

Lock-in Problem for Parallel Rotor-router Walks

Jérémie Chalopin, Shantanu Das, Pawel Gawrychowski, Adrian Kosowski,
Arnaud Labourel, Przemyslaw Uznanski

► **To cite this version:**

Jérémie Chalopin, Shantanu Das, Pawel Gawrychowski, Adrian Kosowski, Arnaud Labourel, et al..
Lock-in Problem for Parallel Rotor-router Walks. 2015. <hal-01021930v3>

HAL Id: hal-01021930

<https://hal.inria.fr/hal-01021930v3>

Submitted on 27 May 2015

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.

Lock-in Problem for Parallel Rotor-router Walks

Jérémie Chalopin¹, Shantanu Das¹, Paweł Gawrychowski², Adrian Kosowski³, Arnaud Labourel¹,
and Przemysław Uznański*⁴

¹LIF, CNRS and Aix-Marseille University, France

²Max-Planck-Institut für Informatik, Saarbrücken, Germany

³Inria – LIAFA – Paris Diderot University, France

⁴Helsinki Institute for Information Technology HIIT, Department of Computer Science, Aalto University, Finland

Abstract

The *rotor-router* model, also called the *Propp machine*, was introduced as a deterministic alternative to the random walk. In this model, a group of identical tokens are initially placed at nodes of the graph. Each node maintains a cyclic ordering of the outgoing arcs, and during consecutive turns the tokens are propagated along arcs chosen according to this ordering in round-robin fashion. The behavior of the model is fully deterministic. Yanovski *et al.* (2003) proved that a single rotor-router walk on any graph with m edges and diameter D stabilizes to a traversal of an Eulerian circuit on the set of all $2m$ directed arcs on the edge set of the graph, and that such periodic behaviour of the system is achieved after an initial transient phase of at most $2mD$ steps.

The case of multiple parallel rotor-routers was studied experimentally, leading Yanovski *et al.* to the conjecture that a system of $k > 1$ parallel walks also stabilizes with a period of length at most $2m$ steps. In this work we disprove this conjecture, showing that the period of parallel rotor-router walks can in fact, be superpolynomial in the size of graph. On the positive side, we provide a characterization of the periodic behavior of parallel router walks, in terms of a structural property of stable states called a *subcycle decomposition*. This property provides us the tools to efficiently detect whether a given system configuration corresponds to the transient or to the limit behavior of the system. Moreover, we provide polynomial upper bounds of $\mathcal{O}(m^4 D^2 + mD \log k)$ and $\mathcal{O}(m^5 k^2)$ on the number of steps it takes for the system to stabilize. Thus, we are able to predict any future behavior of the system using an algorithm that takes polynomial time and space. In addition, we show that there exists a separation between the stabilization time of the single-walk and multiple-walk rotor-router systems, and that for some graphs the latter can be asymptotically larger even for the case of $k = 2$ walks.

1 Introduction

Dynamical processes occurring in nature provide inspiration for simple, yet powerful distributed algorithms. For example, the *heat equation*, which describes real-world processes such as heat and particle diffusion, also proves useful when designing schemes for load-balancing and token rearrangement in a discrete graph scenario. In the diffusive model of load-balancing on a network, each node of the network is initially endowed with a certain load value, and in each step it distributes a fixed proportion of its load evenly among its neighbors. Given that such a balancing operation is performed for load which is infinitely divisible (so-called *continuous diffusion*), in the long term the distribution of load converges on a degree-regular network to uniform over all nodes. When load is composed of indivisible unit tokens, the continuous diffusion process is no longer practicable. It is, however, possible to design randomized schemes in which the *expected* value of load of each node at each moment of time corresponds precisely

*Part of the work was done while the author was affiliated with LIF, CNRS and Aix-Marseille University (supported by the Labex Archimède and by the ANR project MACARON (ANR-13-JS02-0002)).

to the value of its load in the corresponding continuous diffusion process. This may be achieved, for instance, by allowing each token of load to follow an independent random walk on the network, as well as by applying more refined techniques admitting stronger concentration of the load distribution, cf. [24]. Such methods are stochastic in their very nature, and it is natural to ask whether there exist *deterministic* methods which mimic this type of stochastic load balancing behavior? The answer is affirmative, with the natural candidate process being the so-called *rotor-router* model.

Formally, the rotor-router mechanism is represented by an undirected anonymous graph $G = (V, E)$. Initially, a set of identical tokens is released on a vertex of the graph. At discrete, synchronous steps, the tokens are propagated according to the deterministic round robin rule, where after sending each token, the pointer is advanced to the next exit port in the fixed cyclic ordering. Such a mechanism has been proposed as a viable alternative to stochastic and random-walk-based processes in the context of load balancing problems [7, 9, 11], exploration of graphs [1, 10, 15, 17, 20], and stabilization of distributed processes [4, 8, 22, 29].

The resemblance between the rotor-router token distribution mechanism and stochastic balancing processes based on continuous diffusion is at least twofold, in that: (1) the number of tokens on each node for the rotor-router process has a bounded discrepancy with respect to that in the continuous diffusion process [6, 25], and (2) when performing time-averaging of load over sufficiently long time intervals, the observed load averages for all nodes in the rotor-router process converge precisely to their corresponding value for the continuous diffusion process.

By contrast to time-averaged load, for any *fixed* moment of time, the deterministic rotor-router process and the stochastic approaches exhibit important differences. A stochastic load balancing process based on tokens following random walks leads the system towards a “heat death” stochastic state, which is completely independent of the starting configuration. By contrast, the rotor-router process is a deterministic process on a graph and its limit behavior may be much more interesting.* The number of possible configurations of a rotor-router system is finite, hence, after a transient initial phase, the process must stabilize to a cyclic sequence of states which will be repeated ever after. Natural questions arise, concerning the eventual structural behavior observed in this limit cycle of the rotor-router system, the length of the limit cycle, and the duration of the stabilization phase leading to it. So far, the only known answer concerned the case when only a single token is operating in the entire system. Yanovski *et al.* [29] showed that such a single token stabilizes within a polynomial number of steps to periodic behavior, in which it performs a traversal of some Eulerian cycle on the directed version of the network graph.

In this work, we provide a complete structural characterization of the limit behavior of the rotor-router for an *arbitrary* number $k > 1$ of tokens. The obtained characterization shows that the rotor-router mechanism provides a way of self-organizing tokens, initially spread out arbitrarily over a graph, into balanced groups, each of which follows a well-defined walk in some part of the network graph. The practical implications of our result may be seen as twofold. On the one hand, when viewing the rotor-router as a load-balancing process, we obtain a better understanding of its limit behavior. On the other hand, when considering each of the tokens as a walker in the graph, we show that the rotor-router may prove to be a viable strategy for perpetual graph exploration, with possible applications in so-called network patrolling problems.

1.1 Related Work

Load balancing. The rotor-router mechanism of token distribution has been considered in problems of balancing workload among network nodes for specific network topologies. In this context, each token is considered as a unit-length task to be performed by one of the processors in a network of computers. Cooper and Spencer [9] studied load balancing with parallel rotor walks in d -dimensional grid graphs and showed a constant bound on the discrepancy between the number of tokens at a given node v in the rotor-router model and the expected number of tokens at v in the random-walk model. The structural properties of the distribution of tokens for a rotor-router system on the 2-dimensional grid were considered by Doerr and Friedrich [11]. Akbari and Berenbrink [2] proved an upper bound of $\mathcal{O}(\log^{3/2} n)$ on the load-balancing discrepancy for hypercubes, and for tori of constant dimensions, they showed that the discrepancy is bounded by a constant. For general d -regular graphs, a bound of $\mathcal{O}(d \log n / \mu)$ on the discrepancy of the rotor-router mechanism with respect to continuous diffusion follows from the general framework of [23], where

*Perhaps the first work to highlight the importance of differences between limit properties of deterministic and stochastic variants of a token-based discrete process on a graph was that of M. Kac, in the setting of statistical mechanics (cf. [26][Section 1.9] for a comprehensive discussion).

μ is the eigenvalue gap of the graph, under the assumption that a sufficient number of self-loops are present at each node of the graph. This discrepancy bound has recently been improved to $O(d\sqrt{\log n/\mu})$ in [6].

Graph exploration. The walks of tokens following fixed local rules at nodes provide a local mechanism of graph exploration. Each token, starting at a node of the graph, moves at each step to one of the adjacent nodes, until it has explored all the nodes of the graph. Such mechanisms tend to be location-oblivious and robust, displaying resilience to changes in network topology. In some cases, we may require the token to periodically visit all nodes, *e.g.*, with the goal of monitoring the network or distributing updates.

One basic exploration technique relies on multiple tokens, each of which follows an independent random walk: at each step, each token chooses one of the arcs incident to the current node uniformly at random and traverses it. The performance of such parallel random walks have been analyzed by Alon *et al.* [3], Efremenko and Reingold [12], and Elsässer and Sauerwald [13], who have demonstrated that in terms of the expected time until all nodes have been visited by at least one token (i.e., the cover time), parallelization brings about a speedup of between $\Theta(\log k)$ and $\Theta(k)$ when running random walks with k parallel tokens.

The rotor-router mechanism has also been studied in the context of graph exploration, sometimes under the name of *Edge Ant Walks* [27, 29], and in the context of traversing a maze and marking edges with pebbles, *e.g.* in [8]. Cover times of rotor-router systems have been investigated by Wagner *et al.* [27] who showed that starting from an arbitrary initial configuration*, a single token following the rotor-router rule explores all nodes of a graph on n nodes and m edges within $\mathcal{O}(nm)$ steps. Later, Bhatt *et al.* [8] showed that after at most $\mathcal{O}(nm)$ steps, the token continues to move periodically along an Eulerian cycle of the (directed symmetric version of the) graph. Yanovski *et al.* [29] and Bampas *et al.* [4] studied the stabilization time and showed that the token starts circulating in the Eulerian cycle within $\Theta(mD)$ steps, in the worst case, for a graph of diameter D . Studies of the rotor router system for specific classes of graphs were performed in [16]. While all these studies were restricted to static graphs, Bampas *et al.* [5] considered the time required for the rotor-router to stabilize to a new Eulerian cycle after an edge is added or removed from the graph.

Studies of the parallel (*i.e.*, multiple token) rotor-router were performed by Yanovski *et al.* [29] and Klasing *et al.* [19], and the speedup of the system due to parallelization was considered for both worst-case and best-case scenarios. In [10], Dereniowski *et al.* establish bounds on the minimum and maximum possible cover time for a worst-case initialization of a k -rotor-router system in a graph G with m edges and diameter D , as $\Omega(mD/k)$ and $\mathcal{O}(mD/\log k)$ respectively. In [20], Kosowski and Pająk provided a more detailed analysis of the speedup for specific classes of graphs, providing tight bounds of cover-time speed-up for all values of k for degree-restricted expanders, random graphs, and constant-dimensional tori. For hypercubes, they resolve the question precisely, except for values of k much larger than n .

1.2 Our Results

In this work we provide a structural characterization of the limit behavior of the rotor-router model with multiple tokens. Yanovski *et al.* [29] conjectured that the rotor-router system enters a short sequence of states (of length at most $2m$), which repeats cyclically ever after. We start this work by disproving this conjecture. In fact, we display an example of a starting configuration which admits a limit cycle with a period of superpolynomial length ($\exp(\Omega(\sqrt{n \log n}))$) with respect to the size of the graph. Our example is similar to the construction presented by Kiwi *et al.* [18] to prove the existence of super-polynomial periods for chip firing games on graphs (although the rules of chip firing games are only very loosely related to those of the rotor-router).

By contrast, it turns out the fact that the rotor-router admits long limit cycles does not signify that the limit behavior of the rotor-router should be perceived as a “disordered” discrete dynamical system. The long period in our counterexample comes from the system being composed from many smaller parts, each of which exhibits a small (but different) period length. We show that for any limit sequence of states in the rotor-router model, the graph can be partitioned into arc-disjoint directed Eulerian cycles, with each token in the limit periodically traversing arcs of one particular cycle. We name such behavior a *subcycle decomposition*, the exact properties of which are described in Section 3. To

*A configuration is defined by: the cyclic order of outgoing arcs, the initial pointers at the nodes, and the current location of the token.

complement the lower bound, we provide an upper bound of $\exp(\mathcal{O}(\sqrt{m \log m}))$ on the period of parallel rotor walks in its limit behavior. This upper bound asymptotically almost matches the lower bound from our example.

There are several consequences of our structural characterization of the limit behavior of the rotor-router. First, we show that it is possible to determine efficiently whether the system has already stabilized (*i.e.*, reached a configuration that will repeat itself) or not. This detection is based on the analysis of the properties of stable states, that is, of how the tokens arriving at a node are distributed into groups leaving on different outgoing arcs. The main point of this analysis is the observation that the cumulative number of tokens entering a vertex v (over the time period $\{t, (t+1), \dots, (t+\Delta t)\}$) is equal to the cumulative number of tokens leaving vertex v (over time $\{(t+1), (t+2), \dots, (t+\Delta t+1)\}$).

Next, by defining an appropriate potential of a system and showing its monotonicity, we can give a polynomial bound on a number of steps necessary for a system with an arbitrary initialization to reach a periodic configuration. We provide an upper bound of $\mathcal{O}(m^4 D^2 + mD \log k)$, together with examples of graphs with initial configuration having just 2 tokens that require $\Omega(m^2 \log n)$ steps. This analysis is presented in Section 4. The obtained polynomial upper bound means that the rotor-router is an efficient means of self-organizing tokens so as to perform a periodic traversal of the edges of the graph.

Finally, Section 5 is dedicated to showing how the previous results can be applied in a constructive way with regard to efficient simulation of a rotor-router system. We show how the properties of subcycle decomposition can be applied to provide a way to preprocess any starting configuration in a way that makes it possible to answer queries of certain type in a polynomial time. This shows that a structural characterization of the rotor-router system is not only an important as a theoretical tool for understanding the limit behavior of the system, but it also as a practical tool for solving certain problems related to the rotor-router system.

As a complementary result, we show for the single-token rotor-router how to efficiently compute the Eulerian traversal cycle on which the token would be locked-in, faster than by running the process directly. A naive simulation would take $\mathcal{O}(mD)$ time, but by using the structural properties of a single token walk together with application of efficient data structures we show how to preprocess the input graph in time $\mathcal{O}(n+m)$ such that we can answer queries about token position at any given time T , in $\mathcal{O}(\log \log m)$ time per query.

2 Model and Preliminaries

Let $G = (V, E)$ be an undirected connected graph with n nodes, m edges and *diameter* D . Let k be the number of tokens. The digraph $\vec{G} = (V, \vec{E})$ is the directed version of G created by replacing every edge (u, v) with two directed arcs \vec{uv} and \vec{vu} . We will refer to the undirected links in graph G as *edges* and to the directed links in the graph \vec{G} as *arcs*. Given a vertex v , we will denote its set of incoming arcs by $\text{in}(v)$ and outgoing arcs by $\text{out}(v)$. Each vertex v of G is equipped with a fixed ordering of all its outgoing arcs $\rho_v = (e_1, e_2, \dots, e_{\deg(v)})$.

The precise definition of the rotor-router model on the system defined by $(\vec{G}, (\rho_v)_{v \in V})$ is as follows: A *state* at the current time step t is a tuple: $\mathcal{S}_t = ((\text{pointer}_v)_{v \in V}, (\text{tokens}_v)_{v \in V})$, where pointer_v is an arc outgoing from node v , which is referred to as *the current port pointer at node v*, and tokens_v is the number of tokens at any given node. For an arc (\vec{vu}) , let $\text{next}(\vec{vu})$ denote the arc after the arc (\vec{vu}) 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 token to pointer_v ,
2. Set $\text{pointer}_v = \text{next}(\text{pointer}_v)$.

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.

For a given state \mathcal{S}_t , we say that it is *stable* iff there exists $t' > t$ such that $\mathcal{S}_{t'} = \mathcal{S}_t$. *The stabilization time* of state \mathcal{S}_0 , denoted t_s , is the smallest value such that \mathcal{S}_{t_s} is stable. We call *the periodicity* of state \mathcal{S}_0 the smallest $t_p > 0$ such that $\mathcal{S}_{t_s} = \mathcal{S}_{t_s+t_p}$.

Throughout the paper, we denote multisets using $\{\{\}\}$ notation, while for integer ranges, we write $[a..b] \stackrel{\text{def}}{=} \{a, a+1, \dots, b\}$, $[a..b) \stackrel{\text{def}}{=} \{a, a+1, \dots, b-1\}$.




Figure 1: (a) The balloon graph and the initial token distribution. (b) The family of graphs \mathcal{G}_r consisting of r balloons.

3 Periodicity of the Rotor-Router System

We begin with the observation that knowledge of the first $t_s + t_p$ states of the system, that is $\mathcal{S}_0, \dots, \mathcal{S}_{t_s+t_p-1}$, gives us full knowledge of any future state for arbitrarily large time $t \geq t_s$: $\mathcal{S}_t = \mathcal{S}_{t_s + ((t-t_s) \bmod t_p)}$.

So as to be able to efficiently predict the future evolution of any rotor-router state, it would be useful to put a polynomial bound on t_p and t_s (with respect to n, m and k). If $k = 1$, due to results from Yanovski *et al.* [29], we know that $t_p = 2m$ and $t_s = O(mD)$. For arbitrary k , Yanovski *et al.* [29] conjectured that $t_p \leq 2m$ for any graph G regardless of the initial state. However, the following negative result disproves their conjecture and shows that the periodicity cannot be polynomially bounded for parallel rotor-routers.

Theorem 3.1. *There exists a family of graphs and initial states, with $k = 2m$ tokens, having the periodicity $t_p = 2^{\Omega(\sqrt{n \log n})}$.*

Proof. We will construct such a family of graphs \mathcal{G}_r for any sufficiently large integer r , and an appropriate initial configuration of tokens. First consider a balloon graph G_x consisting of a cycle of $x > 3$ vertices $\{v_0, v_1, \dots, v_{x-1}\}$ and an additional vertex v_x (called the *base vertex*) that is joined by an edge to vertex v_{x-1} of the cycle (see Figure 1(a)). Let the initial token distribution at vertices (v_0, \dots, v_x) be $(1, 2, 2, \dots, 2, 4, 1)$. Further let the exit pointers at vertex v_i , $0 \leq i \leq x - 2$ be oriented towards $v_{i-1 \bmod x}$ (in the counter-clockwise direction, in the figure), while at the vertex v_{x-1} the exit pointer is oriented towards v_0 (i.e. in the opposite direction). At the base vertex v_x there is only one outgoing arc and so, the exit pointer at v_x will always point towards this arc.

Observe that for a vertex of out-degree two, the exit pointer remains unchanged if an even number of tokens exit this vertex in the current round, while the exit pointer is rotated if an odd number of tokens exit in the current round.

We will now analyze the movement of tokens along the arcs of the graph in each round. Figure 2 shows an example for a balloon graph G_x , where $x = 5$. During the first round, the number of tokens moving on the arcs $(v_0, v_1), (v_1, v_2), \dots, (v_{x-1}, v_0)$ of the cycle (in the clockwise direction) is given by the sequence $S_0 = (0, 1, 1, \dots, 1, 2)$. During the same round, the number of tokens moving on the arcs in the counter-clockwise direction on the cycle is given by $(1, 1, \dots, 1)$. On the branch edge (v_{x-1}, v_x) there is exactly one token moving in each direction (See Figure 2(a)).

During the second round, the number of tokens moving on the arcs $(v_0, v_1), (v_1, v_2), \dots, (v_{x-1}, v_0)$ of the cycle (in the clockwise direction) is given by the sequence $S_1 = (1, 1, \dots, 1, 2, 0)$ which is a cyclic rotation of the sequence S_0 . The number of tokens moving on the arcs in the counter-clockwise direction on the cycle is still given by $(1, 1, \dots, 1)$. Again, the branch edge (v_{x-1}, v_x) has exactly one token moving in each direction (See Figure 2(b)).

Continuing with the above analysis, it is easy to see that in subsequent rounds, the number of tokens moving on the arcs of the cycle (in the clockwise direction) is given by cyclic rotations of S_0 , i.e., by the sequences $(1, \dots, 1, 2, 0, 1), (1, \dots, 1, 2, 0, 1, 1), (1, \dots, 1, 2, 0, 1, 1, 1)$ and so on (See Figure 2(c-e)). The number of tokens moving along the cycle in the counterclockwise direction is always one token per arc of the cycle. On the branch edge (v_0, v_x) there is




Figure 2: The number of tokens circulating on the arcs of a balloon graph G_5 in each round.

exactly one token moving in each direction in each round. Since the length of the sequence S_0 is $|S_0| = x$, after every x steps the configuration of tokens moving on the arcs of the cycle is the same. In other words, the periodicity of this rotor-router system is x . Notice that the graph G_x has $x + 1$ vertices and $2(x + 1)$ arcs, and there are exactly $2(x + 1)$ tokens in the system.

We will now construct the family of graphs \mathcal{G}_r . For any given r , let p_1, p_2, \dots, p_r be the first r prime numbers starting from $p_1 = 3$. We take r balloon graphs of sizes $(1 + p_1), (1 + p_2), \dots, (1 + p_r)$ respectively and join them by merging all the base vertices into one vertex (see Figure 1(b)). In each balloon graph we place the tokens as before, such that the merged base vertex now contains r tokens. During each step, r tokens will exit the base vertex through the r outgoing arcs and r other tokens will enter the base vertex through the r incoming arcs. Thus, irrespective of the initial state of the exit pointer at the base vertex, the system will behave in the same manner. The behavior of the system in the distinct balloons would be independent of each other and for each balloon of size $(1 + p_i)$ the configuration of the balloon would repeat itself in exactly p_i steps as before. Thus, the global state of the system would repeat in

$\text{lcm}(p_1, \dots, p_r) = \prod_{i=1}^r p_i$ steps. Note that the size of the graph, \mathcal{G}_r , is given by $n = 1 + \sum_{i=1}^r p_i = \Theta(r^2 \log r)$. In general, for any given integer n , we can construct a similar example graph by partitioning the $n - 1$ vertices into balloons of appropriate sizes joined to the n th vertex, such that the period of the system is equal to the *Landau function* [21] $g(n - 1) = 2^{\Omega(\sqrt{n \log n})}$. □

We remark that a similar result exists for parallel chip-firing games [18].

We now present an upper bound on the periodicity of k parallel rotor walks, for arbitrary values of k . First, we will show that even though a stable state can exhibit very long (super-polynomial) periodicity, the underlying graph G can be partitioned into parts, such that each part separately exhibits small (linear) periodicity.

We will use calligraphic large letters (e.g. $\mathcal{L} : \bar{E} \cup V \rightarrow \mathbb{Z}$) to denote *token distributions*. Thus:

- $\mathcal{L}_t(v)$ (*load* of node v) is number of tokens located at node v in time step t ,
- $\mathcal{L}_t(e)$ (*load* of arc e) is number of tokens sent out on arc e at time step t .

Thus, although tokens cannot be located on edges in our model, we can view the tokens in vertex as located already on the outgoing ports that they will be distributed to.

We will use specifically \mathcal{L}_t to denote token distribution associated with state \mathcal{S}_t . (It is important to note, that it is possible for two states to satisfy $\forall_e \mathcal{L}_t(e) = \mathcal{L}_{t'}(e)$ and yet $\mathcal{S}_t \neq \mathcal{S}_{t'}$, as we also require that pointers be in the same positions in identical states.)

We begin with a series of observations on the token distribution process in a rotor-router system.

Observation 3.2. *Since every token is pushed onto some outgoing arc, we have:*

$$\sum_{e \in \text{out}(v)} \mathcal{L}_t(e) = \mathcal{L}_t(v); \sum_{e \in \text{in}(v)} \mathcal{L}_t(e) = \mathcal{L}_{t+1}(v).$$

We also generalize token distribution into the *cumulative token distribution*. Given two time steps $t_1 \leq t_2$, we define $\mathcal{C}_{t_1}^{t_2} \stackrel{\text{def}}{=} \sum_{t \in [t_1 .. t_2]} \mathcal{L}_t$, in particular, for a vertex v and arc e :

$$\mathcal{C}_{t_1}^{t_2}(v) \stackrel{\text{def}}{=} \sum_{t \in [t_1 .. t_2]} \mathcal{L}_t(v), \mathcal{C}_{t_1}^{t_2}(e) \stackrel{\text{def}}{=} \sum_{t \in [t_1 .. t_2]} \mathcal{L}_t(e).$$

Consequently, a natural generalization of Observation 3.2 from load to cumulative load is as follow:

Observation 3.3. $\sum_{e \in \text{out}(v)} \mathcal{C}_{t_1}^{t_2}(e) = \mathcal{C}_{t_1}^{t_2}(v); \sum_{e \in \text{in}(v)} \mathcal{C}_{t_1}^{t_2}(e) = \mathcal{C}_{t_1+1}^{t_2+1}(v)$.

The next observation follows from the fact that rotor-router distributes tokens in a round-robin fashion among all outgoing arcs of a vertex:

Observation 3.4. $\forall_{e_1, e_2 \in \text{out}(v)} |\mathcal{C}_{t_1}^{t_2}(e_1) - \mathcal{C}_{t_1}^{t_2}(e_2)| \leq 1$.

Since arbitrarily large discrepancies (between two incoming edges) of incoming number of tokens are smoothed discretely, we can see the process of token propagation as a load-balancing scheme.

We now define the concept of *potential* of a token distribution system, which will be helpful to derive the necessary and sufficient conditions for a system state to be stable.

Definition 3.5. *Given a token distribution \mathcal{A} over edges, we define its potential as: $\Phi(\mathcal{A}) \stackrel{\text{def}}{=} \sum_{e \in \bar{E}} (\mathcal{A}(e))^2$.*

We also introduce a shorthand notation for the i -th potential of a given rotor-router state \mathcal{S}_t as: $\Phi_i(\mathcal{S}_t) \stackrel{\text{def}}{=} \Phi(\mathcal{C}_t^{t+i}) = \sum_{e \in \bar{E}} (\mathcal{C}_t^{t+i}(e))^2$.

Note that $\Phi_1 \equiv \Phi$. It is important to note that while arbitrary convex function can be used in the potential definition, usage of quadratic function will prove advantageous when analyzing the speed of convergence to a stable state, not only its properties.

The following folklore lemma provides us with a characterization of the minimum of the potential sums.

Lemma 3.6. *Over all partitions of integer S into d integers, the partition $\{\{\lfloor \frac{S}{d} \rfloor, \dots, \lfloor \frac{S}{d} \rfloor, \lceil \frac{S}{d} \rceil, \dots, \lceil \frac{S}{d} \rceil\}\}$ uniquely minimizes the value of sum of squares of elements.*

All omitted proofs are provided in the Appendix.

Lemma 3.7. For arbitrary i and t , the i -th potential is non-increasing: $\Phi_i(\mathcal{S}_{t+1}) \leq \Phi_i(\mathcal{S}_t)$.

Proof. To prove the lemma we have to observe how the round-robin property of the rotor-router acts locally on the groups of tokens (cumulative over the time interval $[t, t+i)$). From Observation 3.3 we know, that:

$$\sum_{e \in \text{out}(v)} \mathcal{C}_{t+1}^{t+i+1}(e) = \mathcal{C}_{t+1}^{t+i+1}(v) = \sum_{e \in \text{in}(v)} \mathcal{C}_t^{t+i}(e).$$

However, from Observation 3.4 and Lemma 3.6 we get that the multiset of values over outgoing arcs *minimizes* the sum of squares. Thus:

$$\sum_{e \in \text{out}(v)} (\mathcal{C}_{t+1}^{t+i+1}(e))^2 \leq \sum_{e \in \text{in}(v)} (\mathcal{C}_t^{t+i}(e))^2,$$

which leads to:

$$\Phi_i(\mathcal{S}_{t+1}) = \sum_v \sum_{e \in \text{out}(v)} (\mathcal{C}_{t+1}^{t+i+1}(e))^2 \leq \sum_v \sum_{e \in \text{in}(v)} (\mathcal{C}_t^{t+i}(e))^2 = \Phi_i(\mathcal{S}_t).$$

□

Observe that Lemma 3.7 implies that if the system is stable, all of the potentials are preserved at every (future) time step. This observation is powerful enough to derive strong characterization of stable states (see Theorem 3.11, equivalence of (i) and (ii)). However, in order to be able to reason about bounds on stabilization time, we need a more powerful notion of being able to characterize even the temporary regularities in token trajectories (for not necessarily stable states).

Definition 3.8. We say that a state \mathcal{S}_T admits a Δt -step subcycle decomposition, if in every vertex v we can define a one-to-one mapping between incoming and outgoing arcs of v $M_v : \text{in}(v) \rightarrow \text{out}(v)$, such that:

$$\forall_{e \in \text{in}(v)} \forall_{t \in [T..T+\Delta t)} \mathcal{L}_t(e) = \mathcal{L}_{t+1}(M_v(e)). \quad (1)$$

The subcycle decomposition has the following equivalent interpretation. We partition $\vec{E} = \vec{E}_1 \cup \dots \cup \vec{E}_c$, such that each \vec{E}_i induces a strongly-connected subgraph of G , and for each \vec{E}_i there exists an Eulerian cycle covering it such that each token traversing arcs of \vec{E}_i follows this particular Eulerian cycle during time steps $T, T+1, \dots, T+\Delta t-1$.

Observe that the mapping M in Definition 3.8 does not need to be necessarily unique. We will call any such mapping M a *valid mapping with respect to Δt subcycle decomposition* if (1) holds for it.

The following lemma gives a series of equivalent characterizations of subcycle decomposition, connecting the existence of such a decomposition during any time interval with lack of potential drop during the time interval, as well as a load-balancing discrepancy criterion over all shorter sub-intervals of time.

Lemma 3.9. The following statements are equivalent:

- (i) \mathcal{S}_T admits a Δt -step subcycle decomposition,
- (ii) $\forall_v \forall_{t, t': [t..t') \subseteq [T..T+\Delta t)}$, $\{\{\mathcal{C}_t^{t'}(e)\}_{e \in \text{in}(v)}\} = \{\{\mathcal{C}_{t+1}^{t'+1}(e)\}_{e \in \text{out}(v)}\}$ (multisets of cumulative loads are preserved locally),
- (iii) $\forall_{t, t': [t..t') \subseteq [T..T+\Delta t)}$ $\{\{\mathcal{C}_t^{t'}(e)\}_{e \in \vec{E}}\} = \{\{\mathcal{C}_{t+1}^{t'+1}(e)\}_{e \in \vec{E}}\}$ (multisets of cumulative loads are preserved globally),
- (iv) $\forall_{0 \leq i \leq \Delta t} \Phi_i(\mathcal{S}_T) = \Phi_i(\mathcal{S}_{T+1}) = \dots = \Phi_i(\mathcal{S}_{T+\Delta t-i+1})$ (potential is constant),
- (v) $\forall_v \forall_{e_1, e_2 \in \text{in}(v)} \forall_{t, t': [t..t') \subseteq [T..T+\Delta t)}$ $|\mathcal{C}_t^{t'}(e_1) - \mathcal{C}_t^{t'}(e_2)| \leq 1$ (incoming discrepancy is at most one).

For a fixed value of $\Delta t = 1$, Lemma 3.9 captures the property that as long as $\Phi(\mathcal{S}_T)$ remains constant, the loads on edges are only permuted between any two consecutive timesteps. Coupled with Lemma 3.7 it immediately implies aforementioned property. Unfortunately, this property is not strong enough for our purposes. However, for arbitrary values of Δt , we can still employ the notion of higher order potentials $\Phi_{\Delta t}$ (as defined previously), and observe the load balancing properties from Observation 3.3. The fact that stable state, by Lemma 3.7 and Lemma 3.9 admits load balancing properties even when collapsing multiple timesteps into a single frame, is our lever which will be used to derive strong properties of such states in the rest of this section.

Definition 3.10. *We say that a state \mathcal{S}_t admits a ∞ -subcycle decomposition if it admits i -steps subcycle decomposition for arbitrarily large i .*

Now we proceed to obtain a more algorithmic characterization of stable states. First, we show that if we do not experience a potential drop during $2m^2$ time steps, then the rotor-router system has reached its limit configuration.

Theorem 3.11. *The following conditions are equivalent:*

- (i) \mathcal{S}_T is stable,
- (ii) \mathcal{S}_T admits a ∞ -subcycle decomposition,
- (iii) \mathcal{S}_T admits a $(2m^2)$ -subcycle decomposition.

As a direct consequence of the proof of Theorem 3.11, we have:

Corollary 3.12. *For a stable state \mathcal{S}_T , any mapping between incoming and outgoing arcs of v denoted $M_v : \text{in}(v) \rightarrow \text{out}(v)$ that is valid with respect to $2m^2$ -subcycle decomposition, is also valid with respect to ∞ -subcycle decomposition.*

We are now ready to provide a stronger characterization of a stable state in Theorem 3.13 (compared to Theorem 3.11), based on refined analysis of the potential behavior.

Theorem 3.13. *State \mathcal{S}_T is stable iff: $\sum_{i=1}^{3m} \Phi_i(\mathcal{S}_T) = \sum_{i=1}^{3m} \Phi_i(\mathcal{S}_{T+2m^2})$.*

Finally, we provide an upper bound on the length of the period for any rotor-router state. It is interesting to see that the upper bound is not far from the period of the example graph given in Theorem 3.1.

Theorem 3.14. *For any stable state \mathcal{S}_T , the period length of the limit cycle is bounded by $t_p = \mathcal{O}(\exp(\sqrt{(m \log m)}))$.*

Proof. Observe, that any arc in G can be part of exactly one cycle in any given subcycle decomposition. As the period length of any stable state is upperbounded by the least common multiple of the length of the cycles, we get the desired upper bound as the value of the *Landau's function* on the total number of arcs in G . \square

4 Stabilization Time of the Rotor-Router System

In this section we provide upper and lower bounds on the stabilization time of parallel rotor-router systems. Since the values of potentials are discrete and non-increasing, in Theorem 3.13 we have a very powerful tool to reason about the stabilization of a state — if the sum of potentials remains unchanged for more than $3m$ time steps, the system has reached a stable state. Thus, we can naively bound each of the potentials by $\mathcal{O}((mk)^2)$, and so also bound the sum of potentials by $\mathcal{O}(m \cdot (mk)^2)$. This gives the following corollary.

Corollary 4.1. *For any initial state \mathcal{S}_0 , there exists $T = \mathcal{O}(m^5 k^2)$ such that \mathcal{S}_T is stable.*

We will now show how to obtain an even stronger bound (in terms of dependence on k), but for this we need a refined upper bound on initial potential. To achieve this, we need to treat the rotor-router system as a load balancing process.

Round-fair processes. As we intend to provide bounds on the values of the i -th potential for rotor-router process (given sufficiently long initialization time), we need to analyze the behavior of the cumulative rotor-router processes, *i.e.*, for a fixed Δt , to observe how the distribution of tokens $\mathcal{C}_t^{t+\Delta t}$ evolves with time. Thus, in the following, we will use the broader concept of *round-fair processes* denoted by \mathcal{W} , as introduced in [23]. Specifically, we will call an algorithm *strictly fair* if, in every step, the number of tokens that are sent out over any two edges incident to a node differs by at most one.

Definition 4.2. A process of token distribution (denoted by \mathcal{W}) is round-fair, if:

$$\forall_{e \in \text{out}(v)} \mathcal{W}_t(e) \in \left\{ \left\lfloor \frac{\mathcal{W}_t(v)}{\deg(v)} \right\rfloor, \left\lceil \frac{\mathcal{W}_t(v)}{\deg(v)} \right\rceil \right\} \quad (2)$$

and no tokens are left in nodes: $\mathcal{W}_t(v) = \sum_{e \in \text{out}(v)} \mathcal{W}_t(e)$.

We observe that any rotor-router process is round-fair. Also, by Observations 3.2, 3.3 and 3.4, for any fixed Δt , cumulative rotor-router in the sense of $\mathcal{W}_t = \mathcal{C}_t^{t+\Delta t}$ is also round-fair.

The round-fairness condition can be strengthened into algorithms which are *cumulatively fair*. We will call an algorithm *cumulatively fair* if for every interval of consecutive time steps, the total number of tokens sent out by a node differs by at most a small constant for any two adjacent edges. It is easy to see that cumulative fair algorithms under the constraint that every token is propagated, are performing the rotor-router distribution (and vice versa, rotor-router distribution is cumulative fair with every token propagated).

In the rest of this section, in order to simplify notation, we will assume that G is not bipartite. For the full formulation of Definition 4.4, Lemma 4.5, $\mathcal{B}(\mathcal{W}_t)$ and Definition 4.7, refer to the Appendix B.

Definition 4.3. A sequence of arcs e_1, e_2, \dots, e_p is called an alternating path (of length $p - 1$), if either every pair of arcs e_{2i}, e_{2i+1} shares starting vertex and every pair of arcs e_{2i-1}, e_{2i} shares ending vertex, or vice-versa: every pair of arcs e_{2i}, e_{2i+1} shares ending vertex and every pair of arcs e_{2i-1}, e_{2i} shares starting vertex.

Definition 4.4. We define the notion of distance between two arcs e, e' , denoted by $d(e, e')$, as a length of shortest alternating path having e and e' as first and last arcs.

In other words, a distance can be treated as a transitive closure of a relation where we define any pair of arcs sharing starting or ending vertex as at distance one.

Lemma 4.5. For any two arcs e, e' in non-bipartite graph G : $d(e, e') \leq 4D + 1$ moreover there exists an alternating path connecting e and e' containing at most $2D$ pairs of arcs sharing ending vertices and at most $2D + 1$ pairs of arcs sharing starting vertices.

Now we will proceed to analyze the behavior of the potential defined as in Definition 3.5, with respect to a round-fair processes.

Recall that $\Phi(\mathcal{W}_t) \stackrel{\text{def}}{=} \sum_{e \in \bar{E}} (\mathcal{W}_t(e))^2$. We will denote the smallest value of the potential achieved by distribution of tokens that preserves sums of loads over load balancing sets of arcs (ignoring the restriction that loads are integers) by:

$$\mathcal{B}(\mathcal{W}_t) \stackrel{\text{def}}{=} 2m \cdot \left(\text{avg}_{e \in \bar{E}} \mathcal{W}_t(e) \right)^2, \quad (3)$$

For non-bipartite graphs, (3) reduces to the following form: $\mathcal{B}(\mathcal{W}_t) = \frac{k^2}{2m} = \text{const}$. The following lemma follows directly from the convexity of quadratic functions.

Lemma 4.6. $\Phi(\mathcal{W}_t) \geq \mathcal{B}(\mathcal{W}_t)$.

Definition 4.7. We say that a configuration of tokens \mathcal{W}_t in non-bipartite graph G has discrepancy over arcs equal to $\max_{e, e' \in \bar{E}} (\mathcal{W}_t(e) - \mathcal{W}_t(e'))$.

The next observation follows directly from (2).

Observation 4.8. *The discrepancy over arcs is non-increasing in time, that is:*

$$\max_{e, e' \in \bar{E}} (\mathcal{W}_t(e) - \mathcal{W}_t(e')) \geq \max_{e, e' \in \bar{E}} (\mathcal{W}_{t+1}(e) - \mathcal{W}_{t+1}(e')).$$

We also put the following bound on the potential drop with respect to the discrepancy of number of tokens over arcs.

Lemma 4.9. *Consider a timestep t such that \mathcal{W}_t has discrepancy over arcs $x > 4D + 1$. Then: $\Phi(\mathcal{W}_t) - \Phi(\mathcal{W}_{t+1}) \geq \frac{(x-4D-1)(x-1)}{4D}$.*

Lemma 4.10. *If \mathcal{W}_t has discrepancy x , then $\Phi(\mathcal{W}_t) \leq \mathcal{B}(\mathcal{W}_0) + \frac{1}{2}mx^2$.*

Theorem 4.11. *If $T \geq 16mD \ln k$, then \mathcal{W}_T has discrepancy over arcs at most $10D$.*

Proof. Observe that for discrepancies $x \geq 10D$ we have, by Lemma 4.9:

$$\Phi(\mathcal{W}_t) - \Phi(\mathcal{W}_{t+1}) \geq \frac{(x-4D-1)(x-1)}{4D} \geq \frac{x^2}{16D}.$$

However, by Lemma 4.10: $x^2 \geq 2 \frac{(\Phi(\mathcal{W}_t) - \mathcal{B}(\mathcal{W}_0))}{m}$. Thus:

$$\Phi(\mathcal{W}_{t+1}) - \mathcal{B}(\mathcal{W}_0) \leq (\Phi(\mathcal{W}_t) - \mathcal{B}(\mathcal{W}_0)) \left(1 - \frac{1}{8mD}\right).$$

Let us assume that after $T \geq 16mD \ln k$ steps the discrepancy is larger than $10D$. Since $\Phi(\mathcal{W}_0) - \mathcal{B}(\mathcal{W}_0) \leq \Phi(\mathcal{W}_0) \leq k^2$, we have:

$$\Phi(\mathcal{W}_T) - \mathcal{B}(\mathcal{W}_0) \leq k^2 \cdot \left(1 - \frac{1}{8mD}\right)^{16mD \ln k} < k^2(1/e)^{2 \ln k} = 1,$$

implying that $\Phi(\mathcal{W}_T) = \lceil \mathcal{B}(\mathcal{W}_0) \rceil$, which implies that $\Phi(\mathcal{W}_T)$ minimizes potential among integer load distribution. Thus \mathcal{W}_T has discrepancy at most 1, a contradiction. \square

We are now ready to prove our main result on the time of stabilization of any rotor-router initial state.

Theorem 4.12. *For any initial state \mathcal{S}_0 , there exists $T = O(m^4 D^2 + mD \log k)$ such that \mathcal{S}_T is stable.*

Proof. Let $t_0 = \lceil 16mD \ln(3km) \rceil$. We observe that the cumulative rotor-router process (taken over Δt rounds) is round-fair, with the number of tokens equal to $\Delta t \cdot k$. For $t \geq t_0$, $\Delta t \leq 3m$, by Theorem 4.11 the token distribution of $\mathcal{C}_t^{t+\Delta t}$ has discrepancy over arcs at most $10D$, thus: $\mathcal{B}(\mathcal{C}_0^{\Delta t}) \leq \Phi_{\Delta t}(\mathcal{S}_t) = \Phi(\mathcal{C}_t^{t+\Delta t}) \stackrel{(4.10)}{\leq} \mathcal{B}(\mathcal{C}_0^{\Delta t}) + 50mD^2$. We next obtain:

$$\sum_{i=1}^{3m} \mathcal{B}(\mathcal{C}_0^i) \leq \sum_{i=1}^{3m} \Phi_i(\mathcal{S}_t) \leq 150m^2 D^2 + \sum_{i=1}^{3m} \mathcal{B}(\mathcal{C}_0^i). \quad (4)$$

Let $T > 300m^4 D^2 + \lceil 16mD \ln(3km) \rceil$. Let us assume that \mathcal{S}_T is not stable. Thus, for all $t \in [t_0 .. T]$, $\sum_{i=1}^{3m} \Phi_i(\mathcal{S}_t) - \sum_{i=1}^{3m} \Phi_i(\mathcal{S}_{t+2m^2}) \geq 1$, and in particular: $\sum_{i=1}^{3m} \Phi_i(\mathcal{S}_{t_0}) - \sum_{i=1}^{3m} \Phi_i(\mathcal{S}_T) \geq \left\lceil \frac{(T-t_0)}{2m^2} \right\rceil > 150m^2 D^2$, which contradicts with (4). \square

4.1 Lower bound

We now give a lower bound on the stabilization time of parallel rotor-router walks.

Theorem 4.13. *For any $N, M > 0$, $N \leq M \leq N^2$, there exists an initialization of the rotor router system in some graph with $\Theta(N)$ nodes and $\Theta(M)$ edges such that the stabilization time is $\Omega(M^2 \log N)$.*

Proof. We start the proof by exhibiting for all $n > 0$ an initial configuration with $k = 2$ tokens on the n -node path, for which the stabilization time is $\Omega(n^2 \log n)$. This construction then extends in a straightforward manner to denser graphs.

We will number the nodes of the path with consecutive integers $1, 2, \dots, n$, starting from its left endpoint. Initially, we assume that both of the tokens are located at node $\lceil n/3 \rceil$. Initially, the pointers of all nodes in the range $[2, \lceil n/3 \rceil]$ are directed towards the left endpoint of the path, and the pointers of all nodes in the range $[\lceil n/3 \rceil + 1, n - 1]$ are directed towards the right endpoint; the pointers of the degree-one nodes 1 and n are fixed in their unique position. We assign the tokens with unique identifiers, 1 and 2 , defined so that the token labeled 1 is always located not further from the left endpoint of the path than token 2 . The identity of a token is persistent over time, and the tokens may only swap identifiers when located at the same node of the path or when simultaneously traversing the same edge of the path in opposite directions.

For the considered initial configuration, the movement of the tokens in the first time steps is such that token 1 is propagated towards the left endpoint of the path (reaching it after $\lceil n/3 \rceil - 1$ steps), while token 2 is propagated towards the right endpoint (reaching it after $t_0 = n - \lceil n/3 \rceil$ steps). After reaching their respective endpoints, the tokens bounce back and begin to move towards the opposite endpoint of the path. All nodes of the path have now been visited by some token. Following the approach introduced in [19], for all subsequent moments of time t , we consider the partition $V = V_1(t) \cup V_2(t)$ of the set of nodes of the path into so called *domains* $V_1(t)$ and $V_2(t)$, where $V_i(t)$ is the set of nodes of the path whose last visit up to time t inclusive was performed by token i , $i = 1, 2$ (ties on the domain boundary are broken in favor of token 1). As observed in [19], each of the domains $V_i(t)$ is a sub-path of the considered path, containing token i which traverses $V_i(t)$ between its two endpoints, enlarging its domain by one node of the neighboring domain each time it reaches the boundary of the two domains. In our case of $k = 2$ tokens, we define the boundary point $b(t)$ so that $V_1(t) = [1, b(t)]$ and $V_2(t) = [b(t) + 1, n]$. Initially, at the time t_0 when token 2 reaches endpoint n for the first time, we have $b(t_0) = \lceil n/3 \rceil$. We will consider the interval of time $[t_0, t_1]$, where t_1 is the first moment of time such that $b(t_1) = \lfloor n/2 \rfloor - 1$. Throughout the interval $[t_0, t_1]$, we have $|V_2(t)| > |V_1(t)|$. Thus, for any node v such that $v = b(t)$ for some $t \in [t_0, t_1]$, if token 2 visits v at some time when heading towards the right endpoint of the path, node v will be subsequently visited by token 1 before token 2 returns to v . It follows that the boundary of the domains can never move to the left by more than one step; formally, for any $t, t' \in [t_0, t_1]$, $t < t'$, we have $b(t') \geq b(t) - 1$.

Now, let $\tau(v)$, for a node $v \in [\lceil n/3 \rceil, \lfloor n/2 \rfloor]$, denote the first moment of time $\tau \in [t_0, t_1]$ such that the domain boundary has reached node v , i.e., $b(\tau) \geq v$. We will show a lower bound on the value of $\tau(\lfloor n/2 \rfloor)$. We begin by lower-bounding the value $\Delta\tau_v = \tau(v+3) - \tau(v)$, for $v \in [\lceil n/3 \rceil, \lfloor n/2 \rfloor - 3]$. Clearly, $\tau(v+3) \geq \tau(v)$. During the considered time interval $(\tau(v), \tau(v+3)]$, the current domain boundary must be visited at least 3 times more by token 1 than by token 2 . Since for any $t \in (\tau(v), \tau(v+3)]$, we have $b(t) \geq v - 1$, and the distance between node $v - 1$ and the left endpoint of the path is $(v - 2)$, the total number of visits of token 1 to the boundary during the considered time interval of length $\Delta\tau_v$ is at most $\lceil \Delta\tau_v / (2(v - 2)) \rceil$. On the other hand, since we also have $b(t) \leq v + 3$ during the considered time interval, and the distance between node $v + 3$ and the right endpoint of the path is $n - v - 3$, we have that the domain boundary is visited by token 2 at least $\lfloor \Delta\tau_v / (2(n - v - 3)) \rfloor$ times during the considered time interval. It follows that:

$$\left\lceil \frac{\Delta\tau_v}{2(v-2)} \right\rceil - \left\lfloor \frac{\Delta\tau_v}{2(n-v-3)} \right\rfloor \geq 3, \quad (5)$$

which directly implies:

$$\begin{aligned} \Delta\tau_v &\geq \left(\frac{1}{2(v-2)} - \frac{1}{2(n-v-3)} \right)^{-1} > 2 \frac{(v-2)(n-v-3)}{n-2v} > \\ &> 2 \frac{(n/3-2)(n/2-3)}{n-2v} = \frac{(n-6)^2}{6} \cdot \frac{1}{n/2-v}. \end{aligned}$$

By summing the time increments $\Delta\tau_v$ over the subpath leading from $\lceil n/3 \rceil$ to $\lfloor n/2 \rfloor$, with a “step” of three nodes, we obtain a lower bound on $\tau(\lfloor n/2 \rfloor)$:

$$\begin{aligned} \tau(\lfloor n/2 \rfloor) &\geq \sum_{a=1}^{n/18-1} \Delta\tau_{\lfloor n/2 \rfloor - 3a} \geq \frac{(n-6)^2}{6} \cdot \sum_{a=1}^{n/18-1} \frac{1}{3a+1} > \\ &> \frac{(n-6)^2}{18} \sum_{a=2}^{n/18} a^{-1} > \frac{(n-6)^2}{18} (\ln n - 5). \end{aligned}$$

Thus, $\tau(\lfloor n/2 \rfloor) = \Omega(n^2 \log n)$. It remains to note that $\tau(\lfloor n/2 \rfloor)$ is a lower bound on the stabilization time of the rotor-router. Indeed, for as long as $t < \tau(\lfloor n/2 \rfloor)$, the domain boundary satisfies $b(t) < \lfloor n/2 \rfloor$, and the left endpoint of the path is visited by token 1 every not more than $2(\lfloor n/2 \rfloor - 2) < n - 1$ steps. On the other hand, in the limit cycle, each node (except endpoints) is visited by a token precisely twice in total during $2m = 2(n - 1)$ steps, a contradiction. Hence, we obtain a lower bound of $\Omega(n^2 \log n)$ on the stabilization time for a configuration of two tokens on the path.

We now extend our construction to obtain graphs with higher edge density admitting a lower bound of $\Omega(m^2 \log n)$ on stabilization time. For given N and M , we build such a graph G starting with an N -node path P_N , with pointers along internal nodes of the path initialized as before. We then attach the path to a pair of identical cliques on $M' = \Theta(M)$ edges by identifying each endpoint of the path with one vertex of the respective clique. The initialization of the pointers within the clique is such that a token entering the clique from the path performs a Eulerian walk of exactly $2M'$ arcs on the clique, returning to the path after traversing each edge of the clique exactly twice and returning all pointers within the clique to their initial position. (Such an initialization of pointers on the clique always exists and corresponds to the stable state pointer arrangement of the single-token rotor-router, for the token located at the distinguished vertex.) In our considerations, the initial location of the tokens is assumed to be on node $\lceil N/3 \rceil$ of the path, as before. The structure of domains within the path is defined analogously as before, and we can disregard the actions of tokens 1 and 2 within their respective cliques completely, replacing them for purposes of analysis by a delay of $2M'$ time steps during which the token is frozen at the left or right endpoint of the path, respectively. Adopting the same notation $\tau(v)$, $v \in [\lceil N/3 \rceil, \lfloor N/2 \rfloor]$, as for the case of the path, we obtain analogously to (5):

$$\left\lceil \frac{\Delta\tau_v}{2M' + 2(v-2)} \right\rceil - \left\lfloor \frac{\Delta\tau_v}{2M' + 2(N-v-3)} \right\rfloor \geq 3,$$

and from there (for $N \geq 9$):

$$\Delta\tau_v \geq \frac{M'^2}{2} \frac{1}{N/2 - v}$$

and finally:

$$\tau(\lfloor N/2 \rfloor) \geq \frac{M'^2}{6} (\ln N - 5).$$

Since $\tau(\lfloor N/2 \rfloor)$ is once again a lower bound on the stabilization time, and since our graph has $\Theta(N)$ nodes and $\Theta(M) = \Theta(M')$ edges, we have obtained the sought example of a rotor router configuration with a stabilization time of $\Omega(M^2 \log N)$. □

5 Simulation of the Rotor-Router

In this section, we answer the question of how to efficiently query for the state of a parallel rotor-router system after a given number of steps. The result below is for an arbitrary number of tokens ($k \geq 1$). For a single token ($k = 1$) rotor-router mechanism we provide a faster simulation algorithm in the Appendix.

Theorem 5.1. *We can preprocess any S_0 , in polynomial time and space (with respect to $n, m, \log k$) so that we can answer queries of state S_τ or queries of $\mathcal{C}_0^\tau(e)$ (the total number of visits until time step τ) both in time $\mathcal{O}(n + m)$.*

Proof. Our first step is to find T such that \mathcal{S}_T is stable. By Theorem 4.12 it is enough to take any $T > 300m^4D^2 + [16mD \ln(3km)]$. We compute and keep states $\mathcal{S}_0, \mathcal{S}_1, \dots, \mathcal{S}_T$, thus answering any queries of \mathcal{S}_τ with $\tau < T$ in $\mathcal{O}(n + m)$ time. We store preprocessed $\mathcal{C}_0^\tau(e)$ for any $\tau \in [0..T]$.

By Corollary 3.12, we can find any valid ∞ -subcycle decomposition of \mathcal{S}_T in polynomial time. By the properties of the subcycle decomposition, we can then find the values of $\mathcal{S}_{T+\tau}(e)$ by finding e' being shifted by τ along the cycle e belongs to. In a similar fashion we find $\mathcal{C}_T^{T+\tau}(e)$ for each arc e , giving us $\mathcal{C}_T^{T+\tau}(v)$ for each vertex v , thus we know the new pointer location for v . Each cycle can be preprocessed with prefix sums such that queries of this type can be answered in $\mathcal{O}(1)$ time, thus giving $\mathcal{O}(n + m)$ time for full $\mathcal{S}_{T+\tau}$ query.

We can preprocess each cycle with prefix sums, thus giving us the access to $\mathcal{C}_T^\tau(e)$ for $\tau \in [T.. \infty)$. By adding the value of $\mathcal{C}_0^\tau(e)$ we get desired $\mathcal{C}_0^\tau(e)$. □

5.1 Simulation of the single token rotor-router

Computational model. In this part of the paper, we assume the standard Word RAM model, where the space is measured in machine words. We assume that all numbers appearing in the input (i.e., n, m , the current time, and so on) fit in one machine word, and all basic arithmetical operations (including indirect addressing) on such words take $\mathcal{O}(1)$ time.

We denote the node where the token starts as v_0 , the node the token is located *after* i steps as v_i and the arc the token traverses during the i -th step as $e_i = (v_{i-1}, v_i)$

We partition the sequence of moves into *phases*, where the walk of the i -th *phase*, denoted \mathcal{E}_i , ends with a move finishing the i -th *full rotation* of starting node (thus \mathcal{E}_i being an Eulerian cycle for a subgraph of \vec{G} containing all outgoing arcs of v_0). Then the walk of the token can be denoted as:

$$\mathcal{E}_\infty = \underbrace{e_1, \dots, e_{a_1}}_{\mathcal{E}_1}, \underbrace{e_{a_1+1}, \dots, e_{a_2}}_{\mathcal{E}_2}, \underbrace{e_{a_2+1}, \dots, e_{a_3}, \dots}_{\mathcal{E}_3}, \dots$$

We recall a consequence of main result from [29].

Proposition 5.2. *The sequence of arcs in a phase $i - 1$ is always a subsequence of sequence from phase i . Moreover, if some arc is not visited during phase i , the inclusion is strict.*

If we denote by p the index of first phase that covers all of the arcs of the graph (thus \mathcal{E}_p being an Eulerian cycle on all arcs), the Proposition 5.2 can be rephrased as $\mathcal{E}_1 \sqsubset \mathcal{E}_2 \sqsubset \dots \sqsubset \mathcal{E}_p = \mathcal{E}_{p+1} = \mathcal{E}_{p+2} = \dots$. For any v , let $\text{expl}(v)$ be the smallest i such that v is visited in the i -th phase. Analogously the smallest i such that $e \in \mathcal{E}_i$ is denoted as $\text{expl}(e)$. We also define $\text{expl}(v_0)$ as 0. We define *the exploratory stage* as maximal consecutive subsequences of \mathcal{E}_i containing arcs not being part of \mathcal{E}_{i-1} . The result of Yanovski *et al.* [29], summarized in the next proposition, characterizes the structure of *exploratory stages*.

Proposition 5.3. *The graph induced in each exploratory stage is Eulerian. In every phase, each arc is traversed at most once, and furthermore in every phase starting with the $(\text{expl}(v) + 1)$ -th one, every arc adjacent to v is traversed exactly once. For every two adjacent nodes v and u we have that $|\text{expl}(v) - \text{expl}(u)| \leq 1$, and actually for every node w there exists a sequence of adjacent nodes $w_0, w_1, w_2, \dots, w_k$, where $w_0 = v_0$ and $w_k = w$, such that $\text{expl}(w_{i+1}) - \text{expl}(w_i) \in \{0, 1\}$.*

To be able to efficiently predict the position of the token at arbitrary time T , we need to efficiently discover the full exploration path $\mathcal{E} = \mathcal{E}_1 \mathcal{E}_2 \dots \mathcal{E}_p$. However, \mathcal{E} might be too large to store it directly (it is of size $\mathcal{O}(m \cdot D)$), therefore we have to store its *compressed form*: the sequence of arcs from \mathcal{E}_p (the Eulerian cycle covering G) together with values of $\text{expl}(e)$ for every arc e .

Consider Algorithm 1. Its correctness follows from the fact that in consecutive iterations of the main loop, we simulate traversals of *exploratory stages* from corresponding $\mathcal{E}_i \setminus \mathcal{E}_{i-1}$. In addition to finding the expl value for every arc, it results in the rotor-router pointers being in the same state as after the token performing a full \mathcal{E} walk. We can

Algorithm 1 Simulation

```
1: reset rotor-router pointers
2:  $\text{expl}(e) \leftarrow -1$  for every edge  $e$ 
3:  $Q \leftarrow [v_0]$  // queue with one element
4:  $i \leftarrow 0$ 
5: while  $Q \neq \emptyset$  do
6:    $i \leftarrow i + 1$ 
7:    $R \leftarrow \emptyset$  // empty queue
8:   for all  $v \in Q$  do
9:      $w \leftarrow \text{port}(v)$  // do a rotor-router walk starting from  $v$  until it reaches an already explored arc
10:    while  $\text{expl}[(v, w)] = -1$  do
11:       $\text{expl}[(v, w)] = i$ 
12:       $\text{advance}(v)$  // advance the pointer of the node  $v$ 
13:       $\text{append } w \text{ to } R$ 
14:       $v \leftarrow w$  // proceed to the next node
15:       $w \leftarrow \text{port}(v)$ 
16:    end while
17:  end for
18:   $Q \leftarrow R$ 
19: end while
```

use this property to find the sequence of arcs from \mathcal{E}_p by simply simulating the rotor-router walk starting from v_0 by another $2 \cdot m$ steps.

Proposition 5.4. *After running Algorithm 1, we can ask queries about position of token after T steps in time $\mathcal{O}(m)$.*

To speed up the above method, we denote $\mathcal{E}_p = e'_1, e'_2, \dots, e'_{2m}$, and let $e'_{2m+1} = e'_1$. Then if we look at the sequence of all values $\text{expl}(e'_1), \text{expl}(e'_2), \dots, \text{expl}(e'_{2m})$, they have the following property.

Proposition 5.5. *For all $i = 1, \dots, 2m$ we have $|\text{expl}(e'_i) - \text{expl}(e'_{i+1})| \leq 1$.*

We want to find the arc traversed by the token in the T -th step. For this we first check if T exceeds $|\mathcal{E}|$, and if so retrieve the $((T - |\mathcal{E}|) \bmod 2m)$ -th element of \mathcal{E}_p . Otherwise we need to locate the appropriate \mathcal{E}_i and its element. This reduces to finding the predecessor of T in the set $\{0, |\mathcal{E}_1|, |\mathcal{E}_1| + |\mathcal{E}_2|, \dots\}$ and then retrieving the $(T - \sum_{j < i} |\mathcal{E}_j|)$ -th element of \mathcal{E}_i . We will show how to perform both operations in $\mathcal{O}(\log \log m)$ time after linear preprocessing.

Lemma 5.6 (*y-fast trees* [28]). *Any set $S \subseteq [1, U]$ can be preprocessed in $\mathcal{O}(|S|)$ space so that we can locate the predecessor of any x in S using $\mathcal{O}(\log \log U)$ time.*

Using the above lemma, the predecessor of T in $\{0, |\mathcal{E}_1|, |\mathcal{E}_1| + |\mathcal{E}_2|, \dots\}$ can be found in $\mathcal{O}(\log \log m)$ time after linear preprocessing, as $U = \sum_i |\mathcal{E}_i| = \mathcal{O}(mn)$. Therefore we focus on preprocessing all \mathcal{E}_i as to allow retrieving the k -th element of any \mathcal{E}_i in $\mathcal{O}(\log \log m)$ time. This is not completely trivial, as we want to keep the preprocessing space linear in $|\mathcal{E}_p|$, not $\sum_i |\mathcal{E}_i|$.

Lemma 5.7. *All \mathcal{E}_i can be preprocessed in $\mathcal{O}(m)$ space so that we can retrieve the k -th element of any \mathcal{E}_i in $\mathcal{O}(\log \log m)$ time.*

Proof. We decompose every \mathcal{E}_i into contiguous parts. Every part is a maximal fragment consisting of arcs which are adjacent in the final \mathcal{E}_p . In other word, every arc such that its predecessor on \mathcal{E}_p is different than on \mathcal{E}_i begins a new part. Because of Proposition 5.5, the number of such arcs on \mathcal{E}_i is bounded by $|\mathcal{E}_{i+1}| - |\mathcal{E}_i|$. This is because if we have two arcs e, e' which are neighbors on \mathcal{E}_i but not \mathcal{E}_p , then there must be at least one arc between them on \mathcal{E}_{i+1} . We create a predecessor structure (implemented using Lemma 5.6) storing the positions on \mathcal{E}_i of all such arcs starting a new fragment. Additionally, the structure keeps for every such arc a pointer to the place where the fragment it begins occurs on \mathcal{E}_p . Then to retrieve the k -th element of \mathcal{E}_i , we use the predecessor structure to find

the part the answer belongs to. Then we return the corresponding element of \mathcal{E}_p , which takes $\mathcal{O}(1)$ time if we store \mathcal{E}_p in an array. Hence the total query time is $\mathcal{O}(\log \log m)$. To bound the total space, observe that it is equal to $\mathcal{O}(|\mathcal{E}_p| + \sum_i |\mathcal{E}_{i+1}| - |\mathcal{E}_i|) = \mathcal{O}(|\mathcal{E}_p|) = \mathcal{O}(m)$. \square

Combining Algorithm 1 and Lemma 5.7, we get following result.

Theorem 5.8. *Graph G can be preprocessed in time $\mathcal{O}(n+m)$ such that we can find in time $\mathcal{O}(\log \log m)$ the position of the token after given time T .*

Another natural problem that we would like to solve is how to preprocess \mathcal{E} using small space so that we can efficiently answer queries of the form: *How many times did the token traversed the node v in the first T steps of exploration?*

We will use similar properties as the ones presented above.

Proposition 5.9. *If $vis_i(v)$ is the number of times v is visited during the i -th phase, then $vis_i(v) = 0$ for all $i < \text{expl}(v)$, $vis_i(v) \in (0, \text{deg}(v)]$ for $i = \text{expl}(v)$, and finally $vis_i(v) = \text{deg}(v)$ for all $i > \text{expl}(v)$.*

Theorem 5.10. *Graph G can be preprocessed in time $\mathcal{O}(n+m)$ so that given a node v and a number T we can compute in time $\mathcal{O}(\log \log m)$ how many times node v was visited during the first T steps.*

Proof. We start with performing the preprocessing giving us the compressed representation of \mathcal{E} and expl value of every edge and every node. Then, for every node v we will create a data structure allowing us to efficiently compute how many times v was visited during the first T steps. To do this, for every node v we will keep a triple $(\text{expl}(v), \mathcal{X}_v, \mathcal{Y}_v)$, where \mathcal{X}_v is a predecessor structure with the timestamps of all visits to v that happened in the $\mathcal{E}_{\text{expl}(v)}$ -th phase, and \mathcal{Y}_v is a predecessor structure with the timestamps of all visits to v in the \mathcal{E}_p -th phase. To process the query, we first apply Theorem 5.8 to determine the position of the token after T steps. Say that in the T -step the token traverses arc e and is in the I -th exploratory phase. If $I < \text{expl}(v)$, then we return 0. If $I = \text{expl}(v)$, then we compute and return the number of elements of \mathcal{X}_v smaller than T , which can be done with a predecessor search. Finally, if $I > \text{expl}(v)$, we know that the total number of times v was visited in all stages up to the $(I-1)$ -th is $|\mathcal{X}_v| + (I - \text{expl}(v) - 1) \cdot \text{deg}(v)$. Hence we only need to count the visits to v in the I -th phase. Because in every further phase the order in which the token visits v and traverses e is the same, this can be done by counting how many times v is visited in the \mathcal{E}_p -th phase before traversing the arc e . After storing for every arc the time it is traversed in the \mathcal{E}_p -th phase, this can be done with a single predecessor search in \mathcal{Y}_v .

The size of predecessor structures kept for each node v is $\mathcal{O}(\text{deg}(v))$ and time necessary to create them is $\mathcal{O}(\text{deg}(v))$, thus giving $\mathcal{O}(n+m)$ total preprocessing time and space. \square

6 Conclusion

The rotor-router process has, in previous work, been identified as an efficient deterministic technique for a number of distributed graph processes, such as graph exploration and load balancing. In these settings, it rivals or outperforms the random walk, in some cases (such as parallel exploration of graphs) providing provable guarantees on performance, the counterparts of which need yet to be shown for the random walk. In this paper, we provide a complete characterization of the long-term behavior of the rotor-router, showing an inherent order in the limit state to which the system rapidly converges. This provides us with a better understanding of, e.g., the long-term load balancing properties of rotor-router-based algorithms, while at the same time opening the area for completely new applications. For instance, in view of our work, the rotor-router becomes a natural candidate for a self-organizing locally coordinated algorithm for the *team patrolling problem* — a task in which the goal is to periodically and regularly traverse all edges of the graph with k agents. This topic, and related questions, such as bounding the maximum distance between tokens on their respective Eulerian cycles in the limit state of the rotor-router, are deserving of future attention.

References

- [1] Y. Afek and E. Gafni. Distributed algorithms for unidirectional networks. *SIAM J. Comput.*, 23(6):1152–1178, 1994.
- [2] H. Akbari and P. Berenbrink. Parallel rotor walks on finite graphs and applications in discrete load balancing. In *SPAA*, pages 186–195. ACM, 2013.
- [3] N. Alon, C. Avin, M. Koucký, G. Kozma, Z. Lotker, and M. R. Tuttle. Many random walks are faster than one. In *SPAA*, pages 119–128, 2008.
- [4] 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. LNCS 5805, 2009.
- [5] E. Bampas, L. Gąsieniec, R. Klasing, A. Kosowski, and T. Radzik. Robustness of the rotor-router mechanism. In *OPODIS*, volume 5923 of *LNCS*, pages 345–358, 2009.
- [6] P. Berenbrink, R. Klasing, A. Kosowski, F. Mallmann-Trenn, and P. Uznanski. Be fair and be selfish! characterizing deterministic diffusive load-balancing schemes with small discrepancy. *CoRR*, abs/1404.4344, 2014.
- [7] P. Berenbrink, R. Klasing, A. Kosowski, F. Mallmann-Trenn, and P. Uznański. Improved analysis of deterministic load-balancing schemes. *CoRR*, abs/1404.4344, 2014.
- [8] S. N. Bhatt, S. Even, D. S. Greenberg, and R. Tayar. Traversing directed eulerian mazes. *J. Graph Algorithms Appl.*, 6(2):157–173, 2002.
- [9] J. N. Cooper and J. Spencer. Simulating a random walk with constant error. *Combinatorics, Probability & Computing*, 15(6):815–822, 2006.
- [10] D. Dereniowski, A. Kosowski, D. Pająk, and P. Uznański. Bounds on the cover time of parallel rotor walks. In *STACS*, volume 25 of *LIPICs*, pages 263–275. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014.
- [11] B. Doerr and T. Friedrich. Deterministic random walks on the two-dimensional grid. *Combinatorics, Probability & Computing*, 18(1-2):123–144, 2009.
- [12] K. Efremenko and O. Reingold. How well do random walks parallelize? In *APPROX-RANDOM*, pages 476–489, 2009.
- [13] R. Elsässer and T. Sauerwald. Tight bounds for the cover time of multiple random walks. *Theor. Comput. Sci.*, 412(24):2623–2641, 2011.
- [14] N. J. Fine and H. S. Wilf. Uniqueness theorems for periodic functions. *Proceedings of the AMS*, 16:109–114, 1965.
- [15] A. S. Fraenkel. Economic traversal of labyrinths. *Mathematics Magazine*, 43:125–130, 1970.
- [16] T. Friedrich and T. Sauerwald. The cover time of deterministic random walks. In *COCOON*, volume 6196 of *LNCS*, pages 130–139, 2010.
- [17] L. Gąsieniec and T. Radzik. Memory efficient anonymous graph exploration. In *WG*, volume 5344 of *LNCS*, pages 14–29, 2008.
- [18] M. Kiwi, R. Ndoundam, M. Tchuente, and E. Goles. No polynomial bound for the period of the parallel chip firing game on graphs. *Theoretical Computer Science*, 136:527–532, 1994.
- [19] 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.

- [20] A. Kosowski and D. Pająk. A case study of cover time for the rotor-router. In *ICALP*, volume 8573 of *LNCS*, pages 544–555. Springer, 2014.
- [21] E. Landau. Über die maximalordnung der permutationen gegebenen grades. *Arch. Math. Phys.*, 5:92–103, 1903.
- [22] 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.
- [23] 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, nov 1998.
- [24] 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 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 341–350. IEEE Computer Society, 2012.
- [25] T. Shiraga, Y. Yamauchi, S. Kijima, and M. Yamashita. L_∞ -discrepancy analysis of polynomial-time deterministic samplers emulating rapidly mixing chains. In Z. Cai, A. Zelikovsky, and A. G. Bourgeois, editors, *Computing and Combinatorics - 20th International Conference, COCOON 2014, Atlanta, GA, USA, August 4-6, 2014. Proceedings*, volume 8591 of *Lecture Notes in Computer Science*, pages 25–36. Springer, 2014.
- [26] C. Thompson. *Mathematical Statistical Mechanics*. Princeton Legacy Library. Princeton University Press, 2015.
- [27] I. A. Wagner, M. Lindenbaum, and A. M. Bruckstein. Distributed covering by ant-robots using evaporating traces. *IEEE Trans. Robotics and Automation*, 15:918–933, 1999.
- [28] D. E. Willard. Log-logarithmic worst-case range queries are possible in space $\theta(n)$. *Information Processing Letters*, 17(2):81–84, 1983.
- [29] V. Yanovski, I. A. Wagner, and A. M. Bruckstein. A distributed ant algorithm for efficiently patrolling a network. *Algorithmica*, 37(3):165–186, 2003.

APPENDIX

A Proofs of Lemmas and Theorems from Section 3

Lemma 3.7. For arbitrary i and t , the i -th potential is nonincreasing: $\Phi_i(\mathcal{S}_{t+1}) \leq \Phi_i(\mathcal{S}_t)$.

Proof. Let us assume on the contrary, that there exist other minimal partition $A = \{a_1, \dots, a_d\}$. As the partition into $\lfloor \frac{S}{d} \rfloor$ and $\lceil \frac{S}{d} \rceil$ is the only one with discrepancy at most 1, there exists pair of elements of A , that is w.l.o.g. a_1 and a_2 such that: $a_1 \geq a_2 + 2$. Thus:

$$a_1^2 + a_2^2 > a_1^2 + a_2^2 - 2(a_1 - a_2 - 1) = (a_1 - 1)^2 + (a_2 + 1)^2,$$

giving that $\{a_1 - 1, a_2 + 1, a_3, \dots, a_d\}$ has smaller sum of squares, which contradicts the assumption. \square

Lemma 3.9. The following statements are equivalent:

(i) \mathcal{S}_T admits a Δt -step subcycle decomposition,

(ii) $\forall v, \forall t, t': [t..t'] \subseteq [T..T+\Delta t] \{\{\mathcal{C}_t^{t'}(e)\}_{e \in \text{in}(v)}\} = \{\{\mathcal{C}_{t+1}^{t'+1}(e)\}_{e \in \text{out}(v)}\}$ (multisets of cumulative loads are preserved locally),

(iii) $\forall t, t': [t..t'] \subseteq [T..T+\Delta t] \{\{\mathcal{C}_t^{t'}(e)\}_{e \in \bar{E}}\} = \{\{\mathcal{C}_{t+1}^{t'+1}(e)\}_{e \in \bar{E}}\}$ (multisets of cumulative loads are preserved globally),

(iv) $\forall 0 \leq i \leq \Delta t \Phi_i(\mathcal{S}_T) = \Phi_i(\mathcal{S}_{T+1}) = \dots = \Phi_i(\mathcal{S}_{T+\Delta t-i+1})$ (potential is constant),

(v) $\forall v, \forall e_1, e_2 \in \text{in}(v) \forall t, t': [t..t'] \subseteq [T..T+\Delta t] |\mathcal{C}_t^{t'}(e_1) - \mathcal{C}_t^{t'}(e_2)| \leq 1$ (incoming discrepancy is at most one).

Proof. (i) \Rightarrow (ii) A straightforward consequence of the Δt -step subcycle decomposition is existence of local bijective mapping $M_v : \text{in}(v) \rightarrow \text{out}(v)$ that:

$$\forall t, t': [t..t'] \subseteq [T..T+\Delta t] \forall e \in \text{in}(v) \mathcal{C}_t^{t'}(e) = \mathcal{C}_{t+1}^{t'+1}(M_v(e)).$$

(ii) \Rightarrow (iii) Taking union over all vertices, we get:

$$\begin{aligned} \forall t, t': [t..t'] \subseteq [T..T+\Delta t] \{\{\mathcal{C}_t^{t'}(e)\}_{e \in \bar{E}}\} &= \bigcup_{v \in V} \{\{\mathcal{C}_t^{t'}(e)\}_{e \in \text{in}(v)}\} = \bigcup_{v \in V} \{\{\mathcal{C}_{t+1}^{t'+1}(e)\}_{e \in \text{out}(v)}\} = \\ &= \{\{\mathcal{C}_{t+1}^{t'+1}(e)\}_{e \in \bar{E}}\}, \end{aligned}$$

from which the claim follows.

(iii) \Rightarrow (iv) Claim follows in a straightforward fashion from the definition of the Φ_i .

(iv) \Rightarrow (v) Assume that, to the contrary, there exists v' and $e_1, e_2 \in \text{in}(v')$ such that:

$$\mathcal{C}_t^{t'}(e_1) - \mathcal{C}_t^{t'}(e_2) \geq 2.$$

Let $i = t' - t$ and let us define:

$$\begin{aligned} c(e_1) &= \mathcal{C}_t^{t'}(e_1) - 1, \\ c(e_2) &= \mathcal{C}_t^{t'}(e_2) + 1, \\ c(e) &= \mathcal{C}_t^{t'}(e) \text{ for every other } e. \end{aligned}$$

We observe, that:

$$\forall v \sum_{e \in \text{in}(v)} \mathcal{C}_t^{t'}(e) = \sum_{e \in \text{in}(v)} c(e).$$

Thus by the same reasoning as in proof of Lemma 3.7:

$$\Phi_i(\mathcal{S}_t) = \sum_v \sum_{e \in \text{out}(v)} \mathcal{C}_t^{t'}(e)^2 > \sum_v \sum_{e \in \text{out}(v)} c(e)^2 = \sum_v \sum_{e \in \text{in}(v)} c(e)^2 \geq \Phi_i(\mathcal{S}_{t+1}),$$

which contradicts the assumption.

(v) \Rightarrow (i) Let us fix v and denote $\text{in}(v) = \{e_1, e_2, \dots, e_d\}$, $\text{out}(v) = \{f_1, f_2, \dots, f_d\}$. Under the assumption, we deduce the existence of permutation (π_1, \dots, π_d) such that:

$$\forall_{0 \leq i \leq \Delta t} \mathcal{C}_T^{T+i}(e_{(\pi_1)}) \geq \mathcal{C}_T^{T+i}(e_{(\pi_2)}) \geq \dots \geq \mathcal{C}_T^{T+i}(e_{(\pi_d)}),$$

because otherwise, we would get that there exists $e_1, e_2, i_1 < i_2 \leq \Delta t$, such that:

$$\mathcal{C}_T^{T+i_1}(e_1) < \mathcal{C}_T^{T+i_1}(e_2), \mathcal{C}_T^{T+i_2}(e_1) > \mathcal{C}_T^{T+i_2}(e_2).$$

implying that

$$\mathcal{C}_{T+i_1}^{T+i_2}(e_1) - \mathcal{C}_{T+i_1}^{T+i_2}(e_2) \geq 2,$$

a contradiction with our initial assumption.

There exists a permutation $(\sigma_1), (\sigma_2), \dots, (\sigma_d)$ such that $f_{(\sigma_1)}, f_{(\sigma_2)}, \dots, f_{(\sigma_d)}$ have the same cyclic ordering as ρ_v , and $f_{(\sigma_1)}$ is the arc pointed to by pointer_v at beginning of time step $T + 1$, along which the first token at time $T + 1$ is propagated. From the cyclic nature of the rotor-router mechanism, we obtain:

$$\forall_{0 \leq i \leq \Delta t} \mathcal{C}_{T+1}^{T+i+1}(f_{(\sigma_1)}) \geq \mathcal{C}_{T+1}^{T+i+1}(f_{(\sigma_2)}) \geq \dots \geq \mathcal{C}_{T+1}^{T+i+1}(f_{(\sigma_d)}).$$

Thus, the sequences $(\mathcal{C}_T^{T+i}(e_{(\pi_1)}), \dots, \mathcal{C}_T^{T+i}(e_{(\pi_d)}))$ and $(\mathcal{C}_{T+1}^{T+i+1}(f_{(\sigma_1)}), \dots, \mathcal{C}_{T+1}^{T+i+1}(f_{(\sigma_d)}))$ are identical since they are sorted sequences of integers, with identical sum and discrepancy not greater than 1 (thus being composed of only two different values).

But $\forall_{0 \leq i \leq \Delta t} \mathcal{C}_T^{T+i}(e_{(\pi_j)}) = \mathcal{C}_{T+1}^{T+i+1}(f_{(\sigma_j)})$ implies $\mathcal{L}_T(e_{(\pi_j)}) = \mathcal{L}_{T+1}(f_{(\sigma_j)})$, $\mathcal{L}_{T+1}(e_{(\pi_j)}) = \mathcal{L}_{T+2}(f_{(\sigma_j)})$, \dots , $\mathcal{L}_{T+\Delta t-1}(e_{(\pi_j)}) = \mathcal{L}_{T+\Delta t}(f_{(\sigma_j)})$, meaning that if we pair each of $e_{(\pi_1)}, e_{(\pi_2)}, \dots$ with corresponding $f_{(\sigma_1)}, f_{(\sigma_2)}, \dots$ we will obtain a subcycle decomposition. □

Theorem 3.11. *The following conditions are equivalent:*

- (i) \mathcal{S}_T is stable,
- (ii) \mathcal{S}_T admits a ∞ -subcycle decomposition,
- (iii) \mathcal{S}_T admits a $(2m^2)$ -subcycle decomposition.

Proof. (ii) \Rightarrow (i) If all tokens traversing arcs follow Eulerian cycles as their routes, then we denote $t_r = \text{lcm}(|\vec{E}_1|, |\vec{E}_2|, \dots, |\vec{E}_p|)$.

Since $\forall_{e \in \vec{E}} \mathcal{L}_T(e) = \mathcal{L}_{T+t_r}(e)$ by assumption, we easily observe that $\forall_{v \in V} \mathcal{L}_T(v) = \mathcal{L}_{T+t_r}(v)$.

To prove that the states are equivalent, we need to show that the rotor-router pointers for each vertex are in the same position at time steps T and $T + t_r$. This is equivalent to: $\forall_{e_1, e_2 \in \text{out}(v)} \mathcal{C}_T^{T+t_r}(e_1) = \mathcal{C}_T^{T+t_r}(e_2)$.

However, if we assume otherwise, then there exist two arcs e_1, e_2 sharing starting points, such that: $|\mathcal{C}_T^{T+t_r}(e_1) - \mathcal{C}_T^{T+t_r}(e_2)| \geq 1$. But that would imply that by taking the cyclic sum twice, $|\mathcal{C}_T^{T+2 \cdot t_r}(e_1) - \mathcal{C}_T^{T+2 \cdot t_r}(e_2)| \geq 2$ which contradicts the rotor-router property.

(i) \Rightarrow (iii) Let us take T' such that $\mathcal{S}_{T'} = \mathcal{S}_T$ and $T' \geq T + 2m^2$. From Lemma 3.7 we deduce: $\forall_{i \geq 0} \Phi_i(\mathcal{S}_T) = \Phi_i(\mathcal{S}_{T+1}) = \Phi_i(\mathcal{S}_{T+2}) = \dots = \Phi_i(\mathcal{S}_{T'})$. By using equivalence of clauses (iv) and (i) in Lemma 3.9, we get the $2m^2$ -subcycle decomposition for \mathcal{S}_T .

(iii) \Rightarrow (ii) Assume that \mathcal{S}_T which has a $(2m^2)$ -subcycle decomposition does not admit a ∞ -subcycle decomposition. We denote by τ the largest integer such that \mathcal{S}_T admits a τ -subcycle decomposition ($\tau \geq 2m^2$). By the

maximality of τ , taking into account the equivalence of clauses (i) and (v) in Lemma 3.9, there exists a vertex v and two incoming arcs $e_1, e_2 \in \text{in}(v)$ such that:

$$\exists_{T' \in [T..T+\tau]} |\mathcal{C}_{T'}^{T'+\tau+1}(e_1) - \mathcal{C}_{T'}^{T'+\tau+1}(e_2)| \geq 2. \quad (6)$$

Let \vec{E}_1, \vec{E}_2 be the Eulerian cycles from the τ -subcycle decomposition (starting at state \mathcal{S}_T), such that $e_1 \in \vec{E}_1, e_2 \in \vec{E}_2$.

Let us denote $\ell = \text{lcm}(|\vec{E}_1|, |\vec{E}_2|)$. Observe that $\ell \leq m^2$, since we either have $\vec{E}_1 \cap \vec{E}_2 = \emptyset$, and then $|\vec{E}_1| + |\vec{E}_2| \leq 2m$ implies that $\ell \leq |\vec{E}_1| \cdot |\vec{E}_2| \leq m^2$, or $\vec{E}_1 = \vec{E}_2$, and then $\ell = |\vec{E}_1| \leq 2m \leq m^2$.[†]

Let us observe, that since the tokens are traversing the arcs over the Eulerian subcycles, we have, for any $[t..t+\ell] \subseteq [T..T+\tau]$:

$$\mathcal{C}_t^{t+\ell}(e_1) = \frac{\ell}{|\vec{E}_1|} \sum_{e \in \vec{E}_1} \mathcal{L}_T(e), \quad (7)$$

$$\mathcal{C}_t^{t+\ell}(e_2) = \frac{\ell}{|\vec{E}_2|} \sum_{e \in \vec{E}_2} \mathcal{L}_T(e). \quad (8)$$

Now we proceed to show that the left-hand sides of the above expressions have to be equal. Indeed, observe that $2\ell \leq 2m^2$ and that condition (v) from Lemma 3.9 also holds for range step $[T..T+2\ell]$, thus we have: $|\mathcal{C}_T^{T+2\ell}(e_1) - \mathcal{C}_T^{T+2\ell}(e_2)| \leq 1$, and moreover, we can write:

$$|\mathcal{C}_T^{T+2\ell}(e_1) - \mathcal{C}_T^{T+2\ell}(e_2)| = \left| \frac{2\ell}{|\vec{E}_1|} \cdot \sum_{e \in \vec{E}_1} \mathcal{L}_t(e) - \frac{2\ell}{|\vec{E}_2|} \cdot \sum_{e \in \vec{E}_2} \mathcal{L}_t(e) \right| = 2 \cdot |\mathcal{C}_T^{T+\ell}(e_1) - \mathcal{C}_T^{T+\ell}(e_2)|.$$

It follows that $\mathcal{C}_T^{T+\ell}(e_1) = \mathcal{C}_T^{T+\ell}(e_2)$, and taking into account (7) and (8) we get that for any $[t..t+\ell] \subseteq [T..T+\tau+1]$:

$$\mathcal{C}_t^{t+\ell}(e_1) = \mathcal{C}_t^{t+\ell}(e_2). \quad (9)$$

Thus, the average number of tokens observed over any multiplicity of the cycle length has to be identical for both arcs e_1, e_2 . We can use this fact to put a bound on discrepancy of number of tokens between both arcs, by analyzing last ℓ steps of the time segment separately. Indeed, by invoking (9) we get (with T' chosen as in (6)), if $T' \leq T + \tau + 1 - \ell$ then: $|\mathcal{C}_{T'}^{T'+\tau+1}(e_1) - \mathcal{C}_{T'}^{T'+\tau+1}(e_2)| \stackrel{(9)}{=} |\mathcal{C}_{T'}^{T'+\tau+1-\ell}(e_1) - \mathcal{C}_{T'}^{T'+\tau+1-\ell}(e_2)| \leq 1$, and otherwise $|\mathcal{C}_{T'}^{T'+\tau+1}(e_1) - \mathcal{C}_{T'}^{T'+\tau+1}(e_2)| \stackrel{(9)}{=} |\mathcal{C}_{T'-\ell}^{T'+\tau+1-\ell}(e_1) - \mathcal{C}_{T'-\ell}^{T'+\tau+1-\ell}(e_2)| \leq 1$, and we see that τ -subcycle decomposition contradicts with (6). □

Theorem 3.13. *State \mathcal{S}_T is stable iff:*

$$\sum_{i=1}^{3m} \Phi_i(\mathcal{S}_T) = \sum_{i=1}^{3m} \Phi_i(\mathcal{S}_{T+2m^2}).$$

Proof. \Rightarrow (only if) This is a straightforward result of characterizations from Theorem 3.11 and preservation of potentials (Lemma 3.9, (i) \Rightarrow (iv)).

\Leftarrow (if) By Theorem 3.11 it is enough to prove that \mathcal{S}_T admits a $2m^2$ -subcycle decomposition.[‡] Let us assume otherwise, that \mathcal{S}_T doesn't admit $2m^2$ -subcycle decomposition, and let τ be the largest integer such that there \mathcal{S}_T admits τ -subcycle decomposition ($\tau < 2m^2$). Similarly to the proof of Theorem 3.11, there exists a vertex v and two incoming arcs $e_1, e_2 \in \text{in}(v)$, such that:

$$\exists_{t \in [T..T+\tau]} |\mathcal{C}_t^{T+\tau+1}(e_1) - \mathcal{C}_t^{T+\tau+1}(e_2)| \geq 2. \quad (10)$$

[†] $2m \leq m^2$ requires $m \geq 2$, however in graphs with $m = 1$ every state is stable.

[‡]Once again we can eliminate case of $m = 1$ and assume that $2m^2 \geq 3m$.

Let us choose a local mapping M valid with respect to the τ -subcycle decomposition (starting at state \mathcal{S}_T), and let \vec{E}_1, \vec{E}_2 be the Eulerian cycles induced by M , such that $e_1 \in \vec{E}_1, e_2 \in \vec{E}_2$. We will look into the structure of sequences

$$\begin{aligned}(a_0, a_1, \dots) &= (\mathcal{L}_T(e_1), \mathcal{L}_{T+1}(e_1), \dots), \\ (b_0, b_1, \dots) &= (\mathcal{L}_T(e_2), \mathcal{L}_{T+1}(e_2), \dots), \\ (c_0, c_1, \dots) &= (a_0 - b_0, a_1 - b_1, \dots).\end{aligned}$$

Since each separate potential is nonincreasing, we deduce that $\forall_{0 \leq i \leq 3m} \Phi_i(\mathcal{S}_T) = \Phi_i(\mathcal{S}_{T+1}) = \dots = \Phi_i(\mathcal{S}_{T+2m^2})$. Thus, all of $\mathcal{S}_T, \mathcal{S}_{T+1}, \dots, \mathcal{S}_{T+2m^2-3m}$ admit a $3m$ -subcycle decomposition. From this we deduce, that:

$$\forall_{[i..j] \subseteq [0..2m^2]} \text{ if } (j - i < 3m) \text{ then } |c_i + c_{i+1} + \dots + c_j| \leq 1, \quad (11)$$

in particular, $c_i \in \{-1, 0, 1\}$ for $0 \leq i < 2m^2$.

From the τ -subcycle decomposition we get that:

$$\forall_{[i..j] \subseteq [0..\tau]} |c_i + c_{i+1} + \dots + c_j| \leq 1$$

and rewriting (10):

$$\exists_{\eta \in [0..\tau]} |c_\eta + c_{\eta+1} + \dots + c_\tau| \geq 2. \quad (12)$$

Before we proceed further with the proof, we need following structural lemma on the words a, b, c :

Lemma A.1. *Let $(a_i), (b_i)$ and (c_i) be defined as in Theorem 3.13. If $c_i = 0$ holds for $2m$ consecutive values of i from $0, \dots, \tau$, then it holds for all of them.*

Proof. $(a_i)_{i=0}^{\leq \tau}$ and $(b_i)_{i=0}^{\leq \tau}$ (a and b for short) are periodic sequences with short periods ($2m$ at most). Let $p_1 = |\vec{E}_1|, p_2 = |\vec{E}_2|$ be the periods lengths of those sequences. We want to prove that if they are identical on corresponding fragments, they are equal everywhere. We look into two cases:

- If $\vec{E}_1 = \vec{E}_2$, then $|p_1| = |p_2| \leq 2m$. By a and b being equal on fragment of length $2m$, we get $p_1 = p_2$, giving $a = b$.
- If $\vec{E}_1 \neq \vec{E}_2$, then we get $|p_1| + |p_2| \leq 2m$. By application of Fine and Wilf's theorem (if two periodic words with period lengths x and y are equal on fragment of length $x + y - \gcd(x, y)$, they are identical, see [14]), we get that $a = b$.

□

By Lemma A.1 we get, that in any sequence that satisfies (12) there cannot be a subsequence of $2m$ consecutive values of 0. However, from (11), we deduce that any two occurrences of 1 (respectively -1) at distance less than $3m$ have to be separated by at least one occurrence of -1 (respectively 1). Thus any occurrences of 1 and -1 have to happen in an alternating fashion (separated by arbitrary number of zeroes), leading to contradiction with (12).

□

B Proofs of Lemmas and Theorems from Section 4

In a bipartite graph G the two groups of tokens, the one traversing at one particular moment from one side to another and the one traversing in the opposite direction, will never meet and exchange tokens. Thus it is possible to have arbitrarily large difference in size of groups traversing arcs at arbitrarily late timesteps. Thus we need to refine our definitions and theorems used in the load balancing analysis.

Balancing sets. Now we introduce the partition of arcs into *balancing sets*. Arcs belonging to the same set will balance the load over time, while no load balancing happens between different sets. We will formalize those notions in the following.

Definition 4.4. We define the notion of distance between two arcs e, e' , denoted by $d(e, e')$, as a length of shortest alternating path having e and e' as first and last arcs. If there is no such path, we say that $d(e, e') = \infty$. We say two arcs e, e' belong to the same balancing set \mathcal{P} iff $d(e, e') < \infty$.

Lemma 4.5. The set of arcs of any bipartite graph is partitioned into exactly two balancing sets, while for every other graph the partition results in one balancing set. Moreover:

$$\forall e, e' \in \mathcal{P} \quad d(e, e') \leq 4D + 1$$

and for any such e, e' there exists an alternating path connecting e and e' containing at most $2D$ pairs of arcs sharing ending vertices and at most $2D + 1$ pairs of arcs sharing starting vertices.

Proof. Since every arc in G has opposite arc, any simple path in G relates to an alternating path. Thus, observe that two arcs are in the same balancing set iff there exists path of even length (in terms of edges) connecting the respective starting vertices v of arc e and v' of arc e' . Moreover, the length of such path increased by one gives an upper bound on $d(e, e')$. Since in a bipartite graph any simple path connecting vertices from the same side is of even length, this instantly gives us bound $d(e, e') \leq D + 1$ for bipartite graphs when e and e' go in the same direction, and $d(e, e') = \infty$ when they go in opposite directions.

To solve the non-bipartite case, let C be the shortest simple cycle of odd length in G (it always exists in non-bipartite graphs). It is easy to show that $|C| \leq 2D + 1$. Let u be the closest vertex on C to v , and u' to v' . Notice that u and u' can be connected by paths of arbitrary parity by using one of two parts of C as a connection. Thus, there exists a path connecting v to v' going through vertices v, u, u', v' (in this order) of even length at most $4D + 1$. Thus the length of the path is at most $4D$.

- If the path doesn't contain both e and e' , we get desired bound $d(e, e') \leq 4D + 1$ with number of pairs sharing starting vertices bounded by $2D + 1$ and pairs sharing ending vertices bounded by $2D$.
- If the path contains at least one of e or e' , then $d(e, e') \leq 4D$ giving also bounds on number of pairs of both types being at most $2D$.

□

We will denote the value of the potential achieved by distribution of tokens that preserves sums of loads over load balancing sets of arcs (ignoring the restriction that loads are integers) by:

$$\mathcal{B}(\mathcal{W}_t) \stackrel{\text{def}}{=} \sum_{\mathcal{P}} |\mathcal{P}| \cdot \left(\text{avg}_{e \in \mathcal{P}} \mathcal{W}_t(e) \right)^2, \quad (13)$$

where the sum is taken over all possible balancing sets \mathcal{P} . Recall that for non-bipartite graphs (13) reduces to:

$$\mathcal{B}(\mathcal{W}_t) = \frac{k^2}{2m} = \text{const},$$

which is consistent with our previous definition, while for bipartite graphs (where k_1 and k_2 are number of tokens in both parts of graph)

$$\mathcal{B}(\mathcal{W}_t) = \frac{k_1^2 + k_2^2}{m} = \text{const}.$$

Definition 4.7. We say that a configuration of tokens \mathcal{W}_t has discrepancy over arcs equal to

$$\max_{\mathcal{P}} \max_{e, e' \in \mathcal{P}} (\mathcal{W}_t(e) - \mathcal{W}_t(e')) = x.$$

Lemma 4.9. Let at time step t a distribution of tokens \mathcal{W}_t has discrepancy over arcs equal to x , for $x > 4D + 1$. Then:

$$\Phi(\mathcal{W}_t) - \Phi(\mathcal{W}_{t+1}) \geq \frac{(x - 4D - 1)(x - 1)}{4D}.$$

Proof. From Lemma 4.5 we know, that there exists an alternating path $e = e_0, e_1, e_2, \dots, e_{h-1}, e_h = e'$ of length at most $4D + 2$, such that each vertex belongs to at most 2 of those arcs, and that e and e' have discrepancy at least x . Thus:

$$\sum_{j=0}^{h-1} |\mathcal{W}_t(e_j) - \mathcal{W}_t(e_{j+1})| \geq x.$$

Moreover, at most $2D$ of pairs (e_j, e_{j+1}) share ending vertices and at most $2D + 1$ share starting vertices. Since by (2) loads of arcs sharing starting vertex can differ by at most 1, we get that there exist $I \leq 2D$ and vertices v_1, \dots, v_I , such that each vertex v_i admits incoming load discrepancy δ_i (two incoming arcs e_{i1} and e_{i2} such that $\mathcal{W}_t(e_{i1}) - \mathcal{W}_t(e_{i2}) \geq \delta_i$), and that:

$$\sum_{i \leq I} \delta_i \geq x - 2D - 1.$$

Observe, that:

$$\begin{aligned} \sum_{e \in \text{in}(v_i)} (\mathcal{W}_t(e))^2 &= (\mathcal{W}_t(e_{i1}))^2 + (\mathcal{W}_t(e_{i2}))^2 + \sum_{e \in \text{in}(v_i) \setminus \{e_{i1}, e_{i2}\}} (\mathcal{W}_t(e))^2 \geq \\ &\geq \left(\mathcal{W}_t(e_{i1}) - \left\lfloor \frac{\delta_i}{2} \right\rfloor \right)^2 + \left(\mathcal{W}_t(e_{i2}) + \left\lfloor \frac{\delta_i}{2} \right\rfloor \right)^2 + \sum_{e \in \text{in}(v_i) \setminus \{e_{i1}, e_{i2}\}} (\mathcal{W}_t(e))^2 \geq \\ &\geq \sum_{e \in \text{out}(v_i)} (\mathcal{W}_{t+1}(e))^2, \end{aligned}$$

which gives

$$\begin{aligned} \sum_{e \in \text{in}(v_i)} (\mathcal{W}_t(e))^2 - \sum_{e \in \text{out}(v_i)} (\mathcal{W}_{t+1}(e))^2 &\geq (\mathcal{W}_t(e_{i1}))^2 + (\mathcal{W}_t(e_{i2}))^2 - \left(\mathcal{W}_t(e_{i1}) - \left\lfloor \frac{\delta_i}{2} \right\rfloor \right)^2 - \\ &- \left(\mathcal{W}_t(e_{i2}) + \left\lfloor \frac{\delta_i}{2} \right\rfloor \right)^2 = 2(\mathcal{W}_t(e_{i1}) - \mathcal{W}_t(e_{i2})) \left\lfloor \frac{\delta_i}{2} \right\rfloor - 2 \left\lfloor \frac{\delta_i}{2} \right\rfloor^2 \geq 2 \left\lfloor \frac{\delta_i}{2} \right\rfloor \left\lfloor \frac{\delta_i}{2} \right\rfloor \geq \frac{\delta_i^2 - 1}{2}, \end{aligned}$$

Taking a sum over all vertices, we get:

$$\begin{aligned} \Phi(\mathcal{W}_t) - \Phi(\mathcal{W}_{t+1}) &\geq \sum_{i \leq I} \frac{\delta_i^2 - 1}{2} \geq \left(\sum_{i \leq I} \frac{\delta_i^2}{2} \right) - \frac{I}{2} \geq \frac{(\sum_{i \leq I} \delta_i)^2}{2I} - \frac{I}{2} \geq \\ &\geq \frac{(x - 2D - 1)^2}{4D} - D = \frac{(x - 4D - 1)(x - 1)}{4D}. \end{aligned}$$

□

Lemma 4.10. If \mathcal{W}_t has discrepancy x , then: $\Phi(\mathcal{W}_t) \leq \mathcal{B}(\mathcal{W}_0) + \frac{1}{2}mx^2$.

Proof. Observe, that by maximizing the sum of squares while preserving the sum and the upper bound on discrepancy, we get that:

$$\sum_{e \in \mathcal{P}} \mathcal{W}_t(e)^2 \leq \frac{|\mathcal{P}|}{2} \cdot (\mathbf{avg}_{e \in \mathcal{P}}(\mathcal{W}_t(e)) + x/2)^2 + \frac{|\mathcal{P}|}{2} \cdot (\mathbf{avg}_{e \in \mathcal{P}}(\mathcal{W}_t(e)) - x/2)^2.$$

Summing above over all \mathcal{P} :

$$\Phi(\mathcal{W}_t) \leq \mathcal{B}(\mathcal{W}_0) + \sum_{\mathcal{P}} |\mathcal{P}| \left(\frac{x}{2}\right)^2.$$

□