

Does Adding More Agents Make a Difference? A Case Study of Cover Time for the Rotor-Router.

Adrian Kosowski, Dominik Pajak

► **To cite this version:**

Adrian Kosowski, Dominik Pajak. Does Adding More Agents Make a Difference? A Case Study of Cover Time for the Rotor-Router.. Journal of Computer and System Sciences, Elsevier, 2019, 106, pp.80-93. 10.1016/j.jcss.2019.07.001 . hal-00950743v3

HAL Id: hal-00950743

<https://hal.inria.fr/hal-00950743v3>

Submitted on 25 Mar 2014

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.

Does Adding More Agents Make a Difference? A Case Study of Cover Time for the Rotor-Router

Adrian Kosowski*

Dominik Pająk†

March 25, 2014

Abstract

We consider the problem of graph exploration by a team of k agents, which follow the so-called rotor router mechanism. Agents move in synchronous rounds, and each node successively propagates agents which visit it along its outgoing arcs in round-robin fashion. It has recently been established by Dereniowski et al. (STACS 2014) that the rotor-router cover time of a graph G , i.e., the number of steps required by the team of agents to visit all of the nodes of G , satisfies a lower bound of $\Omega(mD/k)$ and an upper bound of $O(mD/\log k)$ for any graph with m edges and diameter D . In this paper, we consider the question of how the cover time of the rotor-router depends on k for many important graph classes. We determine the precise asymptotic value of the rotor-router cover time for all values of k for degree-restricted expanders, random graphs, and constant-dimensional tori. For hypercubes, we also resolve the question precisely, except for values of k much larger than n . Our results can be compared to those obtained by Elsässer and Sauerwald (ICALP 2009) in an analogous study of the cover time of k independent parallel random walks in a graph; for the rotor-router, we obtain tight bounds in a slightly broader spectrum of cases. Our proofs take advantage of a relation which we develop, linking the cover time of the rotor-router to the mixing time of the random walk and the local divergence of a discrete diffusion process on the considered graph.

1 Introduction

Graph exploration is a task in which a team of agents is initially placed on a subset of nodes of the graph, and the agents are required to move around the graph so that each node is visited by at least one agent. Exploration with multiple walks is usually studied in a scenario where k agents are placed on some set of starting nodes and deployed in parallel, in synchronous steps. The principal parameter of interest is the *cover time* of the process, i.e., the number of steps until each node of the graph has been visited by at least one agent, for a worst case initial placement of agents in the graph. The agents may be endowed with different capabilities, ranging from a *priory* complete knowledge of the graph topology, through online scenarios in which they need to discover a map of the graph, to the most restrictive, where agents are in some sense passive (oblivious), and their movement is governed by simple local rules within the system.

*Inria Paris-Rocquencourt, France

†Inria Bordeaux Sud-Ouest, France

In this context, a fundamental problem concerns the cover time of k independent parallel random walks in a graph. Alon *et al.* [3], Efremenko and Reingold [9], and Elsässer and Sauerwald [11] have considered the notion of the *speedup* of the random walk for an undirected graph G , defined as the ratio between the cover time of a k -agent walk in G for worst-case initial positions of agents and that of a single-agent walk in G starting from a worst-case initial position, as a function of k . A characterization of the speedup has been achieved for many graph classes, although the question of the minimum and maximal values of speedup attainable in general is still open. The smallest known value of speedup for the random walk is $\Theta(\log k)$, attained e.g. for the cycle, while the largest known value is $\Theta(k)$, attained in a bounded range of values of k for many graph classes, such as expanders, cliques, and stars.

Our focus in this paper is on the deterministic model of walks in graphs known as the *rotor-router*. In the rotor-router model, introduced by Priezzhev *et al.* [17], the behaviour of the agent is fully controlled by the undirected graph in which it operates. The edges outgoing from each node v are arranged in a fixed cyclic order known as a *port ordering*, which does not change during the exploration. Each node v maintains a *pointer* which indicates the edge to be traversed by the agent during its next visit to v . If the agent has not visited node v yet, then the pointer points to an arbitrary edge adjacent to v . The next time when the agent enters node v , it is directed along the edge indicated by the pointer, which is then advanced to the next edge in the cyclic order of the edges adjacent to v . Each agent propagated by the rotor-router is a memoryless entity, but due to the existence of pointers, the rotor-router system as a whole is not Markovian. On the other hand, the system requires no special initialization, and its state at any moment of time is a valid starting state for the process.

State-of-the-art for the rotor-router.

For the case of a single agent, it is known that for any n -node graph of m edges and diameter D , the cover time of the rotor-router in a worst-case initialization in the graph is precisely $\Theta(mD)$ [19, 4]. After $\Theta(mD)$ time, the trajectory of the agent stabilizes to a periodic Eulerian traversal of the set of directed edges of the graph.

For $k > 1$ agents, no similar structural properties are observed, and in particular the rotor-router system may stabilize to a limit cycle of configurations of length exponential in n [16]. Recently, Klasing *et al.* [14] have provided the first evidence of speedup, showing that for the special case when G is a cycle, a k -agent system explores an n -node cycle $\Theta(\log k)$ times more quickly than a single agent system for $k < n^{1/11}$. A result for general graphs has been obtained by Dereniowski *et al.* [7], who show that the cover time of a k -agent system is always between $\Theta(mD/k)$ and $\Theta(mD/\log k)$, for any graph.

There exist interesting similarities between the behavior of the rotor-router and the random walk. For a single agent, the (deterministic) cover time of the rotor-router and the (expected) cover time of the random walk prove to be surprisingly convergent for many special graph classes, such as cycles and constant-degree expanders (same order of cover time for a single agent), cliques and stars (where the rotor-router is faster by a factor of $\Theta(\log n)$), or hypercubes (where the rotor-router is slower by a factor of $\Theta(\log n)$). A larger difference in cover time is observed, e.g., for the 2-dimensional grid, where the cover time of the rotor-router for a single agent is $\Theta(n^{3/2})$, as compared to the $\Theta(n \log n \log \log n)$ cover time for the random walk. For general graphs, the $\Theta(mD)$ bound on the cover time of the rotor-router can be compared to the upper bound of $O(mD \log n)$ on the cover time

of the random walk, although the latter bound is far from tight for many graph classes.

Our results.

In this paper, we ask about how the cover time of the k -agent rotor router depends on the number of agents k for specific graph classes, and investigate whether the similarities in terms of cover time of the rotor-router and the random-walk extend beyond the single agent case. We determine the precise asymptotic value of the cover time for degree-restricted expanders, random graphs, constant-dimensional tori, and hypercubes. Our results can be seen as complementary to those of Elsässer and Sauerwald [11], who studied the cover time of k multiple random walks in the same graph classes. We show that for all of the considered graph classes (except cycles), the cover time of both the rotor-router and the random walk admits a speedup of precisely $\Theta(k)$ for relatively small values of k , but above a certain threshold, the speedup of the rotor-router and random walk become divergent. (For cycles, both processes admit a speedup of $\Theta(\log k)$.) Our results are succinctly presented in Table 1.

We recall that for $k = 1$, the cover time of the rotor-router is $\Theta(mD)$, and note that for sufficiently large k ($k > n\Delta^D$ for a graph of maximum degree Δ), the cover time of the rotor-router is equal to precisely D , since the graph can be flooded with agents starting from a fixed node initially having Δ^D agents. Above this threshold ($k > n\Delta^D$), adding new agents to the system does not speed up exploration. The results we obtain show that for complete graphs, random graphs, and expanders, a cover time of $\Theta(D)$ is attained already for much smaller teams of agents. These graphs also display dichotomous behaviour: up to a certain threshold value of $k_1 = \Theta(m)$, the cover time decreases linearly with the number of agents, and above this threshold, the cover time remains fixed at $\Theta(D)$. We show that the cycle also admits this type of single-threshold behaviour, but with logarithmic speedup in the range of small k , with a cover time of $\Theta(n^2/\log k)$ for $k < k_1 = 2^n$, and a cover time of $\Theta(n)$ for $k \geq k_1$.

Interestingly, we prove that the d -dimensional torus for constant d (with $D = n^{1/d}$) admits precisely two threshold values of k (cf. Table 1). For $k < k_1 = n^{1-1/d}$, the speedup is linear with k ; for $k_1 \leq k < k_2 = 2^{n^{1/d}}$, the cover time further decreases with $\log(k/k_1)$, and above k_2 , the cover time is asymptotically fixed at $\Theta(n^{1/d})$. We remark that the for parallel random walks, the situation appears to be similar, however the question of obtaining a complete characterization remains open. For the hypercube, we also prove threshold behaviour for the speedup of the k -agent rotor-router, showing that there exist at least three threshold values of k (linear speedup for small k , a flat period with no speed-up for k slightly larger than n , a further period of slow growth, and finally a flat period for extremely large k). We also completely characterize the cover time of the hypercube for k up to a point beyond the first threshold.

Intuition of approach: exploration vs. diffusion.

In contrast to the case of parallel random walks, in the rotor-router system multiple agents interact with the same set of pointers at nodes, and the agents cannot be considered independent. However, the link between the multi-agent rotor-router and the parallel random walk processes becomes more apparent when the number of agents is extremely large ($k \gg n$), so that multiple agents are located at each node of the graph. Then, a fixed node v of degree d in the graph, which contains $a_t(v)$ agents at a given moment of time t , will send them out along outgoing links in the next step of the rotor-

Table 1: Cover time of the k -agent rotor-router system for different values of k in a n -node graph with m edges and diameter D . The results for d -dimensional tori are presented for d constant. The result for expanders concerns the case when the ratio of the maximum degree and the minimum degree of the graph is $O(1)$. The result for random graphs holds in the Erdős-Renyi model with edge probability $p > (1 + \varepsilon)\frac{\log n}{n}$, $\varepsilon > 0$, a.s.

<i>Graph</i>	<i>k</i>	<i>Cover time</i>	<i>Reference</i>
General graph	$\leq \text{poly}(n)$	$O\left(\frac{mD}{\log k}\right)$	[7]
		$\Omega\left(\frac{mD}{k}\right)$	[7]
Cycle	$< 2^n$	$\Theta\left(\frac{n^2}{\log k}\right)$	[14] (for $k < n^{1/11}$); Thm. 6.1
	$\geq 2^n$	$\Theta(n)$	
d -dim. torus	$< n^{1-1/d}$	$\Theta\left(\frac{n^{1+1/d}}{k}\right)$	Thm. 5.2
	$\in [n^{1-1/d}, 2^{n^{1/d}}]$	$\Theta\left(\frac{n^{2/d}}{\log(k/n^{1-1/d})}\right)$	Thm. 5.2
	$> 2^{n^{1/d}}$	$\Theta(n^{1/d})$	
Hypercube	$< n\frac{\log n}{\log \log n}$	$\Theta\left(\frac{n \log^2 n}{k}\right)$	Cor. 7.1
	$\in \left[n\frac{\log n}{\log \log n}, n2^{\log^{1-\varepsilon} n}\right]$ (for any $\varepsilon > 0$)	$\Theta(\log n \log \log n)$	Thm. 7.2
	$> n2^{\log^{1-\varepsilon} n}$	$O(\log n \log \log n)$	Thm. 7.2
	$> n^{\log_2 n}$	$\Theta(\log n)$	
Complete	$< n^2$	$\Theta\left(\frac{n^2}{k}\right)$	Thm. 4.5
	$\geq n^2$	$\Theta(1)$	
Expander	$< n \log n$	$\Theta\left(\frac{n \log^2 n}{k}\right)$	Thm. 4.5
	$\geq n \log n$	$\Theta(\log n)$	
Random graph	$< n \log n$	$\Theta\left(\frac{n \log^2 n}{k}\right)$	Thm. 4.5
	$\geq n \log n$	$\Theta(\log n)$	

router process, propagating the pointer at each step, so that each of its neighbours receives either $\lfloor a_t(v)/d \rfloor$ or $\lceil a_t(v)/d \rceil$ agents. In an analogous parallel random walk process, the *expected* number of agents following each of the outgoing links of a node v containing $a_t(v)$ agents will be $a_t(v)/d$. In fact, both the random walk and the rotor-router can be seen as different forms of discretization of the *continuous diffusion process*, in which a node having real-valued load $a_t(v)$ sends out precisely $a_t(v)/d$ load to each of its neighbours in the given time step. Discrete diffusion processes appear in research areas including statistical physics and distributed load balancing problems, and some studies of rotor-router-type systems have also been devoted to their diffusive properties. It is known, in particular, that, at any moment time, the difference of the number of agents located at a node between the rotor-router system and that in continuous diffusion is bounded by $\Theta(d \log n \mu^{-1})$ for d -regular graphs with eigenvalue gap μ [18], given identical initialization. This difference can even be bounded by constant for the case of lines [6] and grids [8]. Some other results in the area can also be found in [13, 1]. In this paper, we observe that, somewhat counter-intuitively, the link between continuous diffusion and the rotor-router can also be exploited for small values of k ($k \ll n$), for which agents as a rule occupy distinct nodes ($a_t(v) = 1$), and rounding $a_t(v)/d$ up or down to the nearest integer makes a major difference.

Organization of results.

In Section 2, we provide a formalization of the rotor-router model and some notation. In Section 3, we outline the technique which we subsequently use to bound the cover time in different graph classes. The main theorem of Section 4 captures the link between the cover time of the k -rotor-router system, the mixing time $\text{MIX}_{1/4}$ of the random walk process in the graph, and a graph parameter known as its discrepancy Ψ [18, 5], in its simplest form. This result directly provides tight bounds on the cover time for most of the considered graph classes, admitting small mixing time. The remaining cases of tori, cycles, and hypercubes are considered in Sections 5, 6, and 7, respectively.

2 Preliminaries

We consider an undirected connected graph $G = (V, E)$ with n nodes, m edges and diameter D . We denote the neighborhood of a node $v \in V$ by $\Gamma(v)$. The degree of a node v is denoted by $\deg(v)$, the maximum degree of the graph is denoted by Δ , and the minimum degree by δ . The directed graph $\vec{G} = (V, \vec{E})$ is the directed symmetric version of G , where the set of arcs $\vec{E} = \{(v, u) : \{v, u\} \in E\}$. We will denote arc (v, u) by $v \rightarrow u$.

Model definition.

We study the rotor-router model (on graph G) with $k \geq 1$ indistinguishable agents, which run in steps, synchronized by a global clock. In each step, each agent moves in discrete steps from node to node along the arcs of graph \vec{G} . A *configuration* at the current step is defined as a triple $((\rho_v)_{v \in V}, (\pi_v)_{v \in V}, \{r_1, \dots, r_k\})$, where ρ_v is a cyclic order of the arcs (in graph \vec{G}) outgoing from node v , π_v is an arc outgoing from node v , which is referred to as *the (current) port pointer at node v* , and $\{r_1, \dots, r_k\}$ is the (multi-)set of nodes currently containing an agent. For each node $v \in V$, the cyclic order ρ_v of the arcs outgoing from v is fixed at the beginning of exploration and does not change in any way

from step to step.

For an arc $v \rightarrow u$, let $next(v \rightarrow u)$ denote the arc next after arc $(v \rightarrow u)$ in the cyclic order ρ_v . The exploration starts from some initial configuration and then keeps running in all future rounds, without ever terminating. During the current step, first each agent i is moved from node r_i traversing the arc π_{r_i} , and then the port pointer π_{r_i} at node r_i is advanced to the next arc outgoing from r_i (that is, π_{r_i} becomes (π_{r_i})). This is performed sequentially for all k agents. Note that the order in which agents are released within the same step is irrelevant from the perspective of the system, since agents are indistinguishable. For example, if a node v contained two agents at the start of a step, then it will send one of the agents along the arc π_v , and the other along the arc $(v, next(\pi_v))$. For simplicity of notation only, we will assume that the ports outgoing from a node of degree d are numbered with consecutive integers $\{0, \dots, d-1\}$, that the function $next$ advances the pointer from port i to port $(i+1) \bmod d$, $0 \leq i < d$, and that the pointers at all nodes of the graph all initially point towards port 0.

Notation.

We will say that a node is *visited* by an agent in round t if the agent is located at this node at the start of round $t+1$. By $\mathbf{n}_t(v)$ we will denote the total number of visits of agents to node v , counting from the initialization of the system to the end of round t of the considered rotor-router process. In particular, $\mathbf{n}_0(v)$ refers to the number of agents at a node directly after initialization (at the start of round 1). Henceforth, we will treat \mathbf{n}_t as a non-negative integer-valued vector of dimension n . The worst-case cover time of the rotor-router for an initialization on graph G with k agents (i.e., for an initialization satisfying $\|\mathbf{n}_0\|_1 = k$), will be denoted by $C_{rr}^k(G)$.

We also introduce some auxiliary notation related to random walks and diffusion on the graph. We will denote by $P_t(v, u)$ the probability that a simple random walk, starting at node v of the graph, is located at u after exactly t steps of the walk, $t \geq 0$. The transition matrix of the random walk will be denoted by \mathbf{M} . For a node $u \in V$, \mathbf{u} will denote a vector of length n with $\mathbf{u}(u) = 1$ and all other entries 0. We recall that the cells of the t -th power of this matrix satisfy the following relation: $\mathbf{u}^T \mathbf{M}^t \mathbf{v} = P_t(v, u)$ [2]. The mixing time after which the random walk on the graph G reaches a total variation distance of at most $1/4$ from its stationary distribution will be denoted by $MIX_{1/4}$ (cf. [15] for further definitions). We will also use the following notation

$$t_{1/4} = \max_{u \in V} \min\{t : \forall v \in V P_t(u, v) \geq \pi_v/4\},$$

which denotes time after which probability of being at any node is at least quarter of the stationary probability regardless of the starting node of the random walk. Value $t_{1/4}$ can be compared to the definition of mixing time used in [10] (the mixing time used in [10] upper bounds $t_{1/4}$).

3 The main technique

To bound the cover time of the rotor-router, for any moment of time t , we will estimate the difference between the number of visits of the rotor-router to a node $x \in V$ up to time t , and the corresponding *expected* number of visits of parallel random walks, starting from the same initial placement of agents in the graph, to the same node x . (The latter notion can be equivalently interpreted as the total amount of load arriving in rounds 1 to t in a similarly initialized continuous diffusion process in load

balancing.) It turns out that the difference (discrepancy) between these two processes is bounded. As soon as the expected total number of visits of parallel random walks to x up to t has exceeded the maximum possible discrepancy with respect to the rotor-router, we can be sure that node x has been visited by the rotor-router at least once up to time t . This is captured by the following lemma.

Lemma 3.1. *Take any graph G . Let t^* be such a time moment that*

$$\forall x \in V \left(\sum_{\tau=0}^{t^*} \mathbf{M}^\tau \mathbf{n}_0 \right) (x) > \Psi_{t^*}$$

where

$$\Psi_t(G) = \max_{v \in V} \sum_{\tau=0}^t \sum_{(u_1, u_2) \in \vec{E}} |P_\tau(u_1, v) - P_\tau(u_2, v)|.$$

Then, the cover time of the k -agent rotor-router with arbitrary initialization on graph G satisfies $C_{rr}^k(G) \leq t^*$.

Before proceeding to prove the lemma, we remark that $\mathbf{M}^\tau \mathbf{n}_0$ is a vector describing the expected number of agents at nodes after τ steps of independent random walks on G , and that the expression $\left(\sum_{\tau=0}^{t^*} \mathbf{M}^\tau \mathbf{n}_0 \right) (x)$ on the left-hand side of the inequality is the before-mentioned expected total number of visits of random walks to x up to time t starting from initial agent placement. The expression Ψ_{t^*} is a generalization of the so-called 1-discrepancy Ψ of the graph, $\Psi = \lim_{t \rightarrow +\infty} \Psi_t$, introduced in [18]. The measure of 1-discrepancy is often applied when comparing a continuous and discrete process at a fixed moment of time t [5, 12], whereas herein we compare the *total* distance of two processes over all steps up to time t .

Proof. Consider the total number of visits $\mathbf{n}_t(u)$ at vertex u until step t by the rotor-router. It may be expressed as the sum of the number of agents initially located in u and the number of agents that entered to u from its neighbors (cf. [19, 14] for details of the argument):

$$\mathbf{n}_t(u) = \mathbf{n}_0(u) + \sum_{v \in \Gamma(u)} \left\lceil \frac{\mathbf{n}_{t-1}(v) - \text{port}(v, u)}{\text{deg}(v)} \right\rceil, \quad (1)$$

where $\text{port}(v, u) \in \{0, 1, \dots, \text{deg}(v) - 1\}$ denotes the label of the port leading from v to u .

We can rewrite Equation (1) as follows

$$\mathbf{n}_t(u) = \sum_{v \in \Gamma(u)} \frac{\mathbf{n}_{t-1}(v)}{\text{deg}(v)} + \mathbf{n}_0(u) + \xi_t(u), \quad (2)$$

where ξ_t is an “error vector” defined as:

$$\xi_t(u) = \sum_{v \in \Gamma(u)} \alpha_t^{(v, u)}, \quad (3)$$

with

$$\alpha_t^{(v, u)} = \left(\left\lceil \frac{\mathbf{n}_{t-1}(v) - \text{port}(v, u)}{\text{deg}(v)} \right\rceil - \frac{\mathbf{n}_{t-1}(v)}{\text{deg}(v)} \right).$$

Note that the values $\alpha_t^{(v,u)}$ are defined over directed arcs of the graph, $(v, u) \in \vec{E}$, satisfying $|\alpha_t^{(v,u)}| \leq 1$ and $\sum_{u \in \Gamma(v)} \alpha_t^{(v,u)} = 0$. Consequently, we have $\sum_{(v,u) \in \vec{E}} \alpha_t^{(v,u)} \mathbf{v} = \mathbf{0}$, and:

$$\xi_t = \sum_{(v,u) \in \vec{E}} \alpha_t^{(v,u)} \mathbf{u} = \sum_{(v,u) \in \vec{E}} \alpha_t^{(v,u)} \cdot (\mathbf{u} - \mathbf{v}).$$

Now, we rewrite (2) as follows:

$$\mathbf{n}_t = \mathbf{M}\mathbf{n}_{t-1} + (\mathbf{n}_0 + \xi_t), \quad (4)$$

where \mathbf{M} is the transition matrix of the random walk on G . Expanding (4) we have:

$$\mathbf{n}_t = \sum_{\tau=0}^t \mathbf{M}^\tau \mathbf{n}_0 + \sum_{\tau=0}^t \mathbf{M}^\tau \xi_{t-\tau}. \quad (5)$$

We will now bound the absolute value of the maximum element of the vector $\sum_{\tau=0}^t \mathbf{M}^\tau \xi_{t-\tau}$.

We have

$$\left\| \sum_{\tau=0}^t \mathbf{M}^\tau \xi_{t-\tau} \right\|_\infty = \left\| \sum_{\tau=0}^t \left(\mathbf{M}^\tau \cdot \sum_{(v,u) \in \vec{E}} \alpha_{t-\tau}^{(v,u)} \cdot (\mathbf{u} - \mathbf{v}) \right) \right\|_\infty \leq \left\| \sum_{\tau=0}^t \sum_{(v,u) \in \vec{E}} \alpha_{t-\tau}^{(v,u)} \mathbf{M}^\tau \cdot (\mathbf{u} - \mathbf{v}) \right\|_\infty$$

Note that since $|\alpha_{t-\tau}^{(u,v)}| \leq 1$

$$\left\| \sum_{\tau=0}^t \mathbf{M}^\tau \xi_{t-\tau} \right\|_\infty \leq \left\| \sum_{\tau=0}^t \sum_{(v,u) \in \vec{E}} |\mathbf{M}^\tau \cdot (\mathbf{u} - \mathbf{v})| \right\|_\infty. \quad (6)$$

We rewrite the above in terms of probability distributions of of random walk on G after τ steps:

$$(\mathbf{M}^\tau \cdot (\mathbf{u} - \mathbf{v})) (w) = P_\tau(u, w) - P_\tau(v, w), \quad (v, u) \in \vec{E}. \quad (7)$$

In this way, we obtain for any $x \in V$:

$$\left| \left(\mathbf{n}_t - \sum_{\tau=0}^t \mathbf{M}^\tau \mathbf{n}_0 \right) (x) \right| = \left| \left(\sum_{\tau=0}^t \mathbf{M}^\tau \xi_{t-\tau} \right) (x) \right| \leq \max_{w \in V} \sum_{\tau=0}^t \sum_{(v,u) \in \vec{E}} |P_\tau(u, w) - P_\tau(v, w)| = \Psi_t(G). \quad (8)$$

Thus, at time t any node the total number of visits in multi-agent rotor-router deviates from expected number of visits by multiple random walks by at most $\Psi_t(G)$. Since at time t^* at any node the expected number of visits by random walk is more than $\Psi_{t^*}(G)$ by assumption, all nodes have been visited at least once by the rotor-router. \square

4 Graphs with small mixing time

Theorem 4.1. *The cover time $C_{rr}^k(G)$ of a k -agent rotor-router with arbitrary initialization on any graph G satisfies*

$$C_{rr}^k(G) \leq t_{1/4}(G) + \frac{4\Delta}{\delta} \frac{n}{k} \Psi(G).$$

Proof. In order to apply Lemma 3.1, we want to find such a time step t that for any $x \in V$, $\sum_{\tau=0}^t (\mathbf{M}^\tau \mathbf{n}_0)(x) > \Psi(G) \geq \Psi_t(G)$.

By the definition of $t_{1/4}(G)$, for any $\tau \geq t_{1/4}(G)$ and for any vertices u, v , the distribution of the random walk satisfies

$$P_\tau(u, v) \geq \frac{\pi_v}{4} = \frac{\deg(v)}{8m} \geq \frac{\delta}{4\Delta n},$$

where π is the stationary distribution of the random walk on G (recall $\pi_v = \deg(v)/2m$).

When considering k independent random walks ($\|\mathbf{n}_0\|_1 = k$), for $\tau \geq t_{1/4}(G)$ we obtain $(\mathbf{M}^\tau \mathbf{n}_0)(x) \geq \frac{k\delta}{4\Delta n}$. Thus

$$\sum_{t=0}^{t_{1/4}(G) + \frac{4\Delta n}{k\delta} \Psi(G)} (\mathbf{M}^t \mathbf{n}_0)(x) > \Psi(G).$$

Thus, by Lemma 3.1, within time $t_{1/4}(G) + \frac{4\Delta n}{k\delta} \Psi(G)$ all nodes of G have been visited by the k -agent rotor-router. \square

In order to apply Theorem 4.1 to special graph classes, we provide convenient bounds on the value of Ψ which hold for regular graphs.

Proposition 4.2. *For any d -regular graph G :*

$$(i) \quad \Psi(G) \leq 4 \sum_{t=0}^{\text{MIX}_{1/4}(G)} \max_{v \in V} \sum_{\{u_1, u_2\} \in E} |P_t(u_1, v) - P_t(u_2, v)|$$

$$(ii) \quad \Psi(G) = O(d \text{MIX}_{1/4}(G)).$$

Proof. We want to approximate value $|P_t(u_1, v) - P_t(u_2, v)|$. Let \mathbf{M} be the transition matrix for the random walk on G and for any node u let \mathbf{u} be a vector of length n with $\mathbf{u}(u) = 1$ and all other entries 0 and let \mathbf{u}^\top be its transposition. We have

$$|P_t(u_1, v) - P_t(u_2, v)| = |\mathbf{u}_1^\top \mathbf{M}^t \mathbf{v} - \mathbf{u}_2^\top \mathbf{M}^t \mathbf{v}| = |(\mathbf{u}_1 - \mathbf{u}_2)^\top \cdot \mathbf{M}^t \mathbf{v}|. \quad (9)$$

Set $t = \tau + a \text{MIX}_{1/4}(G)$ in the Equation (9).

$$|P_t(u_1, v) - P_t(u_2, v)| = |(\mathbf{u}_1 - \mathbf{u}_2)^\top \cdot \mathbf{M}^\tau \cdot (\mathbf{M}^{a \text{MIX}_{1/4}(G)} \cdot \mathbf{v})|. \quad (10)$$

Vector $\mathbf{M}^{a \text{MIX}_{1/4}(G)} \cdot \mathbf{v}$ is the distribution of position of random walk starting from v after $a \text{MIX}_{1/4}(G)$. Let $\frac{1}{n}$ be vector of size n with all values $1/n$. Then

$$\mathbf{M}^{a \text{MIX}_{1/4}(G)} \cdot \mathbf{v} = \frac{1}{n} + \mathbf{e}_a^v,$$

where \mathbf{e}_a^v is the vector of deviations from stationary distribution for random walk of length $a \text{MIX}_{1/4}(G)$ starting at v . Since G is regular then $\frac{1}{n}$ is its stationary distribution. From the properties of mixing time of random walk [15] we have that $\sum_{w \in V} |\mathbf{e}_a^v(w)| \leq 2^{-a+1}$. We transform the Equation (10)

$$\begin{aligned} |P_t(u_1, v) - P_t(u_2, v)| &= |(\mathbf{u}_1 - \mathbf{u}_2)^\top \cdot \mathbf{M}^\tau \cdot (\frac{1}{n} + \mathbf{e}_a^v)| \\ &= |(\mathbf{u}_1 - \mathbf{u}_2)^\top \cdot \mathbf{M}^\tau \cdot \frac{1}{n} + (\mathbf{u}_1 - \mathbf{u}_2)^\top \cdot \mathbf{M}^\tau \cdot \mathbf{e}_a^v|. \end{aligned}$$

Vector $\frac{\mathbf{1}}{n}$ is an eigenvector of matrix \mathbf{M} thus $\mathbf{M}^\tau \cdot \frac{\mathbf{1}}{n} = \frac{\mathbf{1}}{n}$. Clearly $(\mathbf{u}_1 - \mathbf{u}_2)^\top \cdot \frac{\mathbf{1}}{n} = 0$. We have

$$\begin{aligned}
\sum_{\{u_1, u_2\} \in E} |P_t(u_1, v) - P_t(u_2, v)| &= \sum_{\{u_1, u_2\} \in E} |(\mathbf{u}_1 - \mathbf{u}_2)^\top \cdot \mathbf{M}^\tau \cdot \mathbf{e}_a^v| \\
&\leq \left(\sum_{\{u_1, u_2\} \in E} |(\mathbf{u}_1 - \mathbf{u}_2)^\top \cdot \mathbf{M}^\tau| \right) \cdot |\mathbf{e}_a^v| \\
&\leq \left\| \sum_{\{u_1, u_2\} \in E} |(\mathbf{u}_1 - \mathbf{u}_2)^\top \cdot \mathbf{M}^\tau| \right\|_\infty \cdot \sum_{w \in V} |\mathbf{e}_a^v(w)| \\
&\leq 2^{-a+1} \max_{v \in V} \sum_{\{u_1, u_2\} \in E} |P_{u_1, v}(\tau) - P_{u_2, v}(\tau)|,
\end{aligned}$$

where $|\mathbf{e}|$ for a vector \mathbf{e} denotes the vector of absolute values of elements. Now we can bound the value of $\Psi(G)$

$$\begin{aligned}
\Psi(G) &= \max_{v \in V} \sum_{t=0}^{\infty} \sum_{\{u_1, u_2\} \in E} |P_t(u_1, v) - P_t(u_2, v)| \leq \sum_{t=0}^{\infty} \max_{v \in V} \sum_{\{u_1, u_2\} \in E} |P_t(u_1, v) - P_t(u_2, v)| \\
&\leq \sum_{a=0}^{\infty} \sum_{\tau=0}^{\text{MIX}_{1/4}(G)-1} \max_{v \in V} \sum_{\{u_1, u_2\} \in E} |P_{\tau+a\text{MIX}_{1/4}(G)}(u_1, v) - P_{\tau+a\text{MIX}_{1/4}(G)}(u_2, v)| \\
&\leq \sum_{a=0}^{\infty} \sum_{\tau=0}^{\text{MIX}_{1/4}(G)-1} 2^{-a+1} \max_{v \in V} \sum_{\{u_1, u_2\} \in E} |P_\tau(u_1, v) - P_\tau(u_2, v)| \\
&= 4 \sum_{\tau=0}^{\text{MIX}_{1/4}(G)-1} \max_{v \in V} \sum_{\{u_1, u_2\} \in E} |P_\tau(u_1, v) - P_\tau(u_2, v)|,
\end{aligned}$$

which finishes the proof of (i). To prove (ii) observe that

$$\sum_{\{u_1, u_2\} \in E} |P_\tau(u_1, v) - P_\tau(u_2, v)| \leq \sum_{\{u_1, u_2\} \in E} (P_\tau(u_1, v) + P_\tau(u_2, v)) = \sum_{u \in V} dP_\tau(u, v) = d,$$

because for regular graphs, $P_\tau(u, v) = P_\tau(v, u)$. □

By combining Theorem 4.1 and Proposition 4.2, we can obtain upper bounds on the cover time of the rotor-router in regular graphs. At this point we provide an auxiliary result, which allows us to extend all our considerations to almost-regular graphs, as well as to show that our bounds on cover time hold regardless of whether the considered graph has self-loops or not. The proof relies on a variant of the *delayed deployment* technique for the rotor-router, introduced in [14].

Proposition 4.3. *Consider a graph G' constructed from G by adding self-loops to vertices, so that in the port ordering at any vertex there are at most x consecutive self-loops. Then, $C_{rr}^k(G')/(x+1) \leq C_{rr}^k(G) \leq C_{rr}^k(G')$.*

Proof. Before proving the proposition, we recall the notion of *delayed deployments* of the rotor-router [14].

A delayed deployment is a modified executions of the k -agent rotor-router system in which some agents may be stopped at a node, skipping their move for some number of rounds. Formally, a delayed deployment D of k agents is defined as a function $D : V \times \mathbb{N} \rightarrow \mathbb{N}$, where $D(v, t) \geq 0$ represents the number of agents which are stopped in node v in step t of the execution of the system. Delayed deployments may be conveniently viewed as algorithmic procedures for delaying agents, and are introduced for purposes of analysis, only. The following lemma relates the cover time of the rotor-router system to that of its delayed deployment.

The following *slowdown lemma* holds.

Lemma 4.4. [14] *Let R be a k -rotor router system with an arbitrarily chosen initialization, and let D be any delayed deployment of R . Suppose that deployment D covers all the nodes of the graph after T rounds, and in at least τ of these rounds, all k agents were active in D . Then, the cover time t_C of the rotor-router system R can be bounded by: $\tau \leq t_C \leq T$.*

Now we proceed to prove the Proposition. It suffices to compare a pair of rotor-router systems in which all agents are initialized at the same positions in G and G' , and the port orderings of G and G' are identical, when disregarding self-loops of G' , establishing the relation between cover times of such a pair of systems.

The proof that $C_{rr}^k(G') \leq C_{rr}^k(G)$ follows directly from Lemma 4.4, since we can construct a delayed deployment D for graph G which simulates the self-loops as in G' . Movements of agents in D on G will be exactly the same as in undelayed k -agent rotor-router operating on G' .

To prove the bound $C_{rr}^k(G')/(x+1) \leq C_{rr}^k(G)$, we prove by induction, that more generally for any time t and vertex v

$$\mathbf{n}_t(v) \leq \mathbf{n}'_{(x+1)t}(v),$$

where $\mathbf{n}_t(v)$ and $\mathbf{n}'_t(v)$ is the total number of visits at vertex v until time t for rotor-router on graphs G and G' respectively (where visits of an agent coming in from a self-loop of v do not count towards $\mathbf{n}'_t(v)$).

Since agents are initialized at the same positions then the claim is true for $t = 0$. Assume that it is true for some $t \geq 0$. Let for any $e \in \vec{E}$ denote by $r_t(e)$ and $r'_t(e)$ the total number of traversals of edge e for rotor-router in G and G' respectively until time t . In G all agents that entered some node v until step t left v until step $t + 1$. On the other hand in G' all agents that entered v until $(x + 1)t$, left until $(x + 1)(t + 1)$. Since order of pointers when considering only arcs from \vec{E} are the same in both graphs then for any arc $e \in \vec{E}$ we have

$$r_{t+1}(e) \leq r'_{(x+1)(t+1)}(e)$$

Since this holds for every arc $e \in \vec{E}$ then

$$\mathbf{n}_{t+1}(v) \leq \mathbf{n}'_{(x+1)(t+1)}(v).$$

Which completes the inductive proof. The above relation immediately implies that $C_{rr}^k(G) \geq (x+1) \leq C_{rr}^k(G')$, which completes the proof. \square

Taking into account Theorem 4.1 and Propositions 4.2 and 4.3, we obtain an upper bound of $O(mD/k)$ on the cover time of the rotor-router in a wide class of almost regular graphs with small mixing time. The complementary lower bound $C_{rr}^k(G) = \Omega(mD/k)$ is due to [7]. These bounds hold for all k , until the trivial bound $C_{rr}^k(G) = \Omega(D)$ is reached.

Theorem 4.5. *For any graph G such that $t_{1/4}(G) = O(D)$ and $\text{MIX}_{1/4}(G) = O(D)$ and $\Delta/\delta = O(1)$ the cover time of the k -agent rotor-router in the worst-case initialization of the system is:*

$$C_{rr}^k(G) = \Theta \left(\max \left\{ \frac{mD}{k}, D \right\} \right).$$

Proof. Note that we only need to consider the case of $k = O(m)$ and show $C_{rr}^k(G) = \Theta(mD/k)$. The lower bound $C_{rr}^k(G) = \Omega(mD/k)$ was shown in [7]. We will focus on the upper bound. Consider a Δ -regular graph G' constructed from G by adding self-loops to vertices. Since $\Delta/\delta = 1$ then adding self-loops to G increases the mixing time by no more than a constant factor, thus $\text{MIX}_{1/4}(G') = O(D)$ and $t_{1/4}(G') = O(D)$.

By Theorem 4.1 and Proposition 4.2 cover time $C_{rr}^k(G')$ of rotor-router on G' is

$$C_{rr}^k(G') \leq t_{1/4}(G') + \frac{4\Delta n}{k\delta} \Psi(G') \leq O(D) + O\left(\frac{ndD}{k}\right) = O\left(\frac{mD}{k}\right)$$

Thus, taking into account Proposition 4.3, deployment D will also cover graph G in time $O\left(\frac{mD}{k}\right)$. \square

Theorem 4.5 immediately implies the results stated in Table 1 for the case of complete graphs, degree-constrained expanders, and Erdős-Renyi graphs with edge probability $p > (1 + \varepsilon)\frac{\log n}{n}$. For cliques it is easy to see that $t_{1/4}(G) = O(1)$. For degree-constrained expanders, and Erdős-Renyi graphs bound on value $t_{1/4}(G)$ can be found in [10] as $t_{1/4}(G)$ is upper bounded by the mixing time used there.

The classes of tori, cycles, and hypercubes require more careful analysis; we consider them in the following Sections.

5 The torus

For the d -dimensional torus, Theorem 4.5 is not applicable, since the mixing time of the torus is $\text{MIX}_{1/4}(G) = \Theta(n^{2/d})$ [15], for constant d , whereas its diameter is $D = \Theta(n^{1/d})$. In the range of $k \leq n^{1-1/d}$, we can apply Theorem 4.1, taking advantage of a known tight bound on $\Psi(G) = \Theta(n^{1/d})$. In this way, we obtain: $C_{rr}^k(G) = O\left(n^{2/d} + \frac{n^{1+1/d}}{k}\right) = O\left(\frac{mD}{k}\right)$. Moreover, the complementary lower bound $C_{rr}^k(G) = \Omega(mD/k)$ holds for all graphs by [7]. This resolves the case of $k \leq n^{1-1/d}$.

To bound the cover time for $k > n^{1-1/d}$, in view of Proposition 4.3, we can equivalently consider the torus with d self-loops added on each node. We will now rely on Lemma 3.1, taking into account tighter bounds on $\Psi_t(G)$ for small values of t . The following bound can be shown by a straightforward Markovian coupling argument.

Lemma 5.1. *If G' is a d -dimensional torus with d self-loops at each node, then $\Psi_t(G') \leq 24d\sqrt{t}$.*

Proof. We apply the Markovian coupling technique to bound values $|P_\tau(u_1, v) - P_\tau(u_2, v)|$, where u_1 and u_2 are neighbors. To construct a coupling on G' , consider three random walks. Let walk W_1 start from u_1 , and let walks W_2 and W_3 start from u_2 . We will view the random walk on G' in every step as choosing one among $2d$ edges and traversing it with probability $1/2$.

Walk W_1 is a standard random walk on G' and W_2 will be constructed based on W_1 . When W_1 in a step chooses an edge, W_2 in the same step chooses the same edge.

Nodes u_1 and u_2 have the same coordinates in $d - 1$ dimensions and differ by 1 in one dimension. Denote by d^* the dimension on which u_1 and u_2 differ.

If W_1 chooses a dimension different from d^* then W_2 makes the same choice whether to traverse the chosen edge or not. Thus, the positions of walks W_1 and W_2 will never differ on a dimension different from d^* . Consider the distance between W_1 and W_2 in dimension d^* . If W_1 chooses an edge from d^* then when choosing whether to traverse it, W_2 makes an opposite choice (if W_1 traverses it, W_2 does not). Thus, whenever W_1 chooses dimension d^* , the distance between these walks decreases by 1 with probability $1/2$ and increases by 1 with probability $1/2$. Denote by T the random variable denoting the time of meeting of walks W_1 and W_2 . Walk W_3 follows W_2 in steps $0, 1, \dots, T$ and then follows W_1 . The pair (W_1, W_3) forms a coupling. Using the theory of coupling [15, Theorem 5.2], since $P\{T > \tau\}$ is the probability that walks W_1 and W_3 have coupled after time τ , we obtain

$$\frac{1}{2} \|P_\tau(u_1, \cdot) - P_\tau(u_2, \cdot)\|_1 \leq P\{T > \tau\}, \quad (11)$$

where $\|P_\tau(u_1, \cdot) - P_\tau(u_2, \cdot)\|_1 = \sum_{v \in V} |P_\tau(u_1, v) - P_\tau(u_2, v)|$.

Since the initial distance between the walks is 1 then by [15, Theorem 2.17] we obtain for $\tau > 0$

$$P\{T > \tau\} \leq 12/\sqrt{\tau} \quad (12)$$

By Equations (11), (12) we have

$$\sum_{v \in V} |P_\tau(u_1, v) - P_\tau(u_2, v)| \leq \frac{24}{\sqrt{\tau}} \quad (13)$$

Let G' be a d -dimensional torus with n vertices and d self-loops at each node. We have that

$$\frac{12dn}{\sqrt{\tau}} \geq \sum_{\{u_1, u_2\} \in E} \sum_{v \in V} |P_\tau(u_1, v) - P_\tau(u_2, v)| = \sum_{v \in V} \sum_{\{u_1, u_2\} \in E} |P_\tau(u_1, v) - P_\tau(u_2, v)|$$

Observe that by the symmetry of G' , for every v value $\sum_{\{u_1, u_2\} \in E} |P_\tau(u_1, v) - P_\tau(u_2, v)|$ is the same, thus for any v

$$\frac{12dn}{\sqrt{\tau}} \geq n \sum_{\{u_1, u_2\} \in E} |P_\tau(u_1, v) - P_\tau(u_2, v)|$$

We obtain

$$\begin{aligned} \Psi_t(G') &= \max_{v \in V} \sum_{\tau=0}^t \sum_{\{u_1, u_2\} \in E} |P_\tau(u_1, v) - P_\tau(u_2, v)| \leq d/2 + 12d \sum_{\tau=1}^t \tau^{-1/2} \\ &\leq d/2 + 12d \left(1 + \int_1^t x^{-1/2} dx \right) \leq 24d\sqrt{t} \end{aligned}$$

which completes the proof. \square

Introducing the above bound into Lemma 3.1 and taking into account properties of the random walk in the torus, for $k = k'n^{1-1/d}$ ($k' > 1$), we eventually obtain a bound on cover time of the form $O\left(\frac{D^2}{\log k'}\right)$. Somewhat surprisingly, this bound is tight, and we propose an initialization of the rotor-router system which achieves this bound precisely. (The proof of tightness relies on a bound on cover time for the cycle, which is introduced in the following section.) In this way, we obtain a complete characterization of the speed-up of the rotor-router on the torus.

Theorem 5.2. *If G is a torus of constant dimension then cover time of k -agent rotor-router is*

- (i) $C_{rr}^k(G) = \Theta\left(\frac{mD}{k}\right)$, for $k \leq n^{1-1/d}$,
- (ii) $C_{rr}^k(G) = \Theta\left(\max\left\{\frac{D^2}{\log k'}, D\right\}\right)$, for $k = k'n^{1-1/d}$, $k > n^{1-1/d}$.

Proof. We start by proving (i). We want to show that $t_{1/4}(G) = O(n^{2/d})$. For sufficiently large n , distribution of random walk on infinite d -dimensional grid can be approximated by the normal distribution $\mathcal{N}_d(0, \sigma^2)$, where \mathcal{N}_d is a product of d independent normal distributions and $\sigma = \sqrt{t}/d$ is the standard deviation in any dimension. Hitting probabilities on the infinite grid lower bound the probabilities for corresponding pairs u, v on torus G . Thus, it is sufficient to bound the hitting probabilities on the infinite grid for points at distance at most $\Theta(n^{1/d})$ in every dimension from the starting point.

The minimum probability will be achieved in the point that is at the maximum distance in every dimension from the starting point. This probability satisfies

$$P_t(u, v) \geq (1 + o(1)) \left(\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{n^{2/d}}{\sigma^2}} \right)^d = (1 + o(1)) (d/\sqrt{2\pi t})^d e^{-\frac{d^3 n^{2/d}}{t}}, \quad (14)$$

where $P_t(u, v)$ is the probability of being in v after t steps of a random walk on the torus G starting from u . If we set $t = cn^{2/d}d^2/\log(d/\sqrt{2\pi})$, where c is an appropriately chosen constant, we obtain $P_t(u, v) \geq n^{-1/4}$. This means that for tori we have $t_{1/4}(G) = O(n^{2/d})$. In the range of $k \leq n^{1-1/d}$, we can now apply Theorem 4.1, taking advantage of a known tight bound on $\Psi = \Theta(n^{1/d})$ [18]. In this way, we obtain: $C_{rr}^k G = O\left(n^{2/d} + \frac{n^{1+1/d}}{k}\right) = O\left(\frac{mD}{k}\right)$. Moreover, the complementary lower bound $C_{rr}^k(G) = \Omega(mD/k)$ holds for all graphs by [7]. This resolves the case of $k \leq n^{1-1/d}$. Now we want to prove (ii). We first show the lower bound. We want to construct a initialization of rotor-router that will lead to desired complexity. Take a set of all nodes that have coordinate 0 in the first dimension. There are $n^{1-1/d}$ such nodes. Assume that k is divisible by $n^{1-1/d}$ and place all agents evenly on these nodes. In each node we have $k/n^{1-1/d}$ agents. Now we initiate the pointers. The initial position of the pointer is on the edge in the first dimension that is along the shortest path to the closest node with coordinate 0 on the first dimension. Initialization of the cyclic order of the arcs can be arbitrary but the same on every node. With such initialization if we take all nodes with the same coordinate in the first dimension then we will always have the same number of agents on these nodes. Thus we can see the exploration of the torus as the exploration of the cycle. We know that the exploration of a cycle of length $n^{1/d}$ using $k/n^{1-1/d}$ takes time $\Theta\left(\frac{n^{2/d}}{\log(k/n^{1-1/d})}\right)$. If the number of agents k is not divisible by $n^{1-1/d}$ we can again use Lemma 4.4 to observe that exploration with k agents will not be faster than exploration with $k' = n^{1-1/d} \lceil k/n^{1-1/d} \rceil$ agents.

Now we want to prove the upper bound. By Lemma 5.1 we know that $\Psi_t(G) \leq 24d\sqrt{t}$. Thus we can use Lemma 3.1 if we find such t^* , that

$$\forall x \in V \left(\sum_{\tau=0}^{t^*} \mathbf{M}^\tau \mathbf{n}_0 \right) (x) \geq 48d\sqrt{t^*}$$

It is sufficient to find such t^* such that for all $t \in [t^*/2, t^*]$ all elements of vector $\mathbf{M}^t \mathbf{n}_0$ are at least equal to $96d/\sqrt{t^*}$. Vector n_0 is non-negative with sum k thus it is sufficient to find t^* , such that all elements of matrix $\mathbf{M}^{t^*/2}$ are at least $96d/(k\sqrt{t^*})$. We want to find $t' = t^*/2$ such that for any $u, v \in V$

$$P_{t'}(u, v) \geq 96d/(k\sqrt{2t'}),$$

where $P_{t'}(u, v)$ is the probability of being in v after t' steps of a random walk on the torus G starting from u . We can use approximation by normal distribution again. By applying Equation (14) we want to find such t that

$$\Theta(1)t^{-d/2}e^{-\frac{d^3n^{2/d}}{t}} \geq \Theta(1)t^{-1/2}/k$$

If we take $k = n^{1-1/d}k'$ and $t = n^{2/d}/x$ then we obtain

$$x^{(d-1)/2}e^{-d^3x} \geq \Theta(1/k')$$

$$\log(k'/c) \geq d^3x - (d-1)/2 \log x,$$

where c is a constant. Thus, if $x = \log(k'/c)/d$ and $k' > c$, then the inequality is satisfied. Thus, for $k > cn^{1-1/d}$ agents, we have the cover time $\Theta\left(\frac{d^3n^{2/d}}{\log(k/(cn^{1-1/d}))}\right) = \Theta\left(\frac{D^2}{\log(k/n^{1-1/d})}\right)$, which completes the proof. \square

6 The cycle

The general case result from [7] allows us to upper-bound the cover time of the k -rotor-router system on the cycle by $O\left(\max\left\{\frac{n^2}{\log k}, n\right\}\right)$, for any $k \geq 1$. On the other hand, the complementary lower bound of $\Omega\left(\frac{n^2}{\log k}\right)$ was only known to hold for $k < n^{1/11}$ [14]. In the following, we extend this lower bound to arbitrary values of k . The proof relies on a modification of the approach used in the proof of Lemma 3.1: whereas Lemma 3.1 can only be used to upper bound cover time, this time we perform a different transformation of (5) for a specific initialization of agents starting from a single node on the ring, for which we can show that the “error term” associated with vector $\xi_{t-\tau}$ is negative. Intuitively, this behaviour is due to an initialization of pointers which delays progress of the agents going along the path to the most distant node of the ring. We eventually obtain the following result.

Theorem 6.1. *If G is a cycle of size n then cover time of k -agent rotor-router is*

$$C_{rr}^k(G) = \Theta\left(\max\left\{\frac{n^2}{\log k}, n\right\}\right).$$

Proof. The upper bound is a direct consequence of [7] and the fact that adding more agents cannot slow down exploration. The lower bound is also shown by [14] for $k \leq n^{1/11}$.

To prove the lower bound for $k > n^{1/11}$, we consider for simplicity a cycle of even length, with $n' = n/2$, and we divide the cycle into two subpaths of length n' along an axis of symmetry crossing a pair of edges. We then perform an initialization of the rotor-router system which is symmetric with respect to this axis (cf.[14] for an explanation why this argument is correct and sufficient). In all further considerations, we will restrict our attention to only one of the subpaths of the cycle. We will number its nodes $u_1, \dots, u_{n'}$, with $u_1 = v$ and $u_{n'} = w$ being its endpoints. We now initialize the rotor-router on the considered path so that all agents are located at vertex v , all ports along the shortest path to v get label 0, and the port leading away from v gets label 1. We will show that it takes the agents in the considered system a long time to reach node w .

Fix a moment of time t , and suppose that none of the agents has reached w until the end of round $t - 1$ inclusive. We will now show that if the condition $t < n'^2/(12 \log k)$ is satisfied, then none of the agents will reach w at time t either, i.e., $\mathbf{n}_t(w) = 0$.

We rely on some of the techniques from the proof of Lemma 3.1. We have from (5):

$$\mathbf{n}_t(w) = \sum_{\tau=0}^t \mathbf{w}^\top \mathbf{M}^\tau \mathbf{n}_0 + \sum_{\tau=0}^t \mathbf{w}^\top \mathbf{M}^\tau \xi_{t-\tau}. \quad (15)$$

Consider a pair of nodes u_i, u_{i+1} , where we recall that u_{i+1} is the neighbor of u_i that is further from v (and closer to w). Then, due to the chosen port initialization, we have

$$\alpha_{t-\tau}^{(u_i, u_{i+1})} = (\lceil (\mathbf{n}_{t-\tau-1}(u_i) - 1)/2 \rceil - \mathbf{n}_{t-\tau-1}(u_i)/2) \leq 0. \quad (16)$$

Since the considered graph has degree 2, we have $\alpha_{t-\tau}^{(u_i, u_{i-1})} + \alpha_{t-\tau}^{(u_i, u_{i+1})} = 0$, and we may write by rearranging the definition of $\xi_{t-\tau}$:

$$\xi_{t-\tau} = \sum_{i=1}^{n'-1} (\alpha_{t-\tau}^{(u_i, u_{i-1})} \mathbf{u}_{i-1} + \alpha_{t-\tau}^{(u_i, u_{i+1})} \mathbf{u}_{i+1}) = \sum_{i=1}^{n'-1} \alpha_{t-\tau}^{(u_i, u_{i+1})} (\mathbf{u}_{i+1} - \mathbf{u}_{i-1}).$$

In the above sum, u_0 should be interpreted as the mirror reflection of u_1 in the other subpath of G , whereas index n' was discarded from the sum since node $w = u_{n'}$ was not visited before time t by assumption.

Introducing the above into (15), taking into account that $\mathbf{n}_0 = k\mathbf{v}$, and expanding, we obtain:

$$\begin{aligned} \mathbf{n}_t(w) &= \sum_{\tau=0}^t \mathbf{w}^\top \mathbf{M}^\tau k\mathbf{v} + \sum_{\tau=0}^t \left(\mathbf{w}^\top \mathbf{M}^\tau \sum_{i=1}^{n'-1} \alpha_{t-\tau}^{(u_i, u_{i+1})} (\mathbf{u}_{i+1} - \mathbf{u}_{i-1}) \right) = \\ &= k \sum_{\tau=0}^t P_\tau(w, v) + \sum_{\tau=0}^t \sum_{i=1}^{n'-1} \alpha_{t-\tau}^{(u_i, u_{i+1})} (P_{t-\tau}(w, u_{i+1}) - P_{t-\tau}(w, u_{i-1})) \leq k \sum_{\tau=0}^t P_\tau(w, v), \end{aligned} \quad (17)$$

where the last inequality holds because $\alpha_{t-\tau}^{(u_i, u_{i+1})} \leq 0$ by equation (16), whereas $P_T(w, u_{i+1}) \geq P_T(w, u_{i-1})$ holds for any time moment T , by the basic properties of a random walk on the cycle starting from vertex w .

In order to show that $\mathbf{n}_t(w) = 0$, it suffices to show that $\mathbf{n}_t(w) < 1$, since this value is an integer. Taking into account (17) we only need to show that $\sum_{\tau=0}^t P_\tau(w, v) < 1/k$. We apply the following

standard bound based on normal approximation of $P_\tau(w, v)$, recalling that the distance between w and v is $n' - 1$, $t < n'^2/(12 \log k)$, and $k > n'^{1/11}$:

$$\sum_{\tau=0}^t P_\tau(w, v) < t \cdot \frac{1}{\sqrt{t}} e^{-n'^2/t} = \sqrt{t} \cdot e^{-n'^2/t} < n' k^{-12} = k^{-1} (n' k^{-11}) < k^{-1},$$

which completes the proof. \square

7 The hypercube

For the hypercube with $n = 2^d$ vertices, the value of $\Psi(G)$ has been precisely derived in [5]. The corresponding asymptotic formula is $\Psi(G) = \Theta(\log^2 n)$. Using this result in combination with Theorem 4.1, we obtain the following corollary.

Corollary 7.1. *If G is a hypercube with n vertices then $C_{rr}^k(G) = \Theta\left(\frac{n \log^2 n}{k}\right) = \Theta\left(\frac{mD}{k}\right)$, for $k \leq n \frac{\log n}{\log \log n}$.*

The behavior of the rotor-router on the hypercube for $k > k_1 = n \frac{\log n}{\log \log n}$ is not completely understood. For $k = k_1$, the value of cover time is $O(\log n \log \log n)$. Interestingly, we can show that there exists a flat ‘‘plateau’’ region above k_1 in which the asymptotic cover time of the hypercube is precisely $\Theta(\log n \log \log n)$. The proof proceeds along slightly more complex lines than the proof of Theorem 6.1. We show that in the considered range of k , $\Theta(\log n \log \log n)$ time is required for k agents starting at one corner of the hypercube to reach the opposite corner, given an arrangement of ports at each node in which the pointer first traverses all ports leading the agent towards the starting vertex.

Theorem 7.2. *If G is a hypercube of size $n = 2^d$ then the cover time of k -agent rotor-router with $k \leq n \cdot 2^{\log^{1-\varepsilon} n}$ agents is $C_{rr}^k(G) > \frac{\varepsilon}{10} \log n \log \log n$, where $\varepsilon \in (0, 1)$ is an arbitrary fixed constant.*

Proof. The proof proceeds along similar lines as that of Theorem 6.1. We identify each vertex with a $d = \log_2 n$ bit vector of coordinates, with $v = 0^d$ and $w = 1^d$ being antipodal vertices. We partition set V into layers L_0, \dots, L_d , such that all vertices belonging to layer L_i have exactly i ones in their binary representation. Now, for each vertex $u \in L_i$, $0 \leq i \leq d$, the port labeling of u is set so that ports $0, 1, \dots, (i-1)$ point to the i neighbours of u belonging to layer L_{i-1} in arbitrary order, while ports $i, (i+1), \dots, (d-1)$ point to the $d-i$ neighbours of u belonging to layer L_{i+1} in arbitrary order. The system is initialized with k agents placed on node v .

Acting in a similar way as in the proof of Theorem 6.1, we will show that for all $t \leq \frac{\varepsilon}{10} \log n \log \log n$, we have $\mathbf{n}_t(w) = 0$. Once again, we consider equality (15), and we prove that each of the summed expressions $\mathbf{w}^\top \mathbf{M}^\tau \xi_{t-\tau}$ is negative, for all $0 \leq \tau \leq t$. We can represent vector $\xi_{t-\tau}$ as follows:

$$\xi_{t-\tau} = \sum_{i=0}^d \sum_{u \in L_i} \left(\sum_{\substack{(u, u^-) \in \vec{E} \\ u^- \in L_{i-1}}} \alpha_{t-\tau}^{(u, u^-)} \mathbf{u}^- + \sum_{\substack{(u, u^+) \in \vec{E} \\ u^+ \in L_{i+1}}} \alpha_{t-\tau}^{(u, u^+)} \mathbf{u}^+ \right).$$

Due to the fact that for each node u from layer L_i , all the ports pointing to layer L_{i-1} are always visited before those in layer L_{i+1} , and there are no other ports adjacent to u , we have:

$$\sum_{\substack{(u,u^+) \in \vec{E} \\ u^+ \in L_{i+1}}} \alpha_{t-\tau}^{(u,u^+)} = - \sum_{\substack{(u,u^-) \in \vec{E} \\ u^- \in L_{i-1}}} \alpha_{t-\tau}^{(u,u^-)} \leq 0.$$

Moreover, by the symmetry of the random walk with respect to coordinates of the binary vector representation of the node, we have that for all nodes x belonging to a layer j , the probability that a walk starting at w is located at x after τ steps is the same, and denoted as $P_\tau(w, x) = P_\tau^j(w)$.

Combining the above observations, we obtain:

$$\begin{aligned} \mathbf{w}^\top \mathbf{M}^\tau \xi_{t-\tau} &= \sum_{i=0}^d \sum_{u \in L_i} \left(\sum_{\substack{(u,u^-) \in \vec{E} \\ u^- \in L_{i-1}}} \alpha_{t-\tau}^{(u,u^-)} P_\tau(w, u^-) + \sum_{\substack{(u,u^+) \in \vec{E} \\ u^+ \in L_{i+1}}} \alpha_{t-\tau}^{(u,u^+)} P_\tau(w, u^+) \right) = \\ &= \sum_{i=0}^d \sum_{u \in L_i} \left(P_\tau^{(i-1)}(w) \sum_{\substack{(u,u^-) \in \vec{E} \\ u^- \in L_{i-1}}} \alpha_{t-\tau}^{(u,u^-)} + P_\tau^{(i+1)}(w) \sum_{\substack{(u,u^+) \in \vec{E} \\ u^+ \in L_{i+1}}} \alpha_{t-\tau}^{(u,u^+)} \right) = \\ &= \sum_{i=0}^d \sum_{u \in L_i} \left((P_\tau^{(i+1)}(w) - P_\tau^{(i-1)}(w)) \sum_{\substack{(u,u^+) \in \vec{E} \\ u^+ \in L_{i+1}}} \alpha_{t-\tau}^{(u,u^+)} \right) \leq 0, \end{aligned}$$

where in the last inequality we took into account that for all moments of time T , $P_T^{(i+1)}(w) \geq P_T^{(i-1)}(w)$, by the properties of the random walk on the hypercube (recall that $w = 1^d$).

Acting as in the derivation of (17), we obtain:

$$\mathbf{n}_t(w) \leq k \sum_{\tau=0}^t P_\tau(w, v) = k \sum_{\tau=0}^t P_\tau^{(0)}(w).$$

A derivation of an upper bound on $\sum_{\tau=0}^t P_\tau^{(0)}(w)$ is obtained in [11] (they consider the lazy random walk on a hypercube with d self-loops at each node, but the same result can be applied to the hypercube without self-loops, after relaxing time bounds by a constant factor of 4):

$$\mathbf{n}_t(w) \leq k \sum_{\tau=0}^t P_\tau^{(0)}(w) \leq \frac{4tk}{n} \left(1 - \left(1 - \frac{1}{\log n} \right)^{4t} \right)^{\log n}$$

Substituting $k \leq n \cdot 2^{\log^{1-\varepsilon} n}$ and $t \leq \frac{\varepsilon}{10} \log n \log \log n$, we obtain $\mathbf{n}_t(w) < 1$, which completes the proof. \square

We leave the question of the cover time of the rotor-router on the hypercube for $k > n \cdot 2^{\log^{1-\varepsilon} n}$ as open.

References

- [1] H. Akbari and P. Berenbrink. Parallel rotor walks on finite graphs and applications in discrete load balancing. In *SPAA*, pages 186–195, 2013.
- [2] D. Aldous and J. Fill. Reversible markov chains and random walks on graphs. <http://stat-www.berkeley.edu/users/aldous/RWG/book.html>, 2001.
- [3] N. Alon, C. Avin, M. Koucký, G. Kozma, Z. Lotker, and M. R. Tuttle. Many random walks are faster than one. *Combinatorics, Probability & Computing*, 20(4):481–502, 2011.
- [4] E. Bampas, L. Gasieniec, N. Hanusse, D. Ilcinkas, R. Klasing, and A. Kosowski. Euler tour lock-in problem in the rotor-router model. In *DISC*, pages 423–435, 2009.
- [5] P. Berenbrink, C. Cooper, T. Friedetzky, T. Friedrich, and T. Sauerwald. Randomized diffusion for indivisible loads. In D. Randall, editor, *SODA*, pages 429–439. SIAM, 2011.
- [6] J. N. Cooper, B. Doerr, J. H. Spencer, and G. Tardos. Deterministic random walks on the integers. *Eur. J. Comb.*, 28(8):2072–2090, 2007.
- [7] D. Dereniowski, A. Kosowski, D. Pajak, and P. Uznanski. Bounds on the cover time of parallel rotor walks. In E. W. Mayr and N. Portier, editors, *STACS*, volume 25 of *LIPICs*, pages 263–275. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014.
- [8] B. Doerr and T. Friedrich. Deterministic random walks on the two-dimensional grid. *Combinatorics, Probability, and Computing*, 18(1-2):123–144, 2009.
- [9] K. Efremenko and O. Reingold. How well do random walks parallelize? In *APPROX-RANDOM*, pages 476–489, 2009.
- [10] R. Elsässer and T. Sauerwald. Tight bounds for the cover time of multiple random walks. In *ICALP (1)*, volume 5555 of *LNCS*, pages 415–426, 2009.
- [11] R. Elsässer and T. Sauerwald. Tight bounds for the cover time of multiple random walks. *Theor. Comput. Sci.*, 412(24):2623–2641, 2011.
- [12] T. Friedrich, M. Gairing, and T. Sauerwald. Quasirandom load balancing. *SIAM J. Comput.*, 41(4):747–771, 2012.
- [13] S. Kijima, K. Koga, and K. Makino. Deterministic random walks on finite graphs. In C. Martnez and H.-K. Hwang, editors, *ANALCO*, pages 18–27. SIAM, 2012.
- [14] R. Klasing, A. Kosowski, D. Pajak, 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] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2006.
- [16] S. D. P. Uznanski. Personal communication. 2014.
- [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 694–705. IEEE Computer Society, 1998.
- [19] V. Yanovski, I. A. Wagner, and A. M. Bruckstein. A distributed ant algorithm for efficiently patrolling a network. *Algorithmica*, 37(3):165–186, 2003.

Acknowledgment

The authors thank Petra Berenbrink, Ralf Klasing, and Frederik Mallmann-Trenn for productive discussions on closely related topics.