

Limit Behavior of the Multi-Agent Rotor-Router System*

J r mie Chalopin¹, Shantanu Das¹, Paweł Gawrychowski², Adrian Kosowski³,
Arnaud Labourel¹, and Przemysław Uznański⁴

¹ LIF, CNRS and Aix-Marseille University, France

² Institute of Informatics, University of Warsaw, Poland

³ Inria Paris and LIAFA, Paris Diderot University, France

⁴ Helsinki 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 experimental observation 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 observation, 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.

*Research supported by the ANR projects DISPLEXITY (ANR-11-BS02-0014) and MAC-ARON (ANR-13-JS02-0002). Part of the work was done while PU was affiliated with LIF, CNRS and Aix-Marseille University, supported by the Labex Archim de and by the ANR project MAC-ARON. A full version of the paper is available online at <http://arxiv.org/abs/1407.3200>.

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 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. [19]. 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 vertices 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 [5,7,9], exploration of graphs [1,8,10,12,15], and stabilization of distributed processes [3,6,17,22].

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 [5,20], 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. For a rotor-router system, the number of possible configurations 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.* [22] 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 [7] 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 [9]. 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 [18], where 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 $\mathcal{O}(d \sqrt{\log n / \mu})$ in [5].

Graph exploration. The rotor-router mechanism has also been studied in the context of graph exploration, sometimes under the name of *Edge Ant Walks* [21,22], and in the context of traversing a maze and marking edges with pebbles, *e.g.* in [6]. Cover times of rotor-router systems have been investigated by Wagner *et al.* [21] 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.* [6] 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.* [22] and Bampas *et al.* [3] 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 [11]. While all these studies were restricted to static graphs,

Bampas *et al.* [4] 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.* [22] and Klasing *et al.* [14], and the speedup of the system due to parallelization was considered for both worst-case and best-case scenarios. In [8], 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 [15], 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.* [22] experimentally observed 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 observation. 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.* [13] 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 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

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

$\{t, (t + 1), \dots, (t + \Delta t)\}$ is equal to the cumulative number of tokens leaving vertex v (over time $\{(t + 1), (t + 2), \dots, (t + \Delta t + 1)\}$), for arbitrary Δt .

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 important as a theoretical tool for understanding the limit behavior of the system, but 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 $\vec{u}v$ and $\vec{v}u$. 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 $(\vec{G}, (\rho_v)_{v \in V})$ 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 $(\vec{v}u)$, let $\text{next}(\vec{v}u)$ denote the arc after the arc $(\vec{v}u)$ in the cyclic order ρ_v . During each step, each node v distributes in round-robin fashion all of its tokens, using the following algorithm:

While there is a token at node v , do

1. Send token to pointer_v ,

2. Set $pointer_v = next(pointer_v)$.

Note that during a single time step all tokens at a node v are sent out and at exactly the next time step all those tokens arrive at their respective destination nodes.

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 $[a..b] \stackrel{\text{def}}{=} \{a, a+1, \dots, b\}$, $[a..b) \stackrel{\text{def}}{=} \{a, a+1, \dots, b-1\}$.

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.* [22], we know that $t_p = 2m$ and $t_s = O(mD)$. For arbitrary k , Yanovski *et al.* [22] experimentally observed that $t_p \leq 2m$ for any graph G regardless of the initial state. However, the following negative result disproves their observation and shows that the periodicity cannot be polynomially bounded for parallel rotor-routers.

Theorem 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. 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.

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.

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. 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 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.




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

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, with arbitrary port ordering (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* [16] $g(n - 1) = 2^{\Omega(\sqrt{n \log n})}$.

We remark that a similar result exists for parallel chip-firing games [13].

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. $\mathcal{L} : \vec{E} \cup V \rightarrow \mathbb{Z}$) 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 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 2 from load to cumulative load is as follow:

Observation 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 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 5. *Given a token distribution \mathcal{A} over edges, we define its potential as: $\Phi(\mathcal{A}) \stackrel{\text{def}}{=} \sum_{e \in \vec{E}} (\mathcal{A}(e))^2$.*

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

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 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.*

Lemma 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 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 4 and Lemma 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 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 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 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 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 9. *The following statements are equivalent:*

- (i) \mathcal{S}_T admits a Δt -step subcycle decomposition,
- (ii) $\forall v \forall t, t': [t..t'] \subseteq [T..T+\Delta t], \{\{C_t^{t'}(e)\}_{e \in \text{in}(v)}\} = \{\{C_{t+1}^{t'+1}(e)\}_{e \in \text{out}(v)}\}$ (multisets of cumulative loads are preserved locally),
- (iii) $\forall t, t': [t..t'] \subseteq [T..T+\Delta t], \{\{C_t^{t'}(e)\}_{e \in \bar{E}}\} = \{\{C_{t+1}^{t'+1}(e)\}_{e \in \bar{E}}\}$ (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) $\forall v \forall e_1, e_2 \in \text{in}(v) \forall t, t': [t..t'] \subseteq [T..T+\Delta t], |C_t^{t'}(e_1) - C_t^{t'}(e_2)| \leq 1$ (incoming discrepancy is at most one).

For a fixed value of $\Delta t = 1$, Lemma 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 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. The fact that stable state, by Lemma 7 and Lemma 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 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 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 11, we have:

Corollary 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 13 (compared to Theorem 11), based on refined analysis of the potential behavior.

Theorem 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 1.

Theorem 14. *For any stable state 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 .

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 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 15. *For any initial state S_0 , there exists $T = \mathcal{O}(m^5 k^2)$ such that 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 [18]. 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 16. *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 2, 3 and 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 subsequent definitions and lemmas, we refer the reader to the full version of the paper.

Definition 17. 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 18. 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 19. 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 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:

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

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

Definition 21. 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 22. 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 23. *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 24. *If \mathcal{W}_t has discrepancy x , then $\Phi(\mathcal{W}_t) \leq \mathcal{B}(\mathcal{W}_0) + \frac{1}{2}mx^2$.*

Theorem 25. *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 23:

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

However, by Lemma 24: $x^2 \geq 2 \frac{(\Phi(\mathcal{W}_t) - \mathcal{B}(\mathcal{W}_0))}{m}$. 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.

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

Theorem 26. *For any initial state \mathcal{S}_0 , there exists $T = O(m^4D^2 + mD \log k)$ 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 25 the token distribution of $\mathcal{C}_t^{t+\Delta t}$ has discrepancy over arcs at most $10D$, thus: $\mathcal{B}(\mathcal{C}_0^{\Delta t}) \leq \Phi_{\Delta t}(\mathcal{S}_t) = \Phi(\mathcal{C}_t^{t+\Delta t}) \stackrel{(24)}{\leq} \mathcal{B}(\mathcal{C}_0^{\Delta t}) + 50mD^2$. We next obtain:

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

Let $T > 300m^4D^2 + \lceil 16mD \ln(3km) \rceil$. 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: $\sum_{i=1}^{3m} \Phi_i(\mathcal{S}_{t_0}) - \sum_{i=1}^{3m} \Phi_i(\mathcal{S}_T) \geq \left\lceil \frac{(T-t_0)}{2m^2} \right\rceil > 150m^2D^2$, which contradicts with (4).

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

Theorem 27. *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 $\Omega(M^2 \log N)$.*

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 full version of the paper.

Theorem 28. *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 S_T is stable. By Theorem 26 it is enough to take any $T > 300m^4D^2 + \lceil 16mD \ln(3km) \rceil$. We compute and maintain states S_0, S_1, \dots, S_T , thus answering any queries of 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 12, we can find any valid ∞ -subcycle decomposition of S_T in polynomial time. By the properties of the subcycle decomposition, we can then find the values of $S_{T+\tau}(e)$ 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 $S_{T+\tau}$ 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 $\mathcal{C}_0^T(e)$ we get desired $\mathcal{C}_0^\tau(e)$.

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. 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.
4. 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.
5. P. Berenbrink, R. Klasing, A. Kosowski, F. Mallmann-Trenn, and P. Uznański. Improved analysis of deterministic load-balancing schemes. In *PODC*, pages 301–310, 2015.
6. S. N. Bhatt, S. Even, D. S. Greenberg, and R. Tayar. Traversing directed eulerian mazes. *J. Graph Algorithms Appl.*, 6(2):157–173, 2002.
7. J. N. Cooper and J. Spencer. Simulating a random walk with constant error. *Combinatorics, Probability & Computing*, 15(6):815–822, 2006.
8. 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, 2014.
9. B. Doerr and T. Friedrich. Deterministic random walks on the two-dimensional grid. *Combinatorics, Probability & Computing*, 18(1-2):123–144, 2009.
10. A. S. Fraenkel. Economic traversal of labyrinths. *Mathematics Magazine*, 43:125–130, 1970.
11. T. Friedrich and T. Sauerwald. The cover time of deterministic random walks. In *COCOON*, volume 6196 of *LNCS*, pages 130–139, 2010.
12. L. Gąsieniec and T. Radzik. Memory efficient anonymous graph exploration. In *WG*, volume 5344 of *LNCS*, pages 14–29, 2008.
13. M. Kiwi, R. Ndoundam, M. Tchunte, and E. Goles. No polynomial bound for the period of the parallel chip firing game on graphs. *Theoretical Computer Science*, 136:527–532, 1994.
14. 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.
15. A. Kosowski and D. Pająk. A case study of cover time for the rotor-router. In *ICALP*, volume 8573 of *LNCS*, pages 544–555, 2014.
16. E. Landau. Über die maximalordnung der permutationen gegebenen grades. *Arch. Math. Phys.*, 5:92–103, 1903.
17. 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.
18. 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.
19. T. Sauerwald and H. Sun. Tight bounds for randomized load balancing on arbitrary network topologies. In *FOCS*, pages 341–350, 2012.
20. T. Shiraga, Y. Yamauchi, S. Kijima, and M. Yamashita. L_∞ -discrepancy analysis of polynomial-time deterministic samplers emulating rapidly mixing chains. In *COCOON 2014*, volume 8591 of *LNCS*, pages 25–36. Springer, 2014.
21. 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.
22. V. Yanovski, I. A. Wagner, and A. M. Bruckstein. A distributed ant algorithm for efficiently patrolling a network. *Algorithmica*, 37(3):165–186, 2003.