

# Recurrent rotor-router configurations

OMER ANGEL AND ALEXANDER E. HOLROYD

We prove the existence of recurrent initial configurations for the rotor walk on many graphs, including  $\mathbb{Z}^d$ , and planar graphs with locally finite embeddings. We also prove that recurrence and transience of rotor walks are invariant under changes in the starting vertex and finite changes in the initial configuration.

AMS 2000 SUBJECT CLASSIFICATIONS: 05C25, 82C20.

KEYWORDS AND PHRASES: Rotor walk, rotor-router, quasi-random, recurrence.

## 1. Introduction

The **rotor walk** is a derandomized variant of random walk on a graph  $G$ , defined as follows. To each vertex of  $G$  we assign a fixed cyclic order of its neighbours. This collection of orders is called the **rotor mechanism**. At each vertex there is a **rotor**: an arrow which can point to any neighbour. An assignment of directions to all the rotors is called a **rotor configuration**. Starting from some rotor configuration, a **particle** is located at some vertex, and the particle location and rotor configuration evolve together in discrete time as follows. At each time step, the rotor at the particle's current location is incremented to point to the next neighbour in the cyclic order, and then the particle moves to this new neighbour. The rotor walk is obtained by repeatedly applying this rule.

We will assume throughout that  $G$  is an infinite, connected, simple, undirected graph with all degrees finite. Given a rotor mechanism, an initial rotor configuration, and an initial particle location, it is then easy to see that the rotor walk either visits each vertex infinitely many times, or visits each vertex finitely many times (see e.g. [6, Lemma 6]). We call these two cases **recurrent** and **transient** respectively.

In many settings there are striking similarities between the behaviour of the rotor walk and the expected behaviour of *random* walk on the same graph; see e.g. [2, 3, 4, 6, 8]. However, with regard to recurrence and transience there can also be differences. For instance, on any  $G$  (even a recurrent

graph) it is easy to find an initial rotor configuration that is transient (indeed, we can arrange for the particle to trace any self-avoiding path). In the other direction, it was shown in [7] that recurrent initial rotor configurations exist on the infinite binary tree. These matters have been investigated further on trees in [1, 7], but much less is known for more general graphs. Our goal in this article is to prove the existence of recurrent rotor configurations in a broad range of settings, including many where the random walk is transient.

We will also show that recurrence and transience of the rotor walk do not depend on the starting vertex. Therefore we may refer to an initial rotor configuration as recurrent or transient without specifying the starting vertex.

**Theorem 1.** *Fix a graph, a rotor mechanism and an initial rotor configuration. The rotor walk is either recurrent for every starting vertex, or transient for every starting vertex.*

Let  $\mathbb{Z}^d$  denote the graph with vertex set  $\mathbb{Z}^d$  and an edge between each pair of vertices at Euclidean distance 1.

**Theorem 2.** *For any  $d \geq 1$  and any rotor mechanism on  $\mathbb{Z}^d$ , there exists a recurrent rotor configuration.*

An example was given in [6] of a recurrent rotor configuration on  $\mathbb{Z}^2$ . For all  $d \geq 3$ , Theorem 2 provides the first proof of the existence of a recurrent configuration for a translation-invariant rotor mechanism. This answers a question posed by Jim Propp in 2003 (personal communication), and stated as an open problem in [6] and [7]. Our proof is strongly motivated by a recent work of Reddy [9], in which it is argued on the basis of simulation evidence that a certain explicit rotor configuration on  $\mathbb{Z}^3$  is recurrent. Our proof also includes a simple explicit rotor configuration.

Our construction can be generalized to many other settings, and in particular we have the following. A planar embedding of a planar graph is called **locally finite** if every bounded region of the plane contains only finitely many vertices. (There exist planar graphs with no locally finite embedding, such as  $\mathbb{Z}^2$  with a singly-infinite path attached to every vertex.)

**Theorem 3.** *For a planar graph with a locally finite planar embedding, and a rotor mechanism in which each rotor points to the neighbours in clockwise or anticlockwise order (possibly with different directions at different vertices), there exists a recurrent rotor configuration.*

**Theorem 4.** *For any graph, there exists a rotor mechanism that admits a recurrent rotor configuration.*

For some graphs including  $\mathbb{Z}^d$ , we can also give recurrent configurations for which the behaviour of the rotor walk is extremely regular, as follows. For  $x \in \mathbb{Z}^d$ , define

$$\alpha(x) := \begin{cases} 1 & \text{if } |x_1|, \dots, |x_d| \text{ have a unique maximum;} \\ 0 & \text{otherwise.} \end{cases}$$

**Theorem 5.** *For any rotor mechanism on  $\mathbb{Z}^d$ , there exists a recurrent rotor configuration such that for the rotor walk started at 0, just before the  $(k+1)$ st traversal from 0 to  $(1, 0, \dots, 0)$ , vertex  $x$  has been entered exactly*

$$[2d(k - \|x\|_\infty) + \alpha(x)]^+$$

*times.*

For the reader's convenience we present an explicit family of configurations on  $\mathbb{Z}^d$  satisfying the conclusion of Theorem 5 (see Section 4 for a more general result). Let  $x \in \mathbb{Z}^d$ . If  $\alpha(x) = 1$ , set the rotor at  $x$  so that it will *next* point to the unique neighbour  $y$  such that  $\|y\|_\infty < \|x\|_\infty$ . If  $\alpha(x) = 0$  and  $x \neq 0$ , let the rotor at  $x$  point towards any neighbour  $y$  with  $\|y\|_\infty < \|x\|_\infty$ . Set the rotor at 0 to next point towards  $(1, 0, \dots, 0)$ .

## 2. Invariance

In this section we prove Theorem 1, and a corollary thereof. These results will not be needed for the proofs of Theorems 2–5. The proof of Theorem 1 uses the Abelian property of the rotor walk. We will use a truncation argument in order to apply a version of the Abelian property for finite graphs proved in [5]. An alternative approach would be to derive a version of the Abelian property that holds for transfinite-time walks on infinite graphs. Write  $d_x$  for the degree of vertex  $x$ .

*Proof of Theorem 1.* It is clearly sufficient to prove that if a rotor configuration is recurrent for the rotor walk started at some vertex  $x$  then it is also recurrent for the walk started at any neighbour  $y$  of  $x$ . For any  $m \geq 0$ , we show that the rotor walk started at  $y$  visits  $x$  at least  $m$  times.

Let  $S$  be the finite set of vertices visited by the rotor walk started at  $x$  until it has made  $d_x + m$  returns to  $x$ . Let  $F$  be the subgraph of  $G$  induced by  $S$ , modified as follows: add a sink vertex  $z$ , and replace each edge of  $G$  leaving  $S$  by a directed edge to  $z$ ; also split  $x$  into two vertices  $x_+$  and  $x_-$ , and split each edge incident to  $x$  into a directed edge from  $x_+$  and a directed

edge to the sink  $x_-$ . By the Abelian property, [5, Lemma 3.9], if we start  $d_x + m$  rotor particles at  $x_+$ , and let them perform rotor walks until they reach the set of sinks  $\{z, x_-\}$ , then they all in fact reach  $x_-$ , regardless of the order in which they move.

Start  $d_x + m$  particles at  $x_+$  and move them in  $F$  in the following order. First, let  $d_x$  particles each take one step. This leaves the rotor configuration unchanged, and one particle at each neighbour of  $x$  (including  $y$ ). Now let the particle at  $y$  perform rotor walk. It will follow exactly the trajectory of the rotor walk in  $G$  started at  $y$ . Each time it is absorbed at  $x_-$ , continue with one of the  $m$  particles remaining at  $x_+$ . It follows that the rotor walk from  $y$  in  $G$  visits  $x$  at least  $m$  times before leaving  $S$ .  $\square$

One consequence of Theorem 1 is that recurrence and transience are also insensitive to local changes in the configuration.

**Corollary 6.** *Fix a graph and a rotor mechanism. If two rotor configurations differ only at finitely many vertices, then they are either both recurrent or both transient.*

*Proof.* It suffices to consider the case of two rotor configurations  $r$  and  $r'$  that differ only at a single vertex  $x$ , at which the rotor is incremented once in  $r'$  compared with  $r$ . Suppose that  $r$  is recurrent, and start a rotor walk at  $x$ . At the first step, the rotor configuration becomes  $r'$ , and the particle moves to a neighbour  $y$ , say. Hence the rotor walk started at  $y$  is recurrent for  $r'$ , i.e.  $r'$  is recurrent.  $\square$

### 3. Recurrent configurations

Theorems 2–4 are consequences of the more general result below. For a set of vertices  $S$ , let  $\partial S$  denote its outer vertex-boundary:

$$\partial S := \{x \in S^C : x \text{ has a neighbour in } S\}.$$

We say that  $S$  has **reflecting boundary** (for a given rotor mechanism and initial rotor configuration) if for every vertex  $y$  in  $\partial S$ , the rotor at  $y$  will send the particle to each of  $y$ 's neighbours in  $S$  before sending it to any other neighbour of  $y$ .

**Proposition 7.** *If for some rotor configuration on a graph  $G$ , every finite set of vertices is a subset of some finite set with reflecting boundary, then the rotor walk starting from any vertex is recurrent.*

To prove the above result, it will sometimes be convenient to identify  $G$  with the directed graph in which each undirected edge is replaced with two directed edges, one in each direction. We say that the rotor walk traverses the directed edge  $(x, y)$  when the particle makes a step from vertex  $x$  to vertex  $y$ .

**Lemma 8.** *If a rotor walk started at  $x$  traverses some directed edge twice, then the first directed edge to be traversed twice is from  $x$ . In particular such a walk has returned to  $x$ .*

*Proof.* Let  $(y, z)$  be the first directed edge to be traversed twice. At that time  $y$  has sent the particle to all other neighbours exactly once, thus  $y$  has emitted the particle  $d_y + 1$  times. Since no other edge has been traversed twice,  $y$  has received the particle at most  $d_y$  times. Thus  $y = x$ .  $\square$

**Lemma 9.** *If  $S$  is a set with reflecting boundary, then the rotor walk started at a vertex  $x \in S$  will return to  $x$  before leaving  $S \cup \partial S$ .*

*Proof.* We prove the stronger statement that the walk will return to  $x$  before making any step from  $\partial S$  to  $S^C$ . Consider the first such step, from some  $y \in \partial S$  to some  $z \notin S$ . At this time, by the definition of reflecting boundary,  $y$  has previously sent the particle to each of its neighbours in  $S$ . However, by our assumption it has received the particle only from its neighbours in  $S$ , therefore it must have received it twice from some such neighbour. Now apply Lemma 8.  $\square$

*Proof of Proposition 7.* Suppose the particle started at  $x$ , and is currently at  $x$ , and let  $A$  be the (finite) set of vertices that have been visited. Then  $A \subseteq S$  for some  $S$  which had reflecting boundary in the initial rotor configuration. Since  $\partial S$  has not yet been visited,  $S$  still has reflecting boundary. By Lemma 9, the walk will return again to  $x$  (before leaving  $S \cup \partial S$ ). Iterating this shows that the walk is recurrent.  $\square$

*Proof of Theorem 2.* By Proposition 7, it suffices to choose an initial rotor configuration so that every cube of the form  $[-n, n]^d$  ( $n = 0, 1, 2, \dots$ ) has reflecting boundary. Since each vertex  $z$  of  $\partial([-n, n]^d)$  has only one neighbour  $y$  in  $[-n, n]^d$ , this is achieved by setting the rotor at  $z$  so that it will next point to  $y$ . Since the boundaries of different cubes are disjoint, this can be done for all  $z$  and  $n$ .  $\square$

*Proof of Theorem 4.* Let  $S_0$  be any finite non-empty set of vertices, and construct  $S_1, S_2, \dots$  inductively by  $S_{i+1} := S_i \cup \partial S_i$ . By Proposition 7, it suffices to choose the rotor mechanism and configuration so that each  $S_i$  has reflecting boundary. This is clearly possible, since the sets  $\partial S_0, \partial S_1, \dots$  are disjoint.  $\square$

To prove Theorem 3, we need a lemma about planar graphs. Given a planar embedding of a planar graph, the edges incident to a given vertex fall in some cyclic order around it. We say that a set of vertices  $S$  is **pincerless** if for every  $x \in \partial S$ , either all neighbours of  $x$  lie in  $S$ , or the incident edges joining  $x$  to  $S$  lie in one contiguous interval in the cyclic order around  $x$ .

**Lemma 10.** *Let  $G$  be an infinite, connected, planar, simple graph with all degrees finite, and with a locally finite planar embedding. Any finite set of vertices is a subset of some finite pincerless set.*

*Proof.* Let  $A$  be a finite set; we will show that it is a subset of some finite pincerless  $B$ . By enlarging  $A$  if necessary, we may assume that  $A$  is connected (i.e. it induces a connected subgraph of  $G$ ).

Consider any  $x \in \partial A$ . Each neighbour of  $x$  lies either in  $A$ , or in a finite or an infinite component of  $(A \cup \{x\})^C$ . We call the three types  $A$ -neighbours,  $F$ -neighbours, and  $I$ -neighbours, respectively. Note that there is at least one  $A$ -neighbour. We claim that it is impossible for two  $A$ -neighbours  $a, a'$  and two  $I$ -neighbours  $i, i'$  to alternate in the cyclic order of neighbours of  $x$  (i.e. to occur in the order  $aia'i'$  when we delete all other neighbours from the order). This follows from planarity and local finiteness, because  $a$  and  $a'$  are connected by a path in  $A(\not\ni x)$ , while  $i$  and  $i'$  are each connected to infinity off  $A \cup \{x\}$ <sup>1</sup>. Therefore, there exists some interval in the cyclic order of neighbours of  $x$  that contains all the  $A$ -neighbours and no  $I$ -neighbours. Let  $J_x$  be the unique minimal such interval, if there is at least one  $I$ -neighbour, and otherwise the set of all  $x$ 's neighbours. Define  $D_x$  to be the union of all the finite components of  $(A \cup \{x\})^C$  corresponding to  $F$ -neighbours in  $J_x$ . Define

$$B := A \cup \bigcup_{x \in \partial A} D_x.$$

The set  $B$  is clearly finite, and we must check that it is pincerless. First, we claim that if a vertex  $z$  is adjacent to some vertex in  $D_y$  for some  $y \in \partial A$  with  $y \neq z$ , then  $z \in B$ . This follows because either  $z \in A \subseteq B$ , or  $z \in (A \cup \{y\})^C$ , in which case  $z$  lies in the same component of the latter

---

<sup>1</sup>More formally, fix a simple path in  $A$  from  $a$  to  $a'$ , and let  $C$  be the curve formed by the embedding of this path together with the edges  $(x, a)$  and  $(x, a')$ . The Jordan curve theorem states that the complement of  $C$  has two components: a bounded interior and an unbounded exterior, with  $C$  being the boundary of each. The alternation assumption at  $x$  implies that exactly one of  $(x, i), (x, i')$ , say  $(x, i)$ , lies in the interior. By planarity, the component of  $(A \cup \{x\})^C$  containing  $i$  must also lie in the interior, contradicting local finiteness.

set as does its neighbour in  $D_y$ , so since  $D_y$  is a union of such components,  $z \in D_y \subseteq B$ .

Now consider any  $x \in \partial B$ . We must have  $x \in \partial A$ , otherwise the above claim would imply  $x \in B$ . Now by the definition of  $D_x$ , all of  $x$ 's neighbours in the interval  $J_x$  lie in  $B$ . Therefore to check the pincerless condition at  $x$  it suffices to show that no neighbour of  $x$  not in  $J_x$  lies in  $B \setminus A$ . By the definition of  $B$ , such a neighbour would lie in  $D_y$  for some  $y \neq x$ , and hence by the claim again we would have  $x \in B$ , a contradiction.  $\square$

*Proof of Theorem 3.* Let  $S_0$  be any finite non-empty set of vertices, and define  $S_1, S_2, \dots$  inductively by taking  $S_{i+1}$  to be a finite pincerless set containing  $S_i \cup \partial S_i$ , by Lemma 10. For  $i \geq 1$ , since  $S_i$  is pincerless, and the rotors rotate clockwise or anticlockwise, the rotors in  $\partial S_i$  can be given initial directions so that  $S_i$  has reflecting boundary. Since the  $\partial S_i$  are disjoint, this can be done for all  $i \geq 1$  simultaneously.  $\square$

#### 4. Exact number of visits

We now turn to Theorem 5, which is a consequence of the more general result below. Given a rotor configuration and some set of vertices  $U \subseteq V(G)$ , we associate a directed graph with vertex set  $V$ , and a directed edge from  $x$  to  $y$  whenever  $x \in U$  and the rotor at  $x$  points to  $y$ .

**Proposition 11.** *Let  $S_0, S_1, \dots$  be finite sets such that  $S_0 = \{a\}$ , and  $S_i \cup \partial S_i \subseteq S_{i+1}$  for all  $i$ . Suppose that in the initial rotor configuration, for each  $i \geq 0$ :*

1.  $S_i$  has reflecting boundary, and
2. the rotors at vertices of  $S_{i+1} \setminus (S_i \cup \partial S_i)$  form a directed forest pointing towards  $S_i \cup \partial S_i$ .

*Consider the rotor walk started at  $a$ , and let  $b$  be the next vertex it visits. Then for every  $k \geq 1$ , in the time interval from the  $k$ th to the  $(k + 1)$ st traversal from  $a$  to  $b$  (inclusive and exclusive respectively), the rotor walk traverses each edge incident to  $S_k$  exactly once in each direction, and no other edges.*

*Proof.* By a **unicycle rooted at  $x$**  we mean a directed graph comprising a single oriented cycle passing through vertex  $x$ , together with a collection of directed trees rooted on the cycle and pointing towards it. We will use the following fact about the rotor walk on a *finite* undirected graph (or indeed an Eulerian directed graph); see e.g. [5, Lemma 4.9]. Starting from any rotor configuration that forms a unicycle rooted at the current particle location, if

$2m$  rotor steps are performed, where  $m$  is the number of edges of the graph, then each edge is traversed exactly once in each direction, and the rotors and particle finish in their initial positions.

The rotor walk is recurrent by Proposition 7. Let  $t_k$  be the time just before the  $(k + 1)$ st traversal from  $a$  to  $b$ , so  $t_0 = 0$ , and at time  $t_k$  the particle is at  $a$ . We will prove by induction that for all  $k \geq 1$ : at time  $t_k$ , no vertex outside  $S_k$  has been visited, and the rotors in  $S_k$  form a unicycle rooted at  $a$ . In the process of proving this we will establish the claim of the proposition.

First consider the time period from  $t_0$  to  $t_1$ . Since  $S_0 = \{a\}$  has reflecting boundary, the particle simply traverses each edge incident to  $a$  in each direction, so the claim of the proposition holds for  $k = 0$ . At time  $t_1$ , the rotors of  $\partial S_0$  all point towards  $a$ , so the rotors of  $S_0 \cup \partial S_0$  form a unicycle rooted at  $a$ . Since the vertices of  $S_1 \setminus (S_0 \cup \partial S_0)$  have not been visited, condition (ii) of the proposition thus implies that the rotors of  $S_1$  form a unicycle rooted at  $a$ , establishing the inductive hypothesis for  $k = 1$ .

Now suppose the inductive hypothesis holds for some  $k \geq 1$ . Consider the finite subgraph  $F_k$  of  $G$  comprising all edges incident to  $S_k$ , with vertex set  $S_k \cup \partial S_k$ . Consider a rotor configuration on  $F_k$  given as follows. Let the vertices in  $S_k$  inherit their positions from  $G$  at time  $t_k$ . Recalling that  $S_k$  has reflecting boundary in  $G$ , fix the rotor at each  $y \in \partial S_k$  so that it will next send the particle to each of  $y$ 's neighbours in  $S_k$ , in the same order as in  $G$ ; thus, the rotor should initially point to the last of these neighbours. Since the rotors of  $S_k$  form a unicycle rooted at  $a$ , the same applies to  $F_k$ . Therefore, running the rotor walk on  $F_k$ , started at  $a$ , until just before the second traversal from  $a$  to  $b$  results in each edge of  $F_k$  being traversed once in each direction, and the same final rotor configuration. Since the rotors in  $\partial S_k$  have only sent the particle towards  $S_k$ , the behaviour of the walk on  $G$  over this time interval is identical. Hence the claim of the proposition holds for  $k$ . At time  $t_{k+1}$ , the rotors in  $S_k \cup \partial S_k$  form a unicycle rooted at  $a$ , and condition (ii) again implies that the same holds when all rotors in  $S_{k+1}$  are included, establishing the inductive hypothesis for  $k + 1$ .  $\square$

*Proof of Theorem 5.* We apply Proposition 11, with  $S_i = [-i, i]^d$ . The boundary  $\partial S_i$  consists precisely of those vertices  $x$  of  $S_{i+1} \setminus S_i$  that have  $\alpha(x) = 1$  (i.e. those on  $(d - 1)$ -dimensional faces of  $S_{i+1} \setminus S_i$ ). We set the rotors at these vertices to point *next* towards  $S_i$ , and those of  $S_{i+1} \setminus (S_i \cup \partial S_i)$  to form a forest pointing towards  $S_i \cup \partial S_i$ , as required. By Proposition 11, between times  $t_k$  and  $t_{k+1}$ , each vertex of  $S_k$  is visited  $2d$  times, and each vertex of  $\partial S_k$  is visited once. Summing over  $k$  gives the claimed expression.  $\square$



### Open questions

1. Do there exist a graph and rotor mechanism for which every initial rotor configuration is transient?
2. While Theorem 5 provides a detailed description of the behaviour of certain recurrent rotor configurations, some simple transient examples remain mysterious. On  $\mathbb{Z}^2$ , let each rotor initially point East and rotate anticlockwise. Start a rotor walk at 0, and restart it at 0 after each escape to infinity. What is the asymptotic growth rate of the number of escapes to infinity prior to the  $n$ th visit to 0, as  $n \rightarrow \infty$ ? (By a result of Schramm, [6, Theorem 10], it is  $o(n)$ .)

### References

- [1] Angel, O. and Holroyd, A. E. (2011). Rotor walks on general trees. *SIAM J. Discrete Math.* **25** 423–446. [MR2801237](#)
- [2] Cooper, J., Doerr, B., Spencer, J., and Tardos, G. (2006). Deterministic random walks. In: *Proceedings of the Eighth Workshop on Algorithm Engineering and Experiments and the Third Workshop on Analytic Algorithmics and Combinatorics*, pp. 185–197. SIAM, Philadelphia, PA. [MR2498151](#)
- [3] Cooper, J. N. and Spencer, J. (2006). Simulating a random walk with constant error. *Combin. Probab. Comput.* **15** 815–822. [MR2271828](#)
- [4] Doerr, B. and Friedrich, T. (2009). Deterministic random walks on the two-dimensional grid. *Combin. Probab. Comput.* **18** 123–144. [MR2497377](#)
- [5] Holroyd, A. E., Levine, L., Mészáros, K., Peres, Y., Propp, J. and Wilson, D. B. (2008). Chip-firing and rotor-routing on directed graphs. In: *In and out of equilibrium. 2. Progr. Probab.* **60** 331–364. Birkhäuser, Basel. [MR2477390](#)
- [6] Holroyd, A. E. and Propp, J. (2010). Rotor walks and Markov chains. In: *Algorithmic probability and combinatorics. Contemp. Math.* **520** 105–126. Amer. Math. Soc., Providence, RI. [MR2681857](#)
- [7] Landau, I. and Levine, L. (2009). The rotor-router model on regular trees. *J. Combin. Theory Ser. A* **116** 421–433. [MR2475025](#)

- [8] Levine, L. and Peres, Y. (2009). Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile. *Potential Anal.* **30** 1–27. [MR2465710](#)
- [9] Reddy A, T. R. (2010). A recurrent rotor-router configuration in  $Z^3$ . Preprint, arXiv:1005.3962.

OMER ANGEL  
DEPARTMENT OF MATHEMATICS  
UNIVERSITY OF BRITISH COLUMBIA  
VANCOUVER, BC V6T 1Z2  
CANADA  
*E-mail address:* [angel@math.ubc.ca](mailto:angel@math.ubc.ca)

ALEXANDER E. HOLROYD  
MICROSOFT RESEARCH  
1 MICROSOFT WAY  
REDMOND, WA 98052  
USA  
*E-mail address:* [holroyd@microsoft.com](mailto:holroyd@microsoft.com)

RECEIVED JANUARY 24, 2011