



HAL
open science

Ergodic Effects in Token Circulation

Adrian Kosowski, Przemyslaw Uznanski

► **To cite this version:**

Adrian Kosowski, Przemyslaw Uznanski. Ergodic Effects in Token Circulation. SODA '18 - Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, Jan 2018, New Orleans, United States. pp.2668-2682. hal-01963249

HAL Id: hal-01963249

<https://inria.hal.science/hal-01963249v1>

Submitted on 21 Dec 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Ergodic Effects in Token Circulation

Adrian Kosowski¹ and Przemysław Uznański²

¹Inria Paris and IRIF, Université Paris Diderot, France

²Department of Computer Science, ETH Zürich, Switzerland

Abstract

We consider a dynamical process in a network which distributes all particles (tokens) located at a node among its neighbors, in a round-robin manner.

We show that in the recurrent state of this dynamics (i.e., disregarding a polynomially long initialization phase of the system), the number of particles located on a given edge, averaged over an interval of time, is tightly concentrated around the average particle density in the system. Formally, for a system of k particles in a graph of m edges, during any interval of length T , this time-averaged value is $k/m \pm \tilde{O}(1/T)$, whenever $\gcd(m, k) = \tilde{O}(1)$ (and so, e.g., whenever m is a prime number). To achieve these bounds, we link the behavior of the studied dynamics to ergodic properties of traversals based on Eulerian circuits on a symmetric directed graph. These results are proved through sum set methods and are likely to be of independent interest.

As a corollary, we also obtain bounds on the *idleness* of the studied dynamics, i.e., on the longest possible time between two consecutive appearances of a token on an edge, taken over all edges. Designing trajectories for k tokens in a way which minimizes idleness is fundamental to the study of the patrolling problem in networks. Our results immediately imply a bound of $\tilde{O}(m/k)$ on the idleness of the studied process, showing that it is a distributed $\tilde{O}(1)$ -competitive solution to the patrolling task, for all of the covered cases. Our work also provides some further insights that may be interesting in load-balancing applications.

1 Introduction

The notion of *ergodicity* captures the characteristic of a dynamical system which displays averaging behavior in both its state space and in time. For example, a finite Markov chain is ergodic if it is both aperiodic and irreducible, that is, for any sufficiently large integer t , it has positive probability to go between any pair of states in exactly t steps. In this work, we focus on systems consisting of multiple particles (or tokens) moving around sites in a finite graph. In such a setting, an important manifestation of ergodicity is the property of *time-averaging*: a measurement performed at a fixed location in the system over a period of time is representative of the averaged behavior of the system, for any site or location.

As an example, consider a system of k non-interacting random walk particles on a connected non-bipartite symmetric directed graph, moving around the graph in synchronous steps. A direct consequence of the ergodicity of this system is that, in the limit, any multiset of k arcs has the same probability of being traversed by the k walks in any given step. The slightly weaker property of time-averaging for this system is phrased as follows: if the graph has m arcs, then for any given arc, the average number of particles traversing this arc tends to k/m per step, when averaging over sufficiently long time intervals. However, when the number of

particles k in the system is small, the length of the time window required to obtain a reliable time-averaged measurement at any single measurement point becomes long: if we observe any given arc for $\Theta(m^2/k^2)$ consecutive steps, we have constant probability of *not* spotting any particle on the arc, even when the expected number of particles passing through this arc during the considered time interval is $\Theta(m/k)$ (cf. e.g. [19, 28] for related analyses and more general considerations of parallel random walks).

In this paper, we investigate a way to obtain faster time-averaging behavior in a system with k particles, and look at the algorithmic consequences of such a property. We study the fundamental deterministic process of particle propagation in which k tokens circulate in a symmetric directed graph, with each node propagating all the tokens it contains to its neighbors by a round-robin mechanism (*RR dynamics*), in each synchronous round. This process is closely related to (and in fact, a variant of) chip-firing games in the sandpile model (see Section 1.1 for more details), but has the advantage that it displays meaningful behavior even in the regime of a very small number of tokens. At the same time, it bears numerous similarities to the model of non-interacting parallel random walks. These similarities have been studied and exhibited in contexts ranging from parallel cover time [15, 28, 29] to load-balancing properties [3, 8, 16, 29, 36].

As our main result, we show that under the assumption of co-prime values of m and k , the RR dynamics in its recurrent state displays very fast time-averaging on arcs: in a time interval of length T , the number of tokens passing through any arc is always in the range $kT/m \pm \tilde{O}(1)$, i.e., the average number of tokens observed per step during the interval is $k/m \pm \tilde{O}(1/T)$, where the \tilde{O} notation hides poly-logarithmic factors in m . In particular, in the studied scenario, each arc is visited by a token at least once in $\tilde{O}(m/k)$ steps. This effect of quasi-regular traversal of arcs in the RR dynamics' recurrent state is only observed when the considered graph is symmetric. To prove its existence, we build on links between the RR dynamics and the structure of Eulerian circuits in the graph. Our main contribution is to define and analyze a process of moving around the nodes of a graph in time, by performing traversals of fragments (contiguous sequences of arcs) of a given Eulerian circuit, and occasionally switching between different fragments of the circuit which meet at the same vertex. Such a process is shown to have the property that any pair of temporally and spatially separated states can be reached from one another within at most $\tilde{O}(1)$ operations of switching between fragments of the Eulerian circuit. The latter result is obtained by proving, in a framework of additive combinatorics, that iterated sums of the set of lengths of closed sub-circuits of any Eulerian circuit grow rapidly, covering the set of all integers in a small number of iterations.

The rest of the paper is organized as follows. We continue the introduction by describing the context of the RR dynamics in Section 1.1, and some applications of our results in Section 1.2. We then provide an exposition of the technical results in Section 2. The remaining sections are devoted to a proof of the results, starting with an exposition of the main structural lemmas on mixing properties of circulations in Section 3 and 4. We relate these results to RR dynamics in Section 5.

1.1 Overview of RR Dynamics

The RR dynamics, variously referred to as the *round-robin* mechanism, the *rotor-router*, the *Eulerian walker* model, the *ant walk* model, or the *Propp machine*, describes a deterministic token propagation mechanism on a (symmetric) directed graph.¹ In each step, the state of the system is represented by the locations of the k tokens on nodes of the graph, and a pointer for each node, indicating one outgoing arc. In discrete, synchronous steps, all tokens are propagated according to the deterministic round robin rule, where after sending out each token, the node sending the token advances its pointer to the next outgoing arc in some fixed cyclic ordering of its out-neighborhood.

¹A process defined on a symmetric directed graph may also be considered as running on the corresponding undirected graph. We consider the symmetric directed variant for an easier discussion of (Eulerian) circuits.

The RR dynamics was first considered in [32] and is a close cousin of the dynamics of *sandpiles*. The latter model was introduced 30 years ago in the seminal work [6] (cf. e.g. [34] for an overview of its algorithmic properties) and gave rise to the notion of *self-organized criticality* (SOC) — a concept since applied in dozens of fields, ranging from Earth sciences to neurobiology, to model systems whose evolution occurs close to a critical point. Whereas sandpile dynamics were the first synthetic physical system recognized to display SOC behavior, the RR dynamics has been shown to exhibit comparable effects, visible for instance in the almost identical long-distance correlation patterns of the two processes on the grid, cf. e.g. [14, 32]. In the large- k regime, the RR system and the synchronous sandpile model (with chip-firing rules [9, 27]) are comparable in many other respects. For instance, in a network load-balancing context, they both belong to a class of processes with a highly desirable property of cumulative load balancing over arcs outgoing from each node [8, 12, 16]. In the small- k regime, the RR dynamics has been studied in the context of its limit behavior and has often been compared to parallel random walk processes, e.g., in terms of its cover time on graphs [1, 15, 20, 23, 29, 39].

The RR dynamics, like any other finite deterministic system, exhibits an initial transient phase, after which it stabilizes to a cyclic traversal of a set of recurrent states. For $k = 1$ token, multiple authors have observed that on symmetric graphs, the dynamics stabilizes in a polynomial number of steps to a traversal of an Eulerian circuit [1, 32, 39] (for the non-symmetric and infinite settings, see e.g. [26]). The underlying link between the RR dynamics and Eulerian circuits is also displayed by the relation between the number of possible states of the system and the number of Eulerian circuits in the graph, as given by the BEST theorem (cf. e.g. [7] for further discussion). For $k > 1$, a structural characterization of the limit behavior of RR for an *arbitrary* number $k > 1$ of tokens was shown in [10]. The obtained characterization shows that the RR process provides a way of organizing tokens into balanced subsets, each of which follows a well-defined walk in some closed circuit on a subset of arcs of the graph. When the graph is *not symmetric*, such token subsets may be arbitrary and display little regularity: for example, on a one-directional cycle in which each node has out-degree 1, a group of k tokens may perform a cyclic traversal of the graph while always located on the same node. Herein we show how, and (plausibly) why, in any *symmetric* directed graph the RR dynamics will (under almost all parameter configurations) self-organize into recurrent states with an arrangement of tokens in which distances between tokens on their respective circuits are almost uniform.

1.2 Motivation and Consequences of Results

In Section 2 we provide a formal exposition of the obtained structural results for Eulerian circuits, and their consequences for the RR model. Before we do this, we present here two examples of concrete scenarios in which the results of this paper find direct application. The first concerns the deployment of the RR process in the task of regularly patrolling network edges, while the second makes use of the RR rules to balance load in a network fairly over periods of time. We remark that for neither of these tasks were decentralized solutions of comparable quality previously known, whether deterministic or randomized. In particular, as discussed further, the application of random-walk-based techniques fails completely for both tasks.

1.2.1 Network Patrolling

The term *patrolling* refers to an act of surveillance, which involves walking perpetually around an area in order to protect or supervise it. It is a convenient description of numerous settings, such as locating objects or humans that need to be rescued from a disaster, ecological monitoring, or intrusion detection. Network administrators may patrol network nodes and links to detect network failures or to discover web pages which need to be indexed by search engines, cf. [30]. Patrolling has been recently intensively studied in robotics (cf. [5, 11, 17, 18, 25, 30, 39]) where it is often viewed as a version of terrain *coverage*, a central task in robotics.

Scenarios for RR dynamics	Patrolling competitive ratio	Cumulated load discrepancy	Reference
small $\gcd(k, m)$:	$\tilde{\mathcal{O}}(\gcd(k, m))$	$\tilde{\mathcal{O}}(\gcd(k, m))$	Pro. 5.14, 5.12
large k :	$\mathcal{O}(1)$, for $k \geq \frac{3}{4}m$	–	Pro. 5.14
small k :	$\tilde{\mathcal{O}}(\sqrt{k})$	–	Pro. 5.14
small diam:	$\mathcal{O}(\text{diam})$	$\mathcal{O}(\text{diam})$	Pro. 5.14, 5.12
small n :	$\tilde{\mathcal{O}}(\sqrt{n})$	$\tilde{\mathcal{O}}(\sqrt{n})$	Pro. 5.14, 5.12
tree topology:	$\mathcal{O}(1)$, for trees	–	Pro. 5.14
Comparison of strategies (for m prime)			
RR dynamics:	$\tilde{\mathcal{O}}(1)$	$\tilde{\mathcal{O}}(1)$	
k parallel random walks or randomized rounding [22]:	$\Omega(m/k)$ (also in expected sense)	$\Omega(\sqrt{T})$ in expectation (i.e., unbounded)	

Table 1: Summary of obtained performance bounds on RR dynamics. The underlying graph is assumed to be symmetric directed, on n nodes, m arcs, diameter diam , and with k tokens. The cumulated load discrepancy presented in the table holds eventually over any time interval, of arbitrary length T .

Patrolling boundaries of terrains and their area have been studied in [2, 17, 18, 31] with approaches placing more emphasis on experimental results.

The accepted measure of the algorithmic efficiency of patrolling is called *idleness* or *refresh time* and describes the frequency with which the points of the environment are visited (cf. [5, 11, 17, 18, 30, 31]); this criterion was first introduced in [30]. Depending on the requirements, the idleness may sometimes be viewed as the average ([17]), worst-case ([7, 39]), probabilistic ([2]) or experimentally verified ([30]) time elapsed since the last visit to a node or edge (cf. also [5, 11]). A survey of diverse approaches to patrolling based on idleness criteria can be found in [5], whereas a theoretical analysis of approaches to patrolling in graph-based models can be found in [11].

Contribution: In the standard edge-patrolling setting on graphs, the asymptotically optimal value of (worst-case) idleness which can be achieved using a team of k agents on a graph with m edges is $\Theta(m/k)$. This is achieved, e.g., by deploying k agents in a centralized manner with equal spacing, traversing a predefined Eulerian circuit in the graph (cf. e.g. [11]). We show that the RR dynamics provides a completely decentralized (and self-organized) solution to the edge-patrolling problem on graphs, achieving idleness of $\tilde{\Theta}(m/k)$ in a wide variety of settings, in particular, when $\gcd(k, m) = 1$, as well as for trees. (We remark that some empirical studies of patrolling of trees with the RR process in trees were provided in [13], without theoretical analysis.) In Table 1 we present the obtained results on the *patrolling competitive ratio* of the RR process, i.e., the achieved value of the idleness divided by the optimal $\Theta(m/k)$ idleness of the centralized process. To the best of our knowledge, designing a decentralized patrolling strategy with small idleness was a relevant open question in the area, and the RR approach may prove to be a viable practical approach. As pointed out in the discussion of time-averaging processes at the beginning of the paper, even when considering probabilistic measures of idleness and when the considered graph is a cycle, a team of parallel random walks achieve an idleness of $\Omega(m^2/k^2)$, which corresponds to a competitive ratio of $\Omega(m/k)$.

1.2.2 Load Balancing

Process in which nodes exchange tokens have also been extensively studied from the perspective of *load balancing*, where one sees the number of tokens at a node as its load, and aims at minimizing the *discrepancy*, i.e., the difference between the maximal and minimal load of a pair of nodes of a network.

Many distributed load balancing processes are inspired by continuous dynamics from nature. For example, the *heat equation*, which describes real-world processes such as heat and particle diffusion, can also be used to model the following continuous load balancing process, assuming arbitrarily divisible load instead of indivisible tokens. Each node with load $x(u)$ sends to each of its d neighbours load $x(u)/(d+1)$, and keeps $x(u)/(d+1)$ load to itself. In such *diffusive load balancing*, it is well known that loads in d -regular graphs will eventually become perfectly balanced [37]. However, such balancing might be impossible to achieve in a discrete setting, and in general it is impossible to simulate a continuous process in a discrete one. Discrete load balancing schemes, which in principle deal with integer-rounding issues in a continuous process, are also in general much harder to analyze. In [33], a process that sends $\lfloor x(u)/(d+1) \rfloor$ or $\lceil x(u)/(d+1) \rceil$ load over each edge is shown to achieve vertex discrepancy of $\mathcal{O}(d \log n/\mu)$ after $\mathcal{O}(\log(Kn)/\mu)$ steps for d -regular graphs, where K is the initial load discrepancy and μ is the spectral gap of the transition matrix of the underlying Markov chain. Since [33], many variants of discretization of diffusive process have been proposed, see [4, 8, 21, 22, 35].

Contribution: This paper opens the ground for the analysis of a more fine-grained measure in discrete load balancing, namely, balancing the load passing through a node, cumulated (summed) over intervals of time. Thus, for an arbitrary fixed interval of length T , rather than consider the load $x_t(u)$ of a node u at time t , we look at the *cumulated load* $\sum_{\tau=t}^{t+T-1} x_\tau(u)$, and bound the *discrepancy* of this value from its average over all nodes of the graph. The existence of this measure was previously indicated in our work [8], though we managed to obtain relatively weak bounds for a wider family of processes. Here, we show that for the specific load balancing process governed by the RR process, the cumulated discrepancy is very low: the cumulated load of every node is identical up to $\tilde{\mathcal{O}}(1)$ when $\gcd(k, m) = 1$ (see Table 1 for a detailed overview of results). In other words, the total number of operations, such as storing a task, sending or receiving a task from a neighbor, which are performed by every node, is fairly balanced over all nodes of the network, counting over any window of time. Note that the cumulative load balancing problem cannot be addressed using a random-walk-type approach: indeed, even a process propagating a single superfluous token with a random walk among otherwise balanced nodes can cause significant discrepancy of cumulated load, which eventually becomes unbounded. This is, in particular, the case for the elegant randomized rounding approach of Friedrich and Sauerwald [22], which is extremely efficient at bounding load differences at any fixed moment of time, but which allows the differences of cumulated load of nodes over a time interval of length T to increase (in expectation) proportionally to \sqrt{T} . (Such behavior is observed even on simple graph topologies, such as the three-vertex cycle, since intuitively, the best guarantees for such a randomized process follow directly from the law of large numbers. The guarantees we obtain for the RR dynamics have a more subtle explanation which relies on combinatorial and number-theoretic properties. Other deterministic load balancing processes known to us show no similar advantage for cumulated load balancing).

We note that our results provide insights for the behavior of the RR dynamics after it has stabilized, and have no direct bearing on its initial stabilization phase.

2 Exposition of Technical Results

2.1 Notation

In this paper, we will consider graph $G = (V, E)$ to be directed and symmetric, i.e., such that every arc $e = (u, v)$ has a corresponding opposite arc $(v, u) =: -e$. We denote $|E| = m$ and $|V| = n$. We allow the graph to have multi-arcs as well as self-loops, however, to extend our notation to loops, we say that a loop is the opposite of itself, only (even if there are many loops in a single vertex). For an arc $e = (u, v)$, we denote $u = \text{pred}(e)$ and $v = \text{succ}(e)$. Unless otherwise stated, we will assume that G is not bipartite, to avoid obvious issues with the periodicity of 2 in the dynamics on G .

We denote the set of integers by \mathbb{Z} , the set of non-negative integers by \mathbb{Z}_+ , and the set of integers in modulo- η arithmetic as \mathbb{Z}_η . For $A, B \subseteq \mathbb{Z}_\eta$ their *sumset* is $A + B := \{(a + b) \bmod \eta : a \in A, b \in B\}$, and for $k \in \mathbb{Z}_+$, $kA := A + \dots + A$, with the sum spanning over k copies of A .

2.2 Eulerian Mixing in Time and Space

The main technical ingredients of this paper are linked to properties of traversals of Eulerian circuits, and, more broadly, of circulations on symmetric directed graphs. We formulate them in this section in the most intuitive form, in a way which is independent of the RR process.

We will call $\varphi : E \rightarrow E$ a (unit) *circulation* on the set of arcs of a symmetric directed graph (V, E) if φ is a bijection on E such that for all arcs $e \in E$, $\text{succ}(e) = \text{pred}(\varphi(e))$.

For a pair of vertices $u, v \in V$ and $t \in \mathbb{Z}_+$, we will write $u \rightarrow_t v$ if there exists an edge e with $\text{pred}(e) = u$ such that $\text{succ}(\varphi^t(e)) = v$, i.e., it is possible to reach v from u by following a t -step walk along the circulation. We will also write $v \rightarrow_{-t} u$ whenever $u \rightarrow_t v$, for $t \in \mathbb{Z}_+$.

We will call φ *Eulerian* if the orbit of any arc is an Eulerian circuit on (V, E) , i.e., for any $e \in E$, we have $\{\varphi^t(e) : t \in \mathbb{Z}\} = E$. In general, we will consider our circulation φ as a union of $g \geq 1$ orbits (cycles), denoted E_1, \dots, E_g , forming a partition of the arc set.

For a circulation φ on a graph $G = (V, E)$, let G_φ be the graph with vertex set $V \times \mathbb{Z}$ and edge set $\{(u, t_u), (v, t_v) : u \rightarrow_{(t_v - t_u)} v\}$. For the important case of Eulerian circulations, i.e., $g = 1$, the (simpler) version of the main technical lemma of our paper may be stated as follows.

Lemma 2.1. *Let φ be an Eulerian circulation on a non-bipartite symmetric directed graph G . Then, graph G_φ has diameter $\mathcal{O}(\log^2 n)$.*

The assumption that the considered graph is symmetric cannot be omitted in the statement of Lemma 2.1; an example of a non-symmetric graph for which G_φ has larger diameter is presented in Fig. 1.

As a side remark, we note that the above Lemma has an intuitive interpretation in that it bounds the distance between a pair of time-separated points $(u, 0)$ and (u, l) by $\mathcal{O}(\log^2 n)$. We can thus write the following informal corollary.

Corollary 2.2. *For any fixed Eulerian circuit in a non-bipartite symmetric directed graph G and any $l \in \mathbb{Z}_m$, there exists a closed walk over the arcs of G of length congruent to l modulo m , constructed by a successive traversal of $\mathcal{O}(\log^2 n)$ contiguous fragments of the given Eulerian circuit.*

Note that the trajectory between nodes $(u, 0)$ and (u, l) which follows from a shortest path in graph G_φ may contain “positive” traversals along iterated applications of φ , as well as “negative” traversals along iterated applications of φ^{-1} . However, one can ensure that all arcs appear in the considered walk a non-negative (or positive) number of times by augmenting it with a traversal of sufficiently large number of iterations of the complete Eulerian circuit φ^m . This does not change the length of the walk, modulo m .




Figure 1: A non-symmetric directed graph for which Lemma 2.1 does not hold: the unique Eulerian circulation φ covering G results in G_φ that has diameter linear in n . If we consider a configuration of RR dynamics with one token per each blue vertex, symmetrically on both sides, the configuration repeats every $m/2$ steps, and no mixing occurs.

The proof of Lemma 2.3 relies on a characterization of the set of lengths of *self-intersections* of an Eulerian circuit. Intuitively, t belongs to the set of self-intersections if there exists a closed walk $u \rightarrow_t u$ along the Eulerian circuit. In Section 3 we exploit the spectral properties of the set of self-intersections to show that for given η , any value in \mathbb{Z}_η may be obtained as a sum of at most $\mathcal{O}(\log^2 \eta)$ lengths of self-intersections. The discussion in the section is formulated in slightly more general notation, leading to Lemma 3.6, which will also prove useful in the consideration of non-Eulerian circulations. The claim then follows readily (Section 4).

Extending considerations beyond Eulerian circulations, we obtain the following generalization of Lemma 2.1 to the case of $g > 1$.

Lemma 2.3. *Let φ be a unit circulation on a non-bipartite symmetric directed graph G . Then, G_φ has diameter at most $\mathcal{O}(g \log^2 n)$.*

The proof of this lemma relies on the same ingredients (Section 3) as the considerations in the Eulerian case, however, the sum set technique only provides us with a way of showing mixing in time when considering time values modulo the greatest common divisor of the lengths of all cycles in the circulation. The general result for arbitrary integer time values follows from an application of the Chinese Remainder Theorem (Section 4).

The obtained bounds on the diameter of G_φ are tight up to polylogarithmic factors, and we subsequently apply them to the analysis of the RR model.

We close this section with the following side remark, which appears to be interesting in its own right but has no bearing on the main results of the paper. It is possible to obtain a slightly stronger spectral property of G_φ for Eulerian φ than merely a polylogarithmic bound on the diameter. For any Eulerian circuit in G , one can design a random process which will terminate at a well-mixed node in G at a well-mixed time, by traversing a polylogarithmic number of fragments of the Eulerian circuit.

Corollary 2.4. *Let φ be an Eulerian circulation on a non-bipartite symmetric directed graph (V, E) . Then, there exists an assignment of weights $w : E(G_\varphi) \rightarrow [0, 1]$, such that the reversible Markov chain following (G_φ, w) achieves good mixing on $V \times \mathbb{Z}_m$ starting from any vertex from this set, after $\tilde{\mathcal{O}}(1)$ steps.*

The proof follows from a straightforward adaptation of the spectral arguments used in our analysis in Section 3.

2.3 RR Dynamics and Short-Time-Averaging

We start this section by formally defining the RR dynamics. Each vertex v of G is equipped with a fixed ordering of all its outgoing arcs $\rho_v = (e_1, e_2, \dots, e_{\deg(v)})$. A state x at the current time step t is a tuple:

$x_t = ((pointer_v)_{v \in V}, (L(v))_{v \in V})$, where $pointer_v$ is an arc outgoing from node v , which is referred to as *the current port pointer at node v* , and $L(v)$ is the number of tokens at any given node. The state space of the RR dynamics is denoted X , and its cardinality is in general readily seen to be (exponentially) larger than $|V|$. The system is initialized to an arbitrary state of the state space X . For an arc (v, u) , let $next((v, u))$ denote the arc after the arc (v, u) in the cyclic order ρ_v . During each step, each node v distributes in round-robin fashion all of its tokens, using the following algorithm:

While there is a token at node v , **do**:

1. Send the token to $pointer_v$,
2. Set $pointer_v = next(pointer_v)$.

Note that during a single time step all tokens at a node v are sent out and at exactly the next time step all those tokens arrive at their respective destination nodes. The total number of tokens in graph is denoted as k .

When the system is in state $x \in X$, let $(L_x(e) : e \in E) \in \mathbb{Z}_+^m$ be the distribution vector of tokens over arcs of the RR dynamics during its transition from state x to the subsequent state, i.e., let $L_x(e)$ represent the number of particles sent out along arc $e \in E$ in the current time step, starting from state x .

Once the RR dynamics has entered a recurrent state, the trajectory followed by the dynamics involves only a small fraction of its entire state space X , and as such, when viewed in high dimension, it is definitely not an ergodic process. However, we show that it is possible to make use of the time-mixing property of (almost) Eulerian circulations captured by Lemma 2.3 to show that RR achieves the sought time-averaging behavior when considering only one-dimensional probe functions $L(e)$. Intuitively, in the recurrent state of the RR dynamics, individual tokens may be seen as traversing cycles of some unit circulation, with balancing behavior occurring at the meeting points of different cycles in the circulation.

Theorem 2.5. *For any $e \in E$ and $t_0, T \in \mathbb{Z}_+$, such that $x(t_0)$ is recurrent, the RR dynamics satisfies:*

$$\frac{1}{T} \sum_{t=t_0}^{t_0+T} L_{x(t)}(e) = \frac{k}{m} \pm \tilde{O}\left(\frac{\gcd(k, m)}{T}\right).$$

The above Theorem, which we prove in Section 5 after introducing more background properties of the RR process, has immediate consequences to problems of idleness in patrolling and cumulative load balancing. This is also discussed in Section 5.

3 Sum-sets of Self-Intersections on a Circulation

Compared to the more high-level exposition in Section 2.2, in the subsequent analysis it is more convenient to work with measures of distance between pairs of arcs rather than pairs of vertices. We thus introduce all further notation based on the arc-based perspective.

Intersections: We will say that an assignment $\lambda : E \rightarrow \mathbb{Z}_\eta$ along the cycle (cycles) of circulation φ is a η -labeling if the following condition is fulfilled: $\forall e \in E \lambda(\varphi(e)) = \lambda(e) + 1$. (This is only possible when $\eta \mid \gcd(|E_1|, |E_2|, \dots, |E_g|)$.) We say that two arcs e_1 and e_2 outgoing from the same vertex, $\text{pred}(e_1) = \text{pred}(e_2)$, span an *intersection* of size $(\lambda(e_1) - \lambda(e_2)) \bmod \eta \in \mathbb{Z}_\eta$. We consider the *set of all intersections* $A_\lambda \subseteq \mathbb{Z}_\eta$:

$$A_\lambda = \{\lambda(e_1) - \lambda(e_2) : \text{pred}(e_1) = \text{pred}(e_2)\}.$$




Figure 2: (a) Example of a graph G with a circulation with no “small” intersections: for any m -labeling, $A_\lambda \subseteq [m/4, 3m/4]$. Examples of assignments of colors in G for: (b) m -labeling and (c) $m/5$ -labeling. In both cases, all 3 colors meet at least once.

Observe that A_λ does not depend on the particular choice of λ , since all λ are identical up to cyclical shift. We immediately have $0 \in A_\lambda$ and $A_\lambda = -A_\lambda$. Observe also that even when η is even, there is an odd element in A_λ . Otherwise, we could partition vertices into two sets, $\{\text{pred}(e) : 2 \mid \lambda(e)\}$ and $\{\text{pred}(e) : 2 \nmid \lambda(e)\}$, which contradicts the assumption that G is not bipartite.

Set A_λ may have a non-trivial structure, e.g., may contain no small elements other than 0 (see e.g. Fig. 2). Fortunately, it turns out we can take advantage of its spectral properties (which we subsequently formally define).

Spectral properties of A_λ : We start by showing the following key property of the set A_λ of all intersections, expressed by Lemma 3.1. We then briefly recall the formalism of Bohr sets, and use Lemma 3.1 to characterize the structure of the Bohr set of A_λ .

Lemma 3.1. *Let $\xi \in \mathbb{Z}_\eta$ with $1 \leq \xi \leq \eta/6$. Let A_λ be the set of intersections for an arbitrary η -labeling of a circulation. Then, there exists $x \in A_\lambda$ such that $\xi \cdot x \in [\frac{1}{3}\eta, \frac{2}{3}\eta]$.*

Proof. Let us consider a continuous version of set E , that is a metric space \mathcal{E} constructed on the set $E \times [0, 1)$, such that for every edge e , we set $\lim_{t \rightarrow 1}(e, t) = (\varphi(e), 0)$. Observe that \mathcal{E} is composed of disjoint spaces $\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_g$, where each \mathcal{E}_i is a continuous version of cycle E_i , isometric to $\mathbb{R}/(|E_i| \cdot \mathbb{Z})$ (i.e., to a cyclic interval of length $|E_i|$).

We say that elements of \mathcal{E} of the form $((v, \cdot), 0)$ are occurrences of v in \mathcal{E} . Now we can extend the labeling λ of E to a labeling $\mathcal{L} : \mathcal{E} \rightarrow \mathbb{R}/(\eta \cdot \mathbb{Z})$, so that $\mathcal{L}((e, 0)) = \lambda(e)$ and $\mathcal{L}|_{\mathcal{E}_i}$ is a locally distance-preserving mapping of \mathcal{E}_i to $\mathbb{R}/(\eta \cdot \mathbb{Z})$. (\mathcal{L} is an “interpolation” of λ , increasing in a continuous fashion at a constant rate between discrete points on which λ was defined.)

We will say that two points in \mathcal{E} *meet* if they are either occurrences of the same vertex $v \in V$, or if they are of the form $(e, t), (-e, 1 - t)$, for some $e \in E$ and $0 < t < 1$, where we recall that e and $-e$ denote a pair of opposing arcs.

Now, for fixed $\xi \in \mathbb{Z}_\eta$ satisfying the assumption of the lemma, we assign to each point $x \in \mathcal{E}$ one of three colors $\{0, 1, 2\}$ according to the following rule:

$$\text{color}(x) = \begin{cases} 0 & \text{if } \xi \cdot \mathcal{L}(x) \in [0, \frac{1}{3}\eta), \\ 1 & \text{if } \xi \cdot \mathcal{L}(x) \in [\frac{1}{3}\eta, \frac{2}{3}\eta), \\ 2 & \text{if } \xi \cdot \mathcal{L}(x) \in [\frac{2}{3}\eta, \eta). \end{cases}$$

Thus, intuitively we have covered the continuous version of the arc set with closed directed continuous cycles, and each cycle is colored along its orientation with colors 0, 1, 2, with colors successively applied to

contiguous segments of length $\eta/(3\xi)$. See Figure 2 for example. Now, by purely topological reasoning, we observe that there is a point of \mathcal{E} , where all 3 colors meet. To be more precise, one of the following happens:

1. There exist $x_0, x_1, x_2 \in \mathcal{E}$, all being occurrences of the same vertex from V , such that $color(x_i) = i$. It follows that $\xi \cdot (\mathcal{L}(x_i) - \mathcal{L}(x_j)) \in [\frac{1}{3}\eta, \frac{2}{3}\eta]$ for some $i \neq j$.
2. There exist $x, y \in \mathcal{E}$ that meet and satisfy the following constraint (up to permutation of colors $0, 1, 2$): $color(x) = 0$, $color(x - \varepsilon) = 2$ (for arbitrarily small ε), and $color(y) = 1$. It follows that $\xi \cdot (\mathcal{L}(y) - \mathcal{L}(x)) \in [\frac{1}{3}\eta, \frac{2}{3}\eta]$.

Since \mathcal{L} and λ agree on all the occurrences of the vertices (if we identify the values on the edges as taken on the starting points), and in the second case we can round pair x, y to one of the endpoints of the edge they lie on, the claim follows. \square

The above lemma, applied over different ξ , yields a valuable characterisation of the set of elements in A_λ .

Bohr set: We define the *Bohr set* with radius α for a set $S \subseteq \mathbb{Z}_\eta$ as²:

$$Bohr(S, \alpha) = \{\xi \in \mathbb{Z}_\eta : \forall s \in S \|s \cdot \xi\|_\eta \leq \alpha\},$$

where $\|\cdot\|_\eta \in \mathbb{R}$ denotes the ‘‘fractional part’’ norm on \mathbb{Z}_η . (Formally, for $x \in \mathbb{Z}_\eta$, we put $\|x\|_\eta = \alpha$ for the unique real $\alpha \in (-1/2, 1/2]$ such that $x \equiv \alpha\eta \pmod{\eta}$; see also [38][beginning of Section 4.4].)

With this definition, Lemma 3.1 can be rephrased as follows: $Bohr(A_\lambda, 1/3 - \varepsilon) \cap (0, \eta/6] = \emptyset$, for arbitrarily small $\varepsilon > 0$. This is not yet the spectral property we need, but A_λ is easily transformable to a set with such a property. Since $A_\lambda = -A_\lambda$, we immediately have $Bohr(A_\lambda, 1/3 - \varepsilon) \cap [5/6\eta, \eta) = \emptyset$ as well. We will say that a is critical for b , if $\|a \cdot b\|_\eta \geq \alpha$. Since $x \cdot (2 \cdot \xi) = (2 \cdot x) \cdot \xi$, if x is a critical element for $2 \cdot \xi$, $2 \cdot x$ is critical for ξ , thus we have $Bohr(2 \cdot A_\lambda, 1/3 - \varepsilon) \cap ([5/12\eta, \eta/2) \cup (\eta/2, 7/12\eta]) = \emptyset$. Additionally, we observe that for any $\xi \in (\eta/6, 5/12\eta]$, $\|2 \cdot \xi\|_\eta \geq 1/6$, thus $Bohr(A_\lambda, 1/6) \cap ((\eta/6, 5/12\eta] \cup [7/12\eta, 5/6\eta)) = \emptyset$.

By $A_\lambda \cup 2 \cdot A_\lambda \subseteq A_\lambda + A_\lambda$, we conclude that $Bohr(A_\lambda + A_\lambda, 1/6) \subseteq \{0, \eta/2\}$. However, since we know that A_λ contains an odd element, we can eliminate $\eta/2$ from the set. Thus finally, we obtain a spectral characterization of set A_λ .

Lemma 3.2. *Let A_λ be defined as in Lemma 3.1. Then: $Bohr(A_\lambda + A_\lambda, 1/6) = \{0\}$.*

Additive combinatorics: Let $A \subseteq \mathbb{Z}_\eta$ be any set. We denote by $\hat{A}(j)$ the j -th ($j \in [0, \eta - 1]$) Fourier coefficient in the transform of 1_A , normalized so that $\hat{A}(0) = 1$ and $|\hat{A}(j)| \leq 1$, for all j . Formally, $\hat{A}(j) := \frac{1}{|A|} \sum_{a \in A} e(aj/\eta)$, where $e(x) \equiv \exp(2\pi i x)$.

Define the *spectrum* $Spec_\alpha(A)$ as the set of all $j \in [0, \eta - 1]$ such that $|\hat{A}(j)| > \alpha$ (i.e., the set of indices of Fourier coefficients which are sufficiently large). The definition makes sense for $\alpha \in (0, 1]$ and $0 \in Spec_\alpha(A)$. (See [38][Def. 4.34]). The following lemma about annihilation of the spectrum is folklore; we provide a short proof for completeness.

Lemma 3.3. *If $Spec_{1-1/s}(A) = \{0\}$, then $\kappa A - \kappa A = \mathbb{Z}_\eta$, for $\kappa \geq s \ln \eta$.*

²For all undefined notation in additive combinatorics, we refer the reader to [24, 38].

Proof. Let $\kappa = s \ln \eta$, and let $r_{\kappa A - \kappa A}(l)$ denote the number of ways to represent the value $l \in \mathbb{Z}_\eta$ using $\kappa A - \kappa A$. We have (cf. e.g. [24][Sec I.8]):

$$\begin{aligned} r_{\kappa A - \kappa A}(l) &= c_k \sum_j |\hat{A}(j)|^{2\kappa} e(-jl/\eta) \geq \\ &\geq c_k \left(1 - \sum_{j \neq 0} |\hat{A}(j)|^{2\kappa} \right), \end{aligned} \tag{1}$$

where c_k is some positive normalizing constant ($c_k = |A|^k/\eta$).

Next, for a set A satisfying the assumption of the Lemma, we have for $j \neq 0$: $|\hat{A}(j)|^{2\kappa} \leq (1 - 1/s)^{2s \ln \eta} < 1/\eta$. Hence, $\sum_{j \neq 0} |\hat{A}(j)|^{2\kappa} < 1$. It follows that for all $l \in \mathbb{Z}_\eta$, we have $r_{\kappa A - \kappa A}(l) > 0$, and so $\kappa A - \kappa A = \mathbb{Z}_\eta$. \square

In the following, we will also use the following property linking the Bohr set and the spectrum.

Proposition 3.4 ([38], variant of Prop. 4.40). *Let $A, B \subseteq \mathbb{Z}_\eta$ be such that $A \subseteq B$ and $0 \in A$. Let $K \geq 1$ and $0 < \varepsilon < 1$. If $|B - A| \leq K|B|$, then*

$$A - A \subseteq \text{Bohr}(\text{Spec}_{1-\varepsilon}(B - A), \sqrt{8\varepsilon K}).$$

For completeness, the proof of the above Proposition is provided in the Appendix.

We are now ready to show our final Lemma, where we extend considerations from the above Proposition to iterated sums of a set.

Lemma 3.5. *For any set $A_1 \subseteq \mathbb{Z}_\eta$, such that $0 \in A_1$, we have that $\text{Spec}_{1-1/576}(A_1 - \kappa A_1) \subseteq \text{Bohr}(A_1, 1/6)$, for some $\kappa = \mathcal{O}(\log \eta)$.*

Proof. Let us denote for short $B = \text{Bohr}(A_1, 1/6)$. In the proof of the lemma, we will use an equivalent property characterizing A_1 , which follows from symmetry of multiplication in expanding the Bohr set definition (i.e., $\forall x \in \mathbb{Z}_\eta, x \neq 0 : \exists \xi \in A_1 \|x \cdot \xi\|_\eta > 1/6$):

$$\forall X \subseteq \mathbb{Z}_\eta, X \not\subseteq B : A_1 \not\subseteq \text{Bohr}(X, 1/6). \tag{2}$$

We now define sets A_i , starting from A_2 , as follows: $A_{i+1} := A_i - A_1$. We will assume that $0 \in A_1$, so $0 \in A_i$ for all i , and $A_{i+1} \supseteq A_i \supseteq \dots \supseteq A_1$.

We now apply Proposition 3.4 to set A_{i+1} , with $K = 2$ and $\varepsilon = 1/576$. From the Proposition, one of the following two cases must hold:

- Either the assumption of the Proposition is not fulfilled, i.e., $|A_{i+1}| > 2|A_i|$, or:
- $A_1 \subseteq A_1 - A_1 \subseteq \text{Bohr}(\text{Spec}_{1-1/576}(A_{i+1}), 1/6)$.

Since A_i size is bounded by η , it suffices to iterate above construction until at most $i \leq \log_2 \eta$ to have the latter case satisfied. Observe that then, by Eq. (2), we must then have $\text{Spec}_{1-1/576}(A_{i+1}) \subseteq B$. \square

Observe that if we take any A_λ being the set of intersections for an arbitrary η -labeling of a circulation, then we have already established that $\text{Bohr}(A_\lambda, 1/6) = \{0\}$ (Lemma 3.2). By Lemma 3.5, we have that for some $\kappa = \mathcal{O}(\log \eta)$, $\text{Spec}_{1-1/576}(A_\lambda - \kappa A_\lambda) \subseteq \text{Bohr}(A_\lambda, 1/6) = \{0\}$. Applying Lemma 3.3 we obtain the following technical claim.

Lemma 3.6. *Let A_λ be the set of intersections for arbitrary η -labeling of a circulation. Then, $\kappa A_\lambda = \mathbb{Z}_\eta$ for some $\kappa = \mathcal{O}(\log^2 \eta)$.*

4 Bounding the Diameter of Graph G_φ

We now apply the sum-set structure of the set of self-intersections of a circulation, expressed by Lemma 3.6, to bound the diameter of graph G_φ . We start with a direct argument for the case of Eulerian circulations, followed by a more involved analysis for the general case.

4.1 Eulerian Circulations (Proof of Lemma 2.1)

Considerations of Eulerian circulations rely on the definition of a metric δ_t , which expresses an upper bound on the number of contiguous fragments of the Eulerian circuit which it is sufficient to traverse in order to return to the starting arc with a time offset of t .

Intersection distance: Fix $\eta = m$. Let us consider a labeling λ and the corresponding set of intersections A_λ . We will say that a function $\delta : \mathbb{Z} \rightarrow \mathbb{Z}^+$ is an *intersection distance measure* if it satisfies the following constraints:

$$\delta_{im} = 0 \quad \text{for any } i \in \mathbb{Z} \quad (3)$$

$$\delta_x \leq 1 \quad \text{when } x \in A_\lambda \quad (4)$$

$$\delta_{-x} = \delta_x \quad \text{for any } x \in \mathbb{Z} \quad (5)$$

$$\delta_{x+y} \leq \delta_x + \delta_y \quad \text{for any } x, y \in \mathbb{Z} \quad (6)$$

From the definition it follows that for any $x \in \kappa A_\lambda$, $\delta_x \leq \kappa$. Thus by Lemma 3.6 we immediately have:

$$\max_{x \in \mathbb{Z}} \delta_x = \max_{x \in \mathbb{Z}_\eta} \delta_x = \mathcal{O}(\log^2 \eta), \quad (7)$$

where the first equality holds by property (3), corresponding to shifts by a complete tour of the Eulerian circuit.

Phrasing the above in terms of distances in G_φ , by considering for any node v the properties of function δ_x for some arc of the form (v, \cdot) , we conclude that the distance in G_φ between any two nodes of the form (v, t) and $(v, t + x)$, $x, t \in \mathbb{Z}$ is bounded by $\mathcal{O}(\log^2 \eta)$. Since for any pair of vertices $u \neq v$ there exists a path $v \rightarrow_{x_{vu}} u$, for some $x_{vu} \in \mathbb{Z}$ we moreover have that pairs of vertices of the form (v, t) , $(u, t + x_{vu})$, $t \in \mathbb{Z}$, are at distance 1 from each other in G_φ . The claim of Lemma 2.1 follows.

4.2 General Circulations (Proof of Lemma 2.3)

We now extend the definition of intersection distance so that it is useful for arbitrary circulations.

Generalizing intersection distance: Let φ be an arbitrary circulation. We will say that a function $\delta : \mathbb{Z} \times E \times E \rightarrow \mathbb{Z}^+$ is a *generalized intersection distance measure* if it satisfies the following constraints:

$$\delta_0(e, e) = 0 \quad \text{for any } e \in E \quad (8)$$

$$\delta_1(e, \varphi(e)) = 0 \quad \text{for any } e \in E \quad (9)$$

$$\delta_0(e_1, e_2) = 1 \quad \text{when } \text{pred}(e_1) = \text{pred}(e_2) \quad (10)$$

$$\delta_{-x}(e_2, e_1) = \delta_x(e_1, e_2) \quad \text{for any } x \in \mathbb{Z}, e_1, e_2 \in E \quad (11)$$

$$\delta_{x+y}(e_1, e_3) \leq \delta_x(e_1, e_2) + \delta_y(e_2, e_3) \quad \text{for any } x, y \in \mathbb{Z}, e_1, e_2, e_3 \in E \quad (12)$$

The generalization is to be understood in the following sense: if φ is an Euclidean circulation and δ is a generalized intersection distance measure, then for arbitrary $e \in E$, $\delta_x := \delta_x(e, e)$ is an intersection distance measure.

Cycle combinatorics: We now define a few more concepts regarding how different cycles in a (non-Eulerian) circulation relate to each other.

First, given a circulation with cycles (E_1, \dots, E_g) , we introduce the notion of *adjacency* of cycles, saying that two cycles are *adjacent* if they share a common vertex. Naturally, we can use this definition of adjacency to define a graph \mathcal{G} on the set of cycles of the circulation, and the corresponding *distance* $\Delta(E_i, E_j)$ between any two cycles E_i and E_j as the number of adjacency hops needed to go between this pair of cycles.

Additionally, given labeling λ of the circulation, we say that two cycles E_i, E_j are λ -*adjacent* if there are two arcs $e_i \in E_i, e_j \in E_j$, such that $\text{pred}(e_i) = \text{pred}(e_j)$ and $\lambda(e_i) = \lambda(e_j)$. We call the graph \mathcal{G}_λ with vertex set corresponding to the set of cycles the cycles $\{E_1, \dots, E_g\}$ and edges corresponding to the λ -adjacency relation the *cycle graph*. For two cycles E_i, E_j , we denote by $\Delta_\lambda(E_i, E_j)$ the distance in the cycle graph \mathcal{G}_λ between respective cycles E_i and E_j . For compactness of notation, given two arcs $e_1 \in E_1, e_2 \in E_2$, we denote $\Delta_\lambda(e_1, e_2) = \Delta_\lambda(E_1, E_2)$. Finally, we say that labeling λ is *cycle-connected* if its cycle graph \mathcal{G}_λ is connected.

It is easy to see that we can always adjust any indexing λ by cyclical shifts, successively for each cycle, so as to incorporate an arbitrary spanning tree of \mathcal{G} in the cycle graph \mathcal{G}_λ . In particular, by considering the BFS tree of \mathcal{G} , whose diameter is at most $2 \cdot \text{diam}(\mathcal{G})$, we have the following result.

Observation 4.1. *For any $\eta \mid \text{gcd}(|E_1|, \dots, |E_g|)$, there is an η -labeling λ such that for any E_i, E_j : $\Delta_\lambda(E_i, E_j) \leq 2 \cdot \text{diam}(\mathcal{G})$.*

From now on we assume that λ satisfies the property given by Observation 4.1.

Lemma 4.2. *For any pair of arcs e_1, e_2 , there exists $\alpha \in \mathbb{Z}$ such that $\alpha \equiv \lambda(e_2) - \lambda(e_1) \pmod{\eta}$ and $\delta_\alpha(e_1, e_2) \leq \Delta_\lambda(e_1, e_2)$. Moreover, there exists $\alpha' \in \mathbb{Z}$, such that $\delta_{\alpha'}(e_1, e_2) \leq \Delta(e_1, e_2)$.*

Proof. We start by covering the first part of claim. If e_1 and e_2 belong to the same cycle, then there is $x \in \mathbb{Z}$ such that $\delta_x(e_1, e_2) = 0$. However, since λ limited to one particular cycle is just a cyclic numbering modulo n , we have that $\lambda(e_2) - \lambda(e_1) \equiv x \pmod{n}$.

We also observe that the claim is satisfied for two arcs that make corresponding cycles λ -adjacent: if $\lambda(e_1) = \lambda(e_2)$ and e_1 and e_2 are adjacent, then simultaneously $\lambda(e_2) - \lambda(e_1) = 0$ and $\delta_0(e_1, e_2) \leq 1$.

By a simple chaining argument we can extend this to any two arcs on λ -adjacent cycles, and by the induction on $\Delta_\lambda(e_1, e_2)$ to any two arcs.

By repeating the same reasoning using cycle adjacency instead of λ -adjacency, we get the second claim (without getting any constraint on α'). \square

Lemma 4.3. *For any two arcs e, e' and $x \in \mathbb{Z}$, we have $\delta_x(e', e') \leq \delta_x(e, e) + 2\Delta(e, e') \leq \delta_x(e, e) + \mathcal{O}(\text{diam}(\mathcal{G}))$.*

Proof. There is $\alpha \in \mathbb{Z}$ such that $\delta_\alpha(e, e') \leq \Delta(e, e')$. Thus $\delta_x(e', e') \leq \delta_\alpha(e', e) + \delta_{x-\alpha+\alpha}(e, e) + \delta_{-\alpha}(e, e') \leq \delta_x(e, e) + 2\Delta(e, e')$. \square

Lemma 4.4. *For any δ that is a generalized intersection distance measure on a circulation with g orbits, we have:*

$$\max_{e \in E, x \in \mathbb{Z}} \delta_x(e, e) = \tilde{\mathcal{O}}(g).$$

Proof. Let $\eta = \gcd(|E_1|, \dots, |E_g|)$, and consider an indexing λ following Observation 4.1. We will once again be working with the set of all intersections $A_\lambda \subseteq \mathbb{Z}_\eta$.

Fix any $x \in A_\lambda$. Thus, there are $e_1 \in E_1, e_2 \in E_2$ such that $\text{pred}(e_1) = \text{pred}(e_2)$ and $\lambda(e_2) - \lambda(e_1) = x$. The cycles E_1, E_2 are not necessarily λ -adjacent, however by Lemma 4.2 we have that $\delta_\alpha(e_1, e_2) = \mathcal{O}(\text{diam}(\mathcal{G}))$ for some $\alpha \equiv x \pmod{\eta}$. Moreover, since $\text{pred}(e_1) = \text{pred}(e_2)$, we have $\delta_\alpha(e_1, e_1) \leq \delta_\alpha(e_1, e_2) + 1$, and so $\delta_\alpha(e_1, e_1) = \mathcal{O}(\text{diam}(\mathcal{G}))$. By a direct application of Lemma 4.3 we can extend this bound to any arc $e \in E$: $\delta_\alpha(e, e) \leq 2 \cdot \Delta(e, e_1) + \delta_\alpha(e_1, e_1) = \mathcal{O}(\text{diam}(\mathcal{G}))$. Thus we have that for any $x \in \kappa A_\lambda$, there is $x' \in \mathbb{Z}$ such that $\delta_{x'}(e, e) = \mathcal{O}(\kappa \cdot \text{diam}(\mathcal{G}))$, and $x' \equiv x \pmod{\eta}$. By Lemma 3.2 and Lemma 3.6, there exists $\kappa = \mathcal{O}(\log^2 n)$ such that $\kappa A_\lambda = \mathbb{Z}_\eta$.

Denote $\beta = (x' - x)/\eta$. We have shown that skewed distances can be bounded for a set which, under the projection $\text{mod } \eta$, is \mathbb{Z}_η . We will now focus on controlling the term β , describing shifts by multiples of η . Pick $e_1 \in E_1$. We have $\delta_{c \cdot |E_1|}(e_1, e_1) = 0$ for any integer c , thus by Lemma 4.3, $\delta_{c \cdot |E_1|}(e, e) = \mathcal{O}(\text{diam}(\mathcal{G}))$ for any arc e .

Recall that we put $\eta = \gcd(|E_1|, \dots, |E_g|)$. We observe that η has at most $\mathcal{O}(\log \eta) \subseteq \mathcal{O}(\log n)$ different primes in its factorization, thus the gcd is already obtained by taking particular subset of $\log n$ elements. It follows that for any integer β , there exists a sequence c_1, \dots, c_g such that $\sum_i c_i |E_i| = \beta \eta$ and there are $\mathcal{O}(\log n)$ non-zero coefficients in c_i . Thus, in particular, $\delta_{\beta \eta}(e, e) \leq \sum_i \delta_{c_i |E_i|}(e, e) = \mathcal{O}(\text{diam}(\mathcal{G}) \cdot \log n)$.

Combining shifts by multiplicities of η with shifts with proper remainders, we finally have: $\delta_x(e, e) \leq \delta_{x'}(e, e) + \delta_{\beta \eta}(e, e) = \tilde{\mathcal{O}}(\text{diam}(\mathcal{G}))$, from which the claim follows. \square

As in the Eulerian case, we observe that the upper bound on $\delta_x(e, e)$ also gives an upper bound on the path length in graph G_φ , up to an additive factor of g which allows us to move between two spatially separated vertices. This completes the proof of Lemma 2.3.

5 Results for RR Dynamics

5.1 Preliminaries: Properties of RR and Auxiliary Definitions

We recall that the RR dynamics was introduced in Section 1.1. We define edge load L_t and cumulated load $C_t^{\Delta t}$ as follows: $L_t(e)$ is the number of tokens traversing arc e at timestep t , and $C_t^{\Delta t}(e) = \sum_{0 \leq i < \Delta t} L_{t+i}(e)$.

The following notation will be useful when discussing the load-balancing and patrolling properties of the RR dynamics.

Load balancing: A standard notion is to define a discrepancy (over arcs) of a token distribution at a given timestep t as $\max\{|L_t(e_1) - L_t(e_2)| : e_1, e_2 \in E\}$. Similarly, a notion of a cumulative discrepancy over time Δt defined as $\max\{|C_t^{\Delta t}(e_1) - C_t^{\Delta t}(e_2)| : e_1, e_2 \in E\}$.

Patrolling: We say that an arc e has an idleness of t if, starting from some point in time, e is visited at least once in every interval of t consecutive steps, and t is the smallest such value. More precisely, we define *idle arc time* at a given time step t as $\text{idle}(t, e) = \min(\Delta t : C_t^{\Delta t}(e) \geq 1)$. We define idleness for a single arc: $\text{idle}(e) = \limsup_{t \rightarrow \infty} \text{idle}(t, e)$. Finally we define *idleness* as: $\text{idle}(G) = \max\{\text{idle}(e) : e \in E\}$.

Limit behavior: We recall following two results characterizing limit behavior of the RR system, reached in polynomially many steps.

Theorem 5.1 ([10], Theorem 11). *RR dynamics is recurrent, if and only if there is a bijection $\varphi : E \rightarrow E$, such that $\forall_{t \geq 0} \forall_{e \in E} L_t(e) = L_{t+1}(\varphi(e))$. Moreover, φ can always be selected so that $\text{pred}(\varphi(e)) = \text{succ}(e)$ for any arc e .*

From now on we will identify the subcycle decomposition with bijection φ . We will denote the Eulerian cycles induced by φ as E_1, E_2, \dots, E_g , with g denoting number of cycles.

Theorem 5.2 ([10], Theorem 26). *RR dynamics reaches its recurrent state in $\mathcal{O}(m^4 \text{diam}^2 + m \text{diam} \log k)$ steps on a graph G of diameter diam .*

By Theorems 5.1 and 5.2 any configuration after polynomially many steps reaches recurrent state. Thus from now on we assume that system at the timestep 0 is already in a recurrent state.

Averaging and number of cycles: Let us denote by k_i the total number of tokens circulating cycle E_i , that is $k_i = \sum_{e \in E_i} L_t(e)$ (this value does not depend on the choice of k). It is easy to see that in order to maintain cumulative discrepancy over arbitrarily long time, the following property has to hold

$$\frac{k_1}{|E_1|} = \frac{k_2}{|E_2|} = \dots = \frac{k}{m}. \quad (13)$$

As a corollary of (13), the number of cycles g has to satisfy $g \mid \text{gcd}(k, m)$, thus $g \leq \text{gcd}(k, m)$ follows. Thus, in particular, if m is prime, then for any $k < m$ we have exactly one cycle. The same occurs whenever $k \perp m$ in the general case.

5.2 RR Dynamics with Eulerian Circulation

Now we proceed to analyze RR dynamics in a recurrent state that has a circulation with a single cycle. Let us fix circulation φ as in Theorem 5.1.

5.2.1 The Similarity Measure

We use the similarity of the sequence to itself under shifts, where this similarity follows from the small discrepancy of loads on arcs that share a starting (predecessor) vertex.

We define the similarity for time shift $t \in \mathbb{Z}$ as:

$$\delta_t = \max_{e \in E, \tau, \Delta t \geq 0} |C_\tau^{\Delta t}(e) - C_{\tau+t}^{\Delta t}(e)|. \quad (14)$$

Lemma 5.3. *δ defined as in (14) is an intersection distance measure.*

Proof. It is immediate to verify that δ satisfies (3),(5) and (6). To show (4), take $i \in A_\lambda$. By the definition, there are two edges e_1, e_2 , such that $\text{pred}(e_1) = \text{pred}(e_2)$ and $\lambda(e_2) = \lambda(e_1) + i$. Thus, $e_2 = \varphi^i(e_1)$. Fix $e \in E$. Since φ is an Eulerian circulation, let $j \geq 0$ be such that $\varphi^j(e) = e_1$. We have, for any τ and Δt :

$$|C_\tau^{\Delta t}(e) - C_{\tau+i}^{\Delta t}(e)| = |C_{\tau+j}^{\Delta t}(e_1) - C_{\tau+j+i}^{\Delta t}(e_2)| \leq 1$$

thus $\delta_i \leq 1$ and (4) is satisfied. \square

5.2.2 Time-Averaging Property

Lemma 5.4. *For RR dynamics on an Eulerian circulation, for any time step t in the recurrent state, $T > 0$ and $e \in E$, we have $|C_t^T(e) - \frac{k}{m} \cdot T| \leq \delta_T \leq \mathcal{O}(\log^2 \eta)$.*

Proof. Fix time interval $T > 0$ and pick t_1 and t_2 , with $t_2 > t_1$, in the recurrent state arbitrarily. We have the following bounds:

$$|C_{t_1}^T(e) - C_{t_2}^T(e)| = |C_{t_1}^{t_2-t_1}(e) - C_{t_1+T}^{t_2-t_1}(e)| \leq \delta_T,$$

where we used the fact that, for an arbitrary sequence x_i and integers $a \leq b \leq c \leq d$, it holds that: $\sum_{i=a}^{b-1} x_i - \sum_{i=c}^{d-1} x_i = \sum_{i=a}^{c-1} x_i - \sum_{i=b}^{d-1} x_i$, and for the last bound, we used (14).

Since t_1, t_2 were arbitrary, bounds on the discrepancy of the sequence $(C_t^T)_{t \in \mathbb{Z}}$, follow. Since this sequence is periodic and by equation (13) its average value is $T \cdot \frac{k}{m}$, the inequality $|C_t^T(e) - \frac{k}{m} \cdot T| \leq \delta_T$ follows immediately from the previously shown bound. Finally, we have $\delta_T \leq \mathcal{O}(\log^2 \eta)$ by (7), since δ_T is an intersection distance measure following Lemma 5.3. \square

Recalling that by definition $C_{t_0}^T(e) = \sum_{t=t_0}^{t_0+T} L_t(e)$, we obtain directly from Lemma 5.4:

$$\frac{1}{T} \sum_{t=t_0}^{t_0+T} L_t(e) = \frac{k}{m} \pm \tilde{\mathcal{O}}\left(\frac{1}{T}\right),$$

which proves Theorem 2.5 for the special case of Eulerian circulations.

5.3 RR Dynamics with General Circulation

We now apply the results on circulations with multiple cycles to provide bounds for arbitrary RR dynamics.

5.3.1 Generalized Similarity Measure

We extended the notion of δ_x into a load similarity measure between pair of arcs, under time shift. Thus, we define skewed distance function, for $t \in \mathbb{Z}$

$$\delta_t(e_1, e_2) = \max_{\tau, \Delta t \geq 0} |C_\tau^{\Delta t}(e_1) - C_{\tau+t}^{\Delta t}(e_2)|.$$

It is easy to verify that δ defined as above satisfies (8),(9),(10),(11), and (12); δ is thus a generalized intersection distance measure with respect to circulation φ . We obtain the following formulation of Lemma 4.4 for the considered function δ_t .

Lemma 5.5.

$$\max_{e \in E, t \in \mathbb{Z}} \delta_t(e, e) = \tilde{\mathcal{O}}(g) \leq \tilde{\mathcal{O}}(\gcd(k, m)).$$

5.3.2 Time-Averaging Property

The following Lemma is an automatic generalization of Lemma 5.4.

Lemma 5.6. *For RR dynamics on an arbitrary circulation, for any time step t in the recurrent state, $T > 0$ and $e \in E$, we have $|C_t^T(e) - \frac{k}{m} \cdot T| \leq \delta_T(e, e)$.* \square

Recalling once again that by definition $C_{t_0}^T(e) = \sum_{t=t_0}^{t_0+T} L_t(e)$, Lemmas 5.5 and 5.6 directly give:

$$\frac{1}{T} \sum_{t=t_0}^{t_0+T} L_t(e) = \frac{k}{m} \pm \tilde{\mathcal{O}}\left(\frac{\gcd(k, m)}{T}\right).$$

which completes the proof of Theorem 2.5 in the general case.

5.3.3 Other Bounds on the Generalized Similarity Measure

Going beyond the bound of Theorem 2.5, we provide a number of corollaries and auxiliary results, upper-bounding the generalized similarity measure $\delta_T(e, e)$ with respect to other graph parameters. We will then use them in conjunction with Lemma 5.6 to bound the cumulated load of the RR dynamics.

Now first provide a $\tilde{\mathcal{O}}(\sqrt{n})$ bound that uses thresholding: we make use of the fact that in the considered circulation, there are either small cycles (so any cycle must have a short one in its proximity), or there are only large cycles, so there can only be a small number of them. We use the notation developed in Section 4, denoting distance between cycles as $\Delta(E_i, E_j)$.

Lemma 5.7. *For any e and integer x , $\delta_x(e, e) = \tilde{\mathcal{O}}(\sqrt{n})$.*

Proof. We will call a cycle that contains at most \sqrt{n} distinct vertices a *small cycle*.

In the adjacency graph of cycles \mathcal{G} , any shortest path of length x between two cycles contains a cycle using at most $2n/x$ distinct vertices.

First, let us assume that there is at least one small cycle in the graph. Thus, for any cycle E_1 there is a small cycle E_2 such that $\Delta(E_1, E_2) \leq \frac{2n}{\sqrt{n}} = \mathcal{O}(\sqrt{n})$. Fix $e_1 \in E_1, e_2 \in E_2$. Since $\text{diam}(E_2) \leq \sqrt{n}$, by Lemma 5.9, $\delta_T(e_2, e_2) = \mathcal{O}(\sqrt{n})$ for any integer T . By Lemma 4.3 we then have $\delta_T(e_1, e_1) \leq \delta_T(e_2, e_2) + 2\Delta(E_1, E_2) = \mathcal{O}(\sqrt{n})$.

If there are no small cycles in the graph, then every cycle is of size at least \sqrt{n} distinct vertices. So $\text{diam}(\mathcal{G}) = \mathcal{O}(\sqrt{n})$ by the previous observation, and the claim follows from Lemma 4.4. \square

Next, the following bound trivially holds for small values of t , subsequently allowing us to perform analysis of the RR for values of k close to m .

Lemma 5.8. *For any e , $\delta_4(e, e) \leq 3$.*

Proof. Observe that since arcs $\text{pred}(-e) = \text{pred}(\varphi(e))$, $\delta_1(e, -e) \leq 1$. Thus, $\delta_2(e, e) \leq 2$. Similarly, one can reason that $\delta_2(\varphi^2(e), e) \leq 3$, giving $\delta_4(e, e) \leq 3$. \square

For our next bound, let us denote by $d(e_1, e_2)$ the distance (length of shortest path) between starting vertices of e_1 and e_2 , measured in G . We also denote $\text{diam}(G) = \max_{e_1, e_2 \in E} d(e_1, e_2)$ and $\text{diam}(E_i) = \max_{e_1, e_2 \in E_i} d(e_1, e_2)$. (For sanity of notation, we define $\text{diam}(E_i)$ so that $\text{diam}(E_i) \geq 1$ even if E_i is a loop.) The following Lemma relates cycle diameter with our analysis.

Lemma 5.9. *For any integer j and edge $e \in E_i$, we have: $\delta_j(e, e) = \mathcal{O}(\text{diam}(E_i)) \leq \mathcal{O}(\text{diam}(G))$.*

Proof. Assume $j > 0$ and fix an integer $r > 0$. Observe that $d(e, \varphi^r(e)) \leq \text{diam}(E_i)$, thus we have $\delta_{r+d(e, \varphi^r(e))}(e, e) \leq \text{diam}(E_i) + 1$. Since for consecutive values of r the value of $r + d(e, \varphi^r(e))$ increases by 0, 1 or 2, either the claim follows, or $\delta_{j-1}(e, e) \leq \text{diam}(E_i) + 1$. However, since in every vertex there is a loop, we have $\delta_1(e, e) \leq 2$ and the claim follows. \square

The following bound on the generalized similarity measure proves useful when considering sufficiently long ranges of time.

Lemma 5.10. *For any integer $B \leq m$ and for any e , we have: $\min_{x \in [B, 2B]} \delta_x(e, e) = \tilde{\mathcal{O}}(m/B)$.*

Proof. First, let us assume that there is at least one cycle with number of arcs not more than B . Thus, by a reasoning similar to one from Theorem 5.7, for any cycle E_1 and arc $e_1 \in E_1$ there exists a cycle E_2 having at most B arcs and that $\Delta(E_1, E_2) \leq \frac{m}{B}$. For any $e_2 \in E_2$ there is integer $B \leq x \leq 2B$ such that $\delta_x(e_2, e_2) = 0$ (it suffices for x to be a multiplicity of $|E_2|$). Thus for any $e_1 \in E_1$, we have $\delta_x(e_1, e_1) \leq 2m/B$. Since e_1 was picked arbitrarily, the claim follows.

If every cycle has at least B arcs, then the number of cycles satisfies $g \leq m/B$. Thus the claim follows from Lemma 4.4. \square

Finally, we obtain the following $\mathcal{O}(1)$ bound on the generalized similarity measure of the RR dynamics in trees. We remark that this bound holds even though trees are bipartite.

Lemma 5.11. *If G is a tree, then for any e and positive integer B , we have: $\min_{x \in [B, 2B]} \delta_x(e, e) = \mathcal{O}(1)$.*

Proof. Cycles in E are DFS traversals of some subtree of G . We can safely analyse each such subtree separately. Take such subtree T , and fix $B \leq 2|T|$. (The size here is measured in number of edges.) We pick arbitrary root of T , and we consider subtrees denoted T_v , that is rooted in some vertex v containing everything “below” v , plus the edge connecting v to its parent. One of two cases occurs: (1) there is v such that $2|T_v| \in [B, 2B]$, or (2) there is v with children w_1, w_2, \dots such that $2|T_v| > 2B$ and $2|T_{w_1}|, 2|T_{w_2}|, \dots < B$.

In the first case, the cycle that takes part in T_v is of size $2|T_v|$, thus we have for some $x \in [B, 2B]$ that $\delta_x(e', e') = 1$, for some e' in this cycle, thus also for every edge in this cycle.

In the second case, assume that w_1, w_2, \dots are in the order in which the cycle visits those vertices. Since $|T_{w_1}| + |T_{w_2}| + \dots = |T_v| - 1 \geq B$, and each of $|T_{w_i}|$ is smaller than $B/2$, we know there prefix sum such that $2|T_{w_1}| + 2|T_{w_2}| + \dots + 2|T_{w_j}| \in [B, 2B]$. Once again we argue that there is a self-intersection of appropriate length, and the claim follows. \square

5.4 Cumulated Load Discrepancy of RR

Lemma 5.6 provides us with bounds on the cumulated load discrepancy of the RR dynamics in terms of the generalized similarity measure $\delta_T(e, e)$. Introducing Lemmas 5.5, 5.9, and 5.7, we directly obtain the following claims.

Proposition 5.12. *The cumulated load discrepancy of RR dynamics in its recurrent state can be upper-bounded as:*

- $\tilde{\mathcal{O}}(\gcd(k, m))$.
- $\mathcal{O}(\text{diam}(G))$.
- $\tilde{\mathcal{O}}(\sqrt{n})$.

\square

5.5 Idleness of RR

Lemma 5.6 also provides a way of upper-bounding the idleness of the RR dynamics in terms of the generalized similarity measure $\delta_T(e, e)$. Indeed, we have the following claim.

Lemma 5.13. *If $T > \frac{m}{k} \cdot \delta_T(e, e)$ for some particular value of T , then $\text{idle}(e) \leq T$.*

Proof. By Lemma 5.6, we obtain $C_t^T(e) \geq \frac{k}{m} \cdot T - \delta_T(e, e) > 0$. Since this bound holds regardless of the choice of t , the claim follows. \square

The following claims are now obtained directly, by combining Lemma 5.13 for an appropriately chosen value of T with specific bounds on $\delta_T(e, e)$, given by Lemma 5.5 (putting $T = \tilde{\mathcal{O}}(\frac{m}{k} \cdot \gcd(k, m))$), Lemma 5.9 (putting $T = \mathcal{O}(\frac{m}{k} \cdot \text{diam}(G))$), Lemma 5.7 (putting $T = \tilde{\mathcal{O}}(\frac{m}{k} \cdot \sqrt{n})$), Lemma 5.8 (putting $T = 4$), 5.11 (putting $T = \arg \min_{x \in [c \cdot m/k, 2c \cdot m/k]} \delta_x(e, e)$ for arbitrary e and $c = \mathcal{O}(1)$ being the bound from Lemma 5.11), and Lemma 5.10 (putting $B = \tilde{\mathcal{O}}(m/\sqrt{k})$).

Proposition 5.14. *For $k \leq m$, the RR dynamics satisfies the following bounds on idleness:*

- $\text{idle}(G) = \tilde{\mathcal{O}}(\frac{m}{k} \cdot \gcd(k, m))$.
- $\text{idle}(G) = \mathcal{O}(\frac{m}{k} \cdot \text{diam}(G))$.
- $\text{idle}(G) = \tilde{\mathcal{O}}(\frac{m}{k} \cdot \sqrt{n})$.
- $\text{idle}(G) \leq 4$ if $k > \frac{3}{4}m$.
- $\text{idle}(G) = \mathcal{O}(\frac{m}{k})$, on a tree.
- $\text{idle}(G) = \tilde{\mathcal{O}}(\frac{m}{k} \cdot \sqrt{k})$.

\square

6 Conclusion

The obtained results show that the RR dynamics achieves almost optimal time-averaging behavior for a vast majority of parameter values. This has implications for the considered patrolling and load-balancing problems, as shown in Table 1, most notably in the case when m and k have no large common divisors. The case of an underlying Eulerian circulation has been fully resolved, and the mixing properties of the associated walks along Eulerian circuits are well understood. For the non-Eulerian case, it is an open question whether improved generalized δ -distance bounds can be obtained, independent of the number of cycles g of the considered circulation, with consequent improved bounds on the performance of the RR dynamics.

Acknowledgments

We thank Jurek Czyzowicz and Jukka Suomela for fruitful discussions. We also thank an anonymous referee for valuable comments.

References

- [1] Y. Afek and E. Gafni. Distributed algorithms for unidirectional networks. *SIAM J. Comput.*, 23(6):1152–1178, 1994.
- [2] N. Agmon, S. Kraus, and G. A. Kaminka. Multi-robot perimeter patrol in adversarial settings. In *ICRA*, pages 2339–2345, 2008.
- [3] H. Akbari and P. Berenbrink. Parallel rotor walks on finite graphs and applications in discrete load balancing. In *SPAA*, pages 186–195, 2013.
- [4] H. Akbari, P. Berenbrink, and T. Sauerwald. A simple approach for adapting continuous load balancing processes to discrete settings. In *PODC*, pages 271–280, 2012.
- [5] A. Almeida, G. Ramalho, H. Santana, P. A. Tedesco, T. Menezes, V. Corruble, and Y. Chevaleyre. Recent advances on multi-agent patrolling. In *SBIA*, pages 474–483, 2004.
- [6] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: An explanation of $1/f$ noise. *Physical Review Letters*, 59:381–384, 1987.
- [7] E. Bampas, L. Gąsieniec, N. Hanusse, D. Ilcinkas, R. Klasing, and A. Kosowski. Euler tour lock-in problem in the rotor-router model. In *DISC*, pages 423–435, 2009.
- [8] P. Berenbrink, R. Klasing, A. Kosowski, F. Mallmann-Trenn, and P. Uznański. Improved analysis of deterministic load-balancing schemes. In *PODC*, pages 301–310, 2015.
- [9] A. Björner, L. Lovász, and P. W. Shor. Chip-firing games on graphs. *European Journal of Combinatorics*, 12:283–291, 1991.
- [10] J. Chalopin, S. Das, P. Gawrychowski, A. Kosowski, A. Labourel, and P. Uznański. Limit behavior of the multi-agent rotor-router system. In *DISC*, pages 123–139, 2015.
- [11] Y. Chevaleyre. Theoretical analysis of the multi-agent patrolling problem. In *IAT*, pages 302–308, 2004.
- [12] J. Cooper and J. Spencer. Simulating a random walk with constant error. *Combinatorics, Probability and Computing*, 15:815–822, November 2006.
- [13] J. Czyzowicz, A. Kosowski, E. Kranakis, and N. Taleb. Patrolling trees with mobile robots. In *FPS*, volume 10128 of *Lecture Notes in Computer Science*, pages 331–344. Springer, 2016.
- [14] R. Dandekar and D. Dhar. Proportionate growth in patterns formed in the rotor-router model. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(11):P11030, 2014.
- [15] D. Dereniowski, A. Kosowski, D. Pająk, and P. Uznański. Bounds on the cover time of parallel rotor walks. In *STACS*, pages 263–275, 2014.
- [16] B. Doerr and T. Friedrich. Deterministic random walks on the two-dimensional grid. *Combinatorics, Probability and Computing*, 18(1-2):123–144, 2009.
- [17] Y. Elmaliach, N. Agmon, and G. A. Kaminka. Multi-robot area patrol under frequency constraints. *Ann. Math. Artif. Intell.*, 57(3-4):293–320, 2009.

- [18] Y. Elmaliach, A. Shiloni, and G. A. Kaminka. A realistic model of frequency-based multi-robot polyline patrolling. In *AAMAS (1)*, pages 63–70, 2008.
- [19] R. Elsässer and T. Sauerwald. Tight bounds for the cover time of multiple random walks. *Theor. Comput. Sci.*, 412(24):2623–2641, 2011.
- [20] A. S. Fraenkel. Economic traversal of labyrinths. *Mathematics Magazine*, 43:125–130, 1970.
- [21] T. Friedrich, M. Gairing, and T. Sauerwald. Quasirandom load balancing. In *SODA*, pages 1620–1629, 2010.
- [22] T. Friedrich and T. Sauerwald. Near-perfect load balancing by randomized rounding. In *STOC*, pages 121–130, 2009.
- [23] L. Gaşieniec and T. Radzik. Memory efficient anonymous graph exploration. In *WG*, pages 14–29, 2008.
- [24] A. Granville. An introduction to additive combinatorics. In *Additive Combinatorics*, CRM Proceedings and Lecture Notes, pages 1–27. American Mathematical Society, 2007.
- [25] N. Hazon and G. A. Kaminka. On redundancy, efficiency, and robustness in coverage for multiple robots. *Robotics and Autonomous Systems*, 56(12):1102–1114, 2008.
- [26] A. Holroyd and J. Propp. Rotor walks and Markov chains. *Algorithmic Probability and Combinatorics*, pages 105–126, 2010.
- [27] M. A. Kiwi, R. Ndoundam, M. Tchente, and E. Goles. No polynomial bound for the period of the parallel chip firing game on graphs. *Theor. Comput. Sci.*, 136(2):527–532, 1994.
- [28] R. Klasing, A. Kosowski, D. Pająk, and T. Sauerwald. The multi-agent rotor-router on the ring: a deterministic alternative to parallel random walks. In *PODC*, pages 365–374, 2013.
- [29] A. Kosowski and D. Pająk. A case study of cover time for the rotor-router. In *ICALP*, pages 544–555, 2014.
- [30] A. Machado, G. Ramalho, J.-D. Zucker, and A. Drogoul. Multi-agent patrolling: An empirical analysis of alternative architectures. In *MABS*, pages 155–170, 2002.
- [31] F. Pasqualetti, A. Franchi, and F. Bullo. On optimal cooperative patrolling. In *CDC*, pages 7153–7158, 2010.
- [32] V. Priezzhev, D. Dhar, A. Dhar, and S. Krishnamurthy. Eulerian walkers as a model of self-organized criticality. *Phys. Rev. Lett.*, 77(25):5079–5082, Dec 1996.
- [33] Y. Rabani, A. Sinclair, and R. Wanka. Local divergence of markov chains and the analysis of iterative load-balancing schemes. In *FOCS*, pages 694–703, 1998.
- [34] A. Ramachandran and A. Schild. Sandpile prediction on a tree in near linear time. In *SODA*, pages 1115–1131. SIAM, 2017.
- [35] T. Sauerwald and H. Sun. Tight bounds for randomized load balancing on arbitrary network topologies. In *FOCS*, pages 341–350, 2012.

- [36] T. Shiraga, Y. Yamauchi, S. Kijima, and M. Yamashita. Total variation discrepancy of deterministic random walks for ergodic markov chains. In *ANALCO*, pages 138–148, 2016.
- [37] R. Subramanian and I. Scherson. An analysis of diffusive load-balancing. In *SPAA*, pages 220–225, 1994.
- [38] T. Tao and V. H. Vu. *Additive Combinatorics*. Cambridge University Press, 2006. Cambridge Books Online.
- [39] V. Yanovski, I. A. Wagner, and A. M. Bruckstein. A distributed ant algorithm for efficiently patrolling a network. *Algorithmica*, 37(3):165–186, 2003.

A Proof of Proposition 3.4

Let $x, y \in A$ and $\xi \in \text{Spec}_{1-\varepsilon}(B - A)$. Then there exists a phase $\theta \in \mathbb{R}/(\eta \cdot \mathbb{Z})$ such that

$$\operatorname{Re} \sum_{z \in B-A} e(\xi \cdot z + \theta) \geq (1 - \varepsilon)|B - A|,$$

thus

$$\sum_{z \in B-A} (1 - \operatorname{Re} e(\xi \cdot z + \theta)) \leq \varepsilon|B - A| \leq \varepsilon K|B|.$$

Since the summand is non-negative, and $B - A$ contains $b - x$,

$$\sum_{b \in B} |1 - \operatorname{Re} e(\xi \cdot (b - x) + \theta)| \leq \varepsilon K|B|$$

and by Cauchy-Schwarz

$$\sum_{b \in B} |1 - \operatorname{Re} e(\xi \cdot (b - x) + \theta)|^{1/2} \leq \varepsilon^{1/2} K^{1/2} |B|.$$

From the elementary identity $|1 - e(\alpha)| = \sqrt{2}|1 - \operatorname{Re} e(\alpha)|^{1/2}$ we conclude that:

$$\sum_{b \in B} |1 - e(\xi \cdot (b - x) + \theta)| \leq \sqrt{2} \varepsilon^{1/2} K^{1/2} |B|.$$

A similar expression is obtained for x replaced by y . By the triangle inequality we conclude that

$$\sum_{b \in B} |e(\xi \cdot (b - y) + \theta) - e(\xi \cdot (b - x) + \theta)| \leq 2\sqrt{2} \varepsilon^{1/2} K^{1/2} |B|.$$

The left hand side is equal to:

$$\sum_{b \in B} |e(\xi \cdot (b - y) + \theta)| \cdot |e(\xi \cdot (x - y)) - 1| = |B| \cdot |e(\xi \cdot (x - y)) - 1|,$$

thus we have:

$$|e(\xi \cdot (x - y)) - 1| \leq \sqrt{8\varepsilon K}.$$

Since we picked $\xi \in \text{Spec}_{1-\varepsilon}(B - A)$ and $x, y \in A$ arbitrarily, the claim follows.