{Pr}\left[|R_k(t)| \geq \sqrt{\frac{\gamma\log(n)}{\beta}}\right] \leq 2\exp\left(-\frac{2\varepsilon^2}{2/\beta}\right) \leq 2\exp(-\gamma\log(n)) = 2n^{-\gamma}.$$

Since $\operatorname{disc}(R(t)) = \max_{k\in[n]} R_k(t) - \min_{k\in[n]} R_k(t)$, applying a union bound over all nodes $k \in [n]$ we see that

$$\mathbf{Pr}\left[\operatorname{disc}(R(t)) \geq 2\cdot\sqrt{\frac{\gamma\log(n)}{\beta}}\right] \leq 2n^{-\gamma+1},$$

which is the claimed concentration bound. □

Using the tools we have developed, we can now bound the global divergence for the random matching model. To do so, we first apply Lemma 2.16, which shows that the matching distribution $\mathcal{D}_{\mathrm{RM}}(G)$ is (g_G, σ_G^2)-good for specific values of g_G and σ_G^2. We then apply Lemma 2.19, which expresses the global divergence in terms of g_G and σ_G^2, and the remainder follows through straightforward, albeit tedious, calculations.

Lemma 2.24 (Global Divergence for Random Matching model)**.** *Assume G is an arbitrary graph. Let $\mathbf{M}^{[t]} = (\mathbf{M}^{\beta}(\tau))_{\tau=1}^{t}$ be an i.i.d. sequence of matching matrices generated by $\mathcal{D}_{\mathrm{RM}}(G)$. Let*

$$T(G) := \min\left\{\mathrm{t_{hit}(G)}\cdot\log(n)/n, \sqrt{d/\lambda(\mathbf{L}(G))}, 1/\lambda(\mathbf{L}(G))\right\}.$$

Then for all $\gamma > 0$ we get with probability at least $1 - n^{-\gamma}$,

$$\Upsilon(\mathbf{M}^{[t]}) = \mathrm{O}\left(\sqrt{\gamma \log(n) \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n} + \frac{T(G)}{\beta}}\right) = \mathrm{O}\left(\sqrt{\gamma \log(n) \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}} + \sqrt{\frac{T(G)}{\beta}}\right).$$

Proof. Fix a node $k \in [n]$. Define $g_G(x) = \frac{1}{16d} \cdot \max\{d \cdot \lambda(\mathbf{L}(G)) \cdot x, x^2/\mathrm{Res}(G), 4x^3/27\}$ and let $\sigma_G^2 := 32 \cdot (\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})/n) + 5$. Then by Lemma 2.16 the matching distribution $\mathcal{D}_{\mathrm{RM}}(G)$ is (g_G, σ_G^2)-good. By Lemma 2.19 we have for all $t \in \mathbb{N}$

$$\mathbf{Pr}\left[\left(\Upsilon_k(\mathbf{M}^{[t]})\right)^2 \leq 8\sigma_G^2((\gamma+1)\log(n) + \log(8\sigma_G^2)) + \frac{1}{\beta} \cdot \int_0^1 \frac{x}{g_G(x)}\,\mathrm{d}x\right] \geq 1 - n^{-(\gamma+1)}. \tag{28}$$

From Claim 2.25 (presented below) it follows that $\int_0^1 x/g_G(x)\,\mathrm{d}x = \mathrm{O}(T(G))$ and from Claim 2.26 (also presented below) it follows that $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})/n \geq 1/2$. Hence with probability at least $1 - n^{-(\gamma+1)}$ we have

$$\left(\Upsilon_k(\mathbf{M}^{[t]})\right)^2 = \mathrm{O}\left(\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n} \cdot \left(\gamma \log(n) + \log\left(\frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}\right)\right) + \frac{T(G)}{\beta}\right). \tag{29}$$

Since $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) = \mathrm{O}(n^3)$ (Proposition 10.16 in [64]), $\log(\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})/n) = \mathrm{O}(\log(n))$, and $\gamma > 1$, we get that

$$\Upsilon_k(\mathbf{M}^{[t]}) = \mathrm{O}\left(\sqrt{\gamma \log(n) \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n} + \frac{T(G)}{\beta}}\right) = \mathrm{O}\left(\sqrt{\gamma \log(n) \cdot \frac{\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})}{n}} + \sqrt{\frac{T(G)}{\beta}}\right).$$

Applying union bound over all $k \in [n]$ finishes the proof. □

Claim 2.25. *It holds that* $\int_0^1 x/g_G(x)\,\mathrm{d}x = \mathrm{O}(T(G))$.

Proof. First, expanding the definition of $g_G(x)$, pulling out constant factors, and simplifying fractions results in

$$\int_0^1 \frac{x}{g_G(x)}\,\mathrm{d}x = 16d \cdot \int_0^1 \min\left\{\frac{1}{d \cdot \lambda(\mathbf{L}(G))}, \frac{\mathrm{Res}(G)}{x}, \frac{27}{4x^2}\right\}\mathrm{d}x,$$

and we write $f_1(x)$, $f_2(x)$, and $f_3(x)$ for the first, second, and third argument of the minimum. For $x \geq 0$, the indefinite integrals of these functions are

$$\int f_1(x)\,\mathrm{d}x = \frac{x}{d \cdot \lambda(\mathbf{L}(G))}; \quad \int f_2(x)\,\mathrm{d}x = \mathrm{Res}(G) \cdot \log(x); \quad \int f_3(x)\,\mathrm{d}x = -\frac{27}{4}x^{-1}.$$

First, we show that $\int_0^1 x/g_G(x)\,\mathrm{d}x = \mathrm{O}(1/\lambda(\mathbf{L}(G)))$: As $\min\{f_1(x), f_2(x), f_3(x)\} \leq f_1(x)$, we bound the integral in question as

$$\int_0^1 \frac{x}{g_G(x)}\,\mathrm{d}x \leq 16d \cdot \int_0^1 \frac{1}{d \cdot \lambda(\mathbf{L}(G))}\,\mathrm{d}x = 16d \cdot \frac{1}{d \cdot \lambda(\mathbf{L}(G))} = \mathrm{O}\left(\frac{1}{\lambda(\mathbf{L}(G))}\right).$$

Next, we show that $\int_0^1 x/g_G(x)\,\mathrm{d}x = \mathrm{O}(\sqrt{d/\lambda(\mathbf{L}(G))})$: Let $x_{1,3} := \sqrt{\frac{27}{4} d\,\lambda(\mathbf{L}(G))}$ be the x such that

$f_1(x) = f_3(x)$. If $x_{1,3} \le 1$, then

$$\begin{aligned}
\int_0^1 \frac{x}{g_G(x)}\,\mathrm{d}x &\le 16d \cdot \left(\int_0^{x_{1,3}} f_1(x)\,\mathrm{d}x + \int_{x_{1,3}}^1 f_3(x)\,\mathrm{d}x\right) \\
&= 16d \cdot \left(\frac{x_{1,3}}{d \cdot \lambda(\mathbf{L}(G))} + \frac{27}{4} \cdot (-1 + x_{1,3}^{-1})\right) \\
&= 16d \cdot \left(\sqrt{\frac{27}{4 \cdot d \cdot \lambda(\mathbf{L}(G))}} + \sqrt{\frac{27}{4 \cdot d \cdot \lambda(\mathbf{L}(G))}} - \frac{27}{4}\right) = \mathrm{O}\left(\sqrt{\frac{d}{\lambda(\mathbf{L}(G))}}\right).
\end{aligned}$$

But if $x_{1,3} > 1$, the same bound also holds: we showed above that the integral in question is bounded by $\mathrm{O}(1/\lambda(\mathbf{L}(G)))$, so that if $x_{1,3} > 1$, we have an upper bound of

$$\int_0^1 \frac{x}{g_G(x)}\,\mathrm{d}x = \mathrm{O}\left(\frac{1}{\lambda(\mathbf{L}(G))}\right) = \mathrm{O}\left(\frac{x_{1,3}}{\lambda(\mathbf{L}(G))}\right) = \mathrm{O}\left(\sqrt{\frac{d}{\lambda(\mathbf{L}(G))}}\right).$$

Last, we show that $\int_0^1 x/g_G(x)\,\mathrm{d}x = \mathrm{O}(\mathrm{t_{hit}}(\mathrm{G})/n \cdot \log(n))$: Let $x_{1,2} := d \cdot \lambda(\mathbf{L}(G)) \cdot \mathrm{Res}(G)$ be the x such that $f_1(x) = f_2(x)$. If $x_{1,2} \le 1$, then

$$\begin{aligned}
\int_0^1 \frac{x}{g_G(x)}\,\mathrm{d}x &\le 16d \cdot \left(\int_0^{x_{1,2}} f_1(x)\,\mathrm{d}x + \int_{x_{1,2}}^1 f_2(x)\,\mathrm{d}x\right) \\
&= 16d \cdot \left(\frac{x_{1,2}}{d \cdot \lambda(\mathbf{L}(G))} + \mathrm{Res}(G) \cdot (\log(1) - \log(x_{1,2}))\right), \\
&= 16d \cdot \left(\mathrm{Res}(G) + \mathrm{Res}(G) \cdot \log\left(\frac{1}{d \cdot \lambda(\mathbf{L}(G)) \cdot \mathrm{Res}(G)}\right)\right) \\
&= \mathrm{O}\left(d \cdot \mathrm{Res}(G) \cdot \log\left(\frac{1}{d \cdot \lambda(\mathbf{L}(G)) \cdot \mathrm{Res}(G)}\right)\right) \\
&= \mathrm{O}\left(\frac{\mathrm{t_{hit}}(\mathrm{G})}{n} \cdot \log\left(\frac{1}{\lambda(\mathbf{L}(G))} \cdot \frac{n}{\mathrm{t_{hit}}(\mathrm{G})}\right)\right) = \mathrm{O}\left(\frac{\mathrm{t_{hit}}(\mathrm{G})}{n} \log(n)\right),
\end{aligned}$$

where the penultimate bound uses the fact that $\mathrm{Res}(G) \cdot |E(G)| = \mathrm{Res}(G) \cdot dn/2 \le \mathrm{t_{hit}}(\mathrm{G})$ (Lemma 6.21), and the final bound uses the fact that the inverse spectral gap of the normalized Laplacian $1/\lambda(\mathbf{L}(G))$ is bounded from above by $\mathrm{O}(n^3)$ (cf. [5]), and that $\mathrm{t_{hit}}(\mathrm{G}) \ge 1$, so that the argument of the logarithm is polynomial in n.

Otherwise, if $x_{1,2} > 1$, the same bound also holds: we show above that the integral is bounded by $\mathrm{O}(1/\lambda(\mathbf{L}(G)))$, so that if $x_{1,2} > 1$ we have an upper bound of

$$\int_0^1 \frac{x}{g_G(x)}\,\mathrm{d}x = \mathrm{O}\left(\frac{1}{\lambda(\mathbf{L}(G))}\right) = \mathrm{O}\left(\frac{x_{1,2}}{\lambda(\mathbf{L}(G))}\right) = \mathrm{O}(d \cdot \mathrm{Res}(G)) = \mathrm{O}\left(\frac{\mathrm{t_{hit}}(\mathrm{G})}{n} \cdot \log(n)\right).$$

Combining the three bounds, we have, as claimed,

$$\int_0^1 \frac{x}{g_G(x)}\,\mathrm{d}x = \mathrm{O}\left(\min\left\{\frac{1}{\lambda(\mathbf{L}(G))}, \sqrt{\frac{d}{\lambda(\mathbf{L}(G))}}, \frac{\mathrm{t_{hit}}(\mathrm{G})}{n} \cdot \log(n)\right\}\right) = \mathrm{O}(T(G)). \qquad \square$$

Claim 2.26. *For any d-regular graph G it holds that $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})/n \ge 1/2$.*

Proof. By the first inequality of Corollary 3.3 in [67] it holds for any nodes $i, j \in V(G)$ that

$$H(i,j) + H(j,i) \geq |E(G)| \cdot \left(\frac{1}{d(i)} + \frac{1}{d(j)}\right).$$

As G is regular we have $d(i) = d(j) = d$ and $|E(G)| = dn/2$, and since the statement holds in particular for any pair of nodes that is adjacent this entails

$$2\, \mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) \geq \frac{dn}{2} \cdot \frac{2}{d} = n,$$

and the claim follows. □

The following lemma can be viewed as a version of Lemma 2.4 adapted to the process ABal. The proof of Lemma 2.4 relies heavily on the independence between load allocation and the choice of matching edges, which does not hold for ABal. In contrast, Lemma 2.27 carefully analyzes this dependency while employing a stronger concentration inequality.

Lemma 2.27 (Contribution of Dynamically allocated items in asynchronous model)**.** *Let G be a regular graph, and let $t \in \mathbb{N}$. Then in* ABal$(\mathcal{D}_1(G), \beta)$*, for all $\gamma > 0$ and for $\hat{\Upsilon} > 0$ such that $\mathbf{Pr}\Big[\max_{k \in V} \Upsilon_k(\mathbf{M}^{[t]}) + 1 > \hat{\Upsilon}\Big] \leq n^{-\gamma}$, we have*

$$\mathbf{Pr}\left[\mathrm{disc}(D(t)) \geq \frac{2\gamma \log(n)}{3} + 2\sqrt{\frac{\gamma \log(n)}{n}} \cdot \hat{\Upsilon}\right] \leq 8n^{-\gamma+1}.$$

Proof. Fix a node $k \in V$. Let $A(\tau)$ be the vector of allocated loads in round τ and recall that we have

$$D(t) = \sum_{\tau=1}^{t} \mathbf{M}^{[\tau,t]} \cdot A(\tau), \quad \text{so that} \quad D_k(t) = \sum_{\tau=1}^{t} \mathbf{M}^{[\tau,t]}_{k,\cdot} \cdot A(\tau).$$

Using $\mathbf{M}^{[\tau,t]} = \mathbf{M}^{[\tau+1,t]} \cdot \mathbf{M}^{(\tau)}$, we can express the kth coordinate of $D(t)$ as

$$D_k(t) = \sum_{\tau=1}^{t} C_k(\tau), \quad \text{where } C_k(\tau) := \mathbf{M}^{[\tau,t]}_{k,\cdot} \cdot A(\tau) = \mathbf{M}^{[\tau+1,t]}_{k,\cdot} \cdot \Big(\mathbf{M}^{(\tau)} \cdot A(\tau)\Big)$$

is the contribution of the load item allocated in round τ to $D_k(t)$. Note that in the second factorization of the $C_k(\tau)$, the two factors are independent as they concern disjoint rounds.

Now consider the sequence $(Y(l))_{l=0}^{t}$ of partial sums $Y(l) = \sum_{\tau=t-l+1}^{t}(C_k(\tau) - 1/n)$ with respect to the natural filtration $\mathcal{F} = (\mathcal{F}(l))_{l=0}^{t}$ on the sequence of edges $(I(t-l), J(t-l))$. In particular, we have

$$Y(0) = 0, \quad Y(l) - Y(l-1) = C_k(t-l) - 1/n, \quad \text{and } Y(t) = D_k(t) - t/n,$$

and $\mathcal{F}(l)$ determines all edges used in rounds $t-l+1$ up to round t. To apply the martingale tail inequality Corollary 6.5 to $(Y(l))_{l=0}^{t}$, we need to check that $\mathbf{E}[Y(l) - Y(l-1) \mid \mathcal{F}(l-1)] = 0$ and that $|Y(l) - Y(l-1)| \leq 1$. For the first condition, note that both $\mathbf{M}^{[\tau,t]}_{k,\cdot}$ and $A(\tau)$ are stochastic vectors (for the latter, this is because exactly one load item is allocated in each round in the asynchronous model). Thus, their inner product $C_k(\tau)$ has a value in the interval $[0;1]$, so that $|Y(l) - Y(l-1)| = |C_k(t-l) - 1/n| \leq 1 - 1/n \leq 1$, as required.

For the second condition, note that

$$\mathbf{E}[Y(l) - Y(l-1) \mid \mathcal{F}(l-1)] = \mathbf{E}\big[C_k(t-l) \,\big|\, ((I(r), J(r)))_{r=t-l+1}^{t}\big] - 1/n,$$

so that it is enough to show that the expected value of the $C_k(\tau)$ is $1/n$ when conditioned on the matching choices in rounds $\tau+1$ to t. The bound given by Corollary 6.5 also involves the quantity

$$\langle Y \rangle_t \coloneqq \sum_{l=1}^{t} \mathbf{E}\big[(Y(l) - Y(l-1))^2 \mid \mathcal{F}(l-1)\big] = \sum_{l=1}^{t} \mathbf{E}\big[(C_k(t-l) - 1/n)^2 \mid \mathcal{F}(l-1)\big],$$

so we will investigate $C_k(\tau)$ more thoroughly than would be required to compute only its conditional expectation.

To this end, let us first make the dependence between $\mathbf{M}^{(\tau)}$ and $A(\tau)$ more explicit. Let $(I(\tau), J(\tau))$ be the random orientation of the random edge selected in round τ, so that the load item in round τ is allocated to $I(\tau)$, and then the load is balanced across the edge $\{I(\tau), J(\tau)\}$. Then

$$(\mathbf{M}^{(\tau)} \cdot A(\tau))_i = \begin{cases} 1 - \beta/2, & \text{if } i = I(\tau), \\ \beta/2, & \text{if } i = J(\tau), \\ 0, & \text{otherwise.} \end{cases}$$

Using this, we may see that

$$\begin{aligned} C_k(\tau) = \mathbf{M}_{k,\cdot}^{[\tau+1,t]} \cdot \Big(\mathbf{M}^{(\tau)} \cdot A(\tau)\Big) &= \sum_{i \in [n]} \mathbf{M}_{k,i}^{[\tau+1,t]} \cdot \left(\left(1 - \frac{\beta}{2}\right) \cdot \mathbf{1}_{i=I(\tau)} + \frac{\beta}{2} \cdot \mathbf{1}_{i=J(\tau)}\right) \\ &= \left(1 - \frac{\beta}{2}\right) \cdot \left(\sum_{i \in [n]} \mathbf{M}_{k,i}^{[\tau+1,t]} \, \mathbf{1}_{i=I(\tau)}\right) + \frac{\beta}{2} \cdot \left(\sum_{i \in [n]} \mathbf{M}_{k,i}^{[\tau+1,t]} \, \mathbf{1}_{i=J(\tau)}\right). \end{aligned} \tag{30}$$

Now $\mathcal{D}_1(G)$ is the uniform distribution over the edges of G, and the node to which load is allocated is a uniformly random endpoint of the chosen edge. Thus, $(I(\tau), J(\tau))$ is distributed uniformly over the oriented edges $\bigcup_{\{i,j\} \in E(G)} \{(i,j), (j,i)\}$. Since G is d-regular, there are $2 \cdot |E(G)| = 2 \cdot (dn/2) = dn$ such oriented edges. Hence, for all $i \in [n]$,

$$\begin{aligned} \mathbf{Pr}[I(\tau) = i] &= \sum_{j \in [n]} \mathbf{Pr}[(I(\tau), J(\tau)) = (i,j)] = \sum_{j \in [n]} \frac{1}{dn} \cdot \mathbf{1}_{\{i,j\} \in E(G)} \\ &= \frac{1}{dn} \cdot |\{j \in [n] \mid \{i,j\} \in E(G)\}| = \frac{1}{dn} \cdot d = \frac{1}{n}. \end{aligned}$$

By an entirely analogous calculation, $\mathbf{Pr}[J(\tau) = i] = 1/n$ holds as well. So $I(\tau)$ and $J(\tau)$ are identically distributed (but not necessarily independent). Because of this, the two sums over $i \in [n]$ on the right-hand side of Equation (30) are also identically distributed. We can now compute the conditional expectation of $C_k(\tau)$. Using Equation (30) and linearity of expectation we see that

$$\begin{aligned} \mathbf{E}\big[C_k(\tau) \,\big|\, ((I(l), J(l)))_{l=\tau+1}^{t}\big] &= \left(\left(1 - \frac{\beta}{2}\right) + \frac{\beta}{2}\right) \cdot \mathbf{E}\left[\sum_{i \in [n]} \mathbf{M}_{k,i}^{[\tau+1,t]} \, \mathbf{1}_{i=I(\tau)} \,\middle|\, \mathbf{M}^{[\tau+1,t]}\right] \\ &= 1 \cdot \sum_{i \in [n]} \mathbf{Pr}[I(\tau) = i] \cdot \mathbf{M}_{k,i}^{[\tau+1,t]} = \frac{1}{n} \cdot \sum_{i \in [n]} \mathbf{M}_{k,i}^{[\tau+1,t]} = \frac{1}{n}. \end{aligned}$$

So $\mathbf{E}[Y(l) - Y(l-1) \mid \mathcal{F}(l-1)] = 1/n - 1/n = 0$, as required for applying Corollary 6.5. So all preconditions of Corollary 6.5 hold. Applying it with $\varepsilon = \gamma \log(n)$ and $\sigma = \hat{\Upsilon}/\sqrt{n}$ yields

$$\mathbf{Pr}\left[|Y(t) - Y(0)| \geq \frac{\gamma \log(n)}{3} + \sqrt{2\gamma \log(n)/n} \cdot \hat{\Upsilon}\right] \leq 2(n^{-\gamma} + \mathbf{Pr}\left[\langle Y \rangle_t > \hat{\Upsilon}^2/n\right]). \tag{31}$$

Note that from Claim 2.28 (see below) it follows that $\langle Y \rangle_t \leq 1/n \cdot (\Upsilon_k(\mathbf{M}^{[t]}) + 1)^2$, and therefore we get,

$$\mathbf{Pr}\left[\langle Y \rangle_t > \hat{\Upsilon}^2/n\right] \leq \mathbf{Pr}\left[1/n \cdot (\Upsilon_k(\mathbf{M}^{[t]}) + 1)^2 > \hat{\Upsilon}^2/n\right] = \mathbf{Pr}\left[\Upsilon_k(\mathbf{M}^{[t]}) + 1 > \hat{\Upsilon}\right] \leq n^{-\gamma}, \tag{32}$$

with the last inequality using the condition on $\hat{\Upsilon}$ in the statement. From Equation (31) together with Equation (32), union bound over all nodes $k \in V$ and observation 2.14 the statement follows. □

Here we prove the claim used in the last lemma.

Claim 2.28. *For $\langle Y \rangle_t$ defined in the last lemma, it holds that $\langle Y \rangle_t \leq 1/n \cdot (\Upsilon_k(\mathbf{M}^{[t]}) + 1)^2$.*

Proof. Recall that $\langle Y \rangle_t := \sum_{l=1}^{t} \mathbf{E}\left[(Y(l) - Y(l-1))^2 \mid \mathcal{F}(l-1)\right] = \sum_{l=1}^{t} \mathrm{Var}[Y(l) - Y(l-1) \mid \mathcal{F}(l-1)]$, with the latter equality using the fact the expected value of $(Y(l) - Y(l-1))$ conditioned on $\mathcal{F}(l-1)$ is 0. And since $Y(l) - Y(l-1) = C_k(t-l) - 1/n$ and $1/n$ is a constant,

$$\langle Y \rangle_t = \sum_{l=1}^{t} \mathrm{Var}[Y(l) - Y(l-1) \mid \mathcal{F}(l-1)] = \sum_{l=1}^{t} \mathrm{Var}\left[C_k(t-l) \,\middle|\, ((I(r), J(r)))_{r=t-l+1}^{t}\right].$$

By Equation (30), and as for two identically distributed random variables A and B, and $a, b \in \mathbb{R}^+$, we have $\mathrm{Var}[aA + bB] = a^2 \mathrm{Var}[A] + 2ab\mathrm{Cov}\, A, B + b^2 \mathrm{Var}[B] \leq (a^2 + 2ab + b^2)\mathrm{Var}[A] = (a+b)^2 \mathrm{Var}[A]$:

$$\begin{aligned}\mathrm{Var}\left[C_k(\tau) \,\middle|\, ((I(l), J(l)))_{l=\tau+1}^{t}\right] &\leq \left(1 - \frac{\beta}{2} + \frac{\beta}{2}\right)^2 \cdot \mathrm{Var}\left[\sum_{i \in [n]} \mathbf{M}_{k,i}^{[\tau+1,t]} \mathbf{1}_{i=I(\tau)} \,\middle|\, \mathbf{M}^{[\tau+1,t]}\right] \\ &= 1^2 \cdot \sum_{i \in [n]} \mathbf{Pr}[I(\tau) = i] \cdot \left(\mathbf{M}_{k,i}^{[\tau+1,t]} - \frac{1}{n}\right)^2 = \frac{1}{n} \cdot \left|\left|\mathbf{M}_{k,\cdot}^{[\tau+1,t]} - \frac{\vec{1}}{n}\right|\right|_2^2.\end{aligned}$$

And hence we may bound $\langle Y \rangle_t$ from above using the global divergence:

$$\begin{aligned}\langle Y \rangle_t &= \sum_{\tau=1}^{t} \mathrm{Var}\left[C_k(\tau) \,\middle|\, ((I(l), J(l)))_{l=\tau+1}^{t}\right] \leq \frac{1}{n} \cdot \sum_{\tau=1}^{t} \left|\left|\mathbf{M}_{k,\cdot}^{[\tau+1,t]} - \frac{\vec{1}}{n}\right|\right|_2^2 \\ &= \frac{1}{n}\left(\left(\Upsilon_k(\mathbf{M}^{[t]})\right)^2 - \left|\left|\mathbf{M}_{k,\cdot}^{[1,t]} - \frac{\vec{1}}{n}\right|\right|_2^2 + \left|\left|\mathbf{M}_{k,\cdot}^{[t+1,t]} - \frac{\vec{1}}{n}\right|\right|_2^2\right) \\ &\leq \frac{1}{n} \cdot \left(\left(\Upsilon_k(\mathbf{M}^{[t]})\right)^2 + 1\right) \leq \frac{1}{n} \cdot \left(\left(\Upsilon_k(\mathbf{M}^{[t]})\right) + 1\right)^2,\end{aligned}$$

which is all that remained to be shown. □

The next lemma is a version of Lemma 2.24 for the asynchronous model. It provides a bound on global divergence in terms of $T(G)$.

Lemma 2.29 (Global Divergence for Asynchronous Model)**.** *Assume G is an arbitrary graph. Let $\mathbf{M}^{[t]} = (\mathbf{M}^{\beta}(\tau))_{\tau=1}^{t}$ be an i.i.d. sequence of matching matrices generated by $\mathcal{D}_1(G)$. Let $T(G)$ be defined* *as in Theorem 2.10. Then for all $\gamma > 0$ we get with probability at least $1 - n^{-\gamma}$,*

$$\Upsilon(\mathbf{M}^{[t]}) = \mathrm{O}\left(\sqrt{\gamma \log(n) \cdot \mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})} + \sqrt{\frac{n \cdot T(G)}{\beta}}\right).$$

Proof. The proof follows the same approach as that of Lemma 2.24, except that in Equation (28), we replace Lemma 2.16 with Lemma 2.20. □

2.7 Bounds for Specific Graph Classes

In this section, we present bounds on the discrepancy for specific classes of graphs in each of the three models considered. We assume, in all cases, that the system is initially empty.

2.7.1 Balancing Circuit Model

The next two corollaries provide an upper and lower bounds on discrepancy, respectively. Corollary 2.30 and Corollary 2.31 are summarized in Table 1 and Table 2, respectively.

Corollary 2.30. *Let $X(t)$ be the state of process $\mathrm{SBAL}(\mathcal{D}_{\mathrm{BC}}(G), 1, m)$ at time t with $X(0) = \vec{0}$ and assume G has n nodes. For an arbitrary round t it holds w.h.p. and in expectation*

- $\mathrm{disc}(X(t)) = \mathrm{O}(\log(n) + \sqrt{(\Delta \cdot m)/(n \cdot \lambda(\mathbf{R}))} \cdot \sqrt{\log(n)})$ *for arbitrary graphs with round matrix* $\mathbf{R}$.
- $\mathrm{disc}(X(t)) = \mathrm{O}(\log(n) + \sqrt{m} \cdot \sqrt{\log(n)})$ *for cycle and regular graphs with constant* Δ.
- $\mathrm{disc}(X(t)) = \mathrm{O}((1 + \sqrt{m/n}) \cdot \log(n))$ *for the two-dimensional torus or hypercube graphs.*
- $\mathrm{disc}(X(t)) = \mathrm{O}(\log(n) + \sqrt{m/n} \cdot \sqrt{\log(n)})$ *for constant three or more-dimensional torus.*

Proof. The bounds follow from a straight-forward combination of the upper bounds on the local divergence from Lemma 2.32 with Theorem 2.1. □

Corollary 2.31. *Let $X(t)$ be the state of process $\mathrm{SBAL}(\mathcal{D}_{\mathrm{BC}}(G), 1, m)$ at time t with $X(0) = \vec{0}$. It holds with constant probability that*

- $\mathrm{disc}(X(t)) = \Omega(\sqrt{m})$, *for cycle, constant d-regular graphs, $t = \Omega(n^2)$ and $m \geq 4\log(n)$.*
- $\mathrm{disc}(X(t)) = \Omega\big(\sqrt{\frac{m}{n} \cdot \log(n)}\big)$ *for two-dimensional torus, $t = \Omega(n)$, and $m \geq 4n$.*
- $\mathrm{disc}(X(t)) = \Omega\big(\sqrt{\frac{m}{n}}\big)$, *for constant $r \geq 3$-dimensional torus, hypercube graphs, $t \in \mathbb{N}$, and $m \geq 4n \cdot \log(n)$.*

Proof. The bounds follow from a straight-forward combination of the bounds on the local divergence from Lemma 2.32 to Theorem 2.5. □

The two corollaries above demonstrate that our bounds are nearly tight for cycle graphs (see Table 2), constant d-regular graphs, r-dimensional torus graphs with constant r, and hypercube graphs. For example, consider a cycle constructed using the Odd–Even scheme with $m \geq \log(n)$. Corollary 2.30 states that, with high probability, the discrepancy is $\mathrm{O}(\sqrt{m \cdot \log(n)})$, whereas Corollary 2.31 shows that, with constant probability, the discrepancy is at least $\Omega(\sqrt{m})$.

We now compute the global divergence for the following concrete graphs and circuits. For cycles of even length, we consider the "Odd–Even" scheme, in which the first matching $\mathbf{m}^{(1)}$ consists of all edges

Table 2: Asymptotic lower bounds on the discrepancy in specific graph classes.

Graph	$\textsc{SBal}(\mathcal{D}_{\mathrm{BC}}(G), 1, m)$, Corollary 2.31
d-regular graph (const. d)	$\sqrt{m}$
cycle C_n	$\sqrt{m}$
2-D torus	$\sqrt{(m/n) \cdot \log(n)}$
r-D torus (const. $r \geq 3$)	$\sqrt{m/n}$
hypercube	$\sqrt{m/n}$

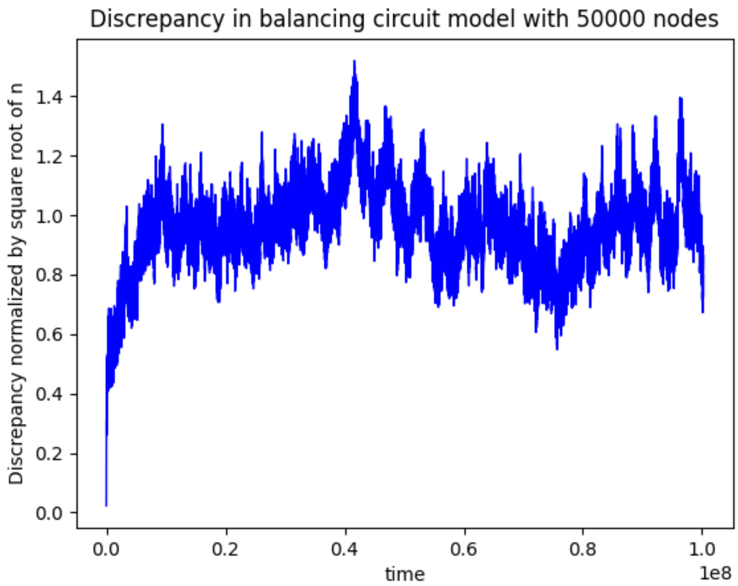

Figure 1: Simulation results for lower bound on discrepancy in the balancing circuit model for cycle with $m = n$. It shows $\Omega(\sqrt{n})$ bound.

$j, (j+1) \bmod n$ for odd $j \in [n]$, and the second matching $\mathbf{m}^{(2)}$ consists of all edges $j, (j+1) \bmod n$ for even $j \in [n]$. More generally, for an r-dimensional torus with node set $[n^{1/r}]^r$, the balancing circuit consists of $2r$ matchings in total, with two matchings for each dimension i, analogous to the cycle construction. For the hypercube, the canonical choice is the dimension-exchange circuit, which consists of $\log_2(n)$ matchings. Nodes u and v are matched in $\mathbf{m}^{(i)}$ if and only if their binary representations differ in bit i only (see, e.g., [33]).

Recall that $\mathbf{m}^{[t]}$ is the sequence of matchings and $\Phi\left(\mathbf{m}_{k,\cdot}^{[\tau,t]}\right) = \left|\left|\mathbf{m}_{k,\cdot}^{[\tau,t]} - \frac{\bar{1}}{n}\right|\right|_2^2$ and $\mathbf{R} := \mathbf{m}^{[1,\Delta]}$. The next lemma bounds the global divergence of some specific graphs for the distribution $\mathcal{D}_{\mathrm{BC}}(G)$.

Lemma 2.32 (Global Divergence)**.** *Let G be a graph and consider $\mathcal{D}_{\mathrm{BC}}(G)$ constructed by Odd-Even scheme such that it produces the round matrix $\mathbf{R}$.*

1. *For each $t \in N$ it holds $(\Upsilon(\mathbf{m}^{[t]}))^2 = \mathrm{O}(\Delta / \lambda(\mathbf{R}))$.*
2. *For a constant Δ and each $t \in \mathbb{N}$ it holds $(\Upsilon(\mathbf{m}^{[t]}))^2 = \mathrm{O}(n)$. It also holds for any $t = \Omega(n^2)$, $(\Upsilon(\mathbf{m}^{[t]}))^2 = \Omega(n)$.*

3. For two-dimensional torus G and for each $t \in \mathbb{N}$ it holds $(\Upsilon(\mathbf{m}^{[t]}))^2 = \mathrm{O}(\log(n))$. It also holds for any $t = \Omega(n)$, $(\Upsilon(\mathbf{m}^{[t]}))^2 = \Omega(\log(n))$.

4. For constant $r \geq 3$-dimensional torus G and each $t \in \mathbb{N}$ it holds $(\Upsilon(\mathbf{m}^{[t]}))^2 = \mathrm{O}(r)$. It also holds for any $t \in \mathbb{N}$, $(\Upsilon(\mathbf{m}^{[t]}))^2 = \Omega(1)$.

5. For hypercube graphs G and each $t \in \mathbb{N}$ it holds $(\Upsilon(\mathbf{m}^{[t]}))^2 = \mathrm{O}(\log(n))$. It also holds for any t, $(\Upsilon(\mathbf{m}^{[t]}))^2 = \Omega(1)$.

Proof. Recall that the sequence of matching matrices $\mathbf{m}^{[t]}$ has global divergence $\Upsilon(\mathbf{m}^{[t]})$, if

$$\forall k \in [n], \sum_{\tau=1}^{t} \left\| \mathbf{m}_{k,\cdot}^{[\tau,t]} - \frac{\vec{1}}{n} \right\|_2^2 \leq \left(\Upsilon(\mathbf{m}^{[t]})\right)^2.$$

Since the matchings are fixed we have $\left(\Upsilon(\mathbf{m}^{[t]})\right)^2 = \max_{w\in[n]} \sum_{\tau=1}^{t} \|\mathbf{m}_{w,\cdot}^{[\tau,t]} - \frac{\vec{1}}{n}\|_2^2$. Consider a node $k \in [n]$ such that $\Upsilon_k(\mathbf{m}^{[t]}) = \Upsilon(\mathbf{m}^{[t]})$. We have seen that

$$\left(\Upsilon_k(\mathbf{m}^{[t]})\right)^2 = \sum_{\tau=1}^{t} \left\| \mathbf{m}_{k,\cdot}^{[\tau,t]} - \frac{\vec{1}}{n} \right\|_2^2 = \sum_{\tau=1}^{t} \Phi\left(\mathbf{m}_{k,\cdot}^{[\tau,t]}\right).$$

Since $\Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right)$ is non increasing in $\tau \in \mathbb{N}$ and $\mathbf{R} := \mathbf{m}^{[1,\Delta]}$, then

$$\left(\Upsilon_k(\mathbf{m}^{[t]})\right)^2 \leq \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{m}_{k,\cdot}^{[1,\tau]}\right) \leq \Delta \cdot \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right).$$

Hence, to bound $\left(\Upsilon(\mathbf{m}^{[t]})\right)^2$, it is enough to bound $\Delta \cdot \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right)$.

General case: Here we get,

$$\Delta \cdot \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) \overset{(a)}{\leq} \Delta \cdot \left(\sum_{\tau=0}^{\infty} (1 - \lambda(\mathbf{R}))^{2\tau}\right) \leq \Delta \cdot \left(\sum_{\tau=0}^{\infty} (1 - \lambda(\mathbf{R}))^{\tau}\right) = \mathrm{O}\left(\frac{\Delta}{\lambda(\mathbf{R})}\right),$$

where (a) follows from [53, Lemma 2] (restated as Lemma 6.22). Note that $\Phi\left(\mathbf{R}_{k,\cdot}^{[1,1]}\right) \leq 1$.

Cycles: Recall that in cycle $\Delta = 2$. It holds that

$$\begin{aligned}\Delta \cdot \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) &= \Delta \cdot \left(\sum_{\tau=1}^{n^2} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) + \sum_{\tau=n^2+1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right)\right) \\ &\overset{(b)}{\leq} \Delta \cdot \left(\sum_{\tau=1}^{n^2} \mathrm{O}\left(\frac{1}{\sqrt{\tau}}\right) + \sum_{\tau=n^2+1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right)\right),\end{aligned}$$

where (b) follows from first statement of Lemma 6.25 and therefore,

$$\begin{aligned}\Delta \cdot \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) &\overset{(c)}{\leq} \Delta \cdot \left(\mathrm{O}(\sqrt{n^2}) + \Phi\left(\mathbf{R}_{k,\cdot}^{[1,n^2]}\right) \cdot \sum_{\tau=1}^{\infty} (1 - \lambda(\mathbf{R}))^{2\tau}\right) \\ &\overset{(d)}{\leq} \Delta \cdot \left(\mathrm{O}(\sqrt{n^2}) + \mathrm{O}\left(\frac{1}{n}\right) \cdot \sum_{\tau=1}^{\infty} (1 - \lambda(\mathbf{R}))^{2\tau}\right) \\ &= \Delta \cdot \mathrm{O}\left(n + \frac{1}{n \cdot \lambda(\mathbf{R})}\right) \overset{(e)}{\leq} \mathrm{O}(\Delta \cdot n) = \mathrm{O}(2n),\end{aligned}$$

where (c) follows from Lemma 6.22 and (d) follows from Lemma 6.25. To see (e), consider that the spectral gap of the round matrix corresponding to a cycle is $\Theta(1/n^2)$, Lemma 6.24. Moreover, for $t = cn^2$ with some constant c, it follows from Lemma 6.25 that

$$\sum_{\tau=1}^{t} \Phi\left(\mathbf{m}_{k,\cdot}^{[\tau,t]}\right) = c_1 \cdot \sum_{\tau=1}^{t} \Phi\left(\mathbf{R}_{k,\cdot}^{[\tau,t/2]}\right) = \sum_{\tau=0}^{t/2-1} \Theta\left(\frac{1}{\sqrt{t/2-\tau}}\right) = \Theta(\sqrt{t/2}) = \Omega(n),$$

for $c_1 \in [1,2]$.

Two-dimensional torus: Note that in r-dimensional torus graphs $\Delta = 2r = 4$, and the spectral gap of the round matrix corresponding to a r-dimensional torus is $\Theta(1/n^{2/r})$, Lemma 6.23. Hence,

$$\Delta \cdot \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) = \Delta \cdot \left(\sum_{\tau=1}^{n^2} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) + \sum_{\tau=n^2+1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) \right) \overset{(f)}{\leq} \Delta \cdot \left(\sum_{\tau=1}^{n^2} \mathrm{O}\left(\frac{1}{\tau}\right) + \sum_{\tau=n^2+1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) \right)$$

where (f) follow from Lemma 6.25. Applying Lemma 6.22 to the right hand side of the last inequality gives us then

$$\begin{aligned}
\Delta \cdot \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) &\leq \Delta \cdot \left(\sum_{\tau=1}^{n^2} \mathrm{O}\left(\frac{1}{\tau}\right) + \Phi\left(\mathbf{R}_{k,\cdot}^{[1,n^2]}\right) \cdot \sum_{\tau=1}^{\infty} (1-\lambda(\mathbf{R}))^{2\tau} \right) \\
&\overset{(h)}{\leq} \Delta \cdot \left(\mathrm{O}(\log(n)) + \mathrm{O}\left(\frac{1}{n^2}\right) \cdot \sum_{\tau=1}^{\infty} (1-\lambda(\mathbf{R}))^{2\tau} \right) \\
&= \mathrm{O}\left(4 \cdot \log(n) + \frac{4 \cdot n^2}{n^2} \right) = \mathrm{O}(4 \log(n)),
\end{aligned}$$

where (h) follow from Lemma 6.25. Moreover, for $t = cn$ with some constant c, it follows from Lemma 6.25 (for $c_1 \in [1,4]$) that

$$\sum_{\tau=1}^{t} \Phi\left(\mathbf{m}_{k,\cdot}^{[\tau,t]}\right) = c_1 \cdot \sum_{\tau=1}^{t} \Phi\left(\mathbf{R}_{k,\cdot}^{[\tau,t/4]}\right) = \sum_{\tau=0}^{t/4-1} \Theta\left(\frac{1}{t/4-\tau}\right) = \Theta(\log(t/4)) = \Omega(\log(n)).$$

Constant three or more-dimensional torus: Let us assume $r = 2(1+\epsilon)$ for some $\epsilon > 0$ then

$$\Delta \cdot \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) \overset{(i)}{\leq} \Delta \cdot \sum_{\tau=1}^{\infty} \tau^{-(1+\epsilon)} \leq \Delta \cdot \left(1 + \int_{1}^{\infty} x^{-(1+\epsilon)}\, dx\right) \leq \Delta \cdot (1 + 1/\epsilon) = \mathrm{O}(2r),$$

where (i) follows form Lemma 6.25.

Hypercubes: Similarly, it holds that

$$\Delta \cdot \sum_{\tau=1}^{\infty} \Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) \overset{(j)}{\leq} \Delta \cdot \left(\sum_{\tau=1}^{\infty} 2^{-\tau} \right) \leq 2 \cdot \Delta = \mathrm{O}(2 \log(n)),$$

where (j) follows from Lemma 6.25. Recall that in hypercube $\Delta \leq \log(n)$.

The lower bound of 1 is trivial. □

2.7.2 Synchronous Model

Corollary 2.33. *Let $X(t)$ be the state of process $\mathrm{SBAL}(\mathcal{D}_{\mathrm{RM}}(G), \beta, m)$ where $X(0) = \vec{0}$. For an arbitrary t it holds w.h.p. and in expectation*

- $\mathrm{disc}(X(t)) = \mathrm{O}(\sqrt{m}\log(n))$ *for any regular graph.*
- $\mathrm{disc}(X(t)) = \mathrm{O}(\log(n) + \sqrt{m\log(n)})$ *for cycles and constant-degree regular graphs.*
- $\mathrm{disc}(X(t)) = \mathrm{O}(\log(n) + \sqrt{m/n} \cdot \log^{3/2}(n))$ *for the two-dimensional torus graphs.*
- $\mathrm{disc}(X(t)) = \mathrm{O}((1 + \sqrt{m/n}) \cdot \log(n))$ *for torus graphs with ≥ 3 dimensions, the hypercube, and all d-regular graphs with $d \geq \lfloor n/2 \rfloor$.*

To show the above corollary we require bounds on $T(G)$ (Lemma 2.34) and bounds on $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})$ (Lemma 2.36). Then the corollary immediately follows from Theorem 2.6.

In the following lemma we provide some bounds on $T(G)$ for several specific graph classes.

Lemma 2.34. *Assume G is a graph with n nodes.*

- *For constant-degree regular graphs G we have $T(G) = \mathrm{O}(n)$.*
- *For a two-dimensional $k \times k$ toroidal mesh G we have $T(G) = \mathrm{O}(\log^2(n))$.*
- *For a r-dimensional $k \times \cdots \times k$ toroidal mesh (with $r \geq 3$) we have $T(G) = \mathrm{O}(\log(n))$.*
- *For a r-dimensional hypercube G we have $T(G) = \mathrm{O}(\log(n))$.*
- *For a d-regular graph G with $d \geq \lfloor \frac{n}{2} \rfloor$ we have $T(G) = \mathrm{O}(\log(n))$.*
- *For an arbitrary d-regular graph G we have $T(G) = \mathrm{O}(n\log(n))$.*

Proof. Recall that $T(G) = \min\Big\{1/\lambda(\mathbf{L}(G)), \sqrt{d/\lambda(\mathbf{L}(G))}, (\mathrm{t_{hit}}(\mathrm{G})/n) \cdot \log(n)\Big\}$, and that $\mathrm{t_{hit}}(\mathrm{G}) \leq 2 \cdot \mathrm{Res}(G) \cdot |E|$ (Lemma 6.21), so that $\mathrm{t_{hit}}(\mathrm{G})/n = \mathrm{O}(d \cdot \mathrm{Res}(G))$.

For d-regular graphs with d being constant, $1/\lambda(\mathbf{L}(G)) = \mathrm{O}(n \cdot d \cdot (\mathrm{diam}(G) + 1))$ by [62], where $\mathrm{diam}(G)$ is the diameter of G. As $\mathrm{diam}(G) \leq n$ and d is constant, $1/\lambda(\mathbf{L}(G)) = \mathrm{O}(n^2)$, so that

$$T(G) = \mathrm{O}(\sqrt{d/\lambda(\mathbf{L}(G))}) = \mathrm{O}(n).$$

For the two-dimensional $k \times k$ toroidal mesh, $d \leq 4$ and $\mathrm{Res}(G) = \mathrm{O}(\log(n))$ by [34, Theorem 6.1], so that

$$T(G) = \mathrm{O}((\mathrm{t_{hit}}(\mathrm{G})/n) \cdot \log(n)) = \mathrm{O}(\log^2(n)).$$

For a r-dimensional $k \times \cdots k$ toroidal mesh with $r \geq 3$, as well as the r-dimensional hypercube, $d \leq 2r$ and $\mathrm{Res}(G) = \mathrm{O}(r^{-1})$ by [34, Theorem 6.1], so that

$$T(G) = \mathrm{O}((\mathrm{t_{hit}}(\mathrm{G})/n) \cdot \log(n)) = \mathrm{O}((d \cdot \mathrm{Res}(G))\log(n)) = \mathrm{O}(r \cdot r^{-1} \cdot \log(n)) = \mathrm{O}(\log(n)).$$

For a d-regular graph G with $d \geq \lfloor \frac{n}{2} \rfloor$, $\mathrm{Res}(G) = \mathrm{O}(d^{-1})$ by [34, Theorem 3.3], so that

$$T(G) = \mathrm{O}((\mathrm{t_{hit}}(\mathrm{G})/n) \cdot \log(n)) = \mathrm{O}((d \cdot \mathrm{Res}(G))\log(n)) = \mathrm{O}(d \cdot d^{-1} \cdot \log(n)) = \mathrm{O}(\log(n)).$$

For general d-regular graphs G, $\mathrm{t_{hit}}(\mathrm{G}) \leq 3n^2 - nd$ by [64, Proposition 10.16], so that

$$T(G) = \mathrm{O}((\mathrm{t_{hit}}(\mathrm{G})/n) \cdot \log(n)) = \mathrm{O}((n^2/n)\log(n)) = \mathrm{O}(n\log(n)).$$

□

To bound $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})$ for many specific graph classes we use the following.

Theorem 2.35 (Theorem 2.10 in [67], [59])**.** *Let G be a graph and $i \in [n]$ be one of its nodes. Then if $J \in [n]$ is chosen uniformly at random from the neighbors of i in G, $\mathbf{E}[H(i,J)] = 2|E|/d(i) - 1$, where $d(i)$ is the degree of i in G.*

This gives us the following bounds.

Lemma 2.36. *Assume G is a graph with n nodes.*

- *For G being a toroidal mesh (including cycles and hypercubes), or being a d-regular graph with $d \geq \lfloor n/2 \rfloor$, we have $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) = \mathrm{O}(n)$*
- *For an arbitrary d-regular graph G we have $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) \leq dn$.*

Proof. Recall that $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) \coloneqq \max_{i,j\in V,\{i,j\}\in E} H(i,j)$. Toroidal meshes are *symmetric* or *arc-transitive* graphs: for every two ordered pairs of adjacent nodes (i_1, j_1) and (i_2, j_2) there is a graph automorphism f such that $f(i_1) = i_2$ and $f(j_1) = j_2$. Hence, for every such two ordered pairs, $H(i_1, i_1) = H(i_2, j_2)$, and thus $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) = H(i,j)$ for any pair of adjacent nodes i, j. So applying Theorem 2.35 shows that $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) = 2|E|/d - 1$. As $|E| = dn/2$ for d-regular graphs, $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) = 2(dn/2)/d - 1 = n - 1 = \mathrm{O}(n)$, as claimed.

For dense graphs we bound $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G})$ as $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) \leq \mathrm{t_{hit}}(\mathrm{G}) \leq 2 \cdot \mathrm{Res}(G) \cdot |E|$ (see Lemma 6.21). As $\mathrm{Res}(G) = \mathrm{O}(1/d)$ by [34, Theorem 3.3], we get since $|E| = dn/2$ that $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) = \mathrm{O}(dn/d) = \mathrm{O}(n)$.

For arbitrary d-regular graphs, $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) \leq 2 \cdot \mathrm{Res}^*(G) \cdot |E|$ by the first statement of Lemma 6.21. As $|E| = dn/2$ for a d-regular graph, and as $\mathrm{Res}^*(G) \leq 1$ (by definition of $\mathrm{Res}^*(G)$ and Lemma 6.16), we thus have $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) \leq 2 \cdot 1 \cdot dn/2 = dn$. □

2.7.3 Asynchronous Model

Combining the bounds on $T(G)$ and on the hitting time with Theorem 2.10 leads us to the following results w.h.p. and in expectation.

Corollary 2.37. *Let $X(t)$ be the state of process* $\mathrm{ABAL}(\mathcal{D}_1(G), 1)$ *where $X(0) = \vec{0}$. For an arbitrary t it holds w.h.p. and in expectation*

1. $\mathrm{disc}(X(t)) = \mathrm{O}(\sqrt{n}\log(n))$ *for any regular graph.*
2. $\mathrm{disc}(X(t)) = \mathrm{O}(\sqrt{n\log(n)})$ *for cycle and constant-degree regular graphs.*
3. $\mathrm{disc}(X(t)) = \mathrm{O}(\log^{3/2}(n))$ *for the two-dimensional torus graph.*
4. $\mathrm{disc}(X(t)) = \mathrm{O}(\log(n))$ *for r-dimensional torus graphs with $r \geq 3$ dimensions, for the hypercube, and for all d-regular graphs with $d \geq \lfloor n/2 \rfloor$.*

2.8 Summary and Open Problems

In this part we analyze discrete load balancing processes on graphs. As our main contribution we bound the discrepancy that arises in dynamic load balancing in three models, the random matching model, the balancing circuit model, and the asynchronous model. Our results for the random matching model and the asynchronous model hold for d-regular graphs, while our analysis for the balancing circuit model applies to arbitrary graphs.

To the best of our knowledge our results constitute the first bounds for discrete, dynamic balancing processes on graphs. Furthermore, our results improve the work by Alistarh et al. [6] who prove that the expected discrepancy is bounded by $\sqrt{n}\log(n)$ in the (arguably simpler) continuousasynchronous process $\text{ABAL}^{(\text{cont})}(\mathcal{D}_1(G), 1)$. We improve their bound to $\sqrt{n\log(n)}$ and additionally show that it holds with high probability. We conjecture that our results are tight up to polylogarithmic factors. However, showing tight upper and lower bounds remains an open problem.

We are confident that our results carry over to arbitrary graphs (as opposed to regular graphs), provided that there exists a lower bound on the probability p_{min} with which an edge is used for balancing. However, to show bounds on the discrepancy one has to overcome fundamental problems such as the bias introduced by high-degree nodes. Another interesting open question is whether the results carry over to a model where the amount of load that may transmitted over an edge in each step is bounded by a constant. If only a single load item can be transferred per edge and step the problem is similar to the token distribution problem (see, for example, [55]).

Finally, we believe that one can also adapt our analysis to variant of a graphical balls-into-bins process. The process works as follows. In each step an edge (i, j) is sampled uniformly at random. W.l.o.g. assume that the load of i is smaller than the load of j by an additive term δ. Then a biased coin is tossed showing heads with probability $p := \min\{1, (1 + \beta \cdot \delta)/2\}$ and tails otherwise, where β is a suitably chosen and non-constant parameter. If the coin hits heads one item is allocated to i and otherwise to j. A formal analysis of this allocation process (as well as of other, related balls-into-bins processes) is beyond the scope of our paper and remains an open problem.

Part Three:

3 Discrete Static Load Balancing on Matchings

3.1 Introduction

In this part, we investigate 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. The 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.

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 [82]. 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 [82], 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 by Ghosh and Muthukrishnan [51], Lovász and Winkler [68], Muthukrishnan et al. [78], and Subramanian and Scherson [86], 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 [14, 47, 82, 85], is as follows: *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 $\tau_S(K, 1)$ denote the number of rounds needed in the idealized (continues) setting to reach discrepancy at most 1. Find a tight bound on the discrepancy after $\mathrm{O}(\tau_S(K, 1))$ rounds in the discrete model.*

Results in a Nutshell. As is commonly done we bound the convergence time in terms of the natural spectral bound $\widetilde{\tau}_{\mathrm{S}}(K) = \log(Kn)/(1 - \lambda)$ on $\tau_{\mathrm{S}}(K, 1)$, instead of $\tau_{\mathrm{S}}(K, 1)$ directly. Furthermore, it was shown in [85, 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 [85] 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 [85] 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 [85].

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 we show that for instances with an *arbitrary* number of tokens, after $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds the number of tokens above the average is $\mathrm{O}(n)$. Hence we can apply the result for a linear number of tokens, which results into a discrepancy of 4 after $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds. Using another $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds we finally reduce the discrepancy down to 3.

Techniques and Comparisons to prior works. 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 [47, 78, 82, 85], which either focus only on one specific matching model or provide tailored analyses for each of them. To unify the three

Table 3: Overview of related results for the balancing circuit model. The stated results from [47, 85], as well as ours, also hold for the random matching model. $\widetilde{\tau}_S(K)$ is the spectral bound for the continuoussetting, and it is defined in Equations (3.62) and (3.63), respectively. All runtime bounds for randomized rounding hold with probability at least $1 - o(1)$.

Reference	Rounds	Discrepancy	Rounding	Graphs
[82]	$O(\widetilde{\tau}_S(K))$	$O(\widetilde{\tau}_S(n))$	det	all
[47]	$O(\widetilde{\tau}_S(K))$	$O\left(\sqrt{\widetilde{\tau}_S(n)}\right)$	rand	all
[47]	$O\big(\widetilde{\tau}_S(K) \cdot (\log\log(n))^3\big)$	$O(1)$	rand	expander
[71]	$2\log_2(n)$	16	rand	hypercube
[70]	$O(\log(n))$	2	rand	hypercube
[85]	$O(\tau_S(K, 1))$	$O(\log^{\varepsilon} n)$	rand	all
[85]	$O(\tau_S(K, 1) \cdot \log\log(n))$	$O(\log\log(n))$	rand	all
[85]	$O(\widetilde{\tau}_S(K))$	$O(1)$	rand	constant degree, $\Delta = O(1)$
Thm. 3.1	$O(\widetilde{\tau}_S(K))$	3	rand	all

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 to [85], the constant factors in both of our running time and discrepancy are explicit and fairly small. In addition to our analysis being tighter and simpler, we also obtain a larger success probability. Finally, the results of [85] 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 [47, 85] 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 uses 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 satisfy a negative association property. This property, together with the Hoeffding-type concentration, is then used in an involved analysis to show that eventually a discrepancy 3 is reached.

3.1.1 Outline

Section 3.2 introduces notations, the standard discrete load-balancing and height-sensitive processes, and defines important concepts. Section 3.3 presents the main theorem and the propositions used in its proof, concluding with an overview of our proof techniques. Section 3.4 lists the technical lemmas with formal analysis, correlation, and concentration results for load vectors, which form the core tools of our analysis. Section 3.5 provides results for specific models: the balancing circuit, random matching, and asynchronous models. Finally, Section 3.6 summarizes the main results and discusses open problems.

3.2 Model and Definitions

We begin by introducing the notation. Subsection 3.2.1 defines the the standard load balancing and the height-sensitive processes, while section 3.2.2 recalls and defines useful definitions.

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. Recall that $\overline{x} = \sum_{i \in V} x_i(0)/n$ is the average load and $\mathrm{disc}(X) = \max_{i \in [n]} X_i - \min_{j \in [n]} X_j$ is the discrepancy of load vector X. We usually assume that the tokens on the nodes are ordered; the height of the token is its number in that order.

3.2.1 Process Definition

Here we first define the standard discrete balancing process which does not specify how tokens are exchanged across the matching edges (similar to [82, 85]). 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 $\mathbf{M}^{[t]} := (\mathbf{M}^{(1)}, \mathbf{M}^{(2)}, \ldots, \mathbf{M}^{(t)})$. The standard load balancing process updates the (discrete) load vector iteratively as follows.

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

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

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* $H_i(t) \in \{1, \ldots, X_{W_i(t)}(t)\}$. Initially tokens are ordered arbitrarily on each node, and the initial height of a token is its position in that order. Two tokens on the two adjacent nodes with the same height are called *sibling*. 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).

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 $\lfloor (X_u(t-1) - X_v(t-1))/2 \rfloor$ tokens from node u to node v, preserving their relative order and adjusting their height accordingly.

2. *Shuffling step:* Swap each token on node u with its sibling on v with probability 1/2, independently from all other tokens. In case where the topmost token at u has no sibling, this token is also moved to v with probability 1/2.

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 3.11, (i) in Section 3.4). Furthermore, *any individual* token performs a random walk with the sequence of matching matrices as transition matrices (Lemma 3.11, (ii)). Crucially, we will prove later that the movements of *different* tokens satisfy a negative association property (see Lemma 3.12).

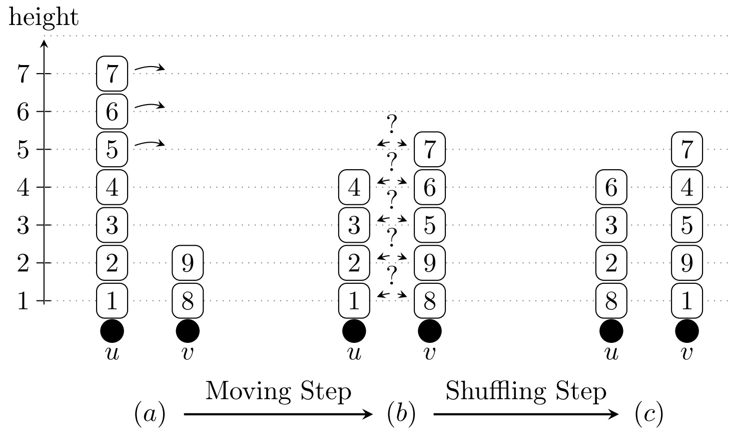

Figure 2: Illustration of the height-sensitive process and the effect on the tokens, which are labeled from 1 to 9. (a): The configuration of tokens before the averaging; (b): the configuration after the moving step and (c): the configuration after the shuffling step.

3.2.2 Properties of Matchings

Assume that nodes are labeled from 1 to n. Recall that the balancing matrix $\mathbf{M}^{(t)} \in [0,1]^{n\times n}$ represents a matching of G in round t as $\mathbf{M}^{(t)}_{u,v} := 1/2$ if $u \neq v$ are matched in round t, $\mathbf{M}^{(t)}_{u,v} := 0$ if $u \neq v$ are not matched in round t, and $\mathbf{M}^{(t)}_{u,u} := 1$ if u is not matched in round t. We define use $\mathbf{M}^{[t_1+1,t_2]}_{u,D} := \sum_{v\in D} \mathbf{M}^{[t_1+1,t_2]}_{u,v}$. By $\vec{1}$ we denote the row vector of length n in which each entry is 1.

Recall that a fixed sequence of matchings $\mathbf{M}^{[t]}$, $t \geq 1$ is called (K,ε)-smoothing if for any $x(0) \in \mathbb{R}^n$ with $\mathrm{disc}(x(0)) \leq K$ we have $\mathrm{disc}(\mathbf{M}^{[1,t]} \cdot x(0)) \leq \varepsilon$. The expression $\mathbf{M}^{[1,t]} \cdot x(0)$ equals the load vector of the continuous load balancing process with initial load $x(0)$ applying the matchings $\mathbf{M}^{[t]}$. It states that the matching sequence is sufficient to balance the load up to ε in the continuous model (see Definition 1.6).

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 3.1. *A sequence of matchings* $\mathbf{M}^{[\infty]} := \left(\mathbf{M}^{(s)}\right)_{s=1}^{\infty}$ *is called* $(\tau_{\mathrm{global}}, \tau_{\mathrm{local}})$-good *if*

$$\inf_{t\in\mathbb{N}_0} \mathbf{Pr}\left[\bigcap_{u\in V} \left\| \mathbf{M}^{[t+1,t+\tau_{\mathrm{global}}]}_{u,\cdot} - \frac{\vec{1}}{n}\right\|_2^2 \leq \frac{1}{n^7}\right] \geq 1 - \frac{1}{n^3}, \tag{3.33}$$

$$\text{and } \inf_{t\in\mathbb{N}_0} \min_{u\in V} \mathbf{Pr}\left[\left\| \mathbf{M}^{[t+1,t+\tau_{\mathrm{local}}]}_{u,\cdot}\right\|_2^2 \leq \frac{1}{\log^{10}(n)}\right] \geq 1 - \frac{1}{\log^{11}(n)}. \tag{3.34}$$

Since each token is performing a random walk according to the matching sequence (Lemma 3.11), 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 Equation (3.33) 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 property of a random walk, since it holds for *any* start node u. With regards to the event in Equation (3.34), it only a 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 $\tau_{\mathrm{global}} = \mathrm{O}(\log(n)/(1-\lambda))$ and $\tau_{\mathrm{local}} = \mathrm{O}(\log\log(n)/(1-\lambda))$ (assuming $p_{\min} = \Omega(1/\Delta)$).

3.3 $(\tau_{\text{global}}\ ,\ \tau_{\text{local}})$-Good Sequence

In this section we analyze $(\tau_{\text{global}}\ ,\ \tau_{\text{local}})$-good sequence of matchings. First we mention and prove the main theorem of this section. In subsection 3.3.1 we show any arbitrary initial load vector reaches a load vector with at most $\mathrm{O}(n)$ tokens above average, after smoothing time (Lemma 3.3). Subsection 3.3.2 shows how to balance a linear number of tokens. It contains the backbone of our analysis. There we show the discrepancy 3 using our developed techniques. Finally we provide a high level of the analysis in subsection 3.3.3. Combining these results, we obtain our main theorem below.

Theorem 3.1. *Let G be any undirected, connected graph on n nodes and consider any initial load vector $x(0) \in \mathbb{N}_0^n$ with $\operatorname{disc}(x(0)) := K > 1$. Assume our process balances the tokens via a $(\tau_{\text{global}}, \tau_{\text{local}})$-good sequence of matchings $\mathbf{M}^{[\infty]}$. Then there exists a time τ with*

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

such that

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

and for any constant $c > 0$,

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

Proof. 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

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

and

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

From Lemma 3.2 (presented below) it follows that $\mathbf{Pr}\left[\overline{\mathcal{G}^\star}\right] \leq n^{-2}$. From Lemma 3.3 (presented in subsection 3.3.1) it follows that

$$\mathbf{Pr}\left[\,\overline{\mathcal{G}_0} \;\middle|\; \mathcal{G}^\star\,\right] \leq 2 \cdot n^{-2}.$$

Note that by the law of total probability we get

$$\begin{aligned}\mathbf{Pr}\left[\overline{\mathcal{G}_0}\right] &= \mathbf{Pr}\left[\,\overline{\mathcal{G}_0} \;\middle|\; \mathcal{G}^\star\,\right] \cdot \mathbf{Pr}[\mathcal{G}^\star] + \mathbf{Pr}\left[\,\overline{\mathcal{G}_0} \;\middle|\; \overline{\mathcal{G}^\star}\,\right] \cdot \mathbf{Pr}\left[\overline{\mathcal{G}^\star}\right] \\ &\leq \mathbf{Pr}\left[\,\overline{\mathcal{G}_0} \;\middle|\; \mathcal{G}^\star\,\right] + \mathbf{Pr}\left[\overline{\mathcal{G}^\star}\right] \leq 3 \cdot n^{-3}. \end{aligned} \tag{3.35}$$

Let $t_1 := t_0 + 2 \cdot \tau_{\text{global}} + 6\ell \cdot \tau_{\text{local}}$. Here we define an event $\mathcal{G}_1 := \{\operatorname{disc}(X(t_1)) \leq 38\}$ and from Lemma 3.5 (presented in subsection 3.3.2) it follows that

$$\mathbf{Pr}\left[\overline{\mathcal{G}_1} \;\middle|\; \mathcal{G}_0\right] \leq \exp(-(1/80) \cdot \ell),$$

and by the law of total probability (Equation (3.35)) we get, $\mathbf{Pr}\left[\overline{\mathcal{G}_1}\right] \leq \exp(-(1/80) \cdot \ell + 3 \cdot n^{-3}$

Let $t_2 := t_1 + 2 \cdot \tau_{\text{global}} + 6\ell \cdot \tau_{\text{local}}$. We define another event $\mathcal{G}_2 := \{\operatorname{disc}(X(t_2)) \leq 4\}$, and from Lemma 3.6 (presented in subsection 3.3.2) it follows that

$$\mathbf{Pr}\left[\,\overline{\mathcal{G}_2} \;\middle|\; \mathcal{G}_1\,\right] \leq \exp(-(1/160) \cdot \ell),$$

and by the law of total probability (eq. (3.35)) we get,

$$\mathbf{Pr}\left[\overline{\mathcal{G}_2}\right] \le \exp(-(1/160) + \exp(-(1/80)\cdot\ell + 3\cdot n^{-3} \le \exp(-(1/200)\cdot\ell).$$

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 (with $\tau := t_2$).

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 $\mathcal{G}_3 := \{\text{disc}(X(\tau)) \le 3\}$. From Lemma 3.7 (presented in subsection 3.3.2) with $\delta := c/2$ it follows that

$$\mathbf{Pr}\left[\overline{\mathcal{G}_3} \;\middle|\; \mathcal{G}_2\right] \le 2\exp(-\log^{1-c}(n)).$$

By the law of total probability (Equation (3.35)) we get,

$$\begin{aligned} \mathbf{Pr}\left[\overline{\mathcal{G}_3}\right] &\le \mathbf{Pr}\left[\overline{\mathcal{G}_3} \;\middle|\; \mathcal{G}_2\right] + \mathbf{Pr}\left[\overline{\mathcal{G}_2}\right] \\ &\le 2\exp(-\log^{1-c}(n)) + \exp(-(1/200)\cdot\ell) \le \exp(-\log^{1-2c}(n)). \end{aligned}$$

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. □

In the rest of this section we prove the propositions used in the theorem. First 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 3.2. *Consider a sequence of $(\tau_{\text{global}}, \tau_{\text{local}})$-good matchings $\mathbf{M}^{[\infty]}$. Let $K \ge 2n^2$ and $0 < \varepsilon \le 1$. Then for $t^* := (3\cdot\log(\frac{K}{\varepsilon})/\log(n))\cdot\tau_{\text{global}}$,*

$$\mathbf{Pr}\left[\ \mathbf{M}^{[t^\star]} \textit{ is } (K,\varepsilon)\textit{-smoothing}\ \right] \ge 1 - n^{-2}.$$

Proof. Recall that $\mathbf{M}^{[t^\star]} = \left(\mathbf{M}^{(s)}\right)_{s=1}^{t^\star}$. We first define $x := 3\left\lceil\log_{2n^2}(\frac{K}{\varepsilon})\right\rceil$. We will consider x subsequent and disjoint subsequences of matchings, $\left(\mathbf{M}^{(s)}\right)_{s=(i-1)\cdot\tau_{\text{global}}+1}^{i\cdot\tau_{\text{global}}}$, 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 $\left(\mathbf{M}^{(s)}\right)_{s=(i-1)\cdot\tau_{\text{global}}+1}^{i\cdot\tau_{\text{global}}}$ is $(1, 1/(2n^2))$-smoothing and one otherwise.

From Definition 3.1 together with observation 6.31 we get

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

We define $Z := \sum_{i=1}^{x} Z_i$ and an event $\Psi := \{Z \le x/2\}$.

By linearity of expectation we get that $\mathbf{E}[Z] \le x/n^3$. From Markov's inequality it follows that $\mathbf{Pr}\left[Z \ge \frac{x}{2}\right] \le \mathbf{Pr}\left[Z \ge \frac{x}{n}\right] \le n^{-2}$ implying that $\mathbf{Pr}[\Psi] \ge 1 - n^{-2}$.

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

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

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

3.3.1 Reducing the Number of Tokens to Linear

We will now present an application of our strong Hoeffding bound ,Lemma 3.15, to bound the number of tokens above the average load (i.e., $\sum_{w \in V} \max\{X_w(t) - \overline{x}, 0\}$); 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 3.15 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 3.3. *Assume our process applies a sequence of matchings* $\mathbf{M}^{[t_1]}$ *which is* $(K, 1/(2n))$*-smoothing on an arbitrary initial load vector with* $\operatorname{disc}(x(0)) \leq K$*. Then it holds that*

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

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

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

From our Hoeffding bound Lemma 3.15 with $\delta := \sqrt{36 \log(n)}$, $a_w := 1$ and $a_u := 0$ for $u \neq w$ and union bound over all nodes $w \in V$ it follows that

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

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,

$$\mathbf{Pr}\left[\overline{\mathcal{E}_i}\right] \leq \sum_{S \subseteq V \colon |S| = n \cdot 2^{-i}} \mathbf{Pr}\left[\bigcap_{u \in S} \{X_u(t_1) \geq \overline{x} + 4 \cdot i\}\right]. \tag{3.36}$$

To calculate this probability we apply Lemma 3.15 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

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

where (a) follows from Lemma 3.15. Plugging this into Equation (3.36) yields

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

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 $i = 0, 1, \ldots, 2\sqrt{\log(n)}$,

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

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

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

This completes the proof. □

3.3.2 Balancing a Linear Number of Tokens

In this subsection 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 3.4, which states that after $\mathrm{O}(\tau_{\text{global}} + \tau_{\text{local}} \cdot \log(n)/\log\log(n))$ rounds the maximum load is $L+1$. An illustration of the proof method can be found in Figure 3.

The next proposition makes uses of Lemma 3.27 and Lemma 3.28 which are stated and proved in the technical lemmas (Section 3.4).

Proposition 3.4. *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 $1 \le L \le \log^7(n)$ and let $\frac{4}{\log^4(n)} \le \varepsilon < L$ (not necessarily constant). Assume our process applies a $(\tau_{\text{global}}, \tau_{\text{local}})$-good sequence of matchings $\mathbf{M}^{[\infty]}$ to an arbitrary initial load vector $x(0)$ with at most $(L-\varepsilon)n$ tokens. Then*

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

Proof. **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

$$\mathcal{G}_1 := \left\{Y(t_1) \le \frac{n}{\log(n)}\right\}.$$

From Lemma 3.27 (Phase 1 Lemma) it follows that

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

Phase 2. For any $t \ge t_1$, we define an auxiliary load vector $\widehat{X}(t)$ with $\widehat{X}_u(t) := \max\{X_u(t)-L, 0\}$. From observation 3.18 it follows that for any $t \ge t_1$, $X_u(t) \le \widehat{X}_u(t) + L$. Hence it is sufficient to show that $\max_{u\in V}\widehat{X}_u(t_2) \le 1$ for some suitable round t_2. Note that from Phase 1 we get $\sum_{u\in V}\widehat{X}_u(t_1) \le n/\log(n)$, i.e. the load vector $\widehat{X}(t_1)$ has at most $n/\log(n)$ tokens. This time we define, for any round $t \ge t_1$, $\widehat{Y}(t) := \sum_{u\in V}\max\{\widehat{X}_u(t)-1, 0\}$, which is equal to the number of tokens with height at least 2 in $\widehat{X}(t)$. 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

$$\mathcal{G}_2 := \{\widehat{Y}(t_2) = 0\}.$$

From Lemma 3.28 (Phase 2 Lemma) with $\varepsilon := 1 - \frac{1}{\log(n)}$ it follows that $\mathbf{Pr}\left[\overline{\mathcal{G}_2} \mid \mathcal{G}_1\right] \le n^{-2}$. And by the law of total probability (stating that for two events A, B we have $\mathbf{Pr}\left[\overline{A}\right] \le \mathbf{Pr}\left[\overline{A} \mid B\right] + \mathbf{Pr}\left[\overline{B}\right]$) we get

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

From the definition of t_2 it follows that

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

finishing the proof. ∎

In the remainder of this Subsection, we show how to reduce the discrepancy to 3 in three steps: first to 38 (Lemma 3.5), then to 4 (Lemma 3.6), and finally to 3 (Lemma 3.7). For these purposes, we frequently apply the techniques developed earlier (Proposition 3.4).

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 3.5). The proof of the result relies on Lemma 3.3 showing that the number of tokens with height at least $\overline{x}+1$ is at most $16n$.

Lemma 3.5. *Let $t^\star := 2 \cdot \tau_{\text{global}} + \frac{6\log(n)}{\log\log(n)} \cdot \tau_{\text{local}}$. Assume our process applies a $(\tau_{\text{global}}, \tau_{\text{local}})$-good sequence of matchings $\mathbf{M}^{[\infty]}$ to a load vector $x(0)$, which satisfies $\sum_{w\in V} \max\{x_w(0) - \overline{x}, 0\} \leq 16 \cdot n$. Then*

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

Proof. For any $t \geq 0$, define an auxiliary load vector $\widetilde{X}_u(t) := \max\{X_u(t) - \lceil \overline{x} \rceil, 0\}$, $u \in V$, that is, we subtract $\lceil \overline{x} \rceil$ tokens from any node (as long as its load is large enough). By assumption, we have $\sum_{w\in V} \widetilde{x}_w(0) \leq 16 \cdot n$. Here we can apply Proposition 3.4 with $L = 17$, $\varepsilon := 1/2$ and $t^\star := 2 \cdot \tau_{\text{global}} + \frac{6\log(n)}{\log\log(n)} \cdot \tau_{\text{local}}$ and obtain for $c := 1/70$,

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

Applying Observation 3.18 we get that the same holds for $X(t^\star)$, resulting in

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

Using a simple symmetry argument (see observation 6.30),

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

A final union bound gives

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

which yields the statement. ∎

Reducing Discrepancy from 38 to 4. Here we show that once the load vector consists of only $O(n)$ tokens and the initial discrepancy is at most 38, then after additional $O(\tau_{\text{global}} + \frac{\log(n)}{\log\log(n)} \cdot \tau_{\text{local}})$ rounds the discrepancy is reduced to 4.

Lemma 3.6. *Assume our process applies a $(\tau_{\text{global}}, \tau_{\text{local}})$-good sequence of matchings $\mathbf{M}^{[\infty]}$ to a load vector $x(0)$, in which $\operatorname{disc}(x(0)) \leq 38$. Then for $t^\star := 2 \cdot \tau_{\text{global}} + 6\tau_{\text{local}} \cdot \frac{\log(n)}{\log\log(n)}$, we have,*

$$\mathbf{Pr}[\operatorname{disc}(X(t^\star)) \leq 4] \geq 1 - \exp\left(-(1/160) \cdot \frac{\log(n)}{\log\log(n)}\right).$$

Proof. Assume without loss of generality that the load values are $\{0, 1, 2, \ldots, 38\}$. Let us pick $L := \lceil \overline{x} + \frac{1}{2} \rceil$ and $\varepsilon := \frac{1}{2}$; clearly $L \leq 39$ since $\overline{x} \leq 38$. Using Proposition 3.4, we obtain that at round $t^\star := 2 \cdot \tau_{\text{global}} + 6\tau_{\text{local}} \cdot \log(n)/\log\log(n)$ the maximum load is at most $L + 1 = \lceil \overline{x} + \frac{1}{2} \rceil + 1$ 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 $\widetilde{x}(0) := 38 - x(0)$. By observation 3.19, $\widetilde{x}(t) = 38 - x(t)$ for all $t \geq 1$. Also, $\overline{(\widetilde{x})} = 38 - \overline{x}$. Repeating the above argument, but now applied to $\widetilde{x}$, yields for any $u \in V$,

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

which implies

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

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 $x(t^\star)$ is at least $\lfloor \overline{x} - \frac{1}{2} \rfloor - 1$. Since all load values at round $t^\star$ are integers in the interval $\left[\lfloor \overline{x} - \frac{1}{2} \rfloor - 1, \lceil \overline{x} + \frac{1}{2} \rceil + 1\right]$, which are at most 5 values, the discrepancy is at most 4. Hence a final union bound gives,

$$\begin{aligned}\mathbf{Pr}[\operatorname{disc}(X(t^\star)) \leq 4] &\geq 1 - 2\exp\left(-\frac{\log(n)}{156 \cdot \log\log(n)} + 8 \cdot \log\log(n)\right) - 4 \cdot n^{-2} \\ &\geq 1 - \exp\left(-\frac{\log(n)}{160 \cdot \log\log(n)}\right).\end{aligned}$$

□

Reducing Discrepancy from 4 to 3. To obtain discrepancy 3, we may not be able to apply Proposition 3.4 directly. For example, if $\overline{x}$ is an integer (or close to an integer), the highest load is $\overline{x} + 2$, and the minimum load is $\overline{x} - 2$, then there is no L and ε such that the application of Proposition 3.4 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 3.4, that most of the primary tokens reduce their height to at most $\overline{x}$. 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 3.28 we obtain that all tokens at height $\overline{x} + 2$ will reduce their height.

Lemma 3.7. *Assume our process applies a* $(\tau_{\text{global}}, \tau_{\text{local}})$*-good sequence of matchings* $\mathbf{M}^{[\infty]}$ *to a load vector* $x(0)$*, in which* $\operatorname{disc}(x(0)) \leq 4$*. Then for any* $0 < \delta < 1/2$ *and* $t^* := 2\tau_{\text{global}} + \frac{5}{\delta \log\log(n)} \cdot \log(n) \cdot \tau_{\text{local}}$*,*

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

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 $\overline{x} \leq 2$. Ideally, we would like to apply Proposition 3.4 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 $\overline{x} \in (2 - \frac{4}{\log^4 n}, 2]$. To overcome this problem, we will first apply the first phase in the proof of Proposition 3.4 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 $\mathbf{M}^{[\infty]}$, 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 3.27 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

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

According to the second statement of observation 6.29, 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

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

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 3.28, by round $t^* := t_1 + \tau_{\text{global}} + \frac{4}{-\log\left(1 - \varepsilon + \frac{2}{\log^4(n)}\right)} \cdot \log(n) \cdot \tau_{\text{local}}$, 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

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

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

3.3.3 An Outline of the Analysis

Before presenting the technical lemmas in detail, we provide a more detailed outline of our analysis. We compile a collection of the most important technical results, which form the core tools of our analysis. The proofs of these tools are significantly more involved than the derivation of the discrepancy bounds using them. Moreover, we believe that several of these tools are of independent interest.

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 3.11). Crucially, we establish that these random walks are negatively correlated:

Lemma 3.8 (simplified version of Lemma 3.12)**.** *Fix a subset of tokens $\mathcal{B} \subseteq \mathcal{T}$ at round 0. Let $t > 0$ be any round, and fix the matchings from round 1 and t. Then for any set $D \subseteq V$, the events $\{W_i(t) \in D\}, i \in \mathcal{B}$ are negatively correlated.*

In comparison to previous work (Lemma 4.2 from [85]), our lemma yields the same statement but here tokens move following the definition of the height-sensitive process, whereas in [85], 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 [85], e.g., we do not need the somewhat unwieldy negative regression condition from [41]. We continue with a Hoeffding-like concentration bound.

Lemma 3.9 (simplified version of Lemma 3.15)**.** *Consider any load vector $x(0)$ with $\operatorname{disc}(x(0)) \leq K$ 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 $\delta > 0$,*

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

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 3.15. 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

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

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 [85], 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 [85, Lemma 3.5] to *arbitrary* stochastic vectors a. The corresponding result in [85] 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 $t_1 \geq t-1$; 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 [47, 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 3.15) is to bound the number of tokens above the average load in Lemma 3.3. This lemma exploits that in Lemma 3.15 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 $\overline{x} + 4 \cdot i$. This is done by choosing an appropriate vector (a_w) and threshold δ in Lemma 3.15. The proof is then concluded by a simple union bound over all possible "bad" subsets $S \subseteq V$.

The second application of our concentration inequality (Lemma 3.15) is Proposition 3.4. It 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 $\mathrm{O}(n)$ tokens. The proof of this proposition constitutes the most challenging part

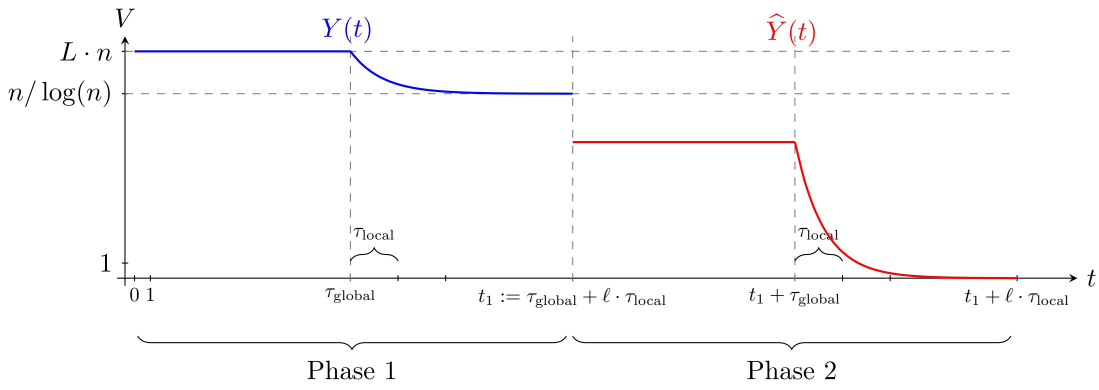

Figure 3: Illustration of phases 1 and 2 in the proof of proposition 3.4. Phase 1 decreases $Y(t)$, the number of tokens with height at least $L+1$, from $n \cdot L$ to $\frac{n}{\log(n)}$. Then Phase 2 decreases $\widehat{Y}(t)$, the number of tokens with height at least $L+2$, to 0. Both phases first use τ_{global} rounds for a "global mixing", and then $\ell := \frac{\log(n)}{\log\log(n)}$ short epochs of length τ_{local}.

of our analysis. Note that in this proposition we consider a generalized statement and instead of $\mathrm{O}(n)$ tokens we balance $(L-\varepsilon)\cdot n$ tokens for a constant $\varepsilon > 0$ and integer $\log^7(n) \geq L \geq 1$. Our goal is to gradually reduce the number of tokens with height $\geq L+1$, until no token with height $\geq L+2$ remains. Proposition 3.4 uses two phases, which are illustrated in Figure 3. In Phase 1 we reduce the number of tokens with height $\geq L+1$ from $(L-\varepsilon)n$ to $n/\log(n)$ (see Lemma 3.27) and in Phase 2 we reduce the number of tokens with height $\geq L+2$ further to 0 (see Lemma 3.28).

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 3.10 (simplified and informal version of Lemma 3.24 (Key Lemma))**.** *Assume that a load vector $x(t)$ has at most $(L-\varepsilon)n$ tokens, where $0 < t$, $0 < \varepsilon < 1$ 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,*

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

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 3.11); 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 correlation lemma (Lemma 3.12), the expected number of tokens which collide with token i at round $t+\tau_{\text{local}}$ can be upper bounded by (see Lemma 3.21),

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

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 our Hoeffding bound for convex combinations of loads (Lemma 3.15) 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 behavior 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.

3.4 Technical lemmas

Here we provide the detailed proofs of our technical lemmas. Our first goal is to prove the two main components of our analysis; first, the negative correlation result about token movements (Lemma 3.12 in subsection 3.4.1) and secondly, the Hoeffding-like concentration inequality (Lemma 3.15 in subsection 3.4.2). Using these two tools we are able to prove the intermediate results.

We start with a simple but useful observation indicating that the height of a token is non-increasing and a single token follows a random walk with transition matrices $\mathbf{M}^{(s)}$ for $s\in[t_1,t_2]$.

Lemma 3.11. *Consider any pair of rounds $t_1\leq t_2$ and let $(\mathbf{M}^{(s)})_{s=t_1}^{t_2}$ 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) $\mathbf{Pr}[\,W_i(t_2)=v\mid w_i(t_1)=u\,]=\mathbf{M}^{[t_1+1,t_2]}_{u,v}$.*

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 [85, 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 $\mathbf{Pr}[\,W_i(t)=v\mid w_i(t_1)=u\,]=\mathbf{M}^{[t_1+1,t]}_{u,v}$. For $t=t_1$, $\mathbf{M}^{[t_1+1,t]}$ is the identity matrix, which means that the induction base holds. For the induction hypothesis, consider $\mathbf{Pr}[\,W_i(t)=v\mid w_i(t_1)=u\,]$. If v is not part of a matching in round t, then

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

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

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

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

3.4.1 Height-Sensitive Negative Association

We next present our Height-Sensitive Negative Association result, which can be interpreted as a form of the negative covariance property introduced in [42]. In essence, this result establishes that the locations of tokens exhibit a negative dependence structure: the probability that a given set of tokens simultaneously lies within a fixed subset D is upper-bounded by the product of their individual probabilities of being in D

Lemma 3.12 (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 $w(t_1) = (w_i(t_1))_{i\in\mathcal{T}}$ 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*

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

Proof. Throughout the proof we always assume that $w(t_1)$ is fixed and we omit the conditioning. Note that by Lemma 3.11, we have $\prod_{i\in\mathcal{B}} \mathbf{Pr}[W_i(t_2) \in D] = \prod_{i\in\mathcal{B}} \mathbf{M}_{w_i(t_1),D}^{[t_1+1,t_2]}$ and therefore it only remains to prove the inequality

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

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

$$Z_i(t) := \sum_{u\in V} \mathbf{1}_{W_i(t)=u} \cdot \mathbf{M}_{u,D}^{[t+1,t_2]} \qquad \text{and} \qquad Z(t) := \prod_{i\in\mathcal{B}} Z_i(t). \tag{3.37}$$

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,

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

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 $W(t) = \big(W_1(t), \ldots, W_{|\mathcal{T}|}(t)\big)$), that is,

$$\mathbf{E}\Big[\, Z(t) \;\Big|\; \mathfrak{F}^{(t-1)} \,\Big] \le Z(t-1), \tag{3.38}$$

which would immediately imply

$$\mathbf{E}[Z(t_2)] \le Z(t_1).$$

It then follows

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

finishing the proof.

It remains to prove Equation (3.38). 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 $t_1 \le t \le t_2$ 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

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

For ease of presentation we define

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

From the definition (3.37) it follows that

$$Z(t-1) = \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i(t-1) \cdot \prod_{i \in T_u \cup T_v} Z_i(t-1) = \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i(t-1) \cdot \left(\frac{p+q}{2}\right)^{|T_u \cup T_v|}. \tag{3.39}$$

We will continue to analyze $\mathbf{E}\big[Z(t) \mid \mathfrak{F}^{(t-1)}\big]$ and eventually upper bound it by the right hand side in Equation (3.39), 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 $\mathbf{E}\Big[Z(t) \;\Big|\; \mathfrak{F}^t\Big]$,

$$\mathbf{E}\Big[Z(t) \;\Big|\; \mathfrak{F}^{(t-1)}\Big] = \mathbf{E}\Bigg[\prod_{i \in B} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg] = \mathbf{E}\Bigg[\prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i(t) \cdot \prod_{i \in T_u \cup T_v} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg],$$

and since for tokens $i \in \mathcal{B} \setminus (T_u \cup T_v)$, $W_i(t)$ is not matched with any node in this round, then $Z_i(t) = Z_i(t-1)$ and the above is

$$= \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i(t-1) \cdot \mathbf{E}\Bigg[\prod_{i \in \mathcal{B}_1} Z_i(t) \cdot \prod_{i \in \mathcal{B}_2} Z_i(t) \cdot \prod_{i \in \mathcal{B}_3} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg]$$

and 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, the above is

$$= \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i(t-1) \cdot \mathbf{E}\Bigg[\prod_{i \in \mathcal{B}_1} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg] \cdot \mathbf{E}\Bigg[\prod_{i \in \mathcal{B}_2} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg] \cdot \mathbf{E}\Bigg[\prod_{i \in \mathcal{B}_3} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg]. \tag{3.40}$$

We will now analyze the three different expectations in Equation (3.40)

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

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

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,

$$\mathbf{E}\Bigg[\prod_{i \in \mathcal{B}_2} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg] = \prod_{i \in \mathcal{B}_2} \mathbf{E}\Big[Z_i(t) \;\Big|\; \mathfrak{F}^{(t-1)}\Big] = \prod_{i \in B_2} \left(\frac{p+q}{2}\right).$$

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,

$$\mathbf{E}\Bigg[\prod_{i \in \mathcal{B}_3} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg] = \prod_{i \in B_3} \left(\frac{p+q}{2}\right).$$

Aggregating the contributions of all tokens in $\mathcal{B}$ in round t and using Equation (3.40) we get

$$\mathbf{E}\Big[Z(t) \;\Big|\; \mathfrak{F}^{(t-1)}\Big]$$
$$= \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i(t-1) \cdot \mathbf{E}\Bigg[\prod_{i \in \mathcal{B}_1} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg] \cdot \mathbf{E}\Bigg[\prod_{i \in \mathcal{B}_2} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg] \cdot \mathbf{E}\Bigg[\prod_{i \in \mathcal{B}_3} Z_i(t) \;\Bigg|\; \mathfrak{F}^{(t-1)}\Bigg]$$

and using Cases 1, 2 and 3 above gives us,

$$\mathbf{E}\Big[\, Z(t) \;\Big|\; \mathfrak{F}^{(t-1)} \,\Big] \overset{(a)}{=} \prod_{i \in \mathcal{B} \setminus (T_u \cup T_v)} Z_i(t-1) \cdot \prod_{i<j \in \mathcal{B}_1 \colon i \sim j} (p \cdot q) \cdot \prod_{i \in \mathcal{B}_2} \frac{p+q}{2} \cdot \prod_{i \in \mathcal{B}_3} \frac{p+q}{2}$$

and using the simple fact the fact that $p \cdot q \leq ((p+q)/2)^2$ (this is a special case of Lemma 6.35, but can be also easily verified by expanding) leads to

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

where the last equality follows form Equation (3.39). Hence $\mathbf{E}\Big[Z(t) \;\Big|\; \mathfrak{F}^{(t-1)}\Big] \leq Z(t-1)$, which is Equation (3.38) and, consequently, completes the proof. □

3.4.2 Concentration Inequality for Convex Combinations of Loads

In this subsection we are to prove our strong Hoeffding bound, Concentration Inequality for Convex Combinations of Loads (Lemma 3.15). This "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. Before proving it, we mention some basic results which will be used in the proof of Lemma 3.15.

We will first state a Hoeffding-like concentration inequality from a related work. Recall that $\varepsilon_{u,v}(s)$ is the rounding error for node u in round s when $[u : v] \in \mathbf{M}^{(s)}$ and from the definition we have $\varepsilon_{u,v}(s) = -\varepsilon_{v,u}(s)$.

Lemma 3.13 (Lemma 3.4 in [85])**.** *Consider an arbitrary but fixed sequence of matchings $\mathbf{M}^{[t]}$ for any rounds $t \geq 1$ and the load vector $x(0)$. For any family of numbers $g^{(s)}_{u,v}, [u : v] \in \mathbf{M}^{(s)}, 1 \leq s \leq t$, define the random variable*

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

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

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

The next lemma can be seen as a generalization of the first statement of Lemma 3.2 from [85], 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 3.15 in its full generality.

Lemma 3.14. *Consider an arbitrary sequence of matchings $\mathbf{M}^{[\infty]}$. Let $(a_k)_{k \in V}$ be any stochastic vector. Then:*

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

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

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

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

Proof. The proof follows closely the proof of Theorem 3.2 from [85]. 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 $1 \le s \le t$. This is a generalization of the potential function in [85], where all $a_k := 1$.

Proof of the first statement. Consider now any round $1 \le s \le t$, 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

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

and rearranging it gives,

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

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

$$
\Psi(s)-\Psi(s-1) = \frac{1}{2}\cdot\sum_{[u:v]\in\mathbf{M}^{(s)}}\left(\sum_{k\in V} a_k\cdot\left(\mathbf{M}^{[s+1,t]}_{u,k} - \mathbf{M}^{[s+1,t]}_{v,k}\right)\right)^2. \tag{3.41}
$$

Therefore,

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

and rearranging gives,

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

where the last inequality holds since the potential $\Psi(s)$ is non-decreasing over s (Equation (3.41)). Moreover, since $\mathbf{M}^{[t+1,t]}$ is the identity matrix $\mathbf{I}$, we have $\mathbf{M}^{[t+1,t]}_{w,k} = 1$ for $k = w$ and $\mathbf{M}^{[t+1,t]}_{w,k} = 0$ for $k \neq w$. Hence

$$
\Psi(t) = \sum_{w\in V}\left(\sum_{k\in V} a_k \mathbf{M}^{[t+1,t]}_{w,k} - \frac{1}{n}\right)^2 = \sum_{w\in V}\left(a_w - \frac{1}{n}\right)^2 = \|a\|_2^2 - \frac{1}{n}. \tag{3.42}
$$

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

Proof of the second statement. Note that since $t_1 \le t$, it follows from Equation (3.41) that $\Psi(t_1) \le \Psi(t)$. Substituting the definition of $\Psi(t_1)$ and $\Psi(t)$ on both sides gives

$$
\sum_{w\in V}\left(\sum_{k\in V} a_k\cdot\mathbf{M}^{[t_1+1,t]}_{w,k} - \frac{1}{n}\right)^2 \le \sum_{w\in V}\left(\sum_{k\in V} a_k\cdot\mathbf{M}^{[t+1,t]}_{w,k} - \frac{1}{n}\right)^2 \stackrel{\text{Eq. (3.42)}}{=} \|a\|_2^2 - \frac{1}{n}. \tag{3.43}
$$

Finally, we have

$$
\sum_{w\in V}\left(\sum_{k\in V} a_k\cdot\mathbf{M}^{[t_1+1,t]}_{w,k}\right)^2 \stackrel{(a)}{=} \sum_{w\in V}\left(\sum_{k\in V} a_k\cdot\mathbf{M}^{[t_1+1,t]}_{w,k} - \frac{1}{n}\right)^2 + \frac{1}{n} \stackrel{\text{Eq. (3.43)}}{=} \|a\|_2^2,
$$

where (a) follows from observation 6.33 (since the matrix $\mathbf{M}^{[t_1+1,t]}$ is doubly stochastic and $\|a\|_1 = 1$). $\square$

Now we prove the main result of this section.

Lemma 3.15 (Hoeffding Bound for Convex Combinations of Loads)**.** *Consider any load vector $x(0)$ with $\mathrm{disc}(x(0)) \leq K$ and any round $t \geq 1$ such that $\mathbf{M}^{[t]}$ 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 $\delta > 0$,*

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

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 3.16. *Consider any load vector $x(0)$ with $\mathrm{disc}(x(0)) \leq K$. Consider any round $t \geq 1$ and let $\mathbf{M}^{[t]}$ be a fixed $(K,1)$-smoothing sequence of matchings. Then,*

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

The exact same bound was shown in [85, Theorem 3.6].

Proof of Lemma 3.15. Recall from Equation (2) we have

$$X(t) = \mathbf{M}^{(t)} \cdot (X(t-1) + A(t)) + \varepsilon(t), \tag{3.44}$$

in which $X(t), \mathbf{M}^{(t)}, \ell(t)$ and $\varepsilon(t)$ are the load vector, the balancing matrix, the newly added load items and rounding error vector in round t. Then from Equation (5) (by repeatedly expanding above equation form up to the beginning of the process) we get

$$X(t) = \underbrace{\mathbf{M}^{[1,t]} \cdot X(0)}_{\substack{I(t)\\ \text{initial load contribution}}} + \underbrace{\sum_{s=1}^{t} \mathbf{M}^{[s,t]} \cdot A(\tau)}_{\substack{D(t)\\ \text{dynamically allocated Load contribution}}} + \underbrace{\sum_{s=1}^{t} \mathbf{M}^{[s+1,t]} \cdot \varepsilon(s)}_{\substack{R(t)\\ \text{rounding error contribution}}}. \tag{3.45}$$

Since in the continuous setting the contribution of rounding error $\mathrm{disc}(R(t))$ and dynamically allocated load $\mathrm{disc}(D(t))$ is zero, then in the idealized (continues) setting

$$\ell_w(t) := \left(\mathbf{M}^{[1,t]} \cdot X(0)\right)_w = \sum_{u \in V} x_u(0) \cdot \mathbf{M}^{[1,t]}_{u,w},$$

is the load of node w after t rounds. Recall that $\varepsilon_{u,v}(s) = \frac{1}{2} \cdot \mathrm{Odd}(X_u(t-1) + X_v(t-1)) \cdot \phi_{u,v}(t)$ is the rounding error; here $\phi_{u,v}(t) \in \{-1,+1\}$ is the random orientation which has *Rademacher distribution* (which takes each value with probability 1/2). Since $\phi_{u,v}(t) = -\phi_{v,u}(t)$ then we get; for any node $w \in V$,

$$X_w(t) - \ell_w(t) = \left(\sum_{s=1}^{t} \mathbf{M}^{[s+1,t]} \cdot \varepsilon(s)\right)_w = \sum_{s=1}^{t} \sum_{[u:v] \in \mathbf{M}^{(s)}} \varepsilon_{u,v}(s) \cdot \left(\mathbf{M}^{[s+1,t]}_{u,w} - \mathbf{M}^{[s+1,t]}_{v,w}\right), \tag{3.46}$$

Note that this equation is a standard formula (e.g., [85, Equation 2.5]). Therefore,

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

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

$$\sum_{w\in V} a_w \cdot \ell_w(t) = \overline{x} + B,$$

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

$$\sum_{w\in V} a_w \cdot X_w(t) - \overline{x} = B + \sum_{s=1}^{t} \sum_{[u:v]\in\mathbf{M}^{(s)}} \varepsilon_{u,v}(s) \cdot \left[\sum_{w\in V} a_w \cdot \left(\mathbf{M}_{u,w}^{[s+1,t]} - \mathbf{M}_{v,w}^{[s+1,t]}\right)\right] \tag{3.47}$$

Following the notation of Lemma 3.13, let us define

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

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 $\mathbf{E}[\varepsilon_{u,v}(s)] = 0$, then $\mathbf{E}[Z] = 0$ and from Lemma 3.13 it follows that,

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

From the first statement of Lemma 3.14 it follows that $\sum_{s=1}^{t}\sum_{[u:v]\in\mathbf{M}^{(s)}}\left(g_{u,v}^{(s)}\right)^2 \le 2\cdot\|a\|_2^2$. Combining these two gives that $\mathbf{Pr}[|Z| \ge \delta] \le 2 \cdot \exp(-\delta^2/(4\cdot\|a\|_2^2))$, which is the first statement. Combining this with Equation (3.47) implies that,

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

and shifting δ by κ completes the proof. □

Here we provide a similar statement with different constant but with a self-contained proof.

Lemma 3.17. *Consider any load vector $x(0)$ with $\operatorname{disc}(x(0)) \le K$ and any round $t \ge 1$ such that $\mathbf{M}^{[t]}$ 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 $\delta > 0$,*

$$\mathbf{Pr}\left[\left|\sum_{w\in V} a_w \cdot X_w(t) - \overline{x}\right| \ge \delta \;\middle|\; \mathbf{M}^{[t]}\right] \le 4 \cdot \exp\left(-\frac{(\delta-\kappa)^2}{\|a\|_2^2}\right).$$

Proof. Let $|B| \leq \kappa$. Recall that Equation (3.47) implies that

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

For clarity, we first define

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

From the definition it follows that

$$Z = \sum_{s=1}^{t} \sum_{(u,v) \in \mathbf{M}^{(s)}} \varepsilon_{u,v}(s) \cdot g_{u,v}^{(s)} = \frac{1}{2} \sum_{s=1}^{t} \sum_{(u,v) \in \mathbf{M}^{(s)}} \mathrm{Odd}(X_u(s-1) + X_v(s-1)) \cdot \phi_{u,v}(s) \cdot g_{u,v}^{(s)}.$$

Since $\mathbf{E}[\varepsilon_{u,v}(s)] = 0$ and $\varepsilon_{u,v}(s)$ is independent of $\mathbf{P}$, the linearity of expectation gives $\mathbf{E}[Z] = 0$. Assume for any $\delta > 0$, we have

$$\Pr[|Z| \geq \delta] \leq 4 \cdot \exp\left(-\frac{\delta^2}{\|a\|_2^2} \right), \tag{3.48}$$

which would complete the proof. Hence, it remains to prove it.

Recall that $\phi_{u,v}(t) \in \{-1, +1\}$ is the random orientation which has Rademacher distribution (that is, it is uniform in $\{-1, +1\}$). Let us define an auxiliary random variable

$$W := \frac{1}{2} \sum_{s=1}^{t} \sum_{(u,v) \in \mathbf{M}^{(s)}} \phi_{u,v}(s) \cdot g_{u,v}^{(s)}.$$

The intuition is to show that an upper bound for W leads to a bound on Z. This part of proof follows along the same lines as the proof of [56, Lemma 1]. For the sake of completes we give the full technical proof here. Note that W is a sum of independent random variables, in particular, a weighted sum of Rademacher variables. We first prove that for any $\delta > 0$

$$\mathbf{Pr}[|Z| \geq \delta] \leq 2 \cdot \mathbf{Pr}[|W| \geq \delta]. \tag{3.49}$$

Consider the edges which for $s \in [1, t]$ satisfy $\mathrm{Odd}(X_u(s-1) + X_v(s-1)) = 0$. These edges make the difference between Z and W. To see this, let $Z > \delta$ for some $\delta > 0$. Then $W \leq \delta$ can happen only if

$$\sum_{s=1}^{t} \sum_{\substack{(u,v) \in \mathbf{M}^{(s)}: \\ \mathrm{Odd}(X_u(s-1) + X_v(s-1)) = 0}} \phi_{u,v}(s) \cdot g_{u,v}^{(s)} < 0.$$

Hence,

$$\mathbf{Pr}[W \leq \delta \mid Z > \delta] \leq \mathbf{Pr}\left[\sum_{s=1}^{t} \sum_{\substack{(u,v) \in \mathbf{M}^{(s)}: \\ \mathrm{Odd}(X_u(s-1) + X_v(s-1)) = 0}} \phi_{u,v}(s) \cdot g_{u,v}^{(s)} < 0 \right].$$

Recall that $\phi_{\cdot,\cdot}(.)$ is a random variable which is 1 or -1 each with probability $1/2$. Moreover, it is independent of $g_{\cdot,\cdot}^{(.)}$. Because of symmetry, the right side happens with probability strictly less than $1/2$.

Therefore,

$$\mathbf{Pr}[W > \delta \mid Z > \delta] \geq 1/2.$$

Consequently it follows that

$$\mathbf{Pr}[W \geq \delta] \geq \mathbf{Pr}[W > \delta \mid Z > \delta] \cdot \mathbf{Pr}[Z > \delta] \geq \frac{1}{2} \cdot \mathbf{Pr}[Z > \delta],$$

and by symmetry

$$\mathbf{Pr}[W \leq -\delta] \geq \frac{1}{2} \cdot \mathbf{Pr}[Z < -\delta].$$

This completes the proof of Equation (3.49). Now we derive a tail bound for

$$W = \frac{1}{2} \sum_{s=1}^{t} \sum_{(u,v) \in \mathbf{M}^{(s)}} \phi_{u,v}(s) \cdot g_{u,v}^{(s)}.$$

We define a sequence of random variables $Y_1, \ldots, Y_t$ such that Y_ℓ for $1 \leq \ell \leq t$ indicates the orientation of the single matching edge (u, v) of round ℓ, i.e., $Y_\ell := \phi_{u,v}(\ell)$. To simplify the notation we define $\mathcal{F}_\ell := (Y_1, Y_2, \ldots, Y_\ell)$. We are to apply Theorem 6.14. Since W is a (weighted) sum of independent random variables, for any $1 \leq \ell \leq t$ (such that edge (u, v) is picked up in round ℓ) it holds that

$$|\mathbf{E}[W \mid \mathcal{F}_\ell] - \mathbf{E}[W \mid \mathcal{F}_{\ell-1}]| \leq \frac{|g_{u,v}^{(\ell)}|}{2}.$$

Note that $\mathbf{E}[W] = 0$. An application of Theorem 6.14 to the martingale $X_\ell := \mathbf{E}[W \mid \mathcal{F}_\ell]$, where $1 \leq \ell \leq t$ gives for any $\delta > 0$,

$$\mathbf{Pr}[|W| \geq \delta] \leq 2 \cdot \exp\left(- \frac{2 \cdot \delta^2}{\sum_{s=1}^{t_1} \sum_{(u,v) \in \mathbf{M}^{(s)}} \left(g_{u,v}^{(s)} \right)^2} \right).$$

From the first statement of Lemma 3.14 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 this with Equation (3.49) gives

$$\mathbf{Pr}[|Z| \geq \delta] \leq 4 \cdot \exp\left(- \frac{\delta^2}{\|a\|_2^2} \right),$$

and shifting δ by κ and recalling the definition of Z concludes the proof. □

Here we collect a few other tools and results that are frequently used only in our analysis. The next two observations help us in coupling arguments.

Observation 3.18. *Fix a sequence of matchings $\mathbf{M}^{[\infty]}$. 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 $\widetilde{x}(0)$ such that for some $\alpha \in \mathbb{N}$ and for all $u \in V$ it holds*

$$\widetilde{x}_u(0) := \max\{x_u(0) - \alpha, 0\}.$$

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

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

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 $\widetilde{X}(t)$ plus α.

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

$$\widehat{x}_u(0) := \max\{x_u(0), \alpha\} \qquad \text{and} \qquad \widetilde{x}_u(0) := \widehat{x}_u(0) - \alpha.$$

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 observation 6.29, it follows that for any round $t \geq 0$ and each node $u \in V$ we have $X_u(t) \leq \widehat{X}_u(t)$. Similarly, from the first statement of observation 6.29 it follows that $\widetilde{X}_u(t) = \widehat{X}_u(t) - \alpha$. A combination of these two implies that

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

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

The next observation uses the concept of orientation, which is defined in Equation (3.46). In our analysis we often need to use a flipped process. The next observation allows us to do it.

Observation 3.19. *Fix a sequence of matchings $\mathbf{M}^{[\infty]}$. Consider two executions of the discrete load balancing protocol with the same matchings but with different initial load vectors $x(0)$ and $\widetilde{x}(0)$, where $\widetilde{x}(0) := K \cdot \vec{1} - x(0)$ for some $K \in \mathbb{N}_0$, and flipped orientations, i.e., $\widetilde{\phi}_{u,v}(t) = -\phi_{u,v}(t)$. Then for any round $t \geq 1$,*

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

Proof. The proof is by induction over $t \in \mathbb{N}_0$. The base case $t = 0$ holds by assumption. Now assuming $\widetilde{x}(t-1) = K \cdot \mathbf{1} - X(t-1)$ 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

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

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

$$X_u(t) = \frac{X_u(t-1) + X_v(t-1)}{2} + \frac{1}{2} \cdot \mathrm{Odd}(X_u(t-1) + X_v(t-1)) \cdot \phi_{u,v}(t)$$

and

$$X_v(t) = \frac{X_u(t-1) + X_v(t-1)}{2} + \frac{1}{2} \cdot \mathrm{Odd}(X_u(t-1) + X_v(t-1)) \cdot \phi_{v,u}(t).$$

Furthermore, we have

$$\begin{aligned}
\widetilde{X}_u(t) &= \frac{\widetilde{X}_u(t-1)+\widetilde{X}_v(t-1)}{2} + \frac{1}{2}\cdot \mathrm{Odd}\Big(\widetilde{X}_u(t-1)+\widetilde{X}_v(t-1)\Big)\cdot \widetilde{\phi}_{u,v}(t)\\
&= \frac{2K-(X_u(t-1)+X_v(t-1))}{2} + \frac{1}{2}\cdot \mathrm{Odd}(2K-(X_u(t-1)+X_v(t-1)))\cdot \widetilde{\phi}_{u,v}(t)\\
&= \frac{2K-(X_u(t-1)+X_v(t-1))}{2} + \frac{1}{2}\cdot \mathrm{Odd}(X_u(t-1)+X_v(t-1))\cdot \widetilde{\phi}_{u,v}(t)\\
&= K - \frac{(X_u(t-1)+X_v(t-1))}{2} - \frac{1}{2}\cdot \mathrm{Odd}(X_u(t-1)+X_v(t-1))\cdot \phi_{u,v}(t)\\
&= K - X_u(t).
\end{aligned}$$

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

3.4.3 Ingredients used in Proposition 3.4

In the rest of this section we prove the main lemmas used in the proof of Proposition 3.4 which are Lemma 3.27 (Phase 1) and Lemma 3.28 (Phase 2). To do so, we have to provide several additional definitions and statements, which ultimately leads towards the key lemma (Lemma 3.24), which proves a drop in the number of tokens with height at least $L+1$ within τ_{local} rounds. From this, Lemma 3.27 and Lemma 3.28 can be derived relatively easily. Note that throughout this section we assume there are initially $\mathrm{O}(L\cdot n)$ tokens for $L \le \log^7(n)$.

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

Definition 3.2. *For any round $t \ge \tau_{\mathrm{global}}$ we define*

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

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

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

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 observation 6.31). We call that property *globally mixing* since it implies that a token starting from any node in round 1 across the entire graph. 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 $[t-\tau_{\mathrm{local}}+1, t]$. We remark that for any round $t \ge 1$, the events $\Gamma_g^{(t)}$ and $\Gamma_\ell^{(t+\tau_{\mathrm{local}})}(u)$ 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 $(\tau_{\mathrm{global}}, \tau_{\mathrm{local}})$-*good* matchings, the proof is a straightforward calculation.

Lemma 3.20. *For any process generating $(\tau_{\mathrm{global}}, \tau_{\mathrm{local}})$-good matchings the following holds.*

1. *For any round $t \ge \tau_{\mathrm{global}}$, $\mathbf{Pr}\left[\Gamma_g^{(t)}\right] \ge 1 - 1/n^3$.*

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

Proof. First note that

$$\mathbf{Pr}\left[\overline{\Gamma_g^{(\tau_{\text{global}})}}\right] = \mathbf{Pr}\left[\bigcup_{u \in V}\left\{\left\|\mathbf{M}_{u,\cdot}^{[1,\tau_{\text{global}}]} - \frac{\vec{1}}{n}\right\|_2^2 > \frac{1}{n^7}\right\}\right]. \tag{3.50}$$

Applying the definition of τ_{global} (Definition 3.1), we get for $t = \tau_{\text{global}}$ that $\mathbf{Pr}\left[\overline{\Gamma_g^{(\tau_{\text{global}})}}\right] \leq 1/n^3$. The first statement follows from the basic fact that $\left\|\mathbf{M}_{u,\cdot}^{[1,t]} - \frac{\vec{1}}{n}\right\|_2^2$ is non-increasing in t (see Observation 6.34). For the second statement, consider an arbitrary round $t \geq \tau_{\text{local}}$ and a fixed node $u \in V$. Here, we get

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

where the last inequality follows from the definition of τ_{local} (Definition 3.1). As above, the proof follows from Observation 6.34. □

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 $e : \mathbb{N}_0 \to \mathbb{N}_0$ 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 $E(k) := [e(k-1)+1, \ldots, e(k)]$ 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.,

$$Y^{(t)} := \sum_{u \in V} \max\{X_u(t) - L, 0\} = \sum_{j \in \mathcal{T}} \mathbf{1}_{H_j(t) \geq L+1}\,. \tag{3.51}$$

We wish to prove that in each epoch $Y^{(t)}$ drops by a constant factor (see Lemma 3.24) 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 3.27). 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

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

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. For the proof we use the negative association lemma after expanding the term $\mathbf{E}[Z_i^{(e(k))}]$.

Lemma 3.21. *Let $\mathbf{M}^{[e(k)]}$ 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*

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

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

$$\begin{aligned}
&\mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)) = w(e(k-1)), \mathbf{M}^{[e(k)]}\Big] \\
&\overset{(a)}{=} \sum_{v \in V} \sum_{j \in \mathcal{T}\colon j \neq i} \mathbf{E}\Big[\, Z_{ij}^{(e(k))}(v) \;\Big|\; W(e(k-1)) = w(e(k-1)), \mathbf{M}^{[e(k)]} \,\Big] \\
&\overset{(b)}{=} \sum_{v \in V} \sum_{j \in \mathcal{T}\colon j \neq i} \mathbf{Pr}\Big[Z_{ij}^{(e(k))}(v) = 1 \;\Big|\; W(e(k-1)) = w(e(k-1)), \mathbf{M}^{[e(k)]}\Big] \\
&= \sum_{v \in V} \sum_{j \in \mathcal{T}\colon j \neq i} \mathbf{Pr}\Big[W_i(e(k)) = v \cap W_j(e(k)) = v \;\Big|\; W(e(k-1)) = w(e(k-1)), \mathbf{M}^{[e(k)]}\Big], \qquad (3.52)
\end{aligned}$$

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 $E(k) = [e(k-1)+1, \ldots, e(k)]$. Crucially, we can now apply our negative association result (Lemma 3.12) to Equation (3.52) to obtain

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

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

$$\mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)), \mathbf{M}^{[e(k)]}\Big] \leq \sum_{w \in V} \left(\sum_{v \in V} \mathbf{M}^{E(k)}_{W_i(e(k-1)),v} \cdot \mathbf{M}^{E(k)}_{w,v} \right) \cdot X_w(e(k-1)). \qquad \square$$

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

Definition 3.3. *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. $\mathcal{L}^{(e(k))}(u) := \left\{ \mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)), \Gamma_\ell^{(e(k))}(u)\Big] \leq L - \varepsilon + \dfrac{1}{\log^4(n)} \right\}.$
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 4 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.

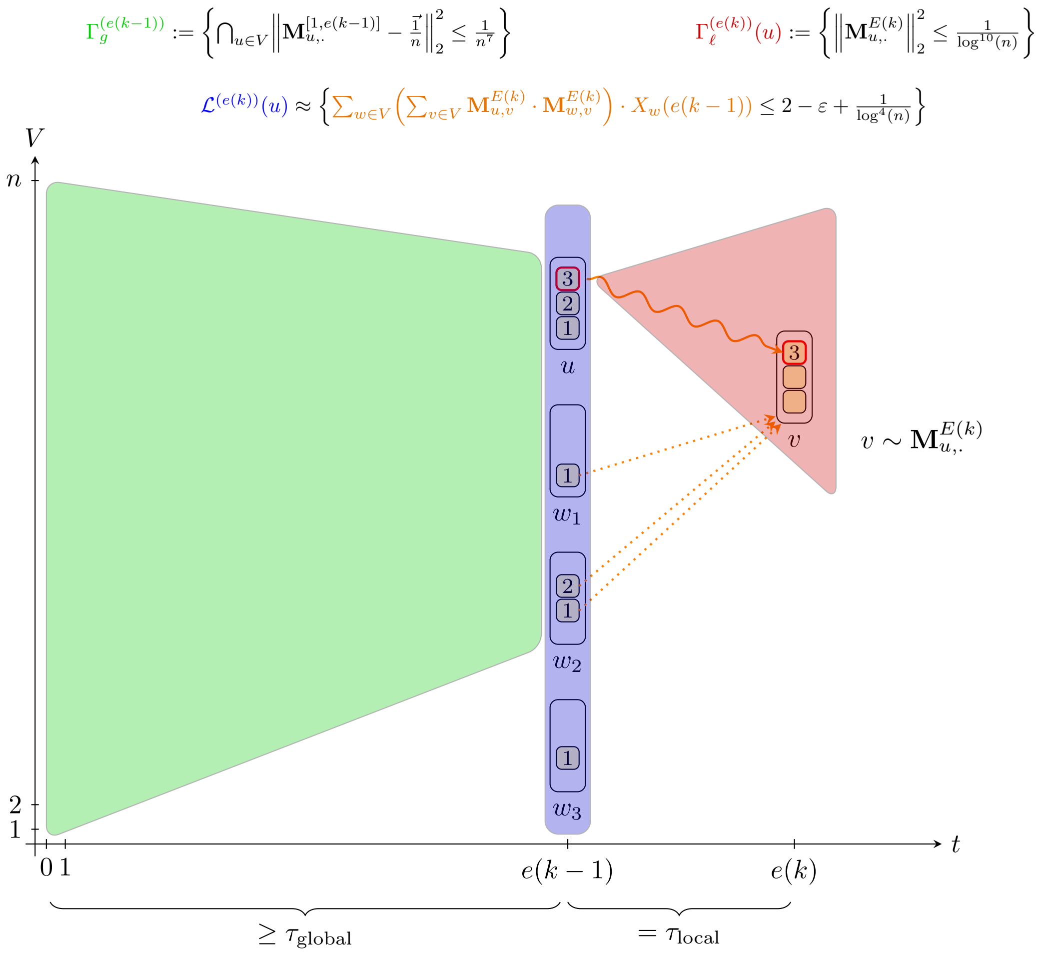

Figure 4: The events $\Gamma_\ell^{(e(k))}(u)$, $\Gamma_g^{(e(k-1))}$ and $\mathcal{L}^{(e(k))}(u)$

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 We prove in Lemma 3.22 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 3.24) that the number of tokens with height $L + 1 = 3$ drops within $E(k)$ by a suitable factor.

Lemma 3.22. *Assume that the load vector $x(0)$ has at most $(L - \varepsilon)n$ tokens for $1 \le L \le \log^7(n)$ being an integer and any $0 < \varepsilon < L$. Then for any milestone $k \ge 1$ we have* $\mathbf{Pr}\left[\mathcal{E}^{(e(k))}\right] \ge 1 - \frac{2}{n^3}$.

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 3.3 which gives

$$\overline{\mathcal{L}^{(e(k))}(u)} = \left\{ \mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)), \Gamma_\ell^{(e(k))}(u) \Big] > L - \varepsilon + \frac{1}{\log^4(n)} \right\} \quad \text{and}$$

$$\overline{\mathcal{E}^{(e(k))}(u)} = \left\{ \overline{\Gamma_g^{(e(k-1))}} \cup \left(\Gamma_\ell^{(e(k))}(u) \cap \overline{\mathcal{L}^{(e(k))}(u)} \right) \right\}.$$

For brevity, let us write $\alpha := \mathbf{Pr}\Big[\overline{\Gamma_g^{(e(k-1))}}\Big]$ and note that $\alpha \le 1/n^3$ by Lemma 3.20. Then,

$$\begin{aligned}
&\mathbf{Pr}\Big[\overline{\mathcal{E}^{(e(k))}(u)}\Big] \\
&\le \alpha + \mathbf{Pr}\Big[\Gamma_\ell^{(e(k))}(u) \cap \overline{\mathcal{L}^{(e(k))}(u)}\Big] \\
&\le \alpha + \mathbf{Pr}\Big[\overline{\mathcal{L}^{(e(k))}(u)}\Big] \\
&= \alpha + \mathbf{Pr}\left[\mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)), \Gamma_\ell^{(e(k))}(u) \Big] > L - \varepsilon + \frac{1}{\log^4(n)} \right] \\
&= \alpha + \mathbf{Pr}\bigg[\mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)), \Gamma_\ell^{(e(k))}(u), \Gamma_g^{(e(k-1))} \Big] \cdot \mathbf{Pr}\Big[\Gamma_g^{(e(k-1))}\Big] + \\
&\qquad\qquad \mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)), \Gamma_\ell^{(e(k))}(u), \overline{\Gamma_g^{(e(k-1))}} \Big] \cdot \mathbf{Pr}\Big[\overline{\Gamma_g^{(e(k-1))}}\Big] > L - \varepsilon + \frac{1}{\log^4(n)} \bigg] \\
&\le \alpha + \mathbf{Pr}\left[\mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)), \Gamma_\ell^{(e(k))}(u), \Gamma_g^{(e(k-1))} \Big] > L - \varepsilon + \frac{1}{\log^4(n)} - \frac{\log^7(n)}{n^2} \right],
\end{aligned} \tag{3.53}$$

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

$$\mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)), \Gamma_\ell^{(e(k))}(u), \overline{\Gamma_g^{(e(k-1))}} \Big] \cdot \mathbf{Pr}\Big[\overline{\Gamma_g^{(e(k-1))}}\Big] \le \frac{\log^7(n)}{n^2}.$$

We thus have

$$\begin{aligned}
&\mathbf{Pr}\Big[\overline{\mathcal{E}^{(e(k))}(u)}\Big] \\
&\le \alpha + \mathbf{Pr}\left[\mathbf{E}\Big[Z_i^{(e(k))} \;\Big|\; W(e(k-1)), \Gamma_\ell^{(e(k))}(u), \Gamma_g^{(e(k-1))} \Big] > L - \varepsilon + \frac{1}{\log^4(n)} - \frac{\log^7(n)}{n^2} \right].
\end{aligned} \tag{3.54}$$

To upper bound Equation (3.53) we apply Claim 3.23 below. Since Claim 3.23 holds for any fixed sequence $\mathbf{M}^{[e(k)]}$ satisfying $\Gamma_\ell^{(e(k))}(u) \cap \Gamma_g^{(e(k-1))}$, we have

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

Applying the union bound twice implies that

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

□

Claim 3.23. *Fix a node u and let $\mathbf{M}^{[e(k)]}$ be a fixed sequence of matchings satisfying $\Gamma_\ell^{(e(k))}(u) \cap \Gamma_g^{(e(k-1))}$.*

Then, it holds that

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

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

$$\begin{aligned}
&\mathbf{Pr}\left[\left\{\mathbf{E}\left[Z_i^{(e(k))} \;\middle|\; W(e(k-1)), \mathbf{M}^{[e(k)]}\right] > L - \varepsilon + \frac{1}{2\log^4(n)}\right\}\right] \\
&\quad \overset{(a)}{\leq} \mathbf{Pr}\left[\sum_{w\in V}\left(\sum_{v\in V} \mathbf{M}_{u,v}^{E(k)} \cdot \mathbf{M}_{w,v}^{E(k)}\right) \cdot X_w(e(k-1)) \geq L - \varepsilon + \frac{1}{2\log^4(n)}\right], \qquad (3.55)
\end{aligned}$$

where (a) follows from Lemma 3.21. To bound the probability of Equation (3.55) we will apply our concentration inequality (Lemma 3.15). Note that, since the matchings are fixed, the only randomness remaining in Lemma 3.15 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 3.15 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 $\mathbf{M}^{[e(k-1)]}$ satisfies the event $\Gamma_g^{(e(k-1))}$, it follows from observation 6.31 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,

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

where (a) follows from the second statement of Lemma 3.14; (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 $\overline{x}$ satisfies $\overline{x} \leq L - \varepsilon$. From Lemma 3.15 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

$$\mathbf{Pr}\left[\sum_{w\in V} a_w \cdot X_w(e(k-1)) \geq L - \varepsilon + \frac{1}{2\log^4(n)}\right] \leq 2 \cdot \exp\left(-\frac{\left(\frac{1}{2\log^4(n)} - \frac{1}{n}\right)^2}{4 \cdot \left(\frac{1}{\log^{10}(n)}\right)}\right) \leq \frac{1}{n^5}. \qquad (3.56)$$

Combining Equations (3.55) and (3.56) finishes the proof of the claim. ☐

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,

$$\widetilde{Y}^{(e(k))} := Y^{(e(k))} \cdot \mathbf{1}_{\cap_{i=0}^{k} \mathcal{E}^{(e(i))}}. \qquad (3.57)$$

Later we will show how to use $Z_i^{(e(k))}$ to bound $\widetilde{Y}^{(e(k))}$, and then eventually $Y^{(e(k))}$. In the next lemma we bound the expected value of $\widetilde{Y}^{(e(k))}$. We will denote by $\left(\mathfrak{F}^{(t)}\right)_{t\geq 0}$ the filtration of the random process;

note that in particular, $\mathfrak{F}^{(t)}$ determines not only the current load vector and previous load vectors $X(t), X(t-1), \ldots, X(1)$, but also all location vectors $W(t), W(t-1), \ldots, W(1)$. In case of randomly generated matchings $\mathbf{M}^{[t]} = \left(\mathbf{M}^{(s)}\right)_{s=1}^{t}$ is determined by $\left(\mathfrak{F}^{(t)}\right)_{t\geq 0}$, 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 $\mathcal{E}$,

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

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

Lemma 3.24 (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 $k \geq 1$,*

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

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

$$\mathbf{E}^{(e(k-1))}\Big[\,\widetilde{Y}^{(e(k))}\,\Big] = \mathbf{E}^{(e(k-1))}\Big[\,Y^{(e(k))} \cdot \mathbf{1}_{\cap_{i=0}^{k}\mathcal{E}^{(e(i))}}\,\Big],$$

applying Lemma 6.36 to it (note that $\mathfrak{F}^{(e(k-1))}$ determines the random variable $\mathbf{1}_{\cap_{i=0}^{k-1}\mathcal{E}^{(e(i))}}$) gives

$$\mathbf{E}^{(e(k-1))}\Big[\,\widetilde{Y}^{(e(k))}\,\Big] = \mathbf{1}_{\cap_{i=0}^{k-1}\mathcal{E}^{(e(i))}} \cdot \mathbf{E}^{(e(k-1))}\Big[\,Y^{(e(k))} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}}\,\Big],$$

and since $\overline{\Gamma_g^{(e(k-1))}}$ implies $\mathbf{1}_{\mathcal{E}^{(e(k))}} = 0$ then,

$$\begin{aligned}
\mathbf{E}^{(e(k-1))}\Big[\,\widetilde{Y}^{(e(k))}\,\Big] &= \mathbf{1}_{\cap_{i=0}^{k-1}\mathcal{E}^{(e(i))}} \cdot \Bigg(\mathbf{Pr}^{(e(k-1))}\Big[\,\Gamma_g^{(e(k-1))}\,\Big] \cdot \mathbf{E}^{(e(k-1))}\Big[\,Y^{(e(k))} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \;\Big|\; \Gamma_g^{(e(k-1))}\,\Big] \\
&\qquad + \mathbf{Pr}^{(e(k-1))}\Big[\,\overline{\Gamma_g^{(e(k-1))}}\,\Big] \cdot 0\Bigg) \\
&\leq \mathbf{1}_{\cap_{i=0}^{k-1}\mathcal{E}^{(e(i))}} \cdot \mathbf{E}^{(e(k-1))}\Big[\,Y^{(e(k))} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \;\Big|\; \Gamma_g^{(e(k-1))}\,\Big],
\end{aligned} \tag{3.58}$$

Further, we have

$$Y^{(e(k))} = \sum_{j\in\mathcal{T}} \mathbf{1}_{H_j(e(k))\geq L+1} \leq \sum_{j\in\mathcal{T}} \mathbf{1}_{H_j(e(k-1))\geq L+1} \cdot \mathbf{1}_{Z_j^{(e(k))}\geq L}, \tag{3.59}$$

where the equality follows from the definition of $Y^{(e(k)}$. To see the inequality, observe that by properties of the height-sensitive processthe 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 Equation (3.59) to Equation (3.58) leads us to

$$
\begin{aligned}
&\mathbf{E}^{(e(k-1))}\left[\widetilde{Y}^{(e(k))}\right] \\
&\quad \leq \mathbf{1}_{\cap_{i=0}^{k-1}\mathcal{E}^{(e(i))}} \cdot \mathbf{E}^{(e(k-1))}\left[\sum_{j\in\mathcal{T}} \mathbf{1}_{H_j(e(k-1))\geq L+1} \cdot \mathbf{1}_{Z_j^{(e(k))}\geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \;\middle|\; \Gamma_g^{(e(k-1))}\right] \\
&\quad \overset{(a)}{=} \mathbf{1}_{\cap_{i=0}^{k-1}\mathcal{E}^{(e(i))}} \cdot \sum_{j\in\mathcal{T}} \mathbf{E}^{(e(k-1))}\left[\mathbf{1}_{H_j(e(k-1))\geq L+1} \cdot \mathbf{1}_{Z_j^{(e(k))}\geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \;\middle|\; \Gamma_g^{(e(k-1))}\right] \\
&\quad \overset{(b)}{=} \mathbf{1}_{\cap_{i=0}^{k-1}\mathcal{E}^{(e(i))}} \cdot \sum_{j\in\mathcal{T}} \mathbf{1}_{H_j(e(k-1))\geq L+1} \cdot \mathbf{E}^{(e(k-1))}\left[\mathbf{1}_{Z_j^{(e(k))}\geq L} \cdot \mathbf{1}_{\mathcal{E}^{(e(k))}} \;\middle|\; \Gamma_g^{(e(k-1))}\right], \qquad (3.60)
\end{aligned}
$$

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 6.36). 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 Equation (3.60) by

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

We upper bound the two remaining terms in the last line separately using Claim 3.25 and Claim 3.26 (see below). Together these two claims yield, using the assumption that $L \leq \log^7(n)$,

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

Applying this to Equation (3.60) gives us

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

□

Here we prove the two claims used in the last lemma.

Claim 3.25. $\mathbf{Pr}^{(e(k-1))}\left[\overline{\Gamma_\ell^{(e(k))}(u)} \;\middle|\; \Gamma_g^{(e(k-1))}\right] \leq \frac{1}{\log^{11}(n)}$.

Proof. As before we define $E(k) = [e(k-1)+1, \ldots, e(k)]$ 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

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

where (a) follows Lemma 3.20 (recall that $e(k) - e(k-1) = \tau_{\text{local}}$). □

Claim 3.26. $\mathbf{E}^{(e(k-1))}\left[\mathbf{1}_{Z_j^{(e(k))}\geq L}\cdot\mathbf{1}_{\mathcal{E}^{(e(k))}}\;\middle|\;\Gamma_g^{(e(k-1))},\Gamma_\ell^{(e(k))}(u)\right]\leq 1-\frac{\varepsilon}{L}+\frac{1}{L\cdot\log^4(n)}.$

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

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

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

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

using the definition of expectations conditional that $\mathbf{E}\left[Z\;\middle|\;\mathcal{E}\right]=\mathbf{E}[Z\cdot\mathbf{1}_{\mathcal{E}}]/\mathbf{Pr}[\mathcal{E}]$, simplifies the above to

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

and using the definition of the expectation we get

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

and applying the so-called "conditional Markov's inequality" (Exercise 8.2.5 in [61]) simplifies the above to,

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

where (c) also uses the definition of expectations conditional on events $\mathcal{E}$, $\mathbf{E}\left[Z\;\middle|\;\mathcal{E}\right]=\frac{\mathbf{E}[Z\cdot\mathbf{1}_{\mathcal{E}}]}{\mathbf{Pr}[\mathcal{E}]}$. Recall that $\mathbf{E}^{(e(k-1))}\left[Z_j^{(e(k))}\;\middle|\;\Gamma_g^{(e(k-1))},\Gamma_\ell^{(e(k))}(u),\mathcal{E}^{(e(k))}\right]$ 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

$$\begin{aligned}
&\mathbf{E}^{(e(k-1))}\Big[\mathbf{1}_{Z_j^{(e(k))}\geq L}\cdot\mathbf{1}_{\mathcal{E}^{(e(k))}}\;\Big|\;\Gamma_g^{(e(k-1))},\Gamma_\ell^{(e(k))}(u)\Big]\\
&\leq \frac{1}{L}\cdot\sup_{\omega}\mathbf{E}\Big[Z_j^{(e(k))}\;\Big|\;\Gamma_g^{(e(k-1))},\Gamma_\ell^{(e(k))}(u),\mathcal{E}^{(e(k))}\Big]\\
&\overset{(a)}{=} \frac{1}{L}\cdot\sup_{\omega}\mathbf{E}\Big[Z_j^{(e(k))}\;\Big|\;\Gamma_g^{(e(k-1))},\Gamma_\ell^{(e(k))}(u),\mathcal{E}^{(e(k))},\mathcal{L}^{(e(k))}(u)\Big]\\
&\leq \frac{1}{L}\cdot\sup_{\omega}\mathbf{E}\Big[Z_j^{(e(k))}\;\Big|\;\Gamma_\ell^{(e(k))}(u),\mathcal{L}^{(e(k))}(u)\Big]\\
&\overset{\text{Def. (3.3)}}{\leq} \frac{1}{L}\cdot\left(L-\varepsilon+\frac{1}{\log^4(n)}\right)=1-\frac{\varepsilon}{L}+\frac{1}{L\cdot\log^4(n)},
\end{aligned}$$

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. □

By repeatedly applying strong Lemma 3.24 over subsequent epochs, we can complete the analysis of Phase 1 and Phase 2.

Lemma 3.27 (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,*

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

Proof. Recall that $e(0)=\tau_{\text{global}}$. Hence $t_1 = e(\log(n)/\log\log(n))$. Note that $\widetilde{Y}^{(e(0))}\leq Y^{(e(0))}\leq n\cdot\log^7(n)$. Applying Lemma 3.24 for $\ell := \frac{\log(n)}{\log\log(n)}$ epochs gives us

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

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

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

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

$$\mathbf{Pr}\left[\widetilde{Y}^{(e(\ell))}\geq\exp\left(-\frac{\varepsilon\cdot\ell}{2L}+\beta\right)\cdot n\cdot\log^7(n)\right]\leq e^{-\beta}. \tag{3.61}$$

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 $\widetilde{Y}^{(e(\ell))}$, the law of total probability and the union bound we get

$$\begin{aligned}
\mathbf{Pr}\left[Y^{(e(\ell))} \geq \frac{n}{\log(n)}\right] &= \mathbf{Pr}\left[Y^{(e(\ell))} \geq \frac{n}{\log(n)} \;\middle|\; \bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}\right] \cdot \mathbf{Pr}\left[\bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}\right] \\
&\quad + \mathbf{Pr}\left[Y^{(e(\ell))} \geq \frac{n}{\log(n)} \;\middle|\; \overline{\bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}}\right] \cdot \mathbf{Pr}\left[\overline{\bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}}\right] \\
&\leq \mathbf{Pr}\left[Y^{(e(\ell))} \geq \frac{n}{\log(n)} \;\middle|\; \bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}\right] + \mathbf{Pr}\left[\overline{\bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}}\right] \\
&= \mathbf{Pr}\left[\widetilde{Y}^{(e(\ell))} \geq \frac{n}{\log(n)}\right] + \mathbf{Pr}\left[\bigcup_{s=0}^{\ell} \overline{\mathcal{E}^{(e(s))}}\right] \\
&\overset{\text{Eq. (3.61)}}{\leq} e^{-\beta} + \sum_{s=0}^{\ell} \mathbf{Pr}\left[\overline{\mathcal{E}^{(e(s))}}\right] \overset{(a)}{\leq} e^{-\beta} + 2 \cdot (\ell+1) \cdot n^{-3},
\end{aligned}$$

where (a) follows from Lemma 3.22. Recalling our choice of β earlier in this proof, and the choice of $\ell = \frac{\log(n)}{\log\log(n)}$, finishes the proof. □

Lemma 3.28 (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 $t_2 := \tau_{\text{global}} + \frac{4}{-\log\left(1-\varepsilon+\frac{2}{\log^4(n)}\right)} \cdot \log(n) \cdot \tau_{\text{local}}$,*

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

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, $\widetilde{Y}^{(e(0))} \leq Y^{(e(0))} \leq n$. Applying Lemma 3.24 with $L = 1$ yields for any epoch $k \geq 1$,

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

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

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

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

$$
\begin{aligned}
&\mathbf{Pr}\Big[Y^{(e(\ell))} \geq 1\Big] \\
&= \mathbf{Pr}\left[Y^{(e(\ell))} \geq 1 \;\middle|\; \bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}\right] \cdot \mathbf{Pr}\left[\bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}\right] + \mathbf{Pr}\left[Y^{(e(\ell))} \geq 1 \;\middle|\; \overline{\bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}}\right] \cdot \mathbf{Pr}\left[\overline{\bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}}\right] \\
&\leq \mathbf{Pr}\left[Y^{(e(\ell))} \geq 1 \;\middle|\; \bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}\right] + \mathbf{Pr}\left[\overline{\bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}}\right] \\
&= \mathbf{Pr}\Big[\widetilde{Y}^{(e(\ell))} \geq 1\Big] + \mathbf{Pr}\left[\overline{\bigcap_{s=0}^{\ell} \mathcal{E}^{(e(s))}}\right] \\
&\leq n^{-3} + \sum_{s=0}^{\ell} \mathbf{Pr}\Big[\overline{\mathcal{E}^{(e(s))}}\Big] \overset{(a)}{\leq} n^{-3} + 2 \cdot (\ell+1) \cdot n^{-3} \leq n^{-2},
\end{aligned}
$$

where the last line used that $\ell = \mathrm{O}(\log(n))$ and (a) follows from Lemma 3.22. □

3.5 Bounds for Specific Models

Note that our main result (Theorem 3.1) 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).

3.5.1 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 $\mathbf{m}^{(1)}, \mathbf{m}^{(2)}, \ldots, \mathbf{m}^{(\Delta)}$. The sequence $\mathbf{m}^{[\infty]} = \big(\mathbf{m}^{(s)}\big)_{s=1}^{\infty}$ 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 [51]. Recall that the round matrix $\mathbf{R}$ is defined as $\mathbf{R} := \prod_{s=1}^{\Delta} \mathbf{m}^{(s)}$ and $\lambda(\mathbf{R})$ is the second (absolute) largest eigenvalue of it. In this model, Theorem 3.1 provides, see Corollary 3.29 below,

$$
\tau = \mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K)), \qquad \text{where} \qquad \widetilde{\tau}_{\mathrm{S}}(K) := \Theta\left(\frac{\Delta \cdot \log(Kn)}{1-\lambda(\mathbf{R})}\right). \tag{3.62}
$$

Corollary 3.29. *Assuming the premise of Theorem 3.1, for the balancing circuit model we have* $\tau = \mathrm{O}(\Delta \cdot \log(Kn)/(1-\lambda(\mathbf{R})))$.

Proof. Using standard spectral arguments (e.g., [82, Theorem 1]) yields that the sequence of matchings $\mathbf{m}^{[t]} = \big(\mathbf{m}^{(s)}\big)_{s=1}^{t}$ for $t := \frac{4\Delta}{1-\lambda(\mathbf{R})} \cdot \log\big(\frac{Kn}{\varepsilon}\big)$ 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 $\tau_{\text{global}} := \frac{7\Delta \cdot \log(n)}{1-\lambda(\mathbf{R})}$ and $\tau_{\text{local}} := \frac{11\Delta \cdot \log\log(n)}{1-\lambda(\mathbf{R})}$. Fix an arbitrary $t \geq 1$. From [85, Lemma 2.4] (restated as Lemma 6.26) it follows for any $u \in V$ and for the choice of τ_{global} as above that we have

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

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

$$\left\|\mathbf{m}_{u,\cdot}^{[t,t+\tau_{\text{local}}]}\right\|_2^2 - \frac{1}{n} \overset{\text{Obs. 6.32}}{=} \left\|\mathbf{m}_{u,\cdot}^{[1,\tau_{\text{local}}]} - \frac{\vec{1}}{n}\right\|_2^2 \overset{\text{Lem. 6.26}}{\leq} \frac{1}{\log^{11} n}.$$

From this we get

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

□

3.5.2 Application to Random Matchings

In this model Δ is maximum degree of G. 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 [32, 51]. To state the results for the random matching model we recall the diffusion matrix $\mathbf{P}$ which is $\mathbf{P}_{u,v} := 1/(2\Delta)$ 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 3.1 provides, see Corollary 3.30 below,

$$\tau = \mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K)), \qquad \text{where} \qquad \widetilde{\tau}_{\mathrm{S}}(K) := \Theta\left(\frac{\log(Kn)}{p_{\min} \cdot \Delta \cdot \big(1 - \lambda(\mathbf{P})\big)}\right) \tag{3.63}$$

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

Corollary 3.30. *Assuming the premise of Theorem 3.1, for the random matching model we have* $\tau = \mathrm{O}\big(\log(Kn)/\big(p_{\min} \cdot \Delta \cdot \big(1 - \lambda(\mathbf{P})\big)\big)\big)$.

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 [85, corollary 2.7] (see Lemma 6.27) it follows for any $u \in V$ and the choice of τ_{global} as above (and sufficiently large n) that we have

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

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

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

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

$$\begin{aligned}\mathbf{Pr}\left[\left\|\mathbf{M}_{u,\cdot}^{[t,t+\tau_{\text{local}}]}\right\|_2^2 \leq \frac{1}{\log^{10} n}\right] &\geq \mathbf{Pr}\left[\left\|\mathbf{M}_{u,\cdot}^{[t,t+\tau_{\text{local}}]}\right\|_2^2 \leq \frac{1}{\log^{11} n} + \frac{1}{n}\right] \\ &\overset{\text{Obs. 6.32}}{=} \mathbf{Pr}\left[\left\|\mathbf{M}_{u,\cdot}^{[t,t+\tau_{\text{local}}]} - \frac{\vec{1}}{n}\right\|_2^2 \leq \frac{1}{\log^{11} n}\right] \overset{\text{Lem. 6.27}}{\geq} 1 - \frac{1}{\log^{11} n}.\end{aligned}$$

□

3.5.3 Application to Asynchronous 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 $\widetilde{\tau}_{\mathrm{S}}(K)$ applies here. Therefore, Theorem 3.1 provides $\tau = \mathrm{O}\Big(\frac{|E|}{\Delta} \cdot \log(Kn)/\big(1 - \lambda(\mathbf{P})\big)\Big) = \mathrm{O}\big(n \cdot \frac{d}{\Delta} \cdot \log(Kn)/\big(1 - \lambda(\mathbf{P})\big)\big)$, where $d = 2|E|/n$ is the average degree of G (see Corollary 3.30).

3.6 Summary and Open Problems

In this work we show that, for the matching model, discrete load balancing is as efficient and effective as continuous load balancing. We show 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 $\mathrm{O}(\log(n)) = \mathrm{O}(\mathrm{diam}(G))$ 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 [85].

At the heart of our analysis lie new correlation and concentration results, which we believe to be of independent interest. As the most involved step, showing that $\mathrm{O}(n)$ tokens can be balanced with maximum load $\mathrm{O}(1)$, 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 [71]). 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 $(\tau_{\mathrm{global}}, \tau_{\mathrm{local}})$-good.

Part Four:

4 Discrete Diffusion on d-Regular Graphs

4.1 Introduction

In this section, we study the vertex-based diffusion process, also referred to as neighborhood-based load balancing, originally introduced in [14] for static setting. We consider a d-regular graph $G = (V, E)$ where each node initially holds an integer number of tokens (also referred to as load items). In each round, every node distributes its tokens as evenly as possible among its neighbors and itself. If a perfectly uniform distribution is not feasible-i.e., some tokens would require splitting-the node distributes the remaining excess tokens randomly among its neighbors and itself, without replacement. The objective is to compute an upper bound on the discrepancy, defined as the difference between the maximum and minimum loads across all nodes in the network.

This protocol is motivated by rotor-router walks, a deterministic alternative to random walks. In a rotor-router walk, each node serves its neighbors in a fixed, cyclic order. Despite its deterministic nature, the resulting walk closely approximates the behavior of a random walk in terms of distribution and coverage. This connection motivated the design of our algorithm. In each round, every node selects a random permutation of its neighbors (including itself) and distributes its tokens sequentially in that order. This round-robin mechanism promotes fairness and facilitates effective load balancing over time.

Results in a nutshell. Building on prior work, we consider the quantity $\widetilde{\tau}_{\mathrm{S}}(K) = \log(Kn)/(1-\lambda)$ as an approximate upper bound on the time required for the discrepancy to fall from an initial value K to 1 in the idealized continuous diffusion, where load items are infinitely divisible. The analysis shows that discrepancy shrinks to polylogarithmic bounds within $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds, with faster convergence for higher-degree graphs, and quantifies the additional discrepancy caused by random load insertions in the dynamic case.

In our setting, we prove that the initial discrepancy K decreases to $\mathrm{O}(d\log(n))$ within $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds. After an additional $i \cdot \widetilde{\tau}_{\mathrm{S}}(n)$ rounds (i.e., at round $t = \mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K) + i \cdot \widetilde{\tau}_{\mathrm{S}}(n))$) the discrepancy becomes $\mathrm{O}(\sqrt[2^i]{d}\log(n) + \sqrt{d\log(n)})$ w.h.p. In particular, at round $t = \mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K) + \widetilde{\tau}_{\mathrm{S}}(n))$ the discrepancy is $\mathrm{O}(\sqrt{d}\log(n))$. Moreover, if $d = \Omega(\log^{1+\varepsilon})$ for any constant $\varepsilon \in [0, 1]$ the discrepancy is already bounded by $\mathrm{O}(\sqrt{d\log(n)})$ after $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds. We then extend our analysis to the dynamic setting, where in each step m new load items are allocated uniformly at random to the nodes before the balancing dynamics are applied. We show that this increases the discrepancy by $\mathrm{O}(\sqrt{m/n} \cdot \sqrt{\widetilde{\tau}_{\mathrm{S}}(n)})$. A summary of the results are provided in Table 4.

Table 4: Results for static discrete diffusion. K is the initial discrepancy and $\varepsilon \in (0, 1)$.

Discrepancy	Rounds	d
$\mathrm{O}(d \cdot \log(n))$	$\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$	-
$\mathrm{O}(\sqrt{d} \cdot \log(n))$	$\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K) + \widetilde{\tau}_{\mathrm{S}}(n))$	-
$\mathrm{O}(\sqrt{d \cdot \log(n)} + \log(n))$	$\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K)) + \log\log(d) \cdot \widetilde{\tau}_{\mathrm{S}}(n)$	-
$\mathrm{O}(\sqrt{d \cdot \log(n)})$	$\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$	$\Omega(\log^{1+\varepsilon} n)$

Techniques and Comparisons to prior works. The best-known result is due to Sauerwald and Sun [85], who show that the discrepancy is $\mathrm{O}(d^2\sqrt{\log(n)})$ within $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds. Their analysis bounds the deviation of a node's load from the average by expressing it as a weighted sum of dependent random variables, where each variable accounts for the error contributions accumulated across edges and rounds. In contrast, we reorganize this sum by grouping the random variables by nodes and their incident edges, which gives a tighter control over deviation. With this method, we prove our first result: the discrepancy is $\mathrm{O}(d\log(n))$ within $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ rounds. This already improves the best-known bound for the regime $d=\omega(\sqrt{\log(n)})$.

Next, we develop a height-sensitive process, which is an alternative realization of the load-balancing dynamics. In this process, each token follows a standard random walk and the token locations are shown to be negatively associated. This property enables a Chernoff-like concentration bound on the load distribution after $\widetilde{\tau}_{\mathrm{S}}(n)$ rounds. Combining these bounds with an invariant, we establish that within $\mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K)+\widetilde{\tau}_{\mathrm{S}}(n))$ rounds, the discrepancy is $\mathrm{O}(\sqrt{d}\log(n))$. This result strictly improves upon existing bounds in the regime $d=\Omega(\sqrt[3]{\log(n)})$.

Outline. The remainder of this part is organized as follows. Section 4.2 introduces the model and definitions. First, we formally define the discrete diffusion and the height-sensitive processes, then recall useful definitions. These are crucial, as they allow us to (a) reason about the convergence time of the discrete process and (b) keep track of the discrepancy. We also introduce the notion of vital tokens, which plays a central role in establishing a bound on the discrepancy after $\widetilde{\tau}_{\mathrm{S}}(K)$ rounds.

Section 4.3 presents our main result for the static setting, along with the key proof techniques. These techniques build on the introduced definitions while also relying on standard tools such as martingales, concentration inequalities for sums of random variables, and the negative association property.

Section 4.4 extends the analysis to the dynamic setting. Here we show that the difference between the dynamic and static settings lies in the presence of newly added load items. In particular, we bound the contribution of these newly allocated items.

Section 4.5 collects the technical lemmas supporting both the intermediate steps and the main results, including a strong tool: the negative association property among the locations of load items.

Section 4.6 derives bounds on the discrepancy at specific rounds in both the static and dynamic settings, following directly from the main theorems. Section 4.7 concludes with a summary and open problems.

4.2 Model and Definitions

We begin by introducing the notation and model. Subsection 4.2.1 formally defines the process, while Subsection 4.2.2 presents the height-sensitive process in detail. Section 4.2.3 defines and recalls crucial definitions for our analysis.

Let G be a d-regular graph with n nodes. We denote by $X(t)$ the discrete load vector at the end of round t where the i-th entry is the load of node i. At round 0, each node starts with an arbitrary load, and the average load is $\overline{x}:=\sum_{i\in V}x_i(0)/n$. Our goal is to bound the discrepancy which for a load vector $X(t)$ is defined as $\operatorname{disc}(X(t)):=\max_{w\in V}X_w(t)-\min_{w\in V}X_w(t)$. Throughout this part, we describe the dynamics using the diffusion matrix $\mathbf{P}$, corresponding to a standard random walk on a G, $\mathbf{P}_{i,j}:=1/(d+1)$ if $(i,j)\in E$ or $i=j$ and 0 otherwise. Here, $\lambda:=\lambda(\mathbf{P})$ represents the second-largest eigenvalue of $\mathbf{P}$ in absolute value. We write $\mathbf{P}^t:=\prod_{s=1}^{t}\mathbf{P}$ and denote by $\mathbf{P}_{i,\cdot}$ the row of $\mathbf{P}$ corresponding to i in matrix $\mathbf{P}$. Due to randomized rounding, $X_i(t)$ is inherently a random variable.

4.2.1 Process Definition

The vertex-based diffusion process proceeds in rounds $1, 2, \ldots$. Fix a round t and node i. Let $X_i(t)$ be the current load of node i. Then, node i sends $\lfloor X_i(t)/(d+1)\rfloor$ tokens to each of its neighbors and keeps the same amount for itself. The remaining $X_i(t) - (d+1)\lfloor X_i(t)/(d+1)\rfloor \in [0, d]$ *excess-tokens* are distributed randomly without replacement among i and its d neighbors. For each edge $[i, j] \in E(G)$, define the random variable $Z_{i,j}(t+1)$ to be 1 if i sends one excess token to j in round $t+1$ and 0 otherwise. Similarly, $Z_{i,i}(t+1) = 1$ if i keeps an excess token for itself and 0 otherwise. Note that $Z_{i,j}(t+1)$ for $j \in N(i) \cup \{i\}$ is a $\{0, 1\}$ random variable with $\Pr[Z_{i,j}(t+1) = 1] = X_i(t)/(d+1) - \lfloor X_i(t)/(d+1)\rfloor$. The number of excess tokens sent out by i satisfies

$$Z_{i,i}(t+1) = X_i(t) - \sum_{j:[i,j]\in E(G)} Z_{i,j}(t+1) - (d+1)\left\lfloor \frac{X_i(t)}{d+1} \right\rfloor. \tag{4.1}$$

The process can thus be described as follows,

$$X_i(t+1) = \left\lfloor \frac{X_i(t)}{d+1} \right\rfloor + Z_{i,i}(t+1) + \sum_{j:[i,j]\in E(G)} \left(\left\lfloor \frac{X_j(t)}{d+1} \right\rfloor + Z_{j,i}(t+1) \right). \tag{4.2}$$

Substituting $Z_{i,i}(t+1)$ from Equation (4.1) in Equation (4.2) yields

$$\begin{aligned} X_i(t+1) &= \left\lfloor \frac{X_i(t)}{d+1} \right\rfloor + X_i(t) - \sum_{j:[i,j]\in E(G)} \left(Z_{i,j}(t+1) - \left\lfloor \frac{X_i(t)}{d+1} \right\rfloor \right) \\ &\qquad + \sum_{j:[i,j]\in E(G)} \left(\left\lfloor \frac{X_j(t)}{d+1} \right\rfloor + Z_{j,i}(t+1) \right) \\ &= X_i(t) + \sum_{j:[i,j]\in E(G)} \left(\left\lfloor \frac{X_j(t)}{d+1} \right\rfloor - \left\lfloor \frac{X_i(t)}{d+1} \right\rfloor + Z_{j,i}(t+1) - Z_{i,j}(t+1) \right). \end{aligned} \tag{4.3}$$

4.2.2 Height-Sensitive Process

We will describe a specific realization of the vertex-based diffusion process. Let $\mathcal{T}$ denote the set of all tokens. Initially tokens on each node $u \in V$ are numbered from 1 to x_u. In each round t, each token i has a *location* $W_i(t) \in V$ and a *height* $H_i(t) \in \{1, \cdots, \lceil X_{W_i(t)}(t)/(d+1)\rceil\}$ as $\lceil i/(d+1)\rceil$.

Each edge $(u, v) \in E$ is equipped with two initially empty *Queues* $Q_{u,v}$ and $Q_{v,u}$ and each node u has a self-loop queue $Q_{u,u}$. Tokens with the same rank in queues $Q_{u,v}$ for node u and $v \in N(u)$ are called *siblings*.

One round of the height sensitive process consists of five steps called *Rounding*, *Queuing*, *Shuffling*, *Swapping* and *De-queuing*. All nodes do each step concurrently. Fix round t and node $u \in V$ and assume u's neighbors are numbered from 1 to $d+1$.

1. **Rounding:** Node u determines how many tokens to send to each neighbor in the vertex-based diffusion process. Specifically, it computes $Z_{u,j}$ (whether an excess token is sent to node j) for all $j = 1, \cdots, d+1$ based on the diffusion process' choices and rounding decisions.

2. **Queuing:** Node u places the tokens into the outgoing queues in a round-robin manner. Token j for j from 1 to $(d+1) \cdot \lfloor x_u/(d+1)\rfloor$ are assigned to $Q_{u,i}$ with $i = j \mod (d+1)$. Then, for each j with $Z_{u,j} = 1$ an excess token is added to $Q_{u,j}$.

3. **Shuffling:** Siblings at each rank $r = 1, 2, \cdots, \lfloor x_u/(d+1) \rfloor$ are shuffled independently among the queues $Q_{u,j}$ for $j = 1, \cdots, d+1$ while retaining their original rank. Excess tokens are not shuffled and remain in their assigned queues.

4. **Swapping:** For each edge, the outgoing queues are exchanged with the corresponding incoming queues: tokens queued to go from $j \in N(u) \setminus u$ to u are placed in u's incoming queue, and vice versa. That is, $Q_{u,j}$ and $Q_{j,u}$ are swapped for all $j \neq u$.

5. **De-queuing:** Each node repeatedly removes one token at a time from each of its incoming queues in neighbor order. This continues until all queues $Q_{u,j}$ are empty. At this points, the tokens on each node can be relabeled.

An illustration of the process is shown in figure 5. During round $t+1$, node u sends $\lfloor X_u(t)/(d+1) \rfloor + Z_{u,j}(t+1)$ tokens to each node $j \in N(u)$. Using the same random choices together with Equation (4.2), we obtain that

Observation 4.1. *Fix a load vector. After one round of Height-Sensitive static diffusion , the load distribution is identical to that of the vertex-based diffusion process.*

4.2.3 Further Definitions

Here, we define and recall key concepts that form the basis of our arguments.

For a round t in the continuous process, each node u sends $X_u(t-1)/(d+1)$ load to each neighbor and itself. This model serves as a benchmark for analyzing the discrete dynamics, allowing discrepancy bounds to be transferred from the continuous to the discrete setting.

For the diffusion matrix $\mathbf{P}$ and an initial load vector with (K, κ)–smoothing time is defined as

$$\tau_{\mathrm{S}}(\mathbf{P}, K, \kappa) = \min\{t \mid \mathrm{disc}(\mathbf{P}^t \cdot X(0)) \leq \kappa\},$$

that is, the minimum number of rounds for the continuous process to reduce the discrepancy from K to κ.

We next define the mixing time of a Markov chain on G with transition matrix $\mathbf{P}$.

Definition 4.1 (Mixing Time)**.** *Let G be a d-regular graph. The mixing time $t_{mix}(\mathbf{P})$ of a Markov chain with state space V and transition matrix $\mathbf{P}$ is*

$$t_{mix}(\mathbf{P}) := \min\left\{t \;\middle|\; \forall u, v \in V \;:\; \left|(\mathbf{P}^{\tau})_{u,v} - \frac{1}{n}\right| \leq \frac{1}{n^3}\right\}.$$

Assume there is only one load item initially located at some node. After $t_{mix}(\mathbf{P})$ rounds, this token is located at each node with probability between $1/n \pm 1/n^3$. Together with the negative association property, this allows us to derive a bound on the discrepancy.

Next observation bounds $\tau_{\mathrm{S}}(\mathbf{P}, K, \kappa)$ and $t_{mix}(\mathbf{P})$ in terms the spectral gap of $\mathbf{P}$. In particular, it shows that $\tau_{\mathrm{S}}(\mathbf{P}, K, 1) = \mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(K))$ and $t_{mix}(\mathbf{P}) = \mathrm{O}(\widetilde{\tau}_{\mathrm{S}}(n))$.

Observation 4.2. *[Theorem 1 in [83] and Lemma 2.1 in [57]] Let G be a d-regular with n nodes and diffusion matrix $\mathbf{P}$. For $K > \kappa$ it holds that $\tau_S(\mathbf{P}, K, \kappa) \leq 2\ln(Kn^2/\kappa)/(1-\lambda)$. Moreover, $t_{mix}(\mathbf{P}) \leq 4\log(n)/(1-\lambda)$.*

When G and $\mathbf{P}$ are clear form context we write $\tau_{\mathrm{S}}(K, \kappa) := \tau_{\mathrm{S}}(\mathbf{P}, K, \kappa)$ and $t_{mix} := t_{mix}(\mathbf{P})$. Moreover, to show the decrease in discrepancy after the continues balancing time $\tau_{\mathrm{S}}(K, 1)$, we model the movement of load items as (dependent) random walks. For this purpose, we define the location vector.

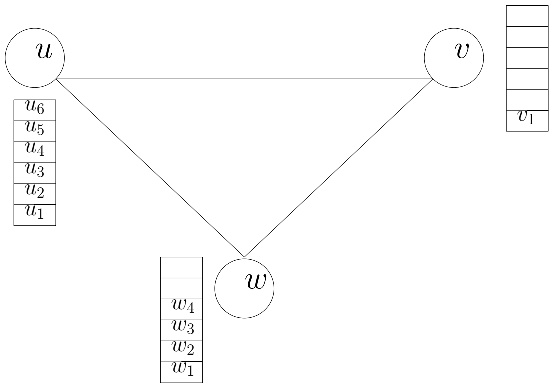

Figure 5: We consider the clock-wise order of neighbors, for instance v is the first neighbor of u. Here $Z_{wu} = 1$ and $Z_{vw} = 1$.

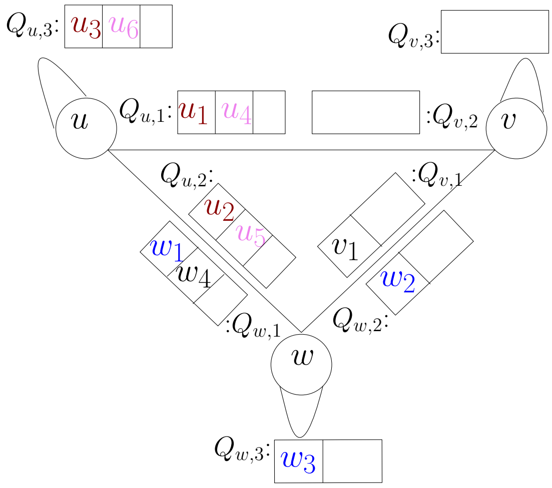

Figure 6: The siblings have the same color. The excess tokens are black.

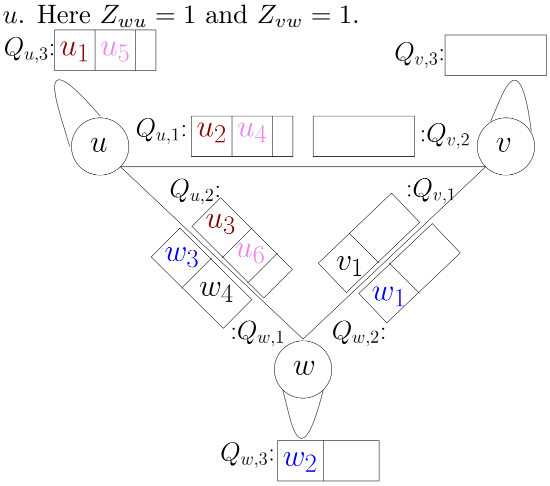

Figure 7: The Siblings are shuffled. For instance, tokens w_1,w_2 and w_3 changed their queues. The rank of tokens does not change.

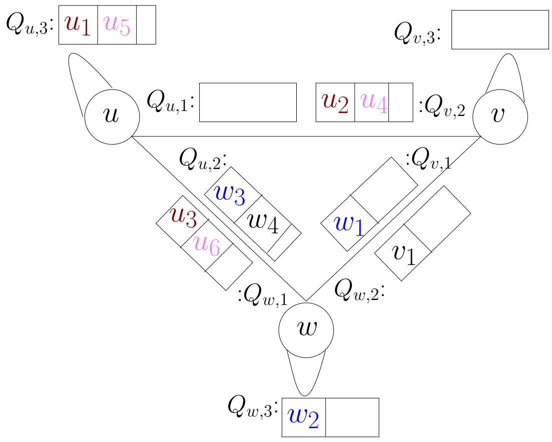

Figure 8: Queues $Q_{u,j}$ and $Q_{j,u}$ for $u \in V$, $j \in N(u)/\{u\}$ are swapped.

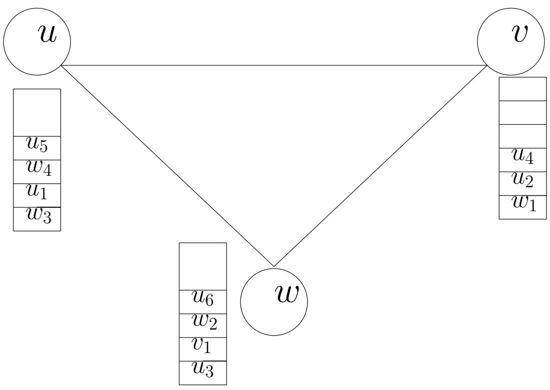

Figure 9: Tokens will be placed on each node.

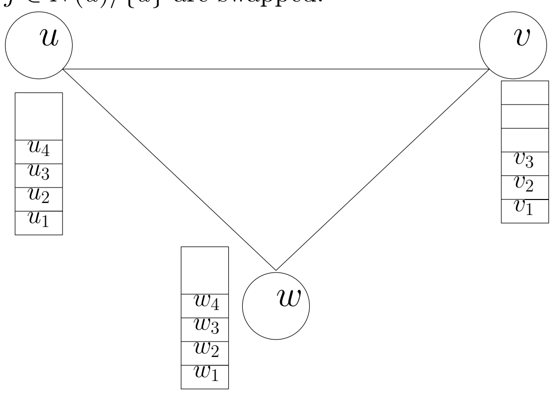

Figure 10: Tokens will be renumbered according to their position on each node.

Definition 4.2 (Location Vector)**.** *The* location vector $W(t)$ *is defined as an* $|\Gamma|$*-dimensional vector where* $W_i(t) = j$ *if token* $i \in \Gamma$ *is located at node* $j \in V$ *at the end of round* t*.*

Here we introduce *vital tokens*, which helps analyze the decrease in the discrepancy after t_{mix} rounds.

Definition 4.3 (Vital Tokens)**.** *For round* t *with load vector* $X(t)$*, define* $\lfloor X(t) \rfloor := (d+1) \cdot \left\lfloor \frac{\min_{u \in V} X_u(t)}{d+1} \right\rfloor$ *the minimum load shared by all nodes that is a multiple of* $d+1$*. The number of vital tokens on node* w *in round* t *is*

$$\Psi_w(t) := X_w(t) - \lfloor X(t) \rfloor, \qquad \textit{and} \qquad \Psi(t) := \sum_{w \in V} \Psi_w(t)$$

is the total number of vital tokens in round t *.*

Observation 4.7 shows that the number of vital tokens is non-increasing over time. Together with the negative association property (lemma 4.10), this ensures that controlling the number of vital tokens received by a node directly translates into a bound on its load.

4.3 Static Diffusion

In this section, we first prove our main result for the static setting (Theorem 4.3). After this proof, the section contains two subsections analyzing the discrepancy at the smoothing time and after the mixing time.

In Subsection 4.3.1, we bound the discrepancy at round $\tau_S(K, \kappa)$, the time required for the continuous process to reduce an (arbitrary) initial discrepancy K down to $\kappa < K$. In particular, at round $\tau_S(K, 1)$, the deviation between the discrete and continuous load is quantified as the contribution of the errors caused by the randomized rounding dynamics. Bounding this carefully, we show that there are $O(nd \log(n))$ vital tokens. In Subsection 4.3.2, we bound the discrepancy after t_{mix} rounds in terms of the (initial) vital tokens, noting that $t_{mix} = O(\tau_S(n, 1))$. Here, using the height-sensitive process, we show that the locations of the tokens are negatively associated, which allows us to derive Chernoff-like tail concentration bounds for the load distribution. Together with the non-increasing property of the vital tokens (Observation 4.7), these results yield a bound on the discrepancy after $\tau_S(K, 1)$ rounds, as stated in Theorem 4.3, while Corollary 4.16 (presented in Section 4.6) provides bounds for specific rounds.

Theorem 4.3. *Consider the discrete static vertex-based diffusion on* d*-regular graph* G *with* n *nodes, diffusion matrix* $\mathbf{P}$ *and any initial load vector with* $\mathrm{disc}(x(0)) := K > 1$*. Let* $i \in [0, \log\log d]$*. For the round* $t^* = \tau_S(\mathbf{P}, K, 1) + i \cdot t_{mix}(\mathbf{P})$*, it holds that*

$$\Pr\left[\mathrm{disc}(X(t^*)) \le 48 \cdot d^{(1/2)^i} \cdot \log(n) + 48\sqrt{d \cdot \log(n)}\right] \ge 1 - 2(i+1) \cdot n^{-3}.$$

Proof. Let $t_0 := \tau_S(\mathbf{P}, K, 1)$. We define *Phase* i for $i \in [1, 1+\log\log d]$ as rounds $[t_0 + (i-1) \cdot t_{mix}(\mathbf{P}) + 1, t_0 + i \cdot t_{mix}(\mathbf{P})]$. With function $e(i)$ we show the last round of phase i i.e., $e(i) := t_0 + i \cdot t_{mix}(\mathbf{P})$ and we let $e(0) := t_0$. For each $i \in [0, \log\log d]$, we define an event

$$\Lambda_i := \left\{\mathrm{disc}(X(e(i))) \le 48 \cdot d^{(1/2)^i} \cdot \log(n) + 48\sqrt{d \cdot \log(n)}\right\} \tag{4.4}$$

and we show by induction that $\mathbf{Pr}\left[\overline{\Lambda_i}\right] \le 2(i+1) \cdot n^{-3}$.

Base case $i = 0$. Here $t = e(0) = \tau_S(\mathbf{P}, K, 1)$. From Lemma 4.4 (presented in Subsection 4.3.1) with $\delta := 8d \cdot \log(n)$ and $\kappa := 1$ it follows that, for a node $u \in V$,

$$\begin{aligned}\Pr[|X_u(t) - \overline{x}| \geq 8d \cdot \log(n) + 1] &\leq \exp\left(-\frac{64d^2 \log^2 n}{2(d+1)^2 + 16d^2 \log(n)/3}\right) \\ &\leq \exp\left(-\frac{64d^2 \log^2 n}{16d^2 \log(n)}\right) \leq n^{-4}.\end{aligned}$$

Note that $\operatorname{disc}(X(t)) \leq 2 \cdot \max_{u \in V} |X_u(t) - \overline{x}|$. An application of union bound over all nodes completes the proof of the base case.

Induction Step. Assume the event Λ_i for some $i \in [0, \log\log(d) - 1]$ holds with probability (at least) $1 - 2(i+1) \cdot n^{-3}$; we show it for $i+1$. The number of vital tokens in round $e(i)$ is

$$\begin{aligned}\Psi(e(i)) &= \sum_{w \in V}\left(X_w(e(i)) - \min_{u \in V} X_u(e(i))\right) + \sum_{w \in V}\left(\min_{u \in V} X_u(e(i)) - \lfloor X(e(i)) \rfloor\right) \\ &\leq n \cdot \operatorname{disc}(X(e(i))) + n \cdot d \leq 48n \cdot \left(d^{(1/2)^i} \log(n) + \sqrt{d \cdot \log(n)}\right) + nd,\end{aligned}$$

where the last inequality follows from the inductive hypothesis (event Λ_i). Then, from Lemma 4.5 (presented in Subsection 4.3.2) it follows that,

$$\mathbf{Pr}\left[\operatorname{disc}(X(e(i+1))) > \frac{2\Psi(e(i))}{n^2} + \sqrt{48\frac{\Psi(e(i))}{n}\log(n)} \;\middle|\; \Lambda_i\right] \leq 2n^{-3}. \tag{4.5}$$

Since $i \geq 0$ and $d \leq n$, then

$$\begin{aligned}\frac{2\Psi(e(i))}{n^2} + \sqrt{48\frac{\Psi(e(i))}{n}\log(n)} \;&\leq\; O(1) + \sqrt{48\frac{48n\left(d^{(1/2)^i}\log(n) + \sqrt{d\log(n)}\right) + nd}{n}\log(n)} \\ &\leq 48d^{(1/2)^{i+1}}\log(n) + 48\sqrt{d\log(n)}.\end{aligned} \tag{4.6}$$

Substituting Equation (4.6) into Equation (4.5) and applying the law of total probability yields

$$\begin{aligned}\mathbf{Pr}\left[\operatorname{disc}(X(e(i+1))) > 48d^{(1/2)^{i+1}}\log(n) + 48\sqrt{d\log(n)}\right] &\leq 2n^{-3} + \mathbf{Pr}\left[\overline{\Lambda_i}\right] \\ &\leq 2(i+2)n^{-3}.\end{aligned}$$

for which we use the inductive hypothesis. It completes the induction step and finishes the proof. □

In the rest of this section we prove the key intermediate results used in the theorem's proof.

4.3.1 Discrepancy at Smoothing Time

The next lemma bounds the load of a node at round $\tau_S(K, \kappa)$ for arbitrary $K > \kappa$. This result improves on the existing bound for $d = \Omega(\sqrt{\log(n)})$, where the previous best was $O(d^2\sqrt{\log(n)})$ ([85]). In particular, it provides our first new discrepancy bound, showing that at round $\tau_S(K, 1)$ the discrepancy is at most $17d \cdot \log(n)$ w.h.p.

Lemma 4.4. *Let G be a d-regular graph with n nodes and diffusion matrix $\mathbf{P}$. Consider an arbitrary initial load vector with $\operatorname{disc}(x(0)) := K > 1$. Then for any node $w \in V$, round $t = \tau_S(\mathbf{P}, K, \kappa)$ with $K > \kappa$*

and any $\delta > 0$, it holds that

$$\Pr[|X_w(t) - \overline{x}| \geq \delta + \kappa] \leq 2 \cdot \exp\left(-\frac{\delta^2}{2(d+1)^2 + \frac{2}{3}d\delta}\right).$$

Proof sketch. The proof proceeds in three main steps. First, we bound the deviation between the discrete and continuous load by introducing a suitable sequence of weighted prefix sums whose terminal value equals the total deviation of interest. Second, we show that this sequence of prefix sums is a martingale with respect to the natural filtration generated by the process, so the total deviation is the terminal value of that martingale. Third, we use negative association to verify the bounded-differences condition required to apply an exponential concentration inequality (Azuma–Hoeffding) to that martingale; this yields that probabilistic deviations are negligible and gives the claimed bound. The remainder of the proof fills in these steps with the precise parameter choices.

Proof. To track of the discrete load, we consider the deviation from the continuous load. Let $\ell(t) := (\ell_u(t))_{t\in\mathbb{N},u\in V}$ denote the continuous load vector, whose dynamics are given by

$$\ell_u(t+1) = \ell_u(t) + \sum_{v:[u,v]\in E(G)} \frac{\ell_v(t) - \ell_u(t)}{d+1}.$$

To obtain a recursion for the discrete process analogous to the continuous dynamics, [14] introduces a random variable $\varepsilon_{u,v}(t+1)$ representing the rounding error at edge $[u,v] \in E$ caused by node $u$at round $t+1$,

$$\varepsilon_{u,v}(t+1) := \frac{X_u(t)}{d+1} - \frac{X_v(t)}{d+1} - \left\lfloor \frac{X_u(t)}{d+1} \right\rfloor + \left\lfloor \frac{X_v(t)}{d+1} \right\rfloor + Z_{v,u}(t+1) - Z_{u,v}(t+1).$$

This allows the discrete process to be expressed as the continuous process plus accumulated rounding errors, enabling precise analysis of the discrepancy. From this together with Equation (4.3), it follows that

$$X_u(t+1) = X_u(t) + \sum_{v:[u,v]\in E(G)} \frac{X_v(t) - X_u(t)}{d+1} + \varepsilon_{u,v}(t+1).$$

Define an error vector $\varepsilon(t)$ with $\varepsilon_u(t) := \sum_{v:[u,v]\in E(G)} \varepsilon_{u,v}(t)$. With this notation the discrete load satisfies $X(t) = \mathbf{P} \cdot X(t-1) + \varepsilon(t)$. Note that this is the same as Equation (3) in which $A(t) := \vec{0}$. Solving the recursion with $\ell(0) = X(0)$ gives

$$X(t) = \mathbf{P}^t \cdot X(0) + \sum_{s=0}^{t-1} \mathbf{P}^s \cdot \varepsilon(t-s) = \ell(t) + \sum_{s=0}^{t-1} \mathbf{P}^s \cdot \varepsilon(t-s)$$

where $\mathbf{P}^0$ is the $n \times n$ identity matrix. For each node $w \in V$, this implies

$$X_w(t) - \ell_w(t) = \sum_{s=0}^{t-1} \sum_{u\in V(G)} \varepsilon(t-s)_u \cdot (\mathbf{P}^s)_{u,w} = \sum_{s=0}^{t-1} \sum_{u\in V} \sum_{v:[u,v]\in E(G)} \varepsilon_{u,v}(t-s) \cdot (\mathbf{P}^s)_{u,w}.$$

Using the antisymmetry $\varepsilon_{u,v}(t-s) = -\varepsilon_{v,u}(t-s)$, we obtain

$$X_w(t) - \ell_w(t) = \sum_{s=0}^{t-1} \sum_{[u,v]\in E(G)} \varepsilon_{u,v}(t-s) \cdot \Big((\mathbf{P}^s)_{u,w} - (\mathbf{P}^s)_{v,w}\Big).$$

Since $t = \tau_{\mathrm{S}}(\mathbf{P}, K, \kappa)$, we have $|\ell_w(t) - \overline{x}| \leq \kappa$ for all $w \in V$ (Definition 1.7), so we can write $\ell_w(t) = \overline{x} + \Theta$, where $|\Theta| \leq \kappa$. It Follows that

$$X_w(t) - \overline{x} = \Theta + \underbrace{\sum_{s=1}^{t} \sum_{[u:v] \in E(G)} \varepsilon_{u,v}(s) \cdot \left(\left(\mathbf{P}^{t-s}\right)_{u,w} - \left(\mathbf{P}^{t-s}\right)_{v,w} \right)}_{:=Z}.$$

Equivalently, define

$$Z := \sum_{s=1}^{t} \sum_{[u:v] \in E(G)} \varepsilon_{u,v}(s) \cdot \left(\left(\mathbf{P}^{t-s}\right)_{u,w} - \left(\mathbf{P}^{t-s}\right)_{v,w} \right). \tag{4.7}$$

Since $\mathbf{E}[\varepsilon_{u,v}(s)] = 0$ and $\varepsilon_{u,v}(s)$ is independent of $\mathbf{P}$, the linearity of expectation gives $\mathbf{E}[Z] = 0$. Assume for any $\delta > 0$, we have

$$\Pr[|Z| \geq \delta] \leq 2 \cdot \exp\left(-\frac{\delta^2}{2(d+1)^2 + 2d\delta/3} \right), \tag{4.8}$$

which would complete the proof. Hence, it remains to prove Equation (4.8). Fix a node $w \in V$. For fixed round t and any $s \leq t$, define

$$G_{u,v}(s) := \left(\mathbf{P}^{t-s}\right)_{u,w} - \left(\mathbf{P}^{t-s}\right)_{v,w}.$$

From Equation (4.7) we can write

$$\begin{aligned}
Z &= \sum_{s=1}^{t} \sum_{[u:v] \in E(G)} \varepsilon_{u,v}(s) \cdot G_{u,v}(s) \\
&= \sum_{s=1}^{t} \sum_{[u:v] \in E(G)} (\mathbf{E}[Z_{u,v}(s)] - Z_{u,v}(s) + Z_{v,u}(s) - \mathbf{E}[Z_{v,u}(s)]) \cdot G_{u,v}(s) \\
&= \sum_{s=1}^{t} \sum_{[u:v] \in E(G)} (\mathbf{E}[Z_{u,v}(s)] - Z_{u,v}(s)) \cdot G_{u,v}(s) - (\mathbf{E}[Z_{v,u}(s)] - Z_{v,u}(s)) \cdot \underbrace{G_{u,v}(s)}_{=-G_{v,u}(s)}
\end{aligned}$$

and since $G_{u,v}(s) = -G_{v,u}(s)$ then we get,

$$\begin{aligned}
Z &= \sum_{s=1}^{t} \sum_{[u:v] \in E(G)} (\mathbf{E}[Z_{u,v}(s)] - Z_{u,v}(s)) \cdot G_{u,v}(s) + (\mathbf{E}[Z_{v,u}(s)] - Z_{v,u}(s)) \cdot G_{v,u}(s) \\
&= \sum_{s=1}^{t} \sum_{u \in V} \sum_{v \in N(u)} (\mathbf{E}[Z_{u,v}(s)] - Z_{u,v}(s)) \cdot G_{u,v}(s),
\end{aligned}$$

Recall that $Z_{u,v}(s) = 1$ with probability $X_u(s-1)/(d+1) - \lfloor X_u(s-1)/d+1 \rfloor$ and 0 otherwise. For clarity, for a fixed node $u \in V$ and round s, define

$$A_u(s) := \sum_{v \in N(u)} (\mathbf{E}[Z_{u,v}(s)] - Z_{u,v}(s)) \cdot G_{u,v}(s).$$

Here, $A_u(s)$ represents the total contribution of node u in round s to the sum Z. Note that $\sum_{v \in N(u)} (\mathbf{E}[Z_{u,v}(s)] - Z_{u,v}(s))$ depends only on u, and the outcome of $Z_{u,v}(s)$ is independent of $G_{u,v}(s)$.

Consider an ordering of nodes from 1 to n. Define a filtration $\mathcal{F}$ where $\mathcal{F}_i$ reveals $Z_{u,v}(s)$ for all

$u \in [n]$, $v \in N(u)$, and rounds $s \in [1, \lfloor i/n \rfloor]$, as well as $Z_{u,v}(s)$ for round $s = \lceil i/n \rceil$ and $u \in [1, i \bmod n]$ and $v \in N(u)$. For $\ell \in [t \cdot n]$, let $1 \le s_1 \le t$ and $1 \le i_1 \le n$ such that $\ell = (s_1 - 1) \cdot n + i_1$ and define

$$Y_\ell := \sum_{s=1}^{s_1-1} \sum_{i=1}^{n} A_i(s) + \sum_{i=1}^{i_1} A_i(s_1).$$

Then $Z = Y_{t \cdot n}$. Since $\mathbf{E}[A_i(s)] = 0$ for all $s \in [t]$ and $i \in [n]$, the sequence of $Y_1, Y_2, \cdots, Y_{t\cdot n}$ forms a martingale with respect to the filtration $\mathcal{F}$, and $\mathbf{E}[Y_{t \cdot n}] = 0$.

We are to apply Theorem 6.12. To do so, we need to bound $|Y_\ell - Y_{\ell-1}|$ and $\mathrm{Var}[Y_\ell \mid \mathcal{F}_{\ell-1}]$ for each $\ell \in [t \cdot n]$. For $\ell = (s_1 - 1) \cdot n + i_1$ with $s_1 \in [t]$ and $i_1 \in [n]$ we have,

$$|Y_\ell - Y_{\ell-1}| = |A_{i_1}(s_1)| \le d.$$

Moreover, we have

$$\mathrm{Var}\Big[Y_\ell \;\Big|\; \mathcal{F}_{\ell-1}\Big] = \mathrm{Var}\Big[Y_{\ell-1} + A_{i_1}(s_1) \;\Big|\; \mathcal{F}_{\ell-1}\Big] = \mathrm{Var}\Big[A_{i_1}(s_1) \;\Big|\; \mathcal{F}_{\ell-1}\Big],$$

where the last inequality holds since $\mathcal{F}_{\ell-1}$ fixes $Y_{\ell-1}$. Using the definition of $A_{i_1}(s_1)$ we get,

$$\begin{aligned}
\mathrm{Var}\Big[Y_\ell \;\Big|\; \mathcal{F}_{\ell-1}\Big] &= \mathrm{Var}\left[\sum_{j \in N(i_1)} (\mathbf{E}[Z_{i_1,j}(s_1)] - Z_{i_1,j}(s_1)) \cdot G_{i_1,j}(s_1) \;\Bigg|\; \mathcal{F}_{\ell-1}\right] \\
&= \mathrm{Var}\left[\sum_{j \in N(i_1)} Z_{i_1,j}(s_1) \cdot |G_{i_1,j}(s_1)| \;\Bigg|\; \mathcal{F}_{\ell-1}\right] \\
&\overset{(a)}{\le} \sum_{j \in N(i_1)} \mathrm{Var}\Big[Z_{i_1,j}(s_1) \cdot |G_{i_1,j}(s_1)| \;\Big|\; \mathcal{F}_{\ell-1}\Big] \\
&= \sum_{j \in N(i_1)} (G_{i_1,j}(s_1))^2 \cdot \mathrm{Var}\Big[Z_{i_1,j}(s_1) \;\Big|\; \mathcal{F}_{\ell-1}\Big] \le \sum_{j \in N(i_1)} (G_{i_1,j}(s_1))^2,
\end{aligned}$$

where (a) follows from Observation 4.8. Applying Theorem 6.12 to the martingale sequence $Y_1, \cdots, Y_{t \cdot n}$ where for $\ell = (s-1) \cdot n + i$ (with $s \in [t]$ and $i \in [n]$) we have $M = d$ and $\sigma_\ell^2 = \sum_{j \in N(i)} (G_{i,j}(s))^2$, gives

$$\Pr[Y_{t \cdot n} \ge \delta] \le \exp\left(-\frac{\delta^2}{2\Big(\sum_{s=1}^{t} \sum_{i=1}^{n} \sum_{j \in N(i)} (G_{i,j}(s))^2 + d\delta/3\Big)}\right).$$

Recall that $\mathbf{E}[Y_{t \cdot n}] = 0$. Furthermore we have

$$\sum_{s=1}^{t} \sum_{i=1}^{n} \sum_{j \in N(i)} (G_{i,j}(s))^2 = 2 \cdot \sum_{s=1}^{t} \sum_{[i,j] \in E} (G_{i,j}(s))^2 \overset{(a)}{\le} (d+1)^2.$$

where (a) follows form [85, Theorem 6.4] with $\gamma = 1 + 1/d$ and $\Delta = d$. Hence, we obtain

$$\Pr[Y_{t \cdot n} \ge \delta] \le \exp\left(-\frac{\delta^2}{2((d+1)^2 + d\delta/3)}\right).$$

Using Theorem 6.13 (with $a_i = 0$) instated of Theorem 6.12 gives

$$\Pr[Y_{t\cdot n} \le -\delta] \le \exp\left(-\frac{\delta^2}{2((d+1)^2 + d\delta/3)}\right).$$

Recalling that $Z = Y_{t\cdot n}$, an union bound over the upper and lower tails yields

$$\Pr[|Z| \ge \delta] \le 2 \cdot \exp\left(-\frac{\delta^2}{2(d+1)^2 + 2d\delta/3}\right).$$

This establishes Equation (4.8) and shifting δ by κ finishes the proof. □

4.3.2 Discrepancy at Mixing Time

The next lemma bounds the discrepancy at the mixing time in terms of the initial number of vital tokens. It shows that the discrepancy after the mixing time, w.h.p., is bounded by $O(1 + \sqrt{\Psi(0) \cdot \log(n)/n})$, provided that $\Psi(0) \le n^2$. The proof leverages the negative association property to derive concentration bounds.

Lemma 4.5. *Let G be a d-regular graph with n nodes. Consider height-sensitive diffusion process with an initial load vector containing $\Psi(0)$ vital tokens. Then for $\tau := t_{mix}(\mathbf{P})$, we have*

$$\mathbf{Pr}\left[\mathrm{disc}(X(\tau)) \le \frac{2\Psi(0)}{n^2} + \sqrt{48\frac{\Psi(0)}{n}\log(n)}\right] \ge 1 - 2n^{-3}.$$

Proof. Let the number of vital tokens received by node u in round τ be

$$Y_u(\tau) := \sum_{i \in [\Psi(0)]} \mathbf{1}_{\{W_i(\tau) = u\}}\,.$$

Then we have

$$\begin{aligned}\mu_u := \mathbf{E}[Y_u(\tau)] &= \sum_{i \in [\Psi(0)]} \mathbf{Pr}[W_i(\tau) = u \mid w_i(0)] \\ &\overset{(a)}{\le} \frac{\Psi(0)}{n}\left(1 + \frac{1}{n^2}\right),\end{aligned}$$

where (a) follows from the first part of the third statement of Observation 4.9. Similarly, from the second part of the same statement, we get

$$\mu_u \ge \frac{\Psi(0)}{n}\left(1 - \frac{1}{n}\right).$$

Although $\mathbf{1}_{\{W_i(\tau)=u\}}$ and $\mathbf{1}_{\{W_j(\tau)=u\}}$ for tokens $i \ne j$ may not be independent, Lemma 4.10 ensures they satisfy the negative association property. Define the events

$$\Lambda_{\max} := \left\{\max_{u \in V} Y_u(\tau) \le \mu_u + \sqrt{12\mu_u \log(n)}\right\}, \quad \text{and} \quad \Lambda_{\min} := \left\{\min_{u \in V} Y_u(\tau) \ge \mu_u - \sqrt{12\mu_u \log(n)}\right\}.$$

By Lemma 6.15 and an union bound over all nodes, the event $\Lambda_{\min} \cap \Lambda_{\max}$ occurs with probability at least $1-2n^{-3}$. Conditioned on this event, we get

$$\begin{aligned}\max_{u\in V} Y_u - \min_{u\in V} Y_u &\le \max_{u\in V}\Big(\mu_u + \sqrt{12\mu_u \log(n)}\Big) - \min_{u\in V}\Big(\mu_u - \sqrt{12\mu_u \log(n)}\Big)\\ &\le \frac{2\Psi(0)}{n^2} + \sqrt{48\frac{\Psi(0)}{n}\log(n)},\end{aligned}$$

where we use the bounds on μ_u from the previous step. From Observation 4.13 it follows that the $\mathrm{disc}(X(\tau)) = \mathrm{disc}(Y(\tau))$ and it finishes the proof.

□

4.4 Dynamic Diffusion

We now extend our analysis to the dynamic discrete vertex-based diffusion process. Let $G=(V,E)$ be a d-regular graph, where each node initially holds an integer number of load items (tokens). At the start of each round, $m \in \mathbb{N}$ new tokens are generated and assigned uniformly at random to the nodes. Afterward, each node redistributes its tokens as evenly as possible among itself and its neighbors. If perfect redistribution is impossible without splitting tokens, any remaining excess tokens are distributed randomly among the neighbors without replacement. This dynamic process introduces additional fluctuations in the load distribution. In the following, we analyze how these newly generated tokens contribute to the overall discrepancy.

Recall that the dynamic load balancing is modeled by a Markov chain $(X(t))_{t\in\mathbb{N}}$, where $X(t) = (X_i(t))_{i\in[n]} \in \mathbb{R}^n$ is the load vector at the end of round t, and $X_i(t)$ is the load of node i at time t. Let $A_i(t)$ denote the number of (newly generated) tokens allocated to node i in step t and define $A(t) = (A_i(t))_{i\in[n]}$. Let $\varepsilon(t) \in \mathbb{R}^n$ be the vector of additive rounding errors in round t, where $\varepsilon_k(t)$ measures the difference between the actual load at node k after step t and the corresponding load in the continuous model with arbitrarily divisible loads.

For the diffusion matrix $\mathbf{P}$ we define the global divergence as

$$\Upsilon(\mathbf{P}) := \max_{k\in[n]} \Upsilon_k(\mathbf{P}), \quad \text{where} \quad \Upsilon_k(\mathbf{P}) := \sqrt{\sum_{s=1}^{\infty}\left\|(\mathbf{P}^s)_{k,\cdot} - \frac{\vec{1}}{n}\right\|_2^2}.$$

The next theorem is the main result of this section.

Theorem 4.6. *Let G be a d-regular with diffusion matrix $\mathbf{P}$. Consider the discrete dynamic vertex-based diffusion on G where m items are added uniformly at random to the nodes in each step, starting from an arbitrary initial load vector such that* $\mathrm{disc}(x(0)) := K > 1$. *Let $i \in [0, \log\log d]$ and define $t^* := \tau_S(\mathbf{P}, K, 1/(2n)) + i \cdot t_{mix}(\mathbf{P})$. Then, with probability at least $1-2(i+1)n^{-3} - 2n^{-\gamma+1}$,*

$$\mathrm{disc}(X(t^*)) \le 48 d^{(1/2)^i}\log(n) + 48\sqrt{d\log(n)} + \frac{8}{3}\gamma\log(n) + \sqrt{32\gamma\log(n)\frac{m}{n}}\Upsilon(\mathbf{P}).$$

Proof. Since the placement of newly added tokens are independent of the matrix $\mathbf{P}$, we can use the same approach as in Equation (5). Expanding the recurrence of Equation (3) gives

$$\begin{aligned} X(t) &= \mathbf{P} \cdot (X(t-1) + A(t)) + \varepsilon(t) \\ &= \mathbf{P} \cdot \Big(\underbrace{(\mathbf{P} \cdot (X(t-2) + A(t-1)) + \varepsilon(t-1))}_{X(t-1)} + A(t) \Big) + \varepsilon(t) \\ &= \mathbf{P}^2 \cdot X(t-2) + \sum_{s=t-1}^{t} \mathbf{P}^{t-s+1} \cdot A(s) + \sum_{s=t-1}^{t} \mathbf{P}^{t-s} \cdot \varepsilon(s) \end{aligned}$$

Iterating this expansion back to the initial round, we obtain

$$X(t) = \underbrace{\mathbf{P}^t \cdot X(0)}_{\substack{I(t) \\ \text{initial load contribution}}} + \underbrace{\sum_{s=1}^{t} \mathbf{P}^{t-s+1} \cdot A(s)}_{\substack{D(t) \\ \text{dynamically allocated Load contribution}}} + \underbrace{\sum_{s=1}^{t} \mathbf{P}^{t-s} \cdot \varepsilon(s)}_{\substack{R(t) \\ \text{rounding error contribution}}} .$$

We denote the three sums in the expansion by $I(t)$, $D(t)$, and $R(t)$ as indicated. By the sub-additivity of discrepancy (Observation 2.13), we have

$$\operatorname{disc}(X(t)) \leq \operatorname{disc}(I(t) + R(t)) + \operatorname{disc}(D(t)), \tag{4.9}$$

showing that it suffices to bound each term individually and sum the results.

Note that $I(t) + R(t)$ corresponds to the static load balancing (without the dynamically added items). Hence, for any $i \in [0, \log\log(d)]$ and $t^\star = \tau_{\mathrm{S}}(\mathbf{P}, K, 1/(2n)) + i \cdot t_{mix}(\mathbf{P})$, Theorem 4.3 implies

$$\mathbf{Pr}\Big[\operatorname{disc}(I(t^\star) + R(t^\star)) > 48 \cdot d^{(1/2)^i} \cdot \log(n) + 48\sqrt{d \cdot \log(n)}\Big] \leq 2(i+1) \cdot n^{-3}. \tag{4.10}$$

To bound the contribution of dynamically allocated items, we relate it to the global divergence of **P**,a s formalized in Lemma 4.14. It states

$$\mathbf{Pr}\left[\operatorname{disc}(D(t^\star)) > \frac{8}{3} \cdot \gamma \log(n) + \sqrt{32\gamma \log(n) \cdot \frac{m}{n}} \cdot \Upsilon(\mathbf{P})\right] \leq 2 \cdot n^{-\gamma+1}. \tag{4.11}$$

Applying an union bound over Equation (4.10) and Equation (4.11) together with Equation (4.9) finishes the proof. Note that the proof of Lemma 4.14 follows directly from the tools developed in Part Two, specifically Lemma 2.4.

□

4.5 Technical Lemmas

In this section, we list the technical lemmas used in our analysis. It is divided into two subsections: one for intermediate results in the static setting, and one for the dynamic setting.

4.5.1 Ingredients Used in Theorem 4.3

We first prove the intermediate results for the static diffusion, highlighting the key component of our analysis: the negative association property Lemma 4.10. We start by showing that the number of vital tokens is non-increasing over time.

Observation 4.7. *The function $\Psi(t)$ is non-increasing over time.*

Proof. Fix round t. Each node has at least $\lfloor X(t) \rfloor$ tokens. During the balancing of round $t+1$ (i.e., the five steps of height-sensitive process), each node receives at least $\lfloor X(t) \rfloor/(d+1)$ tokens from each of its $d+1$ neighbors. Hence, we get

$$\lfloor X(t+1) \rfloor \geq (d+1) \cdot \frac{\lfloor X(t) \rfloor}{d+1}.$$

Using conservation of total load, we get

$$\begin{aligned}\Psi(t+1) &= \sum_{w \in V} (X_w(t+1) - \lfloor X(t+1) \rfloor) = \left(\sum_{w \in V} X_w(t) \right) - n \cdot \lfloor X(t+1) \rfloor \\ &\leq \left(\sum_{w \in V} X_w(t) \right) - n \cdot \lfloor X(t) \rfloor = \Psi(t).\end{aligned}$$

□

The following simple observation is crucial in our analysis. It bounds the variance of a weighted sum of negatively dependent variables by the sum of the individual variances, which is essential for deriving the discrepancy bound after the smoothing time (Lemma 4.4).

Observation 4.8. *Consider the vertex-based diffusion process. Fix a round t and node i. Let a_j for $j \in N(i)$ be non-negative (or non-positive) constants. Then*

$$Var\left[\sum_{j \in N(i)} a_j \cdot Z_{i,j}(t) \right] \leq \sum_{j \in N(i)} a_j^2 \cdot Var[Z_{ij}(t)].$$

Proof. First, note that each node has at most d excess token in each step and has $d+1$ neighbors. For distinct neighbours $j, k \in N(i)$, we have

$$\begin{aligned}\Pr[Z_{i,j}(t) = 1 \wedge Z_{i,k}(t) = 1] &= \Pr[Z_{i,j}(t) = 1 \mid Z_{i,k}(t) = 1] \cdot \Pr[Z_{i,k}(t) = 1] \\ &\leq \Pr[Z_{i,j}(t) = 1] \cdot \Pr[Z_{i,k}(t) = 1], \end{aligned} \tag{4.12}$$

where the inequality holds since conditioning on node k receiving an excess token form i in round t, may decrease the probability that node j receives an excess token form i in round t.

Now using the definition of variance via covariance, we have

$$\begin{aligned}\text{Var}\left[\sum_{j \in N(i)} a_j \cdot Z_{ij}(t) \right] &= \sum_{j \in N(i)} a_j^2 \cdot \text{Var}[Z_{ij}(t)] + \sum_{j \neq k:\ j,k \in N(i)} \text{Cov}[a_j \cdot Z_{ij}(t), a_k \cdot Z_{ik}(t)] \\ &= \sum_{j \in N(i)} a_j^2 \cdot \text{Var}[Z_{i,j}(t)] + \sum_{j \neq k:\ j,k \in N(i)} a_j \cdot a_k \cdot \mathbf{E}[Z_{ij}(t) \cdot Z_{ik}(t)] - a_j \cdot a_k \cdot \mathbf{E}[Z_{ij}(t)] \cdot \mathbf{E}[Z_{ik}(t)] \\ &= \sum_{j \in N(i)} a_j^2 \cdot \text{Var}[Z_{ij}(t)] + \sum_{j \neq k:\ j,k \in N(i)} a_j \cdot a_k \cdot \mathbf{Pr}[Z_{ij}(t) = 1 \wedge Z_{ik}(t) = 1] \\ &\qquad - a_j \cdot a_k \cdot \mathbf{Pr}[Z_{ij}(t) = 1] \cdot \mathbf{Pr}[Z_{ik}(t) = 1],\end{aligned}$$

and applying Equation (4.12) to the right-hand side of the last equation gives

$$
\begin{aligned}
\operatorname{Var}\left[\sum_{j\in N(i)} a_j \cdot Z_{ij}(t)\right] &\le \sum_{j\in N(i)} a_j^2 \cdot \operatorname{Var}[Z_{ij}(t)] \\
&\quad + \sum_{j\neq k:\ j,k\in N(i)} a_j \cdot a_k \cdot \mathbf{Pr}[Z_{ij}(t)=1]\cdot \mathbf{Pr}[Z_{ik}(t)=1] - a_j\cdot a_k \cdot \mathbf{Pr}[Z_{ij}(t)=1]\cdot\mathbf{Pr}[Z_{ik}(t)=1] \\
&= \sum_{j\in N(i)} a_j^2 \cdot \operatorname{Var}[Z_{ij}(t)],
\end{aligned}
$$

finishing the proof. □

We state an useful observation that is frequently used in our analysis. It shows that the height of a token is non-increasing over time and provides a bound on the probability that a single token reaches any node after the mixing time. Recall that $W_i(t_1)$ is the (random) node on which the token i is located in round t_1.

Observation 4.9. *Consider the height-sensitive diffusion process and fix a token $t \in \mathcal{T}$.*

1. *For any pair of rounds $t_1 < t_2$, it holds that $H_i(t_2) \le H_i(t_1)$.*
2. *For the token's location, it holds that*

$$\mathbf{Pr}[W_i(t_2) = w \mid W_i(t_1) = v] = \left(\mathbf{P}^{t_2-t_1}\right)_{v,w}.$$

   Consequently, for any set $D \subseteq V$, $\mathbf{Pr}[\,W_i(t_2) \in D \mid W(t_1) = w(t_1)\,] = (\mathbf{P}^{t_2-t_1})_{w_i(t_1),D}$.

3. *Let $t_{mix}(\mathbf{P})$ be the mixing time of a Markov chain with state space V and transition matrix $\mathbf{P}$. For $t_2 \ge t_1 + t_{mix}(\mathbf{P})$:*
   - *(a) $\mathbf{Pr}[\,W_i(t_2) = v \mid W_i(t_1) = w\,] \le \frac{1}{n} + \frac{1}{n^3}$.*
   - *(b) $\mathbf{Pr}[\,W_i(t_2) = v \mid W_i(t_1) = w\,] \ge \frac{1}{n} - \frac{1}{n^2}$.*

Proof. Assume $H_i(t+1) \le H_i(t)$ holds for all previous rounds: then a simple induction over t completes the proof. It remains to prove that $H_i(t_1+1) \le H_i(t_1)$ for a single round. Let $W_i(t_1) = v$ and $H_i(t_1) = k \in \mathbb{N}$. For each buffer $Q_{u,j}$ and each step $s \in [5]$ of the height-sensitive process, let $Q^s_{i,j}$ be the state of the queue after step s. Initially $Q^0_{i,j} = \emptyset$ for all $u \in V$ and $j \in N(u)$.

The rank of token i does not change during the queuing, shuffling and swapping steps. Now consider the de-queuing step. Let $w := W_i(t_1+1)$ and suppose token i is in queue $Q^4_{w,d+1}$. De-queuing proceeds rank after rank. Since token i has rank k, before it is dequeued, node w receives at most k tokens from each queue $Q^4_{w,j}$ for $j = 1, \cdots, d$ and $k-1$ from $Q^4_{w,d+1}$. Hence, the height of token i after this step satisfies

$$H_i(t_1+1) \le \left\lceil \frac{k\cdot d + (k-1) + 1}{d+1} \right\rceil = k = H_i(t_1).$$

This shows that the height does not increase in a single round, completing the induction.

Now we prove the second statement. Fix nodes $u, v \in V$ and round $t \ge t_1$. We prove by induction on t that for a token i located at node w in round t_1, the probability it is on v at round t equals $(\mathbf{P}^{t-t_1})_{w,v}$.

- Base case: $t = t_1$. Trivially, $\mathbf{P}^0_{w,v} = \mathbf{I}_{w,v}$, so the claim holds.
- Induction step: Assume the claim holds for round $t \ge t_1$. Let $W_i(t) = u$ at the end of round t. Consider round $t+1$ and distinguish cases:

Case 1: Token i is an excess token before the rounding of round $t+1$. At rounding step, it is decided which queue receive token i. During the shuffling, swapping, and de-queuing steps, neither the rank of token i nor the queue containing token i changes until swapping step. Suppose node u has $r \le d$ excess tokens at the start of the round. The excess tokens are distributed among the $d+1$ neighbors (including itself) uniformly without replacement. The probability that v receives i-th token for $i \le r$ is

$$\begin{aligned}
\mathbf{Pr}[\, W_i(t+1) = v \mid W_i(t) = u \,] &= \mathbf{Pr}\big[i \in Q^1_{u,v}\big] \\
&= \underbrace{\frac{d}{d+1} \cdot \frac{d-1}{d} \cdots \frac{(d+1)-(i-1)}{(d+1)-(i-2)}}_{\text{node } v \text{ receives none of the first } i-1 \text{ tokens}} \cdot \frac{1}{(d+1)-(i-1)} \\
&= \frac{1}{d+1}.
\end{aligned}$$

This shows that each neighbor, including v, receives an excess token with uniform probability $1/(d+1)$, matching $\mathbf{P}_{u,v}$.

Case 2: Token i is a non-excess token before the rounding of round $t+1$. If token i has rank k after the queuing step (i.e., height k on node u), it is shuffled with its d siblings. Since $u \in N(u)$, the token is assigned to any queue $Q^3_{u,v}$ for $v \in N(u)$ with equal probability $1/(d+1)$. Hence for any $v \in N(u)$,

$$\mathbf{Pr}[W_i(t+1) = v \mid W_i(t) = u] = \mathbf{Pr}\big[i \in Q^3_{u,v}\big] = \frac{1}{d+1}.$$

Applying the law of total probability over both cases gives

$$\begin{aligned}
\mathbf{Pr}[W_i(t+1) = v] &= \sum_{u \in N(v)} \mathbf{Pr}[W_i(t+1) = v \mid W_i(t) = u] \cdot \mathbf{Pr}[W_i(t) = u] \\
&= \sum_{u \in N(v)} \frac{1}{d+1} \cdot \mathbf{Pr}[W_i(t) = u] = \sum_{u \in N(v)} \mathbf{P}_{u,v} \cdot \mathbf{Pr}[W_i(t) = u].
\end{aligned}$$

Since $w_i(t_1) = w$, by the induction hypothesis it follows that

$$\mathbf{Pr}[W_i(t+1) = v] = \sum_{u \in N(v)} \mathbf{P}_{u,v} \cdot \mathbf{P}^{t-t_1}_{w,u} = \mathbf{P}^{t+1-t_1}_{w,v},$$

which completes the induction and proves the second statement.

The first part of the third statement follows directly from the second statement and the definition of mixing time (Definition 4.1). To prove the second part, observe that

$$\sum_{v \in V} \mathbf{Pr}[\, W_i(t_2) = v \mid W_i(t_1) = w \,] = 1,$$

since token i should be located on some node in any round. Then for a fixed node $v \in V$ we get,

$$\begin{aligned}
\mathbf{Pr}[\, W_i(t_2) = v \mid W_i(t_1) = w \,] &= 1 - \sum_{u \in V, u \neq v} \mathbf{Pr}[\, W_i(t_2) = u \mid W_i(t_1) = w \,] \\
&\ge 1 - (n-1) \cdot \left(\frac{1}{n} + \frac{1}{n^3}\right) \ge \frac{1}{n} - \frac{1}{n^2},
\end{aligned}$$

for which we use the upper bound from the first part of third statement. □

Negative Association. In the next lemma, we establish the negative association property between the locations of tokens in the height-sensitive process (lemma 4.10). This property states that, for any subset of tokens and any subset of nodes, the probability that all tokens end up in the specified nodes is at most the product of their individual probabilities.

Let $\mathcal{S}(t)$ denote the Height-Sensitive diffusion process and $\mathcal{W}(t)$ the random walk process with adjacency matrix $\mathbf{P}$, both restricted to t rounds. For a fixed round t and $t' \leq t$, we define an intermediate process $\mathcal{SW}(t', t)$ that performs t' rounds of process $\mathcal{S}$ and $t - t'$ rounds of process $\mathcal{W}$ afterward. By construction $\mathcal{SW}(0, t)$ is identical to $\mathcal{W}(t)$, while $\mathcal{SW}(t, t)$ coincides with $\mathcal{S}(t)$. Introducing this intermediate process allows us to interpolate between the two dynamics, which is useful for transferring probabilistic bounds from the well-understood random walk to the height-sensitive diffusion process.

We denote the location fo token j after t' rounds by $W_j(t')$, $\widehat{W}_j(t')$ and $\widetilde{W}_j(t')$ for the height-sensitive, intermediate, and random walk process, respectively, assuming all tokens start at the same initial positions: $W_j(0) = \widehat{W}_j(0) = \widetilde{W}_j(0)$.

Lemma 4.10 (Negative Association)**.** *Consider a time $t \geq 1$. Moreover, let $\mathcal{B} \subseteq \mathcal{T}$ and $D \subseteq V$ be an arbitrary subset of tokens and nodes, respectively. Then the following holds:*

$$\mathbf{Pr}\left[\bigcap_{i \in B} \{W_i(t) \in D\}\right] \leq \prod_{i \in B} \mathbf{Pr}\left[\widetilde{W}_i(t) \in D\right].$$

Proof. Note that $\prod_{i \in B} \mathbf{Pr}\left[\widetilde{W}_i(t) \in D\right] = \mathbf{Pr}\left[\bigcap_{i \in B}\{\widetilde{W}_i(t) \in D\}\right]$, since the tokens in random walk process ($\mathcal{W}$) move independent. To prove the statement we consider the intermediate process $\mathcal{SW}(t', t)$ for $t' \in [0, t]$. Recall that $\mathcal{SW}(t', t)$ performs the first t' rounds according to the height-sensitive process($\mathcal{SW}$) and the remaining $t - t'$ rounds as standard random walk ($\mathcal{W}$). For each round $t' \in [0, t]$ and each token $i \in \mathcal{B}$ we define a random variable

$$Z_i(t') := \sum_{u \in V} \mathbf{1}_{\{\widehat{W}_i(t') = u\}} \cdot \left(\mathbf{P}^{t-t'}\right)_{u,D} \qquad \text{and} \qquad Z(t') := \prod_{i \in \mathcal{B}} Z_i(t').$$

Here $Z_i(t')$ represents the probability that token i, starting at node $\widehat{W}_i(t')$, ends up in a node $w \in D$ after $t - t'$ rounds of a standard random walk.

Claim 4.11 (stated below) states that $Z(0) = \prod_{i \in \mathcal{B}} \mathbf{Pr}\left[\widetilde{W}_i(t) \in D\right]$, and $\mathbf{E}[Z(t)] = \mathbf{Pr}\left[\bigcap_{i \in B}\{W_i \in D\}\right]$. The lemma then follows immediately if $\mathbf{E}[Z(t)] \leq Z(0)$. To prove this, we show that for each round $t' \in \{1, 2, \cdots, t\}$ that

$$\mathbf{E}\left[\, Z(t') \;\middle|\; \widehat{W}(t'-1) = \widehat{w}(t'-1) \,\right] \leq Z(t'-1). \tag{4.13}$$

By the law of total probability, this implies $\mathbf{E}[Z(t')] \leq Z(t'-1)$. Iterating this inequality over all rounds completes the proof. Thus, it remain to show Equation (4.13). Since the location vector $W(t'-1)$ is fixed, we can drop the conditioning.

First we calculate $Z(t'-1)$ and then $\mathbf{E}[Z(t')]$ using the intermediate process $\mathcal{SW}(t', t)$.

Let T_u denote the set of tokens from $\mathcal{B}$ located on node u at round $t'-1$. We have

$$\begin{aligned} Z(t'-1) &= \prod_{i \in \mathcal{B}} \left(\sum_{u \in V} \mathbf{1}_{\{\widehat{W}_i(t'-1) = u\}} \cdot \left(\mathbf{P}^{t-t'+1}\right)_{u,D} \right) = \prod_{u \in V} \prod_{i \in T_u} \left(\mathbf{P}^{t-t'+1}\right)_{u,D} \\ &= \prod_{u \in V} \left(\left(\mathbf{P}^{t-t'+1}\right)_{u,D} \right)^{|T_u|} = \prod_{u \in V} \left(\sum_{v \in N(u)} \frac{\left(\mathbf{P}^{t-t'}\right)_{v,D}}{d+1} \right)^{|T_u|}. \end{aligned} \tag{4.14}$$

Next, for the expectation at round t', we have

$$\mathbf{E}[Z(t')] = \mathbf{E}\left[\prod_{i\in\mathcal{B}} Z_i(t')\right] = \mathbf{E}\left[\prod_{u\in V}\prod_{i\in T_u} Z_i(t')\right] = \prod_{u\in V}\mathbf{E}\left[\prod_{i\in T_u} Z_i(t')\right], \tag{4.15}$$

where the last equality holds because tokens on different nodes move independently within a fixed round.

We now focus on a single term corresponding to a fixed node u in Equation (4.15) and analyze the steps of round t' in $\mathcal{SW}(t',t)$ process, which coincide with round t' in $\mathcal{S}(t')$ process.

We define Q^i_{uv} for $i \in [5]$ as the state of queue $Q_{u,v}(t')$ after the i-th step of round t'. Let $h(u) := \lceil x_u(t'-1)/(d+1)\rceil$ denote the maximum height of any token on node u at the end of round $t'-1$. We partition the tokens in T_u into sibling sets $S_1, \cdots, S_{h(u)}$where

$$S_j(u) := \left\{ i \in T_u \;\middle|\; \text{token } i \text{ is in } Q^2_{u,k} \text{ with rank } j \text{ for some } k \in N(u)\right\}.$$

These are the tokens of rank j on outgoing queues after the queuing step. The destinations of these tokens are fixed after the shuffling step, in which siblings of the same rank (except non-excess tokens) are shuffled among outgoing queues independently across ranks. From this independence property follows that we can write the expectation over the tokens on node u as a product over the sibling sets:

$$\mathbf{E}\left[\prod_{i\in T_u} Z_i(t')\right] = \prod_{j=1}^{h(u)}\mathbf{E}\left[\prod_{i\in S_j(u)} Z_i(t')\right].$$

Expanding $Z_i(t')$ gives

$$\mathbf{E}\left[\prod_{i\in T_u} Z_i(t')\right] = \prod_{j=1}^{h(u)}\mathbf{E}\left[\prod_{i\in S_j(u)}\left(\sum_{v\in V}\mathbf{1}_{\widehat{W}_i(t')=v}\cdot\left(\mathbf{P}^{t-t'}\right)_{v,D}\right)\right].$$

Since tokens in $S_j(u)$ can only move to neighbours of u, this reduces to

$$\mathbf{E}\left[\prod_{i\in T_u} Z_i(t')\right] = \prod_{j=1}^{h(u)}\mathbf{E}\left[\prod_{i\in S_j(u)}\left(0+\sum_{v\in N(u)}\mathbf{1}_{\{\widehat{W}_i(t')=v\}}\cdot\left(\mathbf{P}^{t-t'}\right)_{v,D}\right)\right].$$

Finally, noting that token $i \in S_j(u)$ goes to node $v \in N(u)$ if and only if it is placed on queue $Q^3_{u,v}$, we have

$$\mathbf{E}\left[\prod_{i\in T_u} Z_i(t')\right] = \prod_{j=1}^{h(u)}\mathbf{E}\left[\prod_{i\in S_j(u)}\left(\sum_{v\in N(u)}\mathbf{1}_{\{i\in Q^3_{u,v}\}}\cdot\left(\mathbf{P}^{t-t'}\right)_{v,D}\right)\right]$$

This expresses the expected product in terms of the independent placement of sibling tokens into outgoing queues. Applying Claim 4.12 (stated below) to the previous expression gives

$$\mathbf{E}\left[\prod_{i\in T_u} Z_i(t')\right] \leq \prod_{j=1}^{h(u)}\prod_{i\in S_j(u)}\mathbf{E}\left[\sum_{v\in N(u)}\mathbf{1}_{\{i\in Q^3_{u,v}\}}\cdot\left(\mathbf{P}^{t-t'}\right)_{v,D}\right]$$

Since each single token in $S_j(u)$ independently chooses a queue, we can evaluate the expectation:

$$\prod_{j=1}^{h(u)} \prod_{i \in S_j(u)} \mathbf{E}\left[\sum_{v \in N(u)} \mathbf{1}_{\{i \in Q^3_{u,v}\}} \cdot \left(\mathbf{P}^{t-t'}\right)_{v,D}\right] = \prod_{j=1}^{h(u)} \prod_{i \in S_j(u)} \sum_{v \in N(u)} \frac{\left(\mathbf{P}^{t-t'}\right)_{v,D}}{d+1}$$

Combining over all sibling sets and tokens, we get,

$$\mathbf{E}\left[\prod_{i \in T_u} Z_i(t')\right] \leq \left(\sum_{v \in N(u)} \frac{\left(\mathbf{P}^{t-t'}\right)_{v,D}}{d+1}\right)^{|T_u|}.$$

Combining this with Equation (4.15) gives

$$\mathbf{E}[Z(t')] \leq \prod_{u \in V} \left(\sum_{v \in N(u)} \frac{\left(\mathbf{P}^{t-t'}\right)_{v,D}}{d+1}\right)^{|T_u|} \overset{\text{Eq. (4.14)}}{=} Z(t-1'),$$

completing the proof of Equation (4.13) and consequently finishing the proof. □

Here we state and prove the claims used in Lemma 4.10.

Claim 4.11. $Z(0) = \prod_{i \in \mathcal{B}} \mathbf{Pr}\left[\widetilde{W}_i(t) \in D\right]$, *and* $\mathbf{E}[Z(t)] = \mathbf{Pr}\left[\bigcap_{i \in B}\{W_i \in D\}\right]$.

Proof. By definition, for each token i, $Z_i(t') = \sum_{u \in V} \mathbf{1}_{\{\widehat{W}_i(t')=u\}} \cdot \left(\mathbf{P}^{t-t'}\right)_{u,D}$, and $Z(t') = \prod_{i \in \mathcal{B}} Z_i(t')$. At time $t' = 0$, the locations $\widehat{W}_i(0)$ is fixed and equal to the starting locations of the random walk process $\widetilde{W}_i(0)$. Thus,

$$Z(0) = \prod_{i \in \mathcal{B}} (\mathbf{P}^t)_{\widehat{W}_i(0),D} = \prod_{i \in \mathcal{B}} (\mathbf{P}^t)_{\widetilde{W}_i(0),D} = \prod_{i \in \mathcal{B}} \mathbf{Pr}\left[\widetilde{W}_i(t) \in D\right].$$

At time $t' = t$ we have $\mathbf{P}^{t-t'} = \mathbf{P}^0 = \mathbf{I}$, so

$$Z(t) = \prod_{i \in \mathcal{B}} \left(\sum_{u \in V} \mathbf{1}_{\{\widehat{W}_i(t)=u\}} \cdot (\mathbf{P}^0)_{u,D}\right) = \prod_{i \in \mathcal{B}} \left(\sum_{u \in D} \mathbf{1}_{\{\widehat{W}_i(t)=u\}}\right).$$

Hence, $Z(t)$ is 1 if all tokens in $\mathcal{B}$ are located on nodes in D at round t, then it can be regarded as a random variable that can take on the values zero or one. Therefore,

$$\mathbf{E}[Z(t)] = \mathbf{Pr}[Z(t) = 1] = \mathbf{Pr}\left[\bigcap_{i \in \mathcal{B}} \left\{\widehat{W}_i(t) \in D\right\}\right] = \mathbf{Pr}\left[\bigcap_{i \in \mathcal{B}} \{W_i(t) \in D\}\right],$$

because the intermediate process $\mathcal{SW}(t,t)$ coincides with the height-sensitive process $\mathcal{S}(t)$. This proves the claim. □

Claim 4.12. *For a fixed sibling set* $S_j(u)$ *in round* t' *it holds*

$$\mathbf{E}\left[\prod_{i \in S_j(u)} \left(\sum_{v \in N(u)} \mathbf{1}_{\{i \in Q^3_{u,v}\}} \cdot \left(\mathbf{P}^{t-t'}\right)_{v,D}\right)\right] \leq \prod_{i \in S_j(u)} \mathbf{E}\left[\left(\sum_{v \in N(u)} \mathbf{1}_{\{i \in Q^3_{u,v}\}} \cdot \left(\mathbf{P}^{t-t'}\right)_{v,D}\right)\right].$$

Proof. Recall that $S_j(u)$ is the set of sibling tokens with rank j on outgoing queues of node u. After the shuffling step, each token $i \in S_j(u)$ must be assigned to exactly one queue, so

$$\sum_{v \in N(u)} \mathbf{1}_{\{i \in Q^3_{u,v}\}} = 1.$$

If token i was an excess token at the start of the round, then its placement in queue $Q^3_{u,v}$ is the same as in $Q^1_{u,v}$. By the Zero-One Lemma in [41] (restated as Lemma 1.4) it follows that, the indicator variables $\mathbf{1}_{\{i \in Q^3_{u,v}\}}$ for $v \in N(u)$ are *negatively associated* (see the Definition 1.5).

Define for each token $i \in S_j(u)$ the function

$$f_i := \sum_{v \in N(u)} \mathbf{1}_{\{i \in Q^3_{u,v}\}} \cdot \left(\mathbf{P}^{t-t'}\right)_{v,D}.$$

The claim we want is

$$\mathbf{E}\left[\prod_i f_i\right] \leq \prod_i \mathbf{E}[f_i].$$

Since each f_i is a non-decreasing function of negatively associated random variables, we can apply the standard result of negatively associated functions, Lemma 6 in [42] (restated as Lemma 6.37), to conclude that $\mathbf{E}[\prod_i f_i] \leq \prod_i \mathbf{E}[f_i]$, finishing the proof. Note that this inequality is know as *Negative Covariance* (Proposition 3 in [42]). □

Occasionally, we need to add or remove the same number of tokens form each node. The next observation show adding or subtracting a multiple of $d+1$ tokens per node does not affect the distribution of excess tokens. It can be seen as the diffusion analogue of Observation 6.29.

Observation 4.13. *Let G be a d-regular graph with n nodes and diffusion matrix $\mathbf{P}$. Consider two executions of the vertex-based diffusion process on G with initial load vectors $x(0)$ and $\widehat{x}(0)$, using the same random choices for all $Z_{i,j}(t)$, nodes $i \in V$, neighbors $j \in N(u)$, and rounds $t \in \mathbb{N}$. If $x(0) = \widehat{x}(0) + \alpha \cdot \mathbf{1}$ for some $\alpha = k \cdot (d+1)$, $k \in \mathbb{Z}$, then $X(t) = \widehat{X}(t) + \alpha \cdot \mathbf{1}$ for all $t \geq 1$. It particular, $\mathrm{disc}(X(t)) = \mathrm{disc}(\widehat{X}(t))$.*

Proof. We prove this by induction. We show that $X_u(t) = \widehat{X}_u(t) + \alpha$ for $t \geq 0$.

- Base case. For $t = 0$, the claim holds since by assumption $x_u(0) = \widehat{x}_u(0) + \alpha$ for all $u \in V$.
- Induction step. Assume for some round t that $x_u(t) = \widehat{x}_u(t) + \alpha$ for all $u \in V$. By assumption, for all $i, j \in V$, $Z_{i,j}(t+1) = \widehat{Z}_{i,j}(t+1)$. Here we get,

$$\begin{aligned}
X_u(t+1) &= \sum_{v \in N(u)} \left\lfloor \frac{x_v(t)}{d+1} \right\rfloor + Z_{v,u}(t+1) \\
&= \sum_{v \in N(u)} \left\lfloor \frac{\widehat{x}_v(t) + k \cdot (d+1)}{d+1} \right\rfloor + \widehat{Z}_{v,u}(t+1) \\
&= \left(\sum_{v \in N(u)} \left\lfloor \frac{\widehat{x}_v(t)}{d+1} \right\rfloor + \widehat{Z}_{v,u}(t+1) \right) + k \cdot (d+1) \\
&= \widehat{X}_u(t+1) + \alpha.
\end{aligned}$$

Since it holds for each node $u \in V$ then we get $X(t+1) = \widehat{X}(t+1) + \alpha \cdot \mathbf{1}$ and it finishes the proof.

□

4.5.2 Ingredients used in Theorem 4.6

In this subsection, we prove Lemma 4.14 and Observation 4.15 which are used in dynamic setting (Theorem 4.6). Together, these lemmas capture the contribution of newly allocated load items to the discrepancy. Our approach follows the same ideas as in part two of this thesis.

Recall that Lemma 2.4 which bounds the $\mathrm{disc}(D(t))$ in terms of the global divergence of the matching sequence $\mathbf{m}^{[t]}$ requires that (a) the matrix chosen in round t and the newly added items are independent, and (b) all matrices in the sequence are double stochastic. Since the diffusion matrix also satisfies both conditions, then the lemma applies to the dynamic setting of the diffusion model as well.

Next lemma is the diffusion analogue of Lemma 2.4.

Lemma 4.14 (Load concentration via global divergence)**.** *Consider the diffusion matrix* $\mathbf{P}$*. Then for all* $\gamma > 0$ *and* $t \in \mathbb{N}$ *we have*

$$\mathbf{Pr}\left[\mathrm{disc}(D(t)) \geq \frac{8}{3} \cdot \gamma \log(n) + \sqrt{32\gamma \log(n) \cdot \frac{m}{n}} \cdot \Upsilon(\mathbf{P})\right] \leq 2 \cdot n^{-\gamma+1}.$$

Proof. Proof is essentially a repetition of Lemma 2.4, in a way that we replace $\mathbf{m}^{[\tau,t]}$ with $\mathbf{P}^{t-\tau+1}$, since we have the same matrix $\mathbf{P}$ in each step . For the sake of completeness, we provide it.

Fix a node $k \in [n]$. First we establish a concentration inequality on $D_k(t)$ in terms of $\Upsilon_k(\mathbf{P})$. Our goal is to decompose $D_k(t)$ into a sum of independent random variables. For the decomposition observe that $D(t) = \sum_{\tau=1}^{t} \mathbf{P}^{t-\tau-1} \cdot A(\tau)$, where $A(\tau)$ is the random load vector corresponding to the m load items allocated at time τ. So the kth coordinate of $D(t)$ is $D_k(t) = \sum_{\tau=1}^{t} \sum_{w\in[n]} \left(\mathbf{P}^{t-\tau-1}\right)_{k,w} \cdot A_w(\tau)$. We define the indicator random variable $B(\tau, j, w)$ for $\tau \in [t], j \in [m]$ and $w \in [n]$ as , starting from their token locations, $\Pr[B(\tau, j, w) = 1] = 1/n$ and $\mathbf{E}[B(\tau, j, w)] = 1/n$. Observe that $A_w(\tau)$, the load allocated to node w at step τ, can be expressed as $\sum_{j\in[m]} B(\tau, j, w)$. Merging this with the value of $D_k(t)$ gives

$$D_k(t) = \sum_{\tau=1}^{t} \sum_{w\in[n]} \left(\mathbf{P}^{t-\tau-1}\right)_{k,w} \cdot \left(\sum_{j\in[m]} B(\tau, j, w)\right) = \sum_{\tau=1}^{t} \sum_{j\in[m]} \underbrace{\left(\sum_{w\in[n]} \left(\left(\mathbf{P}^{t-\tau-1}\right)_{k,w} \cdot B(\tau, j, w)\right)\right)}_{=:C_k(\tau,j)}.$$

For a fixed $\tau \in [t]$ and $j \in [m]$ we define $C_k(\tau, j) := \sum_{w\in[n]} \left(\mathbf{P}^{t-\tau-1}\right)_{k,w} \cdot B(\tau, j, w)$. This random variable measures the contribution of j-th load item of round τ to $D_k(t)$. Note that the load items are allocated independently from each other. Since $\mathbf{P}^{t-\tau-1}$ are fixed matrices, then $C_k(\tau, j)$ and $C_k(\tau', j')$ are independent for all τ and τ' and $j \neq j'$. To apply the concentration inequality from Theorem 6.9 we need to show that $C_k(\tau, j) \leq 1$ and compute an upper bound on $\mathrm{Var}[C_k(\tau, j)]$. Showing the first condition is easy since exactly one of the indicator random variables $B(\tau, j, w)$ is one and $\left(\mathbf{P}^{t-\tau-1}\right)_{k,w}$ has a value between zero and one.

It remains to consider the variance of $C_k(\tau, j)$. First note that by linearity of expectation

$$\begin{aligned}\mathbf{E}[C_k(\tau, j)] = \mathbf{E}\left[\sum_{w\in[n]} \left(\left(\mathbf{P}^{t-\tau-1}\right)_{k,w} \cdot B(\tau, j, w)\right)\right] &= \sum_{w\in[n]} \left(\mathbf{P}^{t-\tau-1}\right)_{k,w} \cdot \mathbf{E}[B(\tau, j, w)] \\ &= \sum_{w\in[n]} \left(\mathbf{P}^{t-\tau-1}\right)_{k,w} \cdot \frac{1}{n} = \frac{1}{n},\end{aligned}$$

where the last equality follows form the fact that $\mathbf{P}^{t-\tau-1}$ is doubly stochastic. Now we get

$$\begin{aligned}\mathrm{Var}[C_k(\tau,j)] &= \mathbf{E}\Big[(C_k(\tau,j)-\mathbf{E}[C_k(\tau,j)])^2\Big]\\ &= \mathbf{E}\left[\Big(\Big(\sum_{w\in[n]}(\mathbf{P}^{t-\tau-1})_{k,w}\cdot B(\tau,j,w)\Big)-\frac{1}{n}\Big)^2\right]\\ &= \sum_{w'\in[n]}\frac{1}{n}\cdot\left((\mathbf{P}^{t-\tau-1})_{k,w'}-\frac{1}{n}\right)^2 = \frac{1}{n}\cdot\left\|(\mathbf{P}^{t-\tau-1})_{k,\cdot}-\frac{\vec{1}}{n}\right\|_2^2,\end{aligned}$$

where we used that for each τ and each j exactly one of the $B(\tau,j,w)$ is one and all others are zero, and each of the n possible cases has uniform probability. Recall that $C_k(\tau,j)$ and $C_k(\tau',j')$ are independent for all τ,τ' and $j\neq j'$. Hence we get

$$\begin{aligned}\mathrm{Var}\left[\sum_{\tau=1}^{t}\sum_{j\in[m]}C_k(\tau,j)\right] &= \sum_{\tau=1}^{t}\sum_{j\in[m]}\mathrm{Var}[C_k(\tau,j)] = \frac{1}{n}\cdot\sum_{\tau=1}^{t}\sum_{j\in[m]}\left\|(\mathbf{P}^{t-\tau-1})_{k,\cdot}-\frac{\vec{1}}{n}\right\|_2^2\\ &\leq \frac{m}{n}\cdot\sum_{s=0}^{\infty}\left\|(\mathbf{P}^{s})_{k,\cdot}-\frac{\vec{1}}{n}\right\|_2^2 = \frac{m}{n}\cdot(\Upsilon_k(\mathbf{P}))^2,\end{aligned}$$

where the final equality uses the definition of the global divergence $\Upsilon_k(\mathbf{P})$. Applying Theorem 6.9 with $M=1$ and $X=D_k(t)=\sum_{\tau=1}^{t}\sum_{j\in[m]}C_k(\tau,j)$ with $\lambda=2\gamma\log(n)/3+\Upsilon_k(\mathbf{P})\cdot\sqrt{2\gamma m/n}$ results in

$$\Pr\left[D_k(t)-t\cdot\frac{m}{n}\geq\frac{2}{3}\cdot\gamma\log(n)+\sqrt{2\gamma\log(n)\cdot\frac{m}{n}}\cdot\Upsilon_k(\mathbf{P})\right]\leq n^{-\gamma}.$$

The lower bound can be established using Theorem 6.10 (with $a_i=0$ and $M=1$) instead of Theorem 6.9. Via a union bound we get

$$\mathbf{Pr}\left[\left|D_k(t)-t\cdot\frac{m}{n}\right|\geq\frac{4}{3}\cdot\gamma\log(n)+\sqrt{8\gamma\log(n)\cdot\frac{m}{n}}\cdot\Upsilon_k(\mathbf{P})\right]\leq 2\cdot n^{-\gamma}.$$

Applying the union bound over all nodes $k\in[n]$ together with Observation 2.14 (stating that $\mathrm{disc}(D(t))\leq 2|D_k(t)-t\cdot m/n|$) finishes the proof.

□

Next simple observation bounds the global divergence in diffusion model.

Observation 4.15 (Spectral gap bound on global divergence)**.** *Consider the diffusion matrix* $\mathbf{P}$*. Then,*

$$\Upsilon_k(\mathbf{P})\leq\sqrt{\frac{1}{1-\lambda(\mathbf{P})}}$$

Proof. Form the analysis of continues diffusion process or reversible Markov-Chain (e.g., [82], [64]), it is known that $\|(\mathbf{P}^t)_{k,\cdot}-\frac{\vec{1}}{n}\|_2\leq\lambda(\mathbf{P})^2\cdot\|(\mathbf{P}^{t-1})_{k,\cdot}-\frac{\vec{1}}{n}\|_2\leq(\lambda(\mathbf{P}))^{2t}$. So, it follows that

$$(\Upsilon_k(\mathbf{P}))^2=\sum_{t=0}^{\infty}\left\|(\mathbf{P}^{\tau})_{k,\cdot}-\frac{\vec{1}}{n}\right\|_2^2\leq\sum_{t=0}^{\infty}(\lambda(\mathbf{P}))^{2t}=\frac{1}{1-(\lambda(\mathbf{P}))^2}\leq\frac{1}{1-\lambda(\mathbf{P})},$$

as claimed. □

4.6 Bounds for Specific Rounds

In this section, we prove two corollaries that give bounds on discrepancy at specific rounds after smoothing time $\tau_{\mathrm{S}}(K,1)$, in both static and dynamic settings.

Corollary 4.16. *Consider the discrete static vertex-based diffusion on d-regular graph G with n nodes, diffusion matrix* $\mathbf{P}$ *and an initial load vector with* $\mathrm{disc}(x(0)) := K > 1$*. Let* $\varepsilon \in [0,1]$ *be a constant.*

1. *For* $d = \Omega\big(\log^{1+\varepsilon} n\big)$*, there is a round* $t^\star = \mathrm{O}(\ln(Kn)/(1-\lambda))$ *such that*

$$\mathbf{Pr}\Big[\mathrm{disc}(X(t_0) \leq 10\sqrt{d\log(n)}\Big] \geq 1 - n^{-2}.$$

2. *For any d, there is a round* $t_1 = \mathrm{O}(\ln(Kn)/(1-\lambda))$ *such that*

$$\mathbf{Pr}\Big[\mathrm{disc}(X(t_1)) \leq 49\sqrt{d}\log(n)\Big] \geq 1 - 4n^{-3}.$$

3. *For any d, there is a round* $t_2 = \mathrm{O}((\ln(Kn) + \log\log(d)\cdot\log(n))/(1-\lambda))$ *such that*

$$\mathbf{Pr}\Big[\mathrm{disc}(X(t_2)) \leq 48\sqrt{d\log(n)} + 96\log(n)\Big] \geq 1 - n^{-2}.$$

Proof. From Theorem 4.3, it follows that at round $t^\star = \tau_{\mathrm{S}}(\mathbf{P},K,1) + i\cdot t_{mix}(\mathbf{P})$ for $i := \log_2(\frac{2(1+\varepsilon)}{\varepsilon})$, with probability at least $1 - 2(i+1)\cdot n^{-3}$,we have

$$\begin{aligned}
\mathrm{disc}(X(t^\star)) &\leq 48\cdot d^{\,1/2^i}\cdot\log(n) + 48\cdot\sqrt{d\cdot\log(n)} \\
&\overset{(a)}{\leq} 48\cdot c\cdot d^{\,1/2^i+1/(1+\varepsilon)} + 48\cdot\sqrt{c\cdot d^{1+1/(1+\varepsilon)}} \\
&= 48\cdot c\cdot d^{\,\varepsilon/2\cdot(1+\varepsilon)+1/(1+\varepsilon)} + 48\cdot\sqrt{c\cdot d^{(2+\varepsilon)/(1+\varepsilon)}} \\
&\leq 48\cdot c\cdot d^{(2+\varepsilon)/(2+2\varepsilon)} + 48\cdot\sqrt{c\cdot d^{(2+\varepsilon)/(1+\varepsilon)}} \\
&\leq d/2 + d/2 = d
\end{aligned}$$

where (a) follows from the assumption that $\log(n) \leq c\cdot d^{1/(1+\varepsilon)}$ for some constant $c > 0$. Since $i = \log_2(\frac{2(1+\varepsilon)}{\varepsilon})$ is a constant, then we have

$$\begin{aligned}
t^\star &= \tau_{\mathrm{S}}(\mathbf{P},K,1/(2n)) + \log_2\left(\frac{2(1+\varepsilon)}{\varepsilon}\right)\cdot t_{mix}(\mathbf{P}) \\
&\leq \tau_{\mathrm{S}}(\mathbf{P},K,1/(2n))\cdot\left(1 + \log_2\left(\frac{2(1+\varepsilon)}{\varepsilon}\right)\right) \\
&= \mathrm{O}(\tau_{\mathrm{S}}(\mathbf{P},K,1/(2n))).
\end{aligned}$$

Now we define an event $\Gamma^\star := \{\mathrm{disc}(X(t^*)) \leq d\}$, and from above it follows that

$$\mathbf{Pr}\big[\overline{\Gamma^\star}\big] \leq 2\cdot\left(\log_2(\frac{2(1+\varepsilon)}{\varepsilon}) + 1\right)\cdot n^{-3}.$$

Here we assume this event happens and we condition on this. Recall that $\lfloor X(t^\star)\rfloor = (d+1)\cdot$

$\lfloor \min_{u\in V} X_u(t^\star)/(d+1)\rfloor$ and the number of vital tokens at round $t^\star$ is

$$\begin{aligned}\Psi(t^\star) &= \sum_{u\in V}\Big(X_u(t^\star) - \min_{v\in V} X_v(t^\star) + \min_{v\in V} X_v(t^\star) - \lfloor X(t^\star)\rfloor\Big)\\ &\le n\cdot \mathrm{disc}(X(t^\star)) + n\cdot d \le 2n\cdot d.\end{aligned}$$

Let $t_0 := t^\star + t_{mix}(\mathbf{P})$. From Lemma 4.5, it then follows that,

$$\mathbf{Pr}\left[\mathrm{disc}(X(t_0)) \le \frac{4d}{n} + \sqrt{48\cdot 2d\cdot \log(n)} \;\middle|\; \Gamma^\star\right] \ge 1 - 2n^{-3}.$$

From the law of total probability then we get,

$$\begin{aligned}\mathbf{Pr}\left[\mathrm{disc}(X(t_0)) \le 10\sqrt{d\log(n)}\right] &\ge 1 - 2n^{-3} - \mathbf{Pr}\left[\overline{\Gamma^\star}\right]\\ &\ge 1 - 2n^{-3} - 2\cdot\left(\log_2\left(\frac{2(1+\varepsilon)}{\varepsilon}\right)+1\right)\cdot n^{-3} \ge 1 - n^{-2},\end{aligned}$$

where the last inequality holds since $\log_2(\frac{2(1+\varepsilon)}{\varepsilon})$ is a constant and it finishes the proof of the first statement. Recall that $\tau_{\mathrm{S}}(\mathbf{P}, K, 1) \le 4\ln(Kn)/(1-\lambda)$ (Observation 4.2).

Let $t_1 := \tau_{\mathrm{S}}(\mathbf{P}, K, 1) + t_{mix}(\mathbf{P})$. The second statement follows directly from Theorem 4.3 (with $i = 1$) together with Observation 4.2.

Now let $t_2 := \tau_{\mathrm{S}}(\mathbf{P}, K, 1) + \log\log(d)\cdot t_{mix}(\mathbf{P})$. Since $d \le n$, from Theorem 4.3 with $i = \log\log(d)$ we get that

$$\mathbf{Pr}\left[\mathrm{disc}(X(t_2)) \le 48d^{(1/2)^i}\log(n) + 48\sqrt{d\log(n)}\right] \le 2(\log\log(d)+1)n^{-3} \le n^{-2},$$

and note that

$$\begin{aligned}\mathrm{disc}(X(t_2)) &\le 48d^{(1/2)^{\log\log(d)}}\log(n) + 48\sqrt{d\log(n)}\\ &\le 48d^{(1/\log d)}\log(n) + 48\sqrt{d\log(n)}\\ &\overset{(c)}{\le} 96\log(n) + 48\sqrt{d\log(n)},\end{aligned}$$

where (c) holds since $d^{1/\log d} \le 2$. This together with Observation 4.2 finishes the proof. □

The next corollary provides the dynamic counterpart to Corollary 4.17.

Corollary 4.17. *Let G be a d-regular graph with diffusion matrix $\mathbf{P}$. Consider the dynamic discrete vertex-based diffusion on graph G (with m new items added in each step) and an initial load vector with $\mathrm{disc}(x(0)) := K > 1$. In case $d = \omega(\log^{1+\varepsilon}(n)$ for some constant $\varepsilon > 0$, there is a round $t = \mathrm{O}(\log(Kn)/(1-\lambda))$ such that ,w.h.p.,*

$$\mathrm{disc}(X(t)) = \mathrm{O}\left(\sqrt{d\log(n)} + \sqrt{\frac{\log(n)}{1-\lambda}}\sqrt{\frac{m}{n}}\right).$$

For any d, there is a round $t = \mathrm{O}((\log(Kn) + \log\log(d)\log(n))/(1-\lambda)$ such that ,w.h.p.,

$$\mathrm{disc}(X(t)) = \mathrm{O}\left(\sqrt{d\log(n)} + \log(n) + \sqrt{\frac{\log(n)}{1-\lambda}}\sqrt{\frac{m}{n}}\right).$$

Proof. The proof follows directly from Corollary 4.16 and Theorem 4.6, using Observation 4.15 to bound the global divergence in terms of the spectral gap of the diffusion matrix. □

Simulation. We also compared the static discrete diffusion with other protocols through simulations. Although we tested several protocols, here we only present two variants of *Threshold load balancing*. In this protocol, for a matched pair of nodes $[u : v] \in \mathbf{M}^{(t)}$ with $X_u(t) > X_v(t)$, if $X_u(t) > \overline{x}$ then u send $X_u(t) - \overline{x}$ tokens to v, each independently with probability $1/2$. In the second variant, the same rule applies, but v accepts tokens only if its load below the average.

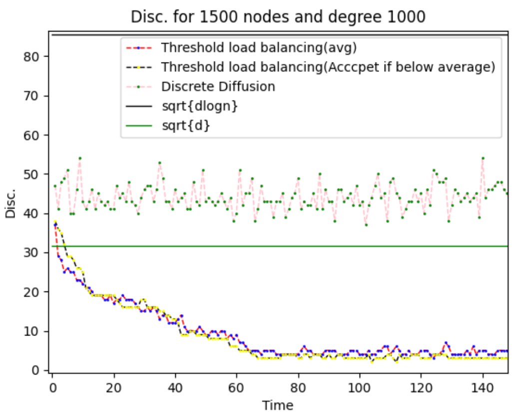

Figure 11: Discrete diffusion after continues balancing time for a regular graph. Time is normalized by the number of nodes.

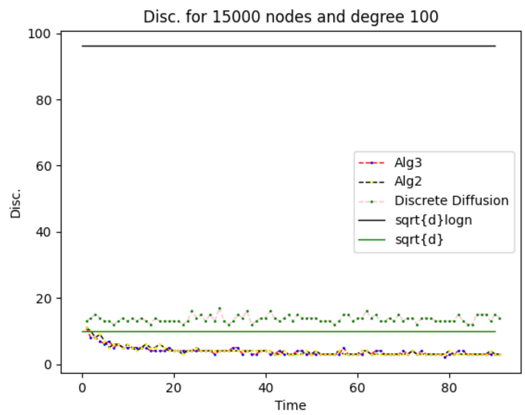

Figure 12: Discrete diffusion after continues balancing time for a regular graph.Time is normalized by the number of nodes.

As the plots show, the discrepancy in static discrete diffusion remains between $\sqrt{\log(n)}$ and $\sqrt{d}\log(n)$.

4.7 Summary and Open Problems

In this work, we show that the discrepancy in the vertex-based discrete diffusion converges to $\mathrm{O}(\sqrt{d\log(n)}+\log(n))$ and extend this result to the dynamic setting. Previously, authors of [85] showed a bound of $\mathrm{O}(d^2\sqrt{\log(n)})$ on the discrepancy for a d-regular graph after continues balancing time $\tau_{\mathrm{S}}(K, 1/(2n))$. Our results improve on this for $d = \Omega(\sqrt[3]{\log(n)})$ and for $d = \Omega(\sqrt[4]{\log(n)})$ when spending $\log\log(d) \cdot t_{mix}$ more rounds (Corollary 4.16). Moreover, we provide the first bound on the discrepancy in the dynamic setting of the discrete vertex-based diffusion.

We conjecture that our bound is tight up to an factor of $\mathrm{O}(\sqrt{\log(n)})$. The intuition is that $\sqrt{d}$ represent a fundamental lower bound, which could potentially be formalized using a variant of the Central Limit theorem, such as the Berry-Essen theorem, for specially constructed graphs and initial configuration similar to our proposed lower bound on discrepancy in part two.

One open question is how tight our dynamic load balancing bounds are for the diffusion model. Improving the bound on the global divergence would be a natural first step, as the other components of our analysis are nearly tight. For constant d-regular graphs, we believe that a constant discrepancy can be shown by adapting the analysis from part four. This raises the question of what the tightest bound is in the regime $d = \mathrm{O}(\sqrt[4]{\log(n)})$ and $d = \Omega(1)$. Finally, an important open problem is whether a tight bound can be established for non-regular graphs. We conjecture this may be possible by designing a height-sensitive protocol that mimics the standard random walk.

Part Five:

5 Token Distribution on Matchings

5.1 Introduction

In this part, we investigate the *token distribution process* on arbitrary connected graphs. Consider a connected graph $G = (V, E)$ with diameter D, where each node initially holds an integer number of load items, or *tokens*. In each round, a matching is selected, and every pair of matched nodes compares their loads. If the load difference between two matched nodes is at least one, the node with the higher load transfers exactly one token to its partner. This simple, local rule gradually balances the load across the network through repeated pairwise exchanges.

Results in a Nutshell. Our main result establishes a *diameter bound* on the final discrepancy and, for the first time, provides an explicit bound on the number of rounds required to achieve it. Specifically, we show that after $\mathrm{O}(K \cdot \widetilde{\tau}_{\mathrm{S}}(K) + (|E| + D^2) \cdot |E| \log(n))$ rounds, the protocol reaches a configuration whose discrepancy is smaller than the graph's diameter D. Here, K denotes the initial discrepancy, and $\widetilde{\tau}_{\mathrm{S}}(K) = \log(Kn)/(1 - \lambda)$ is a spectral parameter reflecting the approximate time it takes for the continuous load balancing to reduce the initial discrepancy K to 1. The proof relies on a simple yet powerful quadratic potential function that captures the evolution of the total load imbalance. Furthermore, we establish a matching lower bound ($\Omega(K/(1 - \lambda))$) on the number of rounds required by the process on regular graphs, matching with the bound established in [49].

Techniques and Comparison to Prior Work. To place our results in context, we compare them with the classical bound presented in [49], which shows that for d-regular graphs, after $\mathrm{O}(dK/\alpha + d^2 \log(n)/\alpha^2)$ rounds (where α denotes the edge expansion of the graph) the discrepancy decreases to $\mathrm{O}(d \log(n)/\alpha)$. Our result improves upon this bound for d-regular graphs satisfying $d \log(n)/\alpha = \Omega(D)$; however, our analysis may exhibit a longer convergence time. Moreover, our bound holds for arbitrary connected graphs.

For instance, in the r-dimensional grid and torus (encompassing the cycle and the hypercube) as well as in complete graphs, our analysis achieves a smaller discrepancy. A concise summary and comparison of the results is presented in Table 5, with a detailed discussion provided in Section 5.5.

Table 5: Overview of related results for the discrepancy in the token distribution process. Note that $r \in [1, \log(n)]$ and $H(r, q)$ is the Hamming (q-ary hypercube) graph. The statements hold w.h.p.

Graph	**Disc. from [49]**	**Disc. from Thm. 5.2**
Cycle	$\mathrm{O}(n \log(n))$	$\theta(n)$
2-dim torus	$\mathrm{O}(\sqrt{n} \cdot \log(n))$	$\theta(\sqrt{n})$
3-dim torus	$\mathrm{O}(\sqrt[3]{n} \cdot \log(n))$	$\theta(\sqrt[3]{n})$
r-dim torus	$\mathrm{O}(rn^{1/r} \cdot \log(n))$	$\theta(rn^{1/r})$
Hypercube	$\mathrm{O}(\log^2(n))$	$\theta(\log(n))$
$H(r, q)$	$\mathrm{O}(r^2 \log(q))$	$\theta(r)$
Complete	$\mathrm{O}(\log(n))$	1
Non-regular	-	D

Overall, our approach offers a unified and elementary analysis based solely on a quadratic potential argument, while simultaneously improving the known bounds on discrepancy across a variety of graph families.

Outline. The remainder of this part is organized as follows. Section 5.2 introduces the notation and model. Section 5.3 presents the main results (upper and lower bound on the discrepancy) for the token distribution problem on arbitrary graphs and the intermediate results used in their proofs. Section 5.4 summarizes the key technical tools. Section 5.5 provides discrepancy bounds for specific graph classes and compares them with related work. Finally, Section 5.6 concludes and outlines open problems.

5.2 Model and Definitions

We begin by introducing the notations, then formally define the process and present an useful observation.

We are given an undirected graph $G(V, E)$ and an arbitrary initial assignment of discrete load to nodes. Time progresses in discrete rounds. As before, the load vector in round $t \in \mathbb{N}$ is shown by $X(t)$, where $X_i(t) \in \mathbb{N}_0$ is the load of node $i \in V$ in round t. The initial load vector is denoted $X(0)$. For a load vector $X(t)$, we let $\mathrm{disc}(X(t))$ denote the discrepancy, defined as $\mathrm{disc}(X(t)) = \max_i X_i(t) - \min_i X_i(t)$. In each round t, a matching $\mathbf{M}^{(t)}$ is given, and the loads are updated according to the following rule: for each edge $(i, j) \in \mathbf{M}^{(t)}$ with $X_i(t) > X_j(t)$, exactly one load item is moved from node i to node j in round t. Formally, the protocol is specified as follows.

For each round $t = 1, 2, 3, \ldots$ and for each edge $(i, j) \in \mathbf{M}^{(t)}$ do:

- $X_i(t) = X_i(t-1) + \mathrm{sgn}(X_j(t-1) - X_i(t-1))$
- $X_j(t) = X_j(t-1) + \mathrm{sgn}(X_i(t-1) - X_j(t-1))$

We assume that G is connected and let d_u denote the degree of node u, $\Delta := \max_u\{d_u\}$ the maximum degree of G, and $D := D(G)$ the diameter of G. Here we recall the diffusion matrix of the graph G. For $\{i, j\} \in E$, we have $\mathbf{P}_{i,j} := 1/2\Delta$, $\mathbf{P}_{i,j} := 1 - d_i/2\Delta$ if $i = j$, and $\mathbf{P}_{i,j} = 0$ otherwise. Recall that $\lambda_1(\mathbf{P}) \geq \ldots \geq \lambda_n(\mathbf{P})$ are the eigenvalues of $\mathbf{P}$ and in this section we let $\lambda = \max\{|\lambda(\mathbf{P})|, |\lambda_n(\mathbf{P})|\}$.

We recall the generating protocol for the matchings from [52].

Observation 5.1 (Lemma 1 in [52])**.** *Consider the following protocol for generating matchings on a graph G with maximum degree Δ. First, Each edge (i, j) is independently put in $\mathbf{M}^{(t)}$ with probability $1/4d_{ij}$ in which $d_{ij} = \max\{d_i, d_j\}$ and d_i is the degree of node i. Then, Each edge (i,j) removes itself from $\mathbf{M}^{(t)}$ if (w, i) or (w, j) is in $\mathbf{M}^{(t)}$ for some $w \in V$. For each round t and edge $(i, j) \in E$ it holds that,*

$$\frac{1}{8\Delta} \leq \mathbf{Pr}\Big[(i, j) \in \mathbf{M}^{(t)}\Big].$$

To analyze the protocol we use the known quadratic potential function. Let $\overline{x}$ be the average load. For each node $u \in V$ and round t, we define

$$\Phi_u(t) := (X_u(t) - \overline{x})^2 \qquad \text{and} \qquad \Phi(t) := \sum_{u \in V} \Phi_u(t).$$

5.3 Token Distribution

In this section, we prove a general upper bound (Theorem 5.2) and a lower bound (Theorem 5.6) on the discrepancy. First we prove the upper bound, then we provide the intermediate results used in its proof and in the end we prove our lower bound.

In Subsection 5.3.1, we show that any initial load vector reaches a load vector with discrepancy at most $2|E|$ after sufficiently many rounds. In Subsection 5.3.2, we demonstrate that, shortly thereafter, the discrepancy decreases to $2D$, where D is the diameter of the graph. Finally, Subsection 5.3.3 reduces it further to D. Together, these results lead to the following theorem.

Theorem 5.2. *Let G be an undirected, connected graph with maximum degree Δ, diffusion matrix $\mathbf{P}$ and diameter D, and let $\mathbf{M}^{[\infty]} := \left(\mathbf{M}^{(s)}\right)_{s=1}^{\infty}$ be a sequence of matchings on G with $p_{\min}$ as the minimum probability of an edge appearing in any matching. Consider token distribution protocol initialized with an arbitrary initial load vector $X(0)$ with $\mathrm{disc}(X(0)) := K > 1$. For round*

$$t^* := \left\lceil \frac{K \cdot (1 + \log(nK^2))}{2\Delta p_{\min} \cdot (1-\lambda)} \right\rceil + \left\lceil \frac{96|E| \cdot n \cdot \log(n)}{p_{\min}} \right\rceil + \left\lceil \frac{24D^2 \cdot n \cdot \log(n)}{p_{\min}} \right\rceil,$$

it holds

$$\mathbf{Pr}[\mathrm{disc}(X(t^*)) \le D] \ge 1 - 6 \cdot n^{-2}.$$

Proof. Let $t_1 := \left\lceil \frac{K \cdot (1+\log(nK^2))}{2\Delta p_{\min} \cdot (1-\lambda)} \right\rceil$ and define an event $\mathcal{E}_1 := \{\mathrm{disc}(X(t_1)) \le 2|E|\}$. From Proposition 5.3 (presented in Subsection 5.3.1) it follows that $\mathbf{Pr}\left[\overline{\mathcal{E}_1}\right] \le 6 \cdot n^{-2}$.

Let $t_2 := t_1 + \left\lceil \frac{96|E| \cdot n \cdot \log(n)}{p_{\min}} \right\rceil$ and define an event $\mathcal{E}_2 := \{\mathrm{disc}(X(t_2) \le 2D\}$. From Proposition 5.4 (presented in Subsection 5.3.2) it follows that $\mathbf{Pr}[\,\mathcal{E}_2 \mid \mathcal{E}_1\,] \ge 1 - n^{-2}$ and by the law of total probability we get, $\mathbf{Pr}[\mathcal{E}_2] \ge 1 - 5 \cdot n^{-2}$.

Let $t_3 := t_2 + \left\lceil \frac{24D^2 \cdot n \cdot \log(n)}{p_{\min}} \right\rceil$ and define an event $\mathcal{E}_3 := \{\mathrm{disc}(X(t_3)) \le D\}$. From Proposition 5.5 (presented in Subsection 5.3.3) it follows that $\mathbf{Pr}[\,\mathcal{E}_3 \mid \mathcal{E}_2\,] \ge 1 - n^{-3}$ and via the law of total probability we get,

$$\mathbf{Pr}[\mathcal{E}_3] \ge 1 - \mathbf{Pr}\left[\overline{\mathcal{E}_2}\right] - n^{-3} \ge 1 - 6n^{-2}.$$

□

In the rest of this subsection we prove the three propositions used in the proof of the Theorem 5.2, and afterward we provide the proof of the lower bound.

5.3.1 Arbitrary Discrepancy to $2|E|$

In the next proposition, we show that any initial load vector reaches a discrepancy of at most $2|E|$ after sufficiently many rounds. To prove this, we argue that as long as the load differences across all edges remain at least $2|E|$, the potential function decreases sufficiently so that, after t rounds (specified in the proposition), the discrepancy falls below $2|E|$ with high probability.

Proposition 5.3. *Consider the token distribution on the sequence $\mathbf{M}^{[\infty]}$ (with $p_{\min}$ as the minimum probability of an edge appearing in any matching) and an arbitrary initial load vector $X(0)$ such that $\mathrm{disc}(X(0)) \le K$. Then, for the round*

$$t := \left\lceil \frac{K \cdot (1 + \log(nK^2))}{2\Delta p_{\min} \cdot (1-\lambda)} \right\rceil,$$

it holds that

$$\mathbf{Pr}[\mathrm{disc}(X(t)) \le 2 \cdot |E|] \ge 1 - 4 \cdot n^{-2}.$$

Proof. We aim to show that there exists a round $t^* \in \mathbb{N}$, $t^* \le t$, such that $\mathrm{disc}(X(t^*)) \le 2|E|$. Since the discrepancy is non-increasing, this suffices to prove the proposition.

Assume, for contradiction, that for all $t^* \le t$ we have

$$\sum_{(i,j)\in E} X_i(t^*) - X_j(t^*) > 2|E|.$$

We apply Lemma 5.7 repeatedly. Define $\beta := 2\Delta p_{\min}/K \cdot (1-\lambda)$. Then

$$\mathbf{E}[\Phi(t)] = \mathbf{E}[\mathbf{E}[\,\Phi(t) \mid X(t-1)\,]] \le \mathbf{E}[\Phi(t-1)] \cdot (1-\beta).$$

Since $\Phi(0) \le nK^2$, the chain rule of conditional expectation gives, for the specified t,

$$\begin{aligned}\mathbf{E}[\Phi(t)] &\le (1-\beta)^t \cdot \Phi(0) \le (1-\beta)^t \cdot nK^2 \le e^{-\beta\cdot t + \log(nK^2)}\\ &= e^{-\beta \cdot \frac{1+\log(nK^2)}{\beta} + \log(nK^2)} = e^{-1}.\end{aligned} \tag{5.1}$$

Define the event $\mathcal{E}(t) := \{\Phi(t) \le |E|^2\}$. From Equation (5.1) and Markov's inequality, we have

$$\mathbf{Pr}\Big[\overline{\mathcal{E}(t)}\Big] = \mathbf{Pr}\big[\Phi(t) > |E|^2\big] \le \frac{\mathbf{E}[\Phi(t)]}{|E|^2} \le \frac{1}{e \cdot |E|^2} \le \frac{4}{n^2}, \tag{5.2}$$

where the last inequality uses $|E| \ge n/2$ since G is connected. By definition, $\operatorname{disc}(X(t)) \le 2 \cdot \sqrt{\Phi(t)}$. Hence, if $\mathcal{E}(t)$ occurs, then

$$\mathbf{Pr}[\operatorname{disc}(X(t)) \le 2 \cdot |E| \mid \mathcal{E}(t)] = 1.$$

Applying the law of total probability gives

$$\mathbf{Pr}[\operatorname{disc}(X(t)) \le 2\cdot|E|] \ge 1 - \mathbf{Pr}\Big[\overline{\mathcal{E}(t)}\Big] \overset{\text{Eq. (5.2)}}{\ge} 1 - 4n^{-2}.$$

This contradicts the assumption that $\sum_{(i,j)\in E} X_i(t^*) - X_j(t^*) > 2|E|$ for all $t^* \le t$. Therefore, the proposition holds. □

5.3.2 Discrepancy from $2|E|$ to $2D$

Recall that D is the diameter of the underlying graph. We define *phases* as consecutive rounds in which, for phase $i \in \mathbb{N}$, the discrepancy lies in the interval $(2|E|/2^{i-1}, 2|E|/2^i]$. The idea is that as long as the discrepancy is sufficiently large ($> 2D$), the quadratic potential function experiences a significant decrease. In particular, for a round in phase r, the expected decrease is at least $p_{\min}|E|/2^r$ (Lemma 5.10). This allows us to bound the length of phase r with high probability (Lemma 5.11), which in turn shows that the discrepancy reaches $2D$ after $\mathrm{O}(|E|n\log(n)/p_{\min})$ rounds, as stated in the next proposition.

Proposition 5.4. *Consider the token distribution on the sequence $\mathbf{M}^{[\infty]}$ (with $p_{\min}$ as the minimum probability of an edge appearing in any matching) and an arbitrary initial load vector $X(0)$ such that $2D < \operatorname{disc}(X(0)) \le 2|E|$. Then, for the round*

$$t_1 := \left\lceil \frac{96|E| \cdot n \cdot \log(n)}{p_{\min}} \right\rceil,$$

it holds that

$$\mathbf{Pr}[\operatorname{disc}(X(t_1)) \le 2D] \ge 1 - n^{-2}.$$

Proof. Let $\ell := \log(|E|/D)$ and that $f(0) := 0$. For each integer $r \in [1, \ell]$ we define function $f(r) := 48|E| \cdot n \cdot \log(n)/p_{\min} \cdot 2^r$ and an event

$$\Gamma_r := \left\{ \mathrm{disc}(X(f(r))) \le \frac{\mathrm{disc}(f(X(r-1)))}{2} \right\} \qquad \text{and} \qquad \Gamma := \bigcap_{r \in [\ell]} \Gamma_r.$$

From Lemma 5.11 it follows that $\mathbf{Pr}\left[\overline{\Gamma_r}\right] \le n^{-3}$. Hence by union bound we get,

$$\mathbf{Pr}\left[\overline{\Gamma}\right] = \mathbf{Pr}\left[\bigcup_{r \in [\ell]} \overline{\Gamma_r}\right] \le \ell \cdot n^{-3} \le n^{-2}. \tag{5.3}$$

Recall that the discrepancy is non-increasing over time. Hence, assuming the event Γ holds we have,

$$\mathrm{disc}(X(f(\ell))) \le \mathrm{disc}(X(f(0))) \cdot \left(\frac{1}{2}\right)^{\log(|E|/D)} \le 2|E| \cdot \left(\frac{1}{2}\right)^{\log(|E|/D)} = 2 \cdot D.$$

Therefore, we have

$$\mathbf{Pr}[\, \mathrm{disc}(X(f(\ell))) \le 2 \cdot D \mid \Gamma \,] = 1,$$

and by the law of total probability we get,

$$\mathbf{Pr}[\mathrm{disc}(X(f(\ell))) \le 2 \cdot D] \ge 1 - \mathbf{Pr}\left[\overline{\Gamma}\right] \overset{\text{Eq. (5.3)}}{\ge} 1 - n^{-2}.$$

Finally the total number of rounds until $f(\ell)$ is,

$$\sum_{r=1}^{\log(|E|/D)} 48 \frac{|E| \cdot n \cdot \log(n)}{p_{\min} \cdot 2^r} \le (48|E| \cdot n \cdot \log(n)/p_{\min}) \cdot \sum_{r=0}^{\infty} 2^{-r} \le 98|E| \cdot n \cdot \log(n)/p_{\min},$$

and setting $t_1 := 98|E| \cdot n \cdot \log(n)/p_{\min}$ completes the proof. □

5.3.3 Discrepancy form $2D$ to D

Here, we show that starting from a load vector with $\mathrm{disc}(X(0)) \le 2D$, the discrepancy drops below D after $\mathrm{O}(D^2 n \log(n)/p_{\min})$ rounds, as stated in Proposition 5.5. For the proof, we first compute the expected change in $\Phi(\cdot)$ after one step, and then we bound the number of rounds required to reach discrepancy D.

Proposition 5.5. *Consider the token distribution on the sequence $\mathbf{M}^{[\infty]}$ (with $p_{\min}$ as the minimum probability of an edge appearing in any matching) and an arbitrary initial load vector $X(0)$ such that $D < \mathrm{disc}(X(0)) \le 2D$. Then, for the round*

$$t_1 := \left\lceil \frac{24D^2 \cdot n \cdot \log(n)}{p_{\min}} \right\rceil,$$

it holds that

$$\mathbf{Pr}[\mathrm{disc}(X(t_1)) \le D] \ge 1 - n^{-3}.$$

Proof. Fix round t and note $\Phi(t) \le 4nD^2$. Since the discrepancy is at least $D + 1$, then there is at least one edge in which its endpoints have load difference of ≥ 2. From Equation (5.7) it follows that $\mathbf{E}[\Phi(t) - \Phi(t+1)] \ge 2 \cdot p_{\min}$ and for $t^* := D^2 \cdot \lceil 4n - 1/2 \rceil / 2 \cdot p_{\min}$ we get $\mathbf{E}[\Phi(t + t^*)] \le D^2/2$. Markov's

Inequality implies $\mathbf{Pr}\big[\Phi(t+t^*) \geq D^2/4\big] \leq 1/2$ and since $\Phi(t)$ is not-increasing over t, then

$$\mathbf{Pr}\left[\bigcup_{j\in[3\log(n)]}\{\Phi(t+j\cdot t^*) < D^2/4\}\right] \geq 1 - \left(\frac{1}{2}\right)^{3\log(n)} \geq 1 - n^{-3}.$$

When this event occurs, then for some $j \in [3\log(n)]$, it holds that $\Phi(t+j\cdot t^*) \leq D^2/4$ and consequently, $\mathrm{disc}(X(t+j\cdot t^*)) \leq D$. Since the discrepancy is non-increasing then setting $t_1 := \lceil 24D^2 \cdot n \cdot \log(n)/p_{\min}\rceil \geq t + 3\log(n)\cdot t^*$ finishes the proof. □

We now establish our lower bound. For the remainder of this section, we assume that the underlying graph G is d-regular and first introduce the relevant variables and definitions. For each round t and any subset $S \subseteq V(G)$, we define

$$\mu_S(t) := \frac{1}{|S|}\sum_{u\in S} X_u(t),$$

as the mean load in the set S. With $\overline{S} := V(G)\setminus S$, we write

$$\Delta_S(t) := \mu_S(t) - \mu_{\overline{S}}(t),$$

for the difference between the mean load in S and its complement. Moreover, let $E(S,\overline{S})$ denote the set of edges crossing the cut between S and $\overline{S}$. Recall that the conductance of a d-regular graph $G=(V,E)$ is defined as

$$\Phi(G) := \min_{\substack{S\subseteq V(G)\\ |S|\leq n/2}} \frac{|E(S,\overline{S})|}{d|S|}.$$

We define T_K to denote the time required for the discrepancy to decrease from K to $K/4$. The next theorem provides a lower bound on T_K, showing that, for sufficiently large K, with constant probability, the process requires at least $\Omega(K/\Phi(G))$ rounds to reduce the initial discrepancy K to $K/4$.

Consider the following d-regular graph $G(V,E)$ with $|V| = n$ in which $n = 2r$ for $r \in \mathbb{N}$ and let $K = 4l$ for $l \in \mathbb{N}$. There is one node with load K, one node with load 0, $(n-2)/2$ nodes with load $3k/4$ and $(n-2)/2$ nodes with load $k/4$. It is clear that $\mathrm{disc}(X(0) = K$ and $\overline{x} = K/2$.

Theorem 5.6. *Consider the randomized token distribution on graph G with the initial load vector described above and assume $K > 8/\Phi(G)$. Then,*

$$\mathbf{Pr}\left[T_K \geq \frac{K}{8\Phi(G)}\right] \geq 1/2.$$

Note that from Cheeger's Inequliaty we get $K/\Phi(G) \leq cK/(1-\lambda)$ for some constant $c > 0$.

Proof. First define

$$t := \min\left\{\frac{K}{8\Phi(G)}, \frac{K^2}{64}\right\}.$$

For $K \geq 8/\Phi(G)$, we have $t = K/(8\Phi(G))$. Recall that $\overline{x} = K/2$. Let $S := \{u \in V \mid X_u(0) \geq \overline{x}\}$ denote the set of overloaded nodes. Note that $|S| = n/2$ and

$$\mu_S(0) = \frac{1}{|S|}\sum_{i\in S} X_i(0) = \frac{2}{n}\cdot\left(K + \frac{n-2}{2}\cdot\frac{3K}{4}\right) \geq \frac{3K}{4},$$

and similarly

$$\mu_{\overline{S}}(0) = \frac{1}{|\overline{S}|}\sum_{i\in\overline{S}} X_i(0) = \frac{2}{n}\cdot\left(0 + \frac{n-2}{2}\cdot\frac{K}{4}\right) \le \frac{K}{4},$$

and therefore $\Delta_S(0) = \mu_S(t) - \mu_{\overline{S}}(t) \ge K/2$. From lemma 5.13 (presented in the technical lemmas' section) with $p := 1/2$ and round t (as specified above), we obtain

$$\mathbf{Pr}[\Delta_S(t) > K/4] \ge 1/2.$$

Observe that, in (each) round t, there exist at least two nodes $u \in S$ and $v \in \overline{S}$ such that $X_u(t) \ge \mu_S(t)$ and $X_v(t) \le \mu_{\overline{S}}(t)$. Combining this with the bound $\Delta_S(t) > K/4$, we obtain

$$X_u(t) - X_v(t) > K/4,$$

which implies $\mathrm{disc}(X(t)) > K/4$. Hence, by definition of T_K, we have $T_K \ge t$, which completes the proof. □

5.4 Technical Lemmas

Here we list the basic and developed tools. Throughout this subsection, for each edge $(i,j) \in E$ and round t, we may assume that $X_i(t) \ge X_j(t)$ (otherwise we consider $(j,i) \in E$). Moreover, we assume a connected graph $G = (V,E)$ with maximum degree Δ, diffusion matrix $\mathbf{P}$ and diameter D together with a sequence of matchings $\mathbf{M}^{[\infty]} := \left(\mathbf{M}^{(s)}\right)_{s=1}^{\infty}$ is given.

5.4.1 Ingredients used in Theorem 5.2

We begin by stating the fundamental results employed in establishing the intermediate lemmas used in the proof of the upper bound, Theorem 5.2.

Discrepancy to $2|E|$. First we prove the intermediate results used in Proposition 5.3. Next lemma is a simple result computing the expected change of the potential as long as the load difference over all edges is at least $2|E|$.

Lemma 5.7. *Consider the Token Distribution on sequence $\mathbf{M}^{[\infty]}$ (with $p_{\min}$ as the minimum probability of an edge appearing in any matching and)with arbitrary initial load vector $X(0)$ such that $\mathrm{disc}(X(0)) \le K$. Assume $\sum_{(i,j)\in E} X_i(t) - X_j(t) > 2\cdot|E|$. Then, for $t \in \mathbb{N}$ with load vector $X(t)$, it holds*

$$\mathbf{E}[\Phi(t+1) \mid X(t)] \le \left(1 - \frac{2p_{min}\cdot\Delta}{K}\cdot(1-\lambda)\right)\cdot\Phi(t).$$

Proof. The proof proceeds along the lines of Theorem 1 in [52]. Fix step t. We assume $\sum_{(i,j)\in E} X_i(t) - X_j(t) \ge 2\cdot|E|$ otherwise we are already done.

Consider an edge $(i,j) \in \mathbf{M}^{(t+1)}$ with $X_i(t) \ge X_j(t)+1$. Then we have,

$$\begin{aligned}\Phi_i(t+1) + \Phi_j(t+1) &= (X_i(t) - 1 - \overline{x})^2 + (X_j(t) + 1 - \overline{x})^2 \\ &= (X_i(t) - \overline{x})^2 - 2(X_i(t) - \overline{x}) + 1 + (X_j(t) - \overline{x})^2 + 2(X_j(t) - \overline{x}) + 1 \\ &= \Phi_i(t) + \Phi_j(t) - 2(X_i(t) - X_j(t) - 1).\end{aligned}$$

Hence for an edge $(i,j) \in \mathbf{M}^{(t+1)}$ with $X_i(t) - X_j(t) \geq 1$, it holds that

$$\Phi_i(t) + \Phi_j(t) - \Phi_i(t+1) - \Phi_j(t+1) = 2(X_i(t) - X_j(t) - 1). \tag{5.4}$$

If a node is not matched in round $t+1$ or is matched with a node sharing the same load, then its contribution to $\Phi(t) - \Phi(t+1)$ equals zero. Accumulating the contribution of all nodes yields,

$$\begin{aligned}
\mathbf{E}[\,\Phi(t) - \Phi(t+1) \mid X(t)\,] &= \sum_{\substack{(i,j)\in E:\\ X_i(t)-X_j(t)\geq 1}} \mathbf{Pr}\Big[(i,j) \in \mathbf{M}^{(t+1)}\Big] \cdot (2(X_i(t) - X_j(t) - 1)) \\
&\geq p_{\min} \cdot \sum_{\substack{(i,j)\in E\\ X_i(t)-X_j(t)\geq 1}} 2(X_i(t) - X_j(t) - 1) \qquad (5.5)\\
&\geq p_{\min} \cdot \left(\sum_{(i,j)\in E} (2(X_i(t) - X_j(t))) - 2|E| \right) \\
&\overset{(a)}{\geq} p_{\min} \cdot \sum_{(i,j)\in E} (X_i(t) - X_j(t)) = p_{\min} \cdot \sum_{(i,j)\in E} \frac{(X_i(t) - X_j(t))^2}{X_i(t) - X_j(t)} \qquad (5.6)
\end{aligned}$$

where (a) follows from the assumption $\sum_{(i,j)\in E} X_i(t) - X_j(t) \geq 2|E|$. Since $X_i(t) - X_j(t) \leq K$, then we have $1/(X_i(t) - X_j(t)) \geq 1/K$ and from Equation (5.6) we get

$$\mathbf{E}[\,\Phi(t) - \Phi(t+1) \mid X(t)\,] \geq p_{\min} \cdot \left(\sum_{(i,j)\in E} \frac{(X_i(t) - X_j(t))^2}{K} \right).$$

Notice that we assume for the reminder of the proof that $X(t) \neq \mathbf{1} \cdot \overline{x}$, otherwise the claim of lemma holds since $\Phi(t+1) = \Phi(t) = 0$. Hence, $\Phi(t) \neq 0$ and

$$\begin{aligned}
\mathbf{E}\left[\frac{\Phi(t) - \Phi(t+1)}{\Phi(t)} \;\middle|\; X(t) \right] &= \frac{\mathbf{E}[\Phi(t) - \Phi(t+1)]}{\Phi(t)} \\
&\geq \frac{\frac{p_{\min}}{K} \cdot \sum_{(i,j)\in E} (X_i(t) - X_j(t))^2}{\sum_{i\in V} (X(t) - \overline{x})^2}.
\end{aligned}$$

By Introducing the vector $y \in \mathbb{R}^n$ where $y_i := X_i(t) - \overline{x}$, we can rewrite the formula above as

$$\mathbf{E}\left[\frac{\Phi(t) - \Phi(t+1)}{\Phi(t)} \;\middle|\; X(t) \right] \geq \frac{p_{\min}}{K} \cdot \frac{y\,\mathbf{L}\,y^T}{y \cdot y^T},$$

where here $\mathbf{L}$ is the Laplacian matrix of G defined by $\mathbf{L}_{i,i} = d_i$, $\mathbf{L}_{i,j} = -1$ for $\{u,v\} \in E$, and $\mathbf{L}_{i,j} = 0$, otherwise. Since $X(t) \neq \mathbf{1} \cdot \overline{x}$ we have that $y \neq \mathbf{0}$ and $y \perp \mathbf{1}$, and the Min-Max characterization of eigenvalue yields

$$\mathbf{E}\left[\frac{\Phi(t) - \Phi(t+1)}{\Phi(t)} \;\middle|\; X(t) \right] \geq \frac{p_{\min}}{K} \cdot \lambda_{n-1}(\mathbf{L}).$$

Since $\mathbf{P} = \mathbf{I} - \frac{1}{2\Delta} \cdot \mathbf{L}$, we have that $\lambda(\mathbf{P}) = 1 - \frac{\lambda_{n-1}(\mathbf{L})}{2\Delta}$ and $\lambda_{n-1}(\mathbf{L}) = 2\Delta \cdot (1 - \lambda(\mathbf{P}))$. This implies,

$$\mathbf{E}[\,\Phi(t+1) \mid X(t)\,] \leq \left(1 - \lambda_{n-1}(\mathbf{L}) \cdot \frac{p_{\min}}{K}\right) \cdot \Phi(t) = \left(1 - \frac{2\Delta \cdot p_{\min}}{K} \cdot (1 - \lambda(\mathbf{P}))\right)\Phi(t).$$

□

Discrepancy from $2|E|$ to $2D$. Here we prove the intermediate results used in Proposition 5.4. Note that since load is moved only from the higher loaded nodes to the lower loaded nodes, therefore:

Observation 5.8. *The discrepancy* $\text{disc}(X(.))$ *in the Token Distribution protocol is non-increasing.*

Moreover, from Equation (5.4) we have that

Observation 5.9. *The potential function* $\Phi(.)$ *in the Token Distribution protocol is non-increasing.*

We partition the time into consecutive intervals of different size. Without loss of generality we assume $\text{disc}(X(0)) \leq 2|E|$. We let $e(0) := 0$. For $r \in \mathbb{N}$, we define *phase* i to be the rounds $[e(r-1)+1, e(r)]$ such that

$$e(r) := \min\left\{ t \mid \text{disc}(X(t)) \leq \frac{2|E|}{2^r} \right\}.$$

Next lemma bounds the decrease in $\Phi(.)$ after each step during phase i. To show this we use the properties of phase r.

Lemma 5.10. *Consider the Token Distribution on sequence* $\mathbf{M}^{[\infty]}$ *(with* $p_{\min}$ *as the minimum probability of an edge appearing in any matching and) with arbitrary initial load vector* $X(0)$ *such that* $\text{disc}(X(0)) \leq 2|E|$*. Consider a phase* $r \in [1, \log(|E|/D)]$ *with* $\text{disc}(X(e(r))) > 2D$*. For any round* t *in this phase, it holds that*

$$\mathbf{E}[\, \Phi(t) - \Phi(t+1) \mid X(t)\,] \geq \frac{p_{\min} \cdot |E|}{2^r}.$$

Proof. From Equation (5.5) it follows that,

$$\begin{aligned} \mathbf{E}[\, \Phi(t) - \Phi(t+1) \mid X(t)\,] &\geq p_{\min} \cdot \sum_{\substack{(i,j)\in E \\ X_i(t)-X_j(t)\geq 1}} 2(X_i(t) - X_j(t) - 1) \\ &\geq p_{\min} \cdot \sum_{\substack{(i,j)\in E \\ X_i(t)-X_j(t)\geq 2}} (X_i(t) - X_j(t)). \end{aligned} \tag{5.7}$$

We are to bound the term in Equation (5.7) using our bounds on the discrepancy in the phase r.

Fix a round t in phase r. Note that by the definition we have $\text{disc}(X(t)) \geq 2|E|/2^r$. Consider two nodes $u, v \in V$ such that $X_u(t) - X_v(t) = \text{disc}(X(t))$. Since there is a path of length (at most) D between u and v in the graph G then there can be at most $D-1$ edges on this path in which the endpoints have load difference of at most 1. Hence,

$$\begin{aligned} \sum_{\substack{(i,j)\in E \\ X_i(t)-X_j(t)\geq 2}} (X_i(t) - X_j(t)) &\geq \text{disc}(X(t)) - (D-1) \\ &\overset{(a)}{\geq} \text{disc}(X(e(r))) - (D) \\ &\overset{(b)}{\geq} \frac{\text{disc}(X(e(r)))}{2} \geq \frac{|E|}{2^r}, \end{aligned} \tag{5.8}$$

where (a) holds since discrepancy is non-increasing over the time and $t \leq e(r)$ and (b) follows from the assumption $\text{disc}(X(e(r))) > 2D$. Therefore, from Equation (5.7) and Equation (5.8) together, it follows that

$$\mathbf{E}[\, \Phi(t) - \Phi(t+1) \mid X(t)\,] \geq \frac{p_{\min} \cdot |E|}{2^r}.$$

□

In the following lemma we bound the number of rounds in phase r. For this we use the last lemma and the Markov's inequality.

Lemma 5.11. *Consider the Token Distribution on sequence* $\mathbf{M}^{[\infty]}$ *(with* $p_{\min}$ *as the minimum probability of an edge being in any matching and) with arbitrary initial load vector* $X(0)$ *such that* $\operatorname{disc}(X(0)) \leq 2|E|$*. Length of phase* $r \in [1, \log(|E|/D)]$ *is at most* $\lceil 48 \cdot |E| \cdot n \cdot \log(n)/(2^r \cdot p_{\min}) \rceil$ *rounds with probability at least* $1 - n^{-3}$*.*

Proof. Fix a phase $r \in [1, \log(|E|/D)]$. Recall that by definition we have $\operatorname{disc}(X(e(r-1))) \leq 2|E|/2^{r-1}$ and $\Phi(X(e(r-1))) \leq n(2|E|/2^{r-1})^2$.

Let us define $t_0 := e(r-1)$ and $t^* := \left\lceil \frac{|E|(16n-1/2)}{p_{\min} \cdot 2^r} \right\rceil$. From Lemma 5.10 (and the law of total expectation) we get,

$$
\begin{aligned}
\mathbf{E}[\, \Phi(t_0 + t^*) \mid X(t_0) \,] &\leq \Phi(t_0) - t^* \cdot \frac{p_{\min}|E|}{2^r} \\
&\leq n \cdot \left(\frac{2|E|}{2^{r-1}} \right)^2 - t^* \cdot \frac{p_{\min}|E|}{2^r} \\
&\leq n \cdot \frac{16|E|^2}{2^{2r}} - \frac{|E| \cdot (16n - 1/2)}{p_{\min} \cdot 2^r} \cdot \frac{p_{\min}|E|}{\cdot 2^r} \\
&= \frac{16n|E^2| - 16n|E^2| + |E|^2/2}{2^{2r}} = \frac{|E|^2}{2^{2r+1}}.
\end{aligned}
\tag{5.9}
$$

Here for integer $j \in [1, 3\log(n)]$, we define *epoch* j to be the rounds $[t_0 + (j-1) \cdot t^* + 1, t_0 + j \cdot t^*]$. For each epoch j we define a good event

$$
\mathcal{E}_j := \left\{ \Phi(t_0 + j \cdot t^*) \leq \frac{|E|^2}{2^{2r}} \right\}.
$$

Note that epochs j and $k \neq j$ consist of disjoint time intervals. From Equation (5.9) and Markov's Inequality we have that $\mathbf{Pr}\left[\overline{\mathcal{E}_j}\right] \leq \frac{1}{2}$. Then we get

$$
\mathbf{Pr}\left[\bigcup_{j \in [3\log(n)]} \mathcal{E}_j \right] = 1 - \mathbf{Pr}\left[\bigcap_{j \in [3\log(n)]} \overline{\mathcal{E}_j} \right] \geq 1 - \left(\frac{1}{2} \right)^{3 \cdot \log(n)} = 1 - \frac{1}{n^3}.
$$

The event $\bigcup_{j \in [3\log(n)]} \mathcal{E}_j$ implies there is at least one epoch $j \in [3\log(n)]$ for which the good event $\mathcal{E}_j$ holds. In the rest of proof we assume the event $\mathcal{E}_j$ for some j holds. Since $\operatorname{disc}(X(t)) \leq 2 \cdot \sqrt{\Phi(t)}$, we have

$$
\operatorname{disc}(X(t_0 + t^* \cdot j)) \leq 2 \cdot \sqrt{\frac{|E|^2}{2^{2r}}} = \frac{2|E|}{2^r},
$$

indicating the round $t_0 + t^* \cdot j$ is the end of phase r (in fact $e(r) \leq t_0 + t^* \cdot j$). Moreover, the number of rounds in phase r is at most

$$
t^* \cdot 3 \cdot \log(n) + t_0 - t_0 \leq 48 \cdot |E| \cdot n \cdot \log(n)/p_{\min} \cdot 2^r
$$

finishing the proof. □

5.4.2 Ingredients used in Theorem 5.6

Here we provide the basic and intermediate results used in the proof of the lower bound stated in Theorem 5.6. To establish the theorem, we prove two auxiliary lemmas. The first lemma shows that the expected change in $\Delta(\cdot)$ after one round is bounded by the conductance $\Phi(G)$.

Recall that for a set $S \subseteq V(G)$ we let $\mu_S(t) = \frac{1}{|S|}\sum_{u\in S} X_u(t)$ denote the mean load in the set S and

$$\Delta_S(t) = \mu_S(t) - \mu_{\overline{S}}(t).$$

Lemma 5.12. *Let $S \subseteq V(G)$ with $|S| \le n/2$. Then, for any round $t \in \mathbb{N}$, it holds that*

$$|\mathbf{E}[\,\Delta_S(t) - \Delta_S(t-1) \mid X(t-1) = x\,]| \le \Phi(G).$$

Proof. Consider an edge $(u, v) \in \mathbf{M}^{(t)}$. If both u and v belong to S (or both to $\overline{S}$), a token may move within S (or $\overline{S}$), so the expected values $\mu_S(t)$ and $\mu_{\overline{S}}(t)$ remain unchanged. In the remaining case, where $u \in S$ and $v \in \overline{S}$ (or vice versa), the sum of loads in S may increase or decrease. Hence,

$$\mathbf{E}[\,\mu_S(t) - \mu_S(t-1) \mid X(t-1) = x\,] = \sum_{(i,j)\in E(S,\overline{S})} \mathbf{Pr}\Big[(i,j) \in \mathbf{M}^{(t)}\Big] \cdot \frac{\operatorname{sgn}(x_j - x_i)}{|S|},$$

$$\mathbf{E}[\,\mu_{\overline{S}}(t) - \mu_{\overline{S}}(t-1) \mid X(t-1) = x\,] = \sum_{(i,j)\in E(S,\overline{S})} \mathbf{Pr}\Big[(i,j) \in \mathbf{M}^{(t)}\Big] \cdot \frac{\operatorname{sgn}(x_i - x_j)}{|\overline{S}|}.$$

Since $\Delta_S(t) = \mu_S(t) - \mu_{\overline{S}}(t)$, linearity of expectation and the fact that $\operatorname{sgn}(-z) = -\operatorname{sgn}(z)$ give

$$\begin{aligned}
&\mathbf{E}[\,\Delta_S(t) - \Delta_S(t-1) \mid X(t-1) = x\,] \\
&\quad= \mathbf{E}[\,\mu_S(t) - \mu_S(t-1) \mid X(t-1) = x\,] - \mathbf{E}[\,\mu_{\overline{S}}(t) - \mu_{\overline{S}}(t-1) \mid X(t-1) = x\,] \\
&\quad= \sum_{(i,j)\in E(S,\overline{S})} \mathbf{Pr}\Big[(i,j) \in \mathbf{M}^{(t)}\Big] \cdot \left(\frac{\operatorname{sgn}(x_j - x_i)}{|S|} - \frac{\operatorname{sgn}(x_i - x_j)}{|\overline{S}|}\right) \\
&\quad\le \sum_{(i,j)\in E(S,\overline{S})} \mathbf{Pr}\Big[(i,j) \in \mathbf{M}^{(t)}\Big] \cdot \left(\frac{1}{|S|} + \frac{1}{|\overline{S}|}\right) \\
&\quad\le \sum_{(i,j)\in E(S,\overline{S})} \frac{1}{d|S|},
\end{aligned}$$

where the last inequality uses Observation 5.1, $\mathbf{Pr}\big[(i,j) \in \mathbf{M}^{(t)}\big] \le 1/(2d)$, and the fact that $|S| \le n/2$ implies $|S| \le |\overline{S}|$.

Applying the triangle inequality, we obtain

$$|\mathbf{E}[\,\Delta_S(t) - \Delta_S(t-1) \mid X(t-1) = x\,]| \le \sum_{(i,j)\in E(S,\overline{S})} \frac{1}{d|S|} \le \frac{|E(S,\overline{S})|}{d|S|} = \Phi(G),$$

where the last inequality holds since the graph G is d-regular. It then completes the proof. □

We next prove the second lemma, which is instrumental for the theorem. The lemma asserts that there exists a round t (in terms of the initial discrepancy and the conductance of the graph) at which, $\Delta_S(t)$ exceeds a threshold determined by the initial discrepancy.

Lemma 5.13. *Let $S \subseteq V(G)$ with $|S| \leq n/2$, and fix an initial load vector with $\Delta_S(0) \geq K$. Then, for any $p \in (0,1)$ and round*

$$t := \min\left\{\frac{K}{4\Phi(G)}, \frac{K^2}{8\log(2/p)}\right\},$$

it holds that

$$\mathbf{Pr}\left[\Delta_S(t) > \frac{K}{2}\right] \geq 1 - p.$$

Proof. For a round $t' \in [1, t]$, define the random variable

$$Z'_t := \Delta_S(t') - t' \cdot \Phi(G).$$

It follows from lemma 5.12 that the sequence $Z_1, Z_2, \ldots, Z_t$ forms a submartingale. Applying lemma 5.12 and using the bound $|\Delta_S(t) - \Delta_S(t-1)| \leq 2$ ($|S \leq n/2$ and since we have matching then a single node $u \in V$ can contribute to this difference by at most $2/|S|$), we obtain via two-sided Azuma's inequality (similar to Theorem 6.6) for $\varepsilon := \sqrt{2t\log(2/p)}$, that

$$\mathbf{Pr}[|\Delta_S(t) - \Delta_S(0)| \geq t \cdot \Phi(G) + \varepsilon] \leq 2\exp\left(-\frac{2\varepsilon^2}{t(2)^2}\right) = 2\exp\left(-\frac{\varepsilon^2}{2t}\right) := p.$$

Since $\Delta_S(0) \geq K$, it follows that $\Delta_S(t) > K/2$ with probability at least $1 - p$, provided

$$t \cdot \Phi(G) + \sqrt{t} \cdot \sqrt{2\log(2/p)} < \frac{K}{2}.$$

This holds if each term is at most $K/4$ individually, i.e., when

$$t \leq \min\left\{\frac{K}{4\Phi(G)}, \left(\frac{K}{2\sqrt{2\log(2/p)}}\right)^2\right\} = \min\left\{\frac{K}{4\Phi(G)}, \frac{K^2}{8\log(2/p)}\right\},$$

which establishes the claim. ☐

5.5 Bounds for Specific Graph Classes

In this section, we present (upper) bounds on the discrepancy for specific classes of graphs and compare them in detail with the most closely related work [49] in which the authors show that in a d-regular graph with n nodes and edge expansion α, after a number of rounds, the discrepancy reaches $O(d\log(n)/\alpha)$. We, in contrast, show that the discrepancy becomes smaller than D after sufficiently many rounds for any connected graph. In the following, we compare our results for concrete graph families.

r-Dimensional Torus. Here $d = 2r$, $\alpha = \Theta(n^{-1/r})$, and $D = \theta(rn^{1/r})$. Note that $r = 1$ and $r = \log(n)$ correspond to the cycle and the hypercube, respectively. It holds that

$$\frac{d\log(n)}{\alpha} = 2rn^{1/r}\log(n) > D,$$

implying that our bound yields a smaller discrepancy. The same holds for r-dimensional grid.

Complete Graphs. Here $\alpha = \Theta(n)$, $d = n-1$, and $D = 1$. Hence,

$$\frac{d\log(n)}{\alpha} = \Theta(\log(n)) > D.$$

Hamming (q-ary hypercube) Graphs $H(r, q)$. Here $n = q^r$, $d = r(q-1)$, and $D = r$. Fixing the first coordinate gives $\alpha(H(r,q)) \le q-1$, and in fact, the boundary-to-volume ratio shows that $\alpha = q-1$. Thus,

$$\frac{d\log(n)}{\alpha} = \frac{r(q-1)\log(q^r)}{q-1} = r^2\log(q) > D,$$

for all $q, r \ge 2$.

Cartesian Powers of a Fixed Connected Regular Graph. Here our analysis again provides a smaller discrepancy. Let G be a d_0-regular graph on n_0 vertices with edge expansion α_0 and diameter D_0. Then $n = n_0^r$, $d = rd_0$, and $D = rD_0$. Taking sets that vary only in one coordinate, one sees that $\alpha \le \alpha_0$ (hence α does not grow with r). Consequently,

$$\frac{d\log(n)}{\alpha} \ge \frac{r^2 d_0 \log(n_0)}{\alpha_0} > rD_0 = D,$$

where the last inequality holds for sufficiently large r. This construction subsumes the hypercube, Hamming, and torus examples, as they are Cartesian powers of small base graphs.

Barbell-type Regular Graphs. Here we also obtain a better bound. Take two d-regular expanders on $n/2$ vertices and connect them by a small number t of matching edges, adjusting interior edges to maintain regularity. If t is small relative to n, the edge expansion satisfies $\alpha = \Theta(t/n)$, while the diameter remains small (at most a small constant greater than the expanders' diameters). Hence,

$$\frac{d\log(n)}{\alpha} > D.$$

Note that a tiny cut (small boundary) makes α very small while keeping the diameter small.

Random d-Regular Graphs. Here we typically have $\alpha = \Theta(1)$ and $D = \Theta(\log(n))$. Then for non-constant d, we have

$$\frac{d\log(n)}{\alpha} = \theta(d\log(n)) = \Omega(D),$$

indicating that for any fixed non-constant d and typical constants, with high probability, our bound provides (asymptotically) smaller discrepancy. Moreover, in case of constant d, our bounds coincide.

For many standard expanders that are locally tree-like, our diameter-based bound also yields a smaller value. Finally, our bound applies to non-regular graphs, whereas the previous bound applies only to regular graphs.

5.6 Summary and Open Problems

In this work, we show that the token distribution process reduces any initial discrepancy $K > D$ to D within $\mathrm{O}\big(K\log(Kn)/(1-\lambda) + (|E| + D^2)|E|\log(n)\big)$ rounds, w.h.p. Moreover, for a d-regular graph with initial discrepancy $\mathrm{disc}(x(0)) = K > 8/\Phi(G)$, it requires $\Omega(K/(1-\lambda))$ rounds to reach a discrepancy of $K/4$ with consonant probability.

Many related works have claimed that the process eventually reaches the diameter of the graph. However, their results (1) do not specify the runtime, and (2) apply only to regular graphs. We show that, for any arbitrary connected graph, the process reaches a load vector with discrepancy at most the diameter. This implies that, for graphs with low expansion-sparse and poorly connected graphs-our bound on the discrepancy improves upon existing results. In particular, for d-regular graphs with n nodes, edge expansion α, and diameter D, when $d\log(n)/\alpha = \Omega(D)$, we show a smaller discrepancy; however, our analysis may require a larger number of rounds.

An interesting question is whether our results can be extended to the multi-port version of the problem, in which nodes of a regular graph may send or receive a token over each edge in their neighborhood whenever the load difference between endpoints is at least $2d$. We believe that a diameter-related discrepancy bound should still be achievable. Another open question is whether one can design a randomized version of the process, in which at most one token is moved over each edge, and show that the discrepancy reaches a small constant, possibly by combining this approach with our results from part three on discrete static load balancing using matchings.

6 Auxiliary Results

In this section we list the known results which we frequently use in our analysis

6.1 Drift Result

In our analysis we use the following tail bound for the sum of a non-increasing sequence of random variables with variable negative drift. The proof uses established methods from drift analysis. In particular, it relies one techniques found in the proof of the Variable Drift Theorem in [63].

Theorem 6.1. *Let $(X(t))_{t\geq 0}$ be a non-increasing sequence of discrete random variables with $X(t) \in \mathbb{R}_0^+$ for all t with fixed $X(0) = x_0$. Assume there exists an increasing function $h\colon \mathbb{R}_0^+ \to \mathbb{R}^+$ and a constant $\sigma > 0$ such that the following holds. For all $t \in \mathbb{N}$ and all $x > 0$ with $\mathbf{Pr}[X(t) = x] > 0$*

1. $\mathbf{E}[X(t+1) \mid X(t) = x] \leq x - h(x)$,
2. *Var*$[X(t+1) \mid X(t) = x] \leq \sigma \cdot (\mathbf{E}[X(t+1) \mid X(t) = x] - x)^2$.

Then the following statements hold.

1. *For all $\delta \in (0,1)$ and any arbitrary but fixed t*

$$\mathbf{Pr}\left[\int_{X(t)}^{x_0} \frac{1}{h(\varphi)}\, \mathrm{d}\varphi \leq (1-\delta)t\right] \leq \exp\left(-\frac{\delta^2 t}{2(\sigma+1)}\right).$$

2. *For all $\delta \in (0,1)$ and $p \in (0,1)$ we define $t_0 := \frac{2(\sigma+1)}{\delta^2}\left(-\log(p) + \log\left(\frac{2(\sigma+1)}{\delta^2}\right)\right)$. Then*

$$\mathbf{Pr}\left[\sum_{t=t_0+1}^{\infty} X(t) \leq \frac{1}{1-\delta} \cdot \int_0^{x_0} \frac{\varphi}{h(\varphi)} \mathrm{d}\varphi\right] \geq 1-p.$$

Proof. Throughout this proof we write

$$f(x) := \int_x^{x_0} \frac{1}{h(\varphi)} \mathrm{d}\varphi.$$

We start by proving the first statement. Let $a, b \in \mathbb{R}^+$ with $a \leq b \leq x_0$ be two arbitrary numbers. Since h is increasing we have $h(a) \leq h(b)$ and $1/h(a) \geq 1/h(b)$. Hence,

$$f(a) - f(b) = \int_a^{x_0} \frac{1}{h(\varphi)} \mathrm{d}\varphi - \int_b^{x_0} \frac{1}{h(\varphi)} \mathrm{d}\varphi = \int_a^b \frac{1}{h(\varphi)} \mathrm{d}\varphi \geq \int_a^b \frac{1}{h(b)} \mathrm{d}\varphi = \frac{b-a}{h(b)}.$$

From condition 1 of the theorem it follows that $\mathbf{E}[X(t+1) \mid X(t) = b] \leq b - h(b)$ and consequently $h(b) \leq b - \mathbf{E}[X(t+1) \mid X(t) = b]$ giving us with $X(t) = b$

$$f(X(t+1)) - f(b) \geq \frac{X(t+1) - b}{\mathbf{E}[X(t+1) - b \mid X(t) = b]}. \tag{6.1}$$

We introduce a new sequence of random variables for which we will derive a lower tail bound, defined as $(Y(t))_{t\in\mathbb{N}}$ given by $Y(0) := 0$ and

$$Y(t+1) := Y(t) + \frac{X(t+1) - X(t)}{\mathbf{E}[X(t+1) - X(t)]}.$$

Comparing this with Equation (6.1) we see that regardless of the value of $X(t)$ it holds that

$$f(X(t+1)) - f(X(t)) \geq \frac{X(t+1) - X(t)}{\mathbf{E}[X(t+1) - X(t)]} = Y(t+1) - Y(t).$$

By induction over t, and since $f(x_0) = \int_{x_0}^{x_0} (1/h(\varphi))\,\mathrm{d}\varphi = 0$ and $Y(0) = 0$, we have for all t

$$f(X(t)) = f(X(t)) - f(x_0) \geq Y(t) - Y(0) = Y(t).$$

From the definition of $(Y_t)_{t\geq 0}$ it follows assuming $X(t) = x$ that

$$\mathbf{E}[Y(t+1) - Y(t) \mid X(t) = x] = \mathbf{E}\left[\frac{X(t+1) - x}{\mathbf{E}[X(t+1) - x]}\right] = 1.$$

Then, from the law of total expectation we get that

$$\begin{aligned}\mathbf{E}[Y(t+1) - Y(t)] &= \sum_x \mathbf{E}[Y(t+1) - Y(t) \mid X(t) = x] \cdot \Pr[X(t) = x] \\ &= \sum_x 1 \cdot \Pr[X(t) = x] = 1.\end{aligned}$$

Since $Y(0) = 0$ it immediately follows that $\mathbf{E}[Y(t)] = t$. Furthermore, we may bound the variance of the change of Y given $X(t) = x$ by

$$\begin{aligned}\mathrm{Var}[Y(t+1) - Y(t) \mid X(t) = x] &= \mathrm{Var}\left[\frac{X(t+1) - x}{\mathbf{E}[X(t+1) - x]}\right] = \frac{\mathrm{Var}[X(t+1) - x]}{(\mathbf{E}[X(t+1) - x])^2} \\ &\overset{(a)}{\leq} \frac{\sigma \cdot (\mathbf{E}[X(t+1)] - x)^2}{(\mathbf{E}[X(t+1) - x])^2} = \sigma,\end{aligned}$$

where (a) follows from Condition 2 of the theorem. The sequence $(Y(t) - \mathbf{E}[Y(t)])_{t\geq 0}$ is a martingale and hence fulfills the preconditions of Theorem 6.6 from [35] (restated as Theorem 6.4) with $a_t \coloneqq 1$ and $\sigma_t^2 \coloneqq \sigma$. Note that $\mathbf{E}[Y(t) - \mathbf{E}[Y(t)]] = 0$. Hence, we obtain

$$\mathbf{Pr}[Y(t) - \mathbf{E}[Y(t)] \leq 0 - \varepsilon] \leq \exp\left(-\frac{\varepsilon^2}{2t(\sigma+1)}\right).$$

Recalling that $f(X(t)) \geq Y(t)$ and $\mathbf{E}[Y(t)] = t$ and setting $\varepsilon = \delta t$ for some $\delta \in (0,1)$ we arrive at the first statement of the theorem;

$$\mathbf{Pr}[f(X(t)) \leq (1-\delta)t] \leq \exp\left(-\frac{\delta^2 t}{2(\sigma+1)}\right).$$

Next we prove the second statement and bound $\sum_{t=t_0+1}^{\infty} X(t)$. Let $T(x) \coloneqq \min\{t \in \mathbb{N} \mid X(t) \leq x\}$ be a hitting time for the event that $X(t) \leq x$. Using $\mathbf{1}_{x<X(t)}$ as the indicator variable (which is one if $x < X(t)$ and zero otherwise) we can write $X(t) = \int_0^{x_0} \mathbf{1}_{X(t)>x}\,\mathrm{d}x$ because x_0 is fixed and $X(t)$ is non-increasing in

t resulting in $X(t) \in [0, x_0]$. As a consequence it holds that

$$\begin{aligned}\sum_{t=t_0+1}^{\infty} X(t) &= \sum_{t=t_0+1}^{\infty} \int_0^{x_0} \mathbf{1}_{X(t)>x} \, \mathrm{d}x = \int_0^{x_0} \left(\sum_{t=t_0+1}^{\infty} \mathbf{1}_{X(t)>x} \right) \mathrm{d}x \\ &= \int_0^{x_0} \left(\sum_{t=t_0+1}^{\infty} \mathbf{1}_{t<T(x)} \right) \mathrm{d}x = \int_0^{x_0} \left(\sum_{t=t_0+1}^{T(x)-1} 1 \right) \mathrm{d}x = \int_0^{x_0} \max\{0, T(x) - (t_0+1)\} \, \mathrm{d}x.\end{aligned}$$

We now proceed to bound the $T(x)$. Using the first statement with a union bound over all t for

$$t > t_0 := \frac{2(\sigma+1)}{\delta^2} \cdot \left(-\log(p) + \log\left(\frac{2(\sigma+1)}{\delta^2} \right) \right)$$

gives us

$$\begin{aligned}\mathbf{Pr}\left[\bigvee_{t=t_0+1}^{\infty} f(X(t)) \leq (1-\delta)t \right] &\leq \sum_{t=t_0+1}^{\infty} \exp\left(-\frac{\delta^2 t}{2(\sigma+1)} \right) \leq \int_{t_0}^{\infty} \exp\left(-\frac{\delta^2 t}{2(\sigma+1)} \right) \mathrm{d}t \\ &= \frac{2(\sigma+1)}{\delta^2} \cdot \exp\left(-\frac{\delta^2 t_0}{2(\sigma+1)} \right) =: p.\end{aligned}$$

As a consequence,

$$\mathbf{Pr}\left[\bigwedge_{t=t_0+1}^{\infty} t \leq \frac{f(X(t))}{1-\delta} \right] \geq 1-p,$$

and

$$\mathbf{Pr}\left[\bigwedge_{t \in \mathbb{N}_0} t \leq \max\left\{ t_0, \frac{f(X(t))}{1-\delta} \right\} \right] \geq 1-p. \tag{6.2}$$

Recalling that $T(x) := \min\{t \in \mathbb{N} \mid X(t) \leq x\}$ Equation (6.2) implies that

$$\mathbf{Pr}\left[\bigwedge_{x<x_0} T(x) - 1 \leq \max\left\{ t_0, \frac{f(X(T(x)-1))}{1-\delta} \right\} \right] \geq 1-p,$$

since $X(T(x)-1) > x$ by the definition of $T(x)$ and f is non-increasing it holds that $f(X(T(x)-1)) \leq f(x)$. It follows that

$$\mathbf{Pr}\left[\bigwedge_{x \leq x_0} T(x) - 1 \leq \max\left\{ t_0, \frac{f(x)}{1-\delta} \right\} \right] \geq 1-p.$$

As a consequence we get that with probability at least $1-p$

$$\int_0^{x_0} \max\{0, T(x) - (t_0+1)\} \, \mathrm{d}x \leq \int_0^{x_0} \max\left\{ 0, \max\left\{ t_0, \frac{f(x)}{1-\delta} \right\} + 1 - (t_0+1) \right\} \mathrm{d}x.$$

Finally, we find that

$$
\begin{aligned}
&\int_0^{x_0} \max\left\{0, \max\left\{t_0, \frac{f(x)}{1-\delta}\right\} + 1 - (t_0+1)\right\} \mathrm{d}x \\
&\quad= \int_0^{x_0} \max\left\{0, \frac{f(x)}{1-\delta} - t_0\right\} \mathrm{d}x \le \frac{1}{1-\delta} \int_0^{x_0} f(x)\,\mathrm{d}x \\
&\quad= \frac{1}{1-\delta} \int_0^{x_0} \int_x^{x_0} \frac{1}{h(\varphi)} \,\mathrm{d}\varphi\,\mathrm{d}x = \frac{1}{1-\delta} \int_0^{x_0} \int_0^{x_0} \frac{\mathbf{1}_{\varphi \ge x}}{h(\varphi)} \,\mathrm{d}\varphi\,\mathrm{d}x \\
&\quad= \frac{1}{1-\delta} \int_0^{x_0} \frac{1}{h(\varphi)} \int_0^{x_0} \mathbf{1}_{x \le \varphi} \,\mathrm{d}x\,\mathrm{d}\varphi = \frac{1}{1-\delta} \cdot \int_0^{x_0} \frac{1}{h(\varphi)} \cdot \varphi\,\mathrm{d}\varphi.
\end{aligned}
$$

Putting everything together we see with probability at least $1-p$ that

$$
\sum_{t=t_0+1}^{\infty} X(t) \le \frac{1}{1-\delta} \cdot \int_0^{x_0} \frac{\varphi}{h(\varphi)} \cdot \mathrm{d}\varphi. \qquad \square
$$

6.2 Concentration Results

The following lemma allows us to turn a high-probability bound into a bound on the expected value. It is a known state.

Lemma 6.2. *Let X be a non-negative real random variable, and let $n \in \mathbb{N}$. Then if there are $c, C > 0$ such that for all $\gamma > 0$,*

$$
\mathbf{Pr}[X \ge (\gamma+1)C] \le cn^{-\gamma},
$$

then

$$
\mathbf{E}[X] \le C \cdot \left(1 + \frac{c}{\log(n)}\right).
$$

Proof. Observe that when $x = (\gamma+1)C$ we have $\gamma = \frac{x}{C} - 1$, so that for all $x \ge C$ we have

$$
\mathbf{Pr}[X \ge x] \le c \cdot n^{-\frac{x}{C}+1}.
$$

Thus,

$$
\begin{aligned}
\mathbf{E}[X] &= \int_0^{\infty} \mathbf{Pr}[X \ge x]\,\mathrm{d}x = \int_0^{C} \mathbf{Pr}[X \ge x]\,\mathrm{d}x + \int_C^{\infty} \mathbf{Pr}[X \ge x]\,\mathrm{d}x \\
&\le \int_0^{C} 1\,\mathrm{d}x + \int_C^{\infty} c \cdot n^{-\frac{x}{C}+1}\,\mathrm{d}x = C + \left[-\frac{cCn^{1-\frac{x}{C}}}{\log(n)}\right]_{x=C}^{\infty} = C + \left[0 + \frac{cCn^{1-1}}{\log(n)}\right] \\
&= C\left(1 + \frac{c}{\log(n)}\right),
\end{aligned}
$$

as claimed. $\square$

Theorem 6.3 (Bhatia-Davis inequality [30])**.** *Let X be a real random variable with $X \in [m, M]$. Then $Var[X] \le (M - \mathbf{E}[X])(\mathbf{E}[X] - m)$.*

Theorem 6.4 (Adapted from Theorem 6.6 in [35])**.** *Let $(X(t))_{t=0}^n$ be a martingale associated with the filter $(\mathcal{F}(t))_{t=0}^n$, where there exist $(a_t)_{t=1}^n$ and $(\sigma_t)_{t=1}^n$ such that for all $t \in [n]$,*

1. *$X(t) - X(t-1) \ge a_t$;*

2. Var$[X(t) \mid \mathcal{F}(t-1)] \leq \sigma_t^2$.

Then for all $\varepsilon > 0$,

$$\mathbf{Pr}[X(n) \leq \mathbf{E}[X(n)] - \varepsilon] \leq \exp\left(-\frac{\varepsilon^2}{2\sum_{i=1}^{n}(a_t^2 + \sigma_t^2)}\right).$$

Corollary 6.5. *Let* $(X(t))_{t=0}^n$ *be a martingale associated with the filter* $(\mathcal{F}(t))_{t=0}^n$*, where*

$$\forall t \in [n], |X(t) - X(t-1)| \leq 1.$$

Then with $\langle X \rangle$ *as in Theorem 6.7, for any* $\varepsilon \geq 0$ *and* $\sigma > 0$,

$$\mathbf{Pr}\Big[|X(n) - X(0)| \geq \frac{\varepsilon}{3} + v\sqrt{2\varepsilon}\Big] \leq 2(e^{-\varepsilon} + \mathbf{Pr}\big[\langle X \rangle_n > v^2\big]).$$

Proof. As $(X(t))_{t=0}^n$ is a martingale, it is also a supermartingale, and it fulfills the conditions of Theorem 6.7 by the assumptions of the claim. So way may use Theorem 6.7 to see that

$$\begin{aligned}\mathbf{Pr}\Big[X(n) - X(0) \geq \frac{\varepsilon}{3} + \sigma\sqrt{2\varepsilon} \wedge \langle X \rangle_n \leq \sigma^2\Big]&\\ \leq \mathbf{Pr}\Big[\exists t \in [n] : X(t) - X(0) \geq \frac{\varepsilon}{3} + \sigma\sqrt{2\varepsilon} \wedge \langle X \rangle_t \leq \sigma^2\Big] &\leq e^{-\varepsilon}.\end{aligned}$$

As $\mathbf{Pr}[A] \leq \mathbf{Pr}[(A \wedge B) \vee B] \leq \mathbf{Pr}[A \wedge B] + \mathbf{Pr}[B]$, this implies that

$$\mathbf{Pr}\Big[X(n) - X(0) \geq \frac{\varepsilon}{3} + \sigma\sqrt{2\varepsilon}\Big] \leq e^{-\varepsilon} + \mathbf{Pr}\big[\langle X \rangle_n \leq \sigma^2\big].$$

The claim follows from applying the same argument to the supermartingale $(-X(t))_{t=0}^n$ and a union bound. ☐

Theorem 6.6 (Azuma–Hoeffding inequality Theorem 13.6 in [75])**.** *Let* $(X(t))_{t=0}^n$ *be a martingale associated with the filter* $(\mathcal{F}(t))_{t=0}^n$*, where there exist non-negative sequences* $(a_t)_{t=1}^n$*,* $(b_t)_{t=1}^n$ *and* $(\sigma_t)_{t=1}^n$ *such that for all* $t \in [n]$,

$$-b_t \leq X(t) - X(t-1) \leq a_t.$$

Then for all $\varepsilon > 0$,

$$\mathbf{Pr}[|X(n) - \mathbf{E}[X(n)]| \geq \varepsilon] \leq 2\exp\left(-\frac{2\varepsilon^2}{\sum_{i=1}^{n}(a_t + b_t)^2}\right).$$

Theorem 6.7 (Adapted from Theorem 2.1 and combined with Remark 2.1 and Equation 18 in [44])**.** *Let* $(X(t))_{t=0}^n$ *be a supermartingale associated with the filter* $(\mathcal{F}(t))_{t=0}^n$*, where* $X(t) - X(t-1) \leq 1$ *for all* $t \in [n]$*. Let* $\langle X \rangle$ *be the quadratic characteristic of* X*, i.e., let*

$$\langle X \rangle_0 = 0, \quad \langle X \rangle_t = \sum_{\tau=1}^{t} \mathbf{E}\Big[(X(\tau) - X(\tau-1))^2 \;\Big|\; \mathcal{F}(\tau-1)\Big], \quad \forall t \in [n].$$

Then, for any $\varepsilon \geq 0$ *and* $\sigma > 0$,

$$\mathbf{Pr}\Big[\exists t \in [n] : X(t) - X(0) \geq \frac{\varepsilon}{3} + v\sqrt{2\varepsilon} \wedge \langle X \rangle_t \leq \sigma^2\Big] \leq e^{-\varepsilon}.$$

Theorem 6.8 (Berry-Esseen Theorem [28, 43] for Non-identical Random Variables). *Let $Y_1, Y_2, \cdots, Y_k$ be independently distributed with $\mathbf{E}[Y_i] = 0$, $\mathbf{E}\left[Y_i^2\right] = Var[Y_i] = \sigma_i^2$ and $\mathbf{E}\left[|Y_i|^3\right] = \rho_i < \infty$. If $F_k(x)$ is the distribution of $\frac{Y_1+Y_2+\cdots+Y_k}{\sqrt{\sigma_1^2+\sigma_2^2+\cdots+\sigma_k^2}}$ and $\Phi_N(x)$ is the standard normal distribution, then*

$$|F_k(x) - \Phi_N(x)| \leq C_0 \cdot \psi_0,$$

where $\psi_0 = \frac{\sum_{i=1}^k \rho_i}{\left(\sum_{i=1}^k \sigma_i^2\right)^{3/2}}$ and C_0 is a constant.

Theorem 6.9 (Theorem 3.4 in [35], [72]). *let X_i $(1 \leq i \leq n)$ be independent random variables satisfying $X_i \leq \mathbf{E}[X_i] + M$, for $1 \leq i \leq n$. We consider the sum $X = \sum_{i=1}^n X_i$ with expectation $\mathbf{E}[X] = \sum_{i=1}^n \mathbf{E}[X_i]$ and variance $Var[X] = \sum_{i=1}^n Var[X_i]$. Then for any $\delta > 0$ we have*

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

Theorem 6.10 (Theorem 4.1 in [35]). *Let X_i denote independent random variable satisfying $X_i \geq \mathbf{E}[X_i] - a_i - M$ for $0 \leq i \leq n$. For $X = \sum_{i=1}^n X_i$ and any $\delta > 0$ we have*

$$\Pr[X \leq \mathbf{E}[X] - \delta] \leq \exp\left(-\frac{\delta^2}{2 \cdot (Var[X] + \sum_{i=1}^n a_i^2 + M\delta/3)}\right).$$

Theorem 6.11 (Theorem 3.4 in [35]). *Let Y_i denote independent random variables satisfying $Y_i \leq \mathbf{E}[Y_i] + a_i + M$, for $1 \leq i \leq n$. For $Y = \sum_{i=1}^n Y_i$and any $\delta > 0$, we have*

$$\Pr[Y \geq \mathbf{E}[Y] + \delta] \leq \exp\left(-\frac{\delta^2}{2(Var[Y] + \sum_{i=1}^n a_i^2 + M\delta/3)}\right).$$

Theorem 6.12 (Theorem 6.1 in [35]). *Let X_i be the martingale associated with a filter F satisfying*

1. *$Var[X_i \mid F_{i-1})] \leq \sigma^2$ for $i \in [n]$,*
2. *$|X_i - X_{i-1}| \leq M$ for $i \in [n]$.*

Then, we have

$$\Pr[X_n - \mathbf{E}[X_n] \geq \delta] \leq \exp\left(-\frac{\delta^2}{2(\sum_{i=1}^n \sigma^2 + M\delta/3)}\right).$$

Theorem 6.13 (Theorem 6.5 in [35]). *Let X_i be the martingale associated with a filter F satisfying*

1. *$Var[X_i \mid F_{i-1})] \leq \sigma^2$ for $i \in [n]$,*
2. *$X_{i-1} - X_i \leq a_i + M$ for $i \in [n]$.*

Then, we have

$$\Pr[X_n - \mathbf{E}[X_n] \leq -\delta] \leq \exp\left(-\frac{\delta^2}{2(\sum_{i=1}^n (\sigma^2 + a_i^2) + M\delta/3)}\right).$$

Theorem 6.14 ([77, page 92, Theorem 4.16, Azuma's Inequality]). *Let $X_0, X_1, \ldots$ 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 $\delta > 0$,*

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

Sum of $\{0,1\}$ random variables which satisfy negative association condition has Chernoff-like tail concentration.

Lemma 6.15 (Proposition 29 in [42]). *The Chernoff-Hoeffding bounds apply to sum of variables that satisfy the negative association condition. In other words, let $Y_1, Y_2, \cdots, Y_n$ be a sequence of $\{0,1\}$ random variables satisfying the negative association condition and $Y = \sum_{i=1}^{n} Y_i$ and $\mu = \mathbf{E}[Y]$. Then it holds,*

1. *For all $\delta > 0$ that*
$$\Pr[Y \geq (1+\delta)\mu] \leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu} \leq \exp\left(-\frac{\delta^2 \cdot \mu}{2+\delta}\right),$$
2. *For all $1 > \delta > 0$ that*
$$\Pr[Y \leq (1-\delta)\mu] \leq \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{\mu} \leq \exp\left(-\frac{\delta^2 \cdot \mu}{2}\right).$$

6.3 Random Walks, Hitting Times and Effective Resistance

In this subsection we present for completeness fundamental definitions and relations concerning random walks, hitting times, and the effective resistance. We start with a definition of the effective resistance of a network in Definition 6.1. For a motivation of the definition see [64, Chapter 9]. Further details and properties can also be found in [40] and [67, Section 4].

Definition 6.1 (Harmonic Functions and Effective Resistance). *Let G be a graph and let $i, j \in [n]$ be nodes of the graph. Then a* harmonic function on G with the poles i and j *(for unit edge weights) is a function $f : [n] \to \mathbb{R}$ such that for all $k \in [n] \setminus \{i, j\}$ we have $f(k) = \frac{1}{d(k)} \cdot \sum_{l \in N_G(k)} f(l)$, where $N_G(k)$ is the set of k's neighbors in G. Given a harmonic function f on G with the poles i and j (with arbitrary boundary values $f(i) \neq f(j)$), the* effective resistance *(or* resistive distance *between i and j in G is given by*

$$\mathrm{Res}(i,j) := \frac{f(i) - f(j)}{\sum_{k \in N_G(i)} |f(k) - f(i)|}.$$

Note that the value is not dependent on the boundary values of the harmonic function.

Note that for boundary values $f(i)$ and $f(j)$ the harmonic function is unique [64, Proposition 9.1].

The following is a well-known property of effective resistances; it is a direct consequence of, e.g., Corollary 9.13 in [64].

Lemma 6.16. *Let G be a graph, and write $\mathrm{d}(i,j)$ for the (standard) distance between i and j in G. Then $\mathrm{Res}(i,j) \leq d(i,j)$.*

For a graph G, and nodes $i, j \in V(G)$, let $H(i,j)$ be the *hitting time from i to j*, i.e., the expected time for a random walk on G starting at i to reach j for the first time.

Theorem 6.17 (Theorem 4.1 (i) in [67]). *Let G be a graph. Then for any $i, j \in V(G)$,*

$$H(i,j) + H(j,i) = 2 \cdot |E| \cdot \mathrm{Res}(i,j).$$

Corollary 6.18. *Let G be a graph. Then for any $i, j \in V(G)$,*

$$\max\{H(i,j), H(j,i)\} \leq 2 \cdot |E(G)| \cdot \mathrm{Res}(i,j) \leq 2 \cdot \max\{H(i,j), H(j,i)\}.$$

Proof. For the first inequality, since one of $H(i,j)$ and $H(j,i)$ is at least the maximum of the two, we have, by Theorem 6.17:

$$\max\{H(i,j), H(j,i)\} \leq H(i,j) + H(j,i) = 2 \cdot |E(G)| \cdot \mathrm{Res}(i,j).$$

And for the second inequality, since both $H(i,j)$ and $H(j,i)$ are at most the maximum of the two, we have, again by Theorem 6.17

$$2 \cdot |E(G)| \cdot \mathrm{Res}(i,j) = H(i,j) + H(j,i) \leq 2 \cdot \max\{H(i,j), H(j,i)\},$$

as claimed. □

Theorem 6.19 (Dirichlet's principle, see Exercise 2.13 in [69]; or Exercise 9.9 in [64], referencing Theorem 6.1 in [65]). *Let u, v be distinct nodes of a graph G. Then*

$$\min_{\substack{\vec{a} \in \mathbb{R}^n \\ a_v = 1 \\ a_u = 0}} \Psi_G(\vec{a}) = \frac{1}{\mathrm{Res}(u,v)}.$$

Theorem 6.20 (Corollary 3.3 in [67], applied to d-regular graphs). *Let G be an arbitrary graph on n nodes. Then*

$$n \leq H(i,j) + H(j,i) \leq \frac{n}{\lambda(\mathbf{L}(G))}.$$

The following lemma is well-known; we state it for completeness. It relates the hitting time of a graph G to its resistive diameter and the edge hitting time of G to the $\mathrm{Res}^*(G)$.

Lemma 6.21. *For any graph $G = (V, E)$*

1. $\mathrm{Res}^*(G) \cdot |E| \leq \mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) \leq 2 \cdot \mathrm{Res}^*(G) \cdot |E|$, *and*
2. $\mathrm{Res}(G) \cdot |E| \leq \mathrm{t}_{\mathrm{hit}}(\mathrm{G}) \leq 2 \cdot \mathrm{Res}(G) \cdot |E|$.

Proof. Recall that

$$\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) \coloneqq \max_{i,j \in V, \{i,j\} \in E} H(i,j),$$

and that

$$\mathrm{Res}^*(G) \coloneqq \max_{i,j \in V, \{i,j\} \in E} \mathrm{Res}(i,j).$$

For the first inequality, let $i, j \in V$ be adjacent nodes for which $\mathrm{Res}(i,j) = \mathrm{Res}^*(G)$. Then, by Corollary 6.18,

$$2 \cdot |E| \cdot \mathrm{Res}^*(G) \leq 2 \cdot |E| \cdot \mathrm{Res}(i,j) \leq 2 \cdot \max H(i,j), H(j,i) \leq 2 \cdot \mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}),$$

which becomes the first inequality after dividing by 2 on both sides. For the second inequality, let $i, j \in V$ be adjacent nodes for which $\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) = H(i,j)$. Then, again by Corollary 6.18,

$$\mathrm{t}^*_{\mathrm{hit}}(\mathrm{G}) = H(i,j) \leq 2 \cdot |E| \cdot \mathrm{Res}(i,j) \leq 2 \cdot |E| \cdot \mathrm{Res}^*(G).$$

The second statement is entirely analogous, except that the $i, j \in V$ are no longer required to be adjacent, and that they are chosen such that $\mathrm{Res}(i,j) = \mathrm{Res}(G)$ for the first inequality, or, for the second inequality, that $H(i,j) = \mathrm{t}_{\mathrm{hit}}(\mathrm{G})$. □

6.4 Basic Results for Load Balancing

In this subsection we mention the useful lemmas from related works which are used in our analysis of the first part. Recall that $\Phi\left(\mathbf{m}_{k,\cdot}^{[\tau,t]}\right) = \left|\left|\mathbf{m}_{k,\cdot}^{[\tau,t]} - \frac{\vec{1}}{n}\right|\right|_2^2$ and $\mathbf{R} := \mathbf{m}^{[1,\Delta]}$ is the round matrix in the balancing circuit model.

Lemma 6.22 (Lemma 2 in [53])**.** *It holds that* $\Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) \leq (1 - \lambda(\mathbf{R}))^{2\tau}$*. More generally,*

$$\Phi\left(\mathbf{R}_{k,\cdot}^{[1,t+\tau]}\right) \leq (1 - \lambda(\mathbf{R}))^{2\tau} \cdot \Phi\left(\mathbf{R}_{k,\cdot}^{[1,t]}\right).$$

Lemma 6.23 ([82])**.** *Let G be a r-dimensional torus graph with the round matrix* $\mathbf{R}$*. Then it holds,*

$$1 - \lambda(\mathbf{R}) = \theta\left(\frac{1}{n^{2/r}}\right).$$

Lemma 6.24 ([82])**.** *Let G be a cycle graph with the round matrix* $\mathbf{R}$*. Then it holds,*

$$1 - \lambda(\mathbf{R}) = \theta\left(\frac{1}{n^2}\right).$$

Lemma 6.25 ([33])**.** *Let G be a r-dimensional torus graph with the round matrix* $\mathbf{R}$*. Then if graph G is*

1. *a cycle, then it holds* $\Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) = \mathrm{O}\left(\frac{1}{\sqrt{\tau}}\right)$.
2. *a 2-dimensional torus, then it holds that* $\Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) = \mathrm{O}\left(\frac{1}{\tau}\right)$.
3. *a r-dimensional with $r \geq 3$, then it holds* $\Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) \leq \tau^{-r/2}$.
4. *a hypercube, then it holds that* $\Phi\left(\mathbf{R}_{k,\cdot}^{[1,\tau]}\right) \leq 2^{-\tau}$.

Here we list basic results from [85] that we use in our analysis.

Lemma 6.26 ([85, Lemma 2.4])**.** *Consider the balancing circuit model with sequence of matchings* $\mathbf{m}^{[\infty]} := \left(\mathbf{m}^{(s)}\right)_{s=1}^{\infty}$ *with round matrix* $\mathbf{R} := \prod_{s=1}^{\Delta} \mathbf{m}^{(s)}$*. Then for any node $u \in V$ it holds*

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

Lemma 6.27 (cf. [85, Corollary 2.7])**.** *Consider the random matching model with sequence of matchings* $\mathbf{M}^{[\infty]}$*. Then for any node $u \in V$ and any round $t \geq 1$ we have*

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

We remark that in contrast to [85, 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 6.28 ([85, 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 $\mathbf{M}^{[t]}$ is $(K, 1/(2n))$-smoothing for*

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

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

Observation 6.29 ([85, Observation 2.11])**.** *Fix a sequence of matchings $\mathbf{M}^{[\infty]}$. 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 $\widetilde{x}(0)$. Then the following two statements hold:*

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

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

Observation 6.30 (cf. [85, Observation 2.12])**.** *Assume* $\operatorname{disc}(x(0)) = K$*. Fix a sequence of matchings $\mathbf{M}^{[\infty]}$. Then for arbitrary positive integers α and t it holds that*

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

The following observation shows an important relationship.

Observation 6.31 (cf. [85, Lemma 2.2])**.** *Assume the sequence of matchings $\mathbf{M}^{[t]}$ satisfies for all $u \in V$ that $\left\|\mathbf{M}^{[1,t]}_{u,\cdot} - \frac{\vec{1}}{n}\right\|_2^2 \leq \left(\frac{\varepsilon}{2K \cdot n}\right)^2$. Then this sequence is (K, ε)-smoothing.*

We will omit the proof, since the result follows immediately by noting that $\max_{v \in V} \left|\mathbf{M}^{[1,t]}_{u,v} - \frac{1}{n}\right| \leq \left\|\mathbf{M}^{[1,t]}_{u,\cdot} - \frac{\vec{1}}{n}\right\|_2$ and then applying the third statement of [85, Lemma 2.2].

The next three lemmas are simple and known results.

Observation 6.32. *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

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

□

Observation 6.33. *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.

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

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

The next observation is a well-known statement and in fact, it is implied by Equation (3.41).

Observation 6.34. *For each node* $w \in V$ *the expression* $\left\|\mathbf{M}^{[1,t]}_{w,\cdot} - \frac{\vec{1}}{n}\right\|_2^2$ *is non-increasing over* t.

Proof. The observation follows from the fact that $\mathbf{1}$ is an eigenvector of any matching matrix with eigenvalue 1, and all its eigenvalues are between $[-1, 1]$. Here we give a more intuitive proof, which uses the balancing process. Let u, v be nodes with $[u : v] \in \mathbf{M}^{(t+1)}$. From the balancing process we have,

$$\mathbf{M}^{[1,t+1]}_{w,u} = \mathbf{M}^{[1,t+1]}_{w,v} = \frac{\mathbf{M}^{[1,t]}_{w,u} + \mathbf{M}^{[1,t]}_{w,v}}{2}.$$

Note that if a node $u \in V$ is not matched in $\mathbf{M}^{(t+1)}$, then

$$\mathbf{M}^{[1,t+1]}_{w,u} = \mathbf{M}^{[1,t]}_{w,u}.$$

We have

$$\begin{aligned}
&\left\|\mathbf{M}^{[1,t+1]}_{w,\cdot} - \frac{\vec{1}}{n}\right\|_2^2 - \left\|\mathbf{M}^{[1,t]}_{w,\cdot} - \frac{\vec{1}}{n}\right\|_2^2 \\
&= \sum_{[u:v]\in\mathbf{M}^{(t+1)}} 2\cdot\left(\frac{\mathbf{M}^{[1,t]}_{w,u} + \mathbf{M}^{[1,t]}_{w,v}}{2} - \frac{1}{n}\right)^2 - \left(\mathbf{M}^{[1,t]}_{w,u} - \frac{1}{n}\right)^2 - \left(\mathbf{M}^{[1,t]}_{w,v} - \frac{1}{n}\right)^2 \\
&= \sum_{[u:v]\in\mathbf{M}^{(t+1)}} \frac{1}{2}\left(\left(\mathbf{M}^{[1,t]}_{w,u} - \frac{1}{n}\right) + \left(\mathbf{M}^{[1,t]}_{w,v} - \frac{1}{n}\right)\right)^2 - \left(\mathbf{M}^{[1,t]}_{w,u} - \frac{1}{n}\right)^2 - \left(\mathbf{M}^{[1,t]}_{w,v} - \frac{1}{n}\right)^2 \\
&\overset{(a)}{=} -\frac{1}{2}\sum_{[u:v]\in\mathbf{M}^{(t+1)}}\left(\mathbf{M}^{[1,t]}_{w,u} - \mathbf{M}^{[1,t]}_{w,v}\right)^2 \le 0,
\end{aligned}$$

where in (a) we use the fact that $(a+b)^2/2 - a^2 - b^2 = -(a-b)^2/2$. □

The next lemma is elementary.

Lemma 6.35. *The following statements hold.*

- *For any $p, q \in [0,1]$ and any even integer k we have,*

$$\left(\frac{p+q}{2}\right)^k \geq p^{k/2} \cdot q^{k/2}.$$

- *For any $p, q \in [0,1]$ and any odd integer k we have,*

$$\left(\frac{p+q}{2}\right)^k \geq \frac{1}{2}p^{\lceil k/2 \rceil} \cdot q^{\lfloor k/2 \rfloor} + \frac{1}{2}p^{\lfloor k/2 \rfloor} \cdot q^{\lceil k/2 \rceil}.$$

Proof. For even k, this is equivalent to the arithmetic-geometric mean inequality.

Consider now the case of odd $k = 2\ell + 1$ for even ℓ. Starting from the arithmetic-geometric mean inequality, we have

$$\left(\frac{p+q}{2}\right) \geq \sqrt{p \cdot q},$$

and raising both sides to the power of 2ℓ,

$$\left(\frac{p+q}{2}\right)^{2\ell} \geq p^\ell q^\ell.$$

Now multiplying both sides by $\frac{p+q}{2}$,

$$\left(\frac{p+q}{2}\right)^{2\ell+1} \geq p^\ell q^\ell \cdot \left(\frac{p+q}{2}\right) = \frac{1}{2}p^{\ell+1}q^\ell + \frac{1}{2}p^\ell q^{\ell+1}.$$

□

Lemma 6.36 (Proposition 13.2.6 in [84])**.** *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*,*

$$\mathbf{E}\Big[X \cdot Y \;\Big|\; \mathfrak{F}\Big] = X \cdot \mathbf{E}\Big[Y \;\Big|\; \mathfrak{F}\Big].$$

Lemma 6.37 (Lemma 2 in [42])**.** *Let $X_1, \cdots, X_n$ satisfy the negative association condition. Then for any non-decreasing function $f_i, i \in [n]$,*

$$\mathbf{E}\left[\prod_{i \in [n]} f_i(X_i)\right] \leq \prod_{i \in [n]} \mathbf{E}[f_i(X_i).]$$

References

[1] M. Abramowitz and I. A. Stegun: *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, ninth Dover printing, tenth GPO printing. Dover, 1964.

[2] H. Ackermann, S. Fischer, M. Hoefer, and M. Schöngens: *Distributed algorithms for QoS load balancing*. In *Distributed Comput.* volume 23 (5-6), 2011, pages 321–330. DOI: 10.1007/s00446-010-0125-1.

[3] W. Aiello, B. Awerbuch, B. M. Maggs, and S. Rao: *Approximate load balancing on dynamic and asynchronous networks*. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, STOC*. ACM, 1993, pages 632–641. DOI: 10.1145/167088.167250.

[4] H. Akbari, P. Berenbrink, and T. Sauerwald: *A simple approach for adapting continuous load balancing processes to discrete settings*. In *Distributed Comput.* volume 29 (2), 2016, pages 143–161.

[5] S. G. Aksoy, F. Chung, M. Tait, and J. Tobin: *The maximum relaxation time of a random walk*. In *Adv. Appl. Math.* volume 101, 2018, pages 1–14. DOI: 10.1016/j.aam.2018.07.002.

[6] D. Alistarh, G. Nadiradze, and A. Sabour: *Dynamic Averaging Load Balancing on Cycles*. In *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020*. LIPIcs, volume 168, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, pages 7:1–7:16. DOI: 10.4230/LIPIcs.ICALP.2020.7.

[7] A. Anagnostopoulos, A. Kirsch, and E. Upfal: *Load Balancing in Arbitrary Network Topologies with Stochastic Adversarial Input*. In *SIAM Journal on Computing*, volume 34 (3), 2005, pages 616–639. DOI: 10.1137/S0097539703437831.

[8] E. Anshelevich, D. Kempe, and J. M. Kleinberg: *Stability of Load Balancing Algorithms in Dynamic Adversarial Systems*. In *SIAM J. Comput.* volume 37 (5), 2008, pages 1656–1673. DOI: 10.1137/050639272.

[9] G. Bankhamer, P. Berenbrink, F. Biermeier, R. Elsässer, H. Hosseinpour, D. Kaaser, and P. Kling: *Fast Consensus via the Unconstrained Undecided State Dynamics*. 2021. arXiv: 2103.10366 [cs.DC].

[10] G. Bankhamer, P. Berenbrink, F. Biermeier, R. Elsässer, H. Hosseinpour, D. Kaaser, and P. Kling: *Fast Consensus via the Unconstrained Undecided State Dynamics*. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*. SIAM, 2022, pages 3417–3429. DOI: 10.1137/1.9781611977073.135.

[11] G. Bankhamer, P. Berenbrink, F. Biermeier, R. Elsässer, H. Hosseinpour, D. Kaaser, and P. Kling: *Population Protocols for Exact Plurality Consensus – How a small chance of failure helps to eliminate insignificant opinions*. 2024. arXiv: 2402.06471 [cs.DC].

[12] G. Bankhamer, P. Berenbrink, F. Biermeier, R. Elsässer, H. Hosseinpour, D. Kaaser, and P. Kling: *Population Protocols for Exact Plurality Consensus: How a small chance of failure helps to eliminate insignificant opinions*. In *PODC '22: ACM Symposium on Principles of Distributed Computing*. ACM, 2022, pages 224–234. DOI: 10.1145/3519270.3538447.

[13] N. Bansal and O. N. Feldheim: *The power of two choices in graphical allocation*. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*. ACM, 2022, pages 52–63. DOI: 10.1145/3519935.3519995.

[14] P. Berenbrink, C. Cooper, T. Friedetzky, T. Friedrich, and T. Sauerwald: *Randomized diffusion for indivisible loads*. In *J. Comput. Syst. Sci.* volume 81 (1), 2015, pages 159–185. DOI: 10.1016/J.JCSS.2014.04.027.

[15] P. Berenbrink, A. Czumaj, A. Steger, and B. Vöcking: *Balanced Allocations: The Heavily Loaded Case*. In *SIAM J. Comput.* volume 35 (6), 2006, pages 1350–1385. DOI: 10.1137/S009753970444435X.

[16] P. Berenbrink, R. Elsässer, T. Friedetzky, H. Hosseinpour, D. Kaaser, P. Kling, and T. Sauerwald: *(Almost) Perfect Discrete Iterative Load Balancing*. 2025. arXiv: 2510.15473 [cs.DC].

[17] P. Berenbrink, R. Elsässer, T. Friedetzky, H. Hosseinpour, D. Kaaser, P. Kling, and T. Sauerwald: *(Almost) Perfect Discrete Iterative Load Balancing*. In, Jan. 2026, pages 4224–4237. DOI: 10.1137/1.9781611978971.156.

[18] P. Berenbrink, T. Friedetzky, and L. A. Goldberg: *The Natural Work-Stealing Algorithm is Stable*. In *SIAM J. Comput.* volume 32 (5), 2003, pages 1260–1279.

[19] P. Berenbrink, T. Friedetzky, and Z. Hu: *A new analytical method for parallel, diffusion-type load balancing*. In *J. Parallel Distributed Comput.* volume 69 (1), 2009, pages 54–61. DOI: 10.1016/J.JPDC.2008.05.005.

[20] P. Berenbrink, T. Friedetzky, D. Kaaser, and P. Kling: *Tight & Simple Load Balancing*. In *2019 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2019, Rio de Janeiro, Brazil, May 20-24, 2019*. IEEE, 2019, pages 718–726. DOI: 10.1109/IPDPS.2019.00080.

[21] 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.

[22] P. Berenbrink, T. Friedetzky, and R. A. Martin: *On the Stability of Dynamic Diffusion Load Balancing*. In *Algorithmica*, volume 50 (3), 2008, pages 329–350. DOI: 10.1007/S00453-007-9081-Y.

[23] P. Berenbrink, L. Hintze, H. Hosseinpour, D. Kaaser, and M. Rau: *Dynamic Averaging Load Balancing on Arbitrary Graphs*. 2023. arXiv: 2302.12201 [cs.DC].

[24] 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*. LIPIcs, volume 261, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, pages 18:1–18:18. DOI: 10.4230/LIPICS.ICALP.2023.18.

[25] P. Berenbrink, M. Hoefer, and T. Sauerwald: *Distributed Selfish Load Balancing on Networks*. In *ACM Trans. Algorithms*, volume 11 (1), 2014, pages 2:1–2:29. DOI: 10.1145/2629671.

[26] P. Berenbrink, R. Klasing, A. Kosowski, F. Mallmann-Trenn, and P. Uznanski: *Improved Analysis of Deterministic Load-Balancing Schemes*. In *ACM Trans. Algorithms*, volume 15 (1), 2019, pages 10:1–10:22. DOI: 10.1145/3282435.

[27] P. Berenbrink, P. Kling, C. Liaw, and A. Mehrabian: *Tight Load Balancing Via Randomized Local Search*. In *2017 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2017*. IEEE Computer Society, 2017, pages 192–201. DOI: 10.1109/IPDPS.2017.52.

[28] A. C. Berry: *The Accuracy of the Gaussian Approximation to the Sum of Independent Variates*. In *Transactions of the American Mathematical Society*, volume 49 (1), 1941, pages 122–136.

[29] D. P. Bertsekas and J. N. Tsitsiklis: *Parallel and Distributed Computation: Numerical Methods*. Prentice Hall, 1989.

[30] R. Bhatia and C. Davis: *A Better Bound on the Variance*. In *Am. Math. Mon.* volume 107 (4), 2000, pages 353–357.

[31] J. E. Boillat: *Load Balancing and Poisson Equation in a Graph.* In *Concurr. Pract. Exp.* volume 2 (4), 1990, pages 289–314. DOI: 10.1002/CPE.4330020403.

[32] S. P. Boyd, A. Ghosh, B. Prabhakar, and D. Shah: *Randomized gossip algorithms.* In *IEEE Trans. Inf. Theory*, volume 52 (6), 2006, pages 2508–2530. DOI: 10.1109/TIT.2006.874516.

[33] L. Cai and T. Sauerwald: *Randomized Load Balancing on Networks with Stochastic Inputs.* In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017.* LIPIcs, volume 80, 2017, pages 139:1–139:14. DOI: 10.4230/LIPIcs.ICALP.2017.139.

[34] A. K. Chandra, P. Raghavan, W. L. Ruzzo, R. Smolensky, and P. Tiwari: *The Electrical Resistance of a Graph Captures its Commute and Cover Times.* In *Comput. Complex.* volume 6 (4), 1997, pages 312–340. DOI: 10.1007/BF01270385.

[35] F. R. K. Chung and L. Lu: *Survey: Concentration Inequalities and Martingale Inequalities: A Survey.* In *Internet Math.* volume 3 (1), 2006, pages 79–127. DOI: 10.1080/15427951.2006.10129115.

[36] J. N. Cooper, B. Doerr, T. Friedrich, and J. Spencer: *Deterministic random walks on regular trees.* In *Random Struct. Algorithms*, volume 37 (3), 2010, pages 353–366. DOI: 10.1002/RSA.20314.

[37] J. N. Cooper and J. Spencer: *Simulating a Random Walk with Constant Error.* In *Comb. Probab. Comput.* volume 15 (6), 2006, pages 815–822. DOI: 10.1017/S0963548306007565.

[38] G. Cybenko: *Dynamic Load Balancing for Distributed Memory Multiprocessors.* In *J. Parallel Distributed Comput.* volume 7 (2), 1989, pages 279–301. DOI: 10.1016/0743-7315(89)90021-X.

[39] R. Diekmann, A. Frommer, and B. Monien: *Efficient schemes for nearest neighbor load balancing.* In *Parallel Comput.* volume 25 (7), 1999, pages 789–812. DOI: 10.1016/S0167-8191(99)00018-6.

[40] P. G. Doyle and J. L. Snell: *Random Walks and Electric Networks.* Carus Mathematical Monographs. Mathematical Association of America, 1984.

[41] D. P. Dubhashi and D. Ranjan: *Balls and bins: A study in negative dependence.* In *Random Struct. Algorithms*, volume 13 (2), 1998, pages 99–124. DOI: 10.1002/(SICI)1098-2418(199809)13:2\<99::AID-RSA1\>3.0.CO;2-M.

[42] D. Dubhashi and D. Ranjan: *Balls and Bins: A Study in Negative Dependence.* In *Random Struct. Algorithms*, volume 13 (2), Sept. 1998, pages 99–124.

[43] C.-G. Esseen: *On the Liapounoff Limit of Error in the Theory of Probability.* Arkiv för matematik, astronomi och fysik. Almqvist & Wiksell, 1942.

[44] X. Fan, I. Grama, and Q. Liu: *Hoeffding's Inequality for Supermartingales.* In *Stochastic Processes and their Applications*, volume 122 (10), Oct. 2012, pages 3545–3559. DOI: 10.1016/j.spa.2012.06.009.

[45] S. Fischer, H. Räcke, and B. Vöcking: *Fast Convergence to Wardrop Equilibria by Adaptive Sampling Methods.* In *SIAM J. Comput.* volume 39 (8), 2010, pages 3700–3735. DOI: 10.1137/090746720.

[46] T. Friedrich, M. Gairing, and T. Sauerwald: *Quasirandom Load Balancing.* In *SIAM J. Comput.* volume 41 (4), 2012, pages 747–771. DOI: 10.1137/100799216.

[47] 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 2009, Bethesda, MD, USA, May 31 - June 2, 2009.* ACM, 2009, pages 121–130. DOI: 10.1145/1536414.1536433.

[48] J. Gehrke, C. G. Plaxton, and R. Rajaraman: *Rapid Convergence of a Local Load Balancing Algorithm for Asynchronous Rings.* In *Theor. Comput. Sci.* volume 220 (1), 1999, pages 247–265. DOI: 10.1016/S0304-3975(98)00243-6.

[49] B. Ghosh, F. T. Leighton, B. M. Maggs, S. Muthukrishnan, C. G. Plaxton, R. Rajaraman, A. W. Richa, R. E. Tarjan, and D. Zuckerman: *Tight analyses of two local load balancing algorithms*. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing*. STOC '95. Association for Computing Machinery, 1995, pages 548–558. DOI: 10.1145/225058.225272.

[50] B. Ghosh, F. T. Leighton, B. M. Maggs, S. Muthukrishnan, C. G. Plaxton, R. Rajaraman, A. W. Richa, R. E. Tarjan, and D. Zuckerman: *Tight Analyses of Two Local Load Balancing Algorithms*. In *SIAM J. Comput.* volume 29 (1), 1999, pages 29–64. DOI: 10.1137/S0097539795292208.

[51] B. Ghosh and S. Muthukrishnan: *Dynamic Load Balancing by Random Matchings*. In *J. Comput. Syst. Sci.* volume 53 (3), 1996, pages 357–370. DOI: 10.1006/jcss.1996.0075.

[52] B. Ghosh and S. Muthukrishnan: *Dynamic Load Balancing by Random Matchings*. In *Journal of Computer and System Sciences*, volume 53 (3), 1996, pages 357–370. DOI: https://doi.org/10.1006/jcss.1996.0075.

[53] B. Ghosh, S. Muthukrishnan, and M. H. Schultz: *First and Second Order Diffusive Methods for Rapid, Coarse, Distributed Load Balancing (Extended Abstract)*. In *Proceedings of the 8th Annual ACM Symposium on Parallel Algorithms and Architectures, SPAA '96*. ACM, 1996, pages 72–81. DOI: 10.1145/237502.237509.

[54] V. Guruswami: *Rapidly Mixing Markov Chains: A Comparison of Techniques (A Survey)*. 2016. arXiv: 1603.01512 [cs.DS].

[55] F. M. auf der Heide, B. Oesterdiekhoff, and R. Wanka: *Strongly Adaptive Token Distribution*. In *Algorithmica*, volume 15 (5), 1996, pages 413–427. DOI: 10.1007/BF01955042.

[56] M. Herlihy and S. Tirthapura: *Randomized smoothing networks*. In *J. Parallel Distributed Comput.* volume 66 (5), 2006, pages 626–632. DOI: 10.1016/J.JPDC.2005.06.009.

[57] M. Hoefer and T. Sauerwald: *Threshold Load Balancing in Networks*. In *CoRR*, volume abs/1306.1402, 2013. arXiv: 1306.1402.

[58] R. M. Karp, M. Luby, and F. M. auf der Heide: *Efficient PRAM Simulation on a Distributed Memory Machine*. In *Algorithmica*, volume 16 (4/5), 1996, pages 517–542. DOI: 10.1007/BF01940878.

[59] J. Keilson: *Markov Chain Models — Rarity and Exponentiality*. Springer New York, 1979. DOI: 10.1007/978-1-4612-6200-8_1.

[60] 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, pages 482–491. DOI: 10.1109/SFCS.2003.1238221.

[61] A. Klenke: *Probability theory - a comprehensive course*. Universitext. Springer, 2008.

[62] H. Landau and A. Odlyzko: *Bounds for eigenvalues of certain stochastic matrices*. In *Linear Algebra and its Applications*, volume 38, 1981, pages 5–15. DOI: 10.1016/0024-3795(81)90003-3.

[63] J. Lengler: *Drift Analysis*. In *Theory of Evolutionary Computation - Recent Developments in Discrete Optimization*. Natural Computing Series. Springer, 2020, pages 89–131. DOI: 10.1007/978-3-030-29414-4_2.

[64] D. Levin and Y. Peres: *Markov Chains and Mixing Times*. AMS, 2017. DOI: 10.1090/mbk/107.

[65] T. M. Liggett: *Interacting Particle Systems*. Springer, 1985. DOI: 10.1007/b138374.

[66] D. Los, T. Sauerwald, and J. Sylvester: *Balanced Allocations with Heterogeneous Bins: The Power of Memory*. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*. SIAM, 2023, pages 4448–4477. DOI: 10.1137/1.9781611977554.CH169.

[67] L. Lovász: *Random walks on graphs*. In *Combinatorics, Paul Erdős is Eighty*, volume 2, 1993, pages 1–46.

[68] L. Lovász and P. Winkler: *Mixing of random walks and other diffusions on a graph*. In *Surveys in Combinatorics*, volume 218, London Mathematical Society Lecture Note Series. Cambridge University Press, 1995, pages 119–154. DOI: 10.1017/CBO9780511662096.007.

[69] R. Lyons and Y. Peres: *Probability on Trees and Networks*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2017. DOI: 10.1017/9781316672815.

[70] 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*. ACM, 2009, pages 178–187. DOI: 10.1145/1583991.1584043.

[71] M. Mavronicolas and T. Sauerwald: *The impact of randomization in smoothing networks*. In *Distributed Comput.* volume 22 (5-6), 2010, pages 381–411. DOI: 10.1007/S00446-009-0087-3.

[72] C. McDiarmid: *Concentration*. In *Probabilistic Methods for Algorithmic Discrete Mathematics*. Springer Berlin Heidelberg, 1998, pages 195–248. DOI: 10.1007/978-3-662-12788-9_6.

[73] H. Meyerhenke: *Shape optimizing load balancing for MPI-parallel adaptive numerical simulations*. In *Graph Partitioning and Graph Clustering, 10th DIMACS Implementation Challenge Workshop*. Contemporary Mathematics, volume 588, American Mathematical Society, 2012, pages 67–82.

[74] M. Mitzenmacher, B. Prabhakar, and D. Shah: *Load Balancing with Memory*. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*. IEEE Computer Society, 2002, pages 799–808. DOI: 10.1109/SFCS.2002.1182005.

[75] M. Mitzenmacher and E. Upfal: *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005. DOI: 10.1017/CBO9780511813603.

[76] V. Mohammadian, N. J. Navimipour, M. Hosseinzadeh, and A. M. Darwesh: *Fault-Tolerant Load Balancing in Cloud Computing: A Systematic Literature Review*. In *IEEE Access*, volume 10, 2022, pages 12714–12731. DOI: 10.1109/ACCESS.2021.3139730.

[77] R. Motwani and P. Raghavan: *Randomized Algorithms*. Cambridge University Press, 1995. DOI: 10.1017/CBO9780511814075.

[78] S. Muthukrishnan, B. Ghosh, and M. H. Schultz: *First- and Second-Order Diffusive Methods for Rapid, Coarse, Distributed Load Balancing*. In *Theory Comput. Syst.* volume 31 (4), 1998, pages 331–354. DOI: 10.1007/s002240000092.

[79] B. Patzák and D. Rypl: *Object-oriented, parallel finite element framework with dynamic load balancing*. In *Adv. Eng. Softw.* volume 47 (1), 2012, pages 35–50. DOI: 10.1016/j.advengsoft.2011.12.008.

[80] D. Peleg and E. Upfal: *The Token Distribution Problem*. In *SIAM J. Comput.* volume 18 (2), 1989, pages 229–243. DOI: 10.1137/0218015.

[81] Y. Peres, K. Talwar, and U. Wieder: *Graphical balanced allocations and the (1 + β)-choice process*. In *Random Struct. Algorithms*, volume 47 (4), 2015, pages 760–775. DOI: 10.1002/rsa.20558.

[82] 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 '98*. IEEE Computer Society, 1998, pages 694–705. DOI: 10.1109/SFCS.1998.743520.

[83] Y. Rabani, A. Sinclair, and R. Wanka: *Local Divergence of Markov Chains and the Analysis of Iterative Load-Balancing Schemes*. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science*. FOCS '98. IEEE Computer Society, 1998, pages 694.

[84] J. S. Rosenthal: *A First Look at Rigorous Probability Theory*, Second Edition. World Scientific, 2006. DOI: 10.1142/6300.

[85] 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*. Full version at https://arxiv.org/abs/1201.2715. IEEE Computer Society, 2012, pages 341–350. DOI: 10.1109/FOCS.2012.86.

[86] 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*. ACM, 1994, pages 220–225. DOI: 10.1145/181014.181361.

[87] G. Zheng, A. Bhatele, E. Meneses, and L. V. Kalé: *Periodic hierarchical load balancing for large supercomputers*. In *Int. J. High Perform. Comput. Appl.* volume 25 (4), 2011, pages 371–385. DOI: 10.1177/1094342010394383.