Local-to-global principles for the hitting sequence of a rotor walk

Giuliano Pezzolo Giacaglia
Department of Mathematics
Massachusetts Institute of Technology
Cambridge, Massachusetts USA
giu@mit.edu

Lionel Levine
Department of Mathematics
Cornell University
Ithaca, New York USA
http://www.math.cornell.edu/~levine

James Propp
Department of Mathematical Sciences
University of Massachusetts Lowell
Lowell, Massachusetts USA
http://jamespropp.org

Linda Zayas-Palmer
Department of Mathematics
Massachusetts Institute of Technology
Cambridge, Massachusetts USA
lzayasp@mit.edu

Submitted: Aug 8, 2011; Accepted: Dec 16, 2011; Published: Jan 6, 2012
Mathematics Subject Classification: 05C25, 05C38, 05C81, 90B10

Abstract

In rotor walk on a finite directed graph, the exits from each vertex follow a prescribed periodic sequence. Here we consider the case of rotor walk where a particle starts from a designated source vertex and continues until it hits a designated target set, at which point the walk is restarted from the source. We show that the sequence of successively hit targets, which is easily seen to be *eventually* periodic, is in fact periodic. We show moreover that reversing the periodic patterns of all rotor sequences causes the periodic pattern of the hitting sequence to be reversed as well. The proofs involve a new notion of equivalence of rotor configurations, and an extension of rotor walk incorporating time-reversed particles.

Keywords: Cycle popping, hitting sequence, monoid action, rotor-router model, sandpile group, sandpile monoid.

1 Introduction

A rotor walk in a graph G is a walk in which the sequence of exits from each vertex is periodic. The sequence of exits from a vertex v is called the *rotor mechanism* at v. Rotor walks have been studied in combinatorics as deterministic analogues of random walks, in computer science as a means of load-balancing and territory exploration, and in statistical

physics as a model of self-organized criticality (see the end of the introduction for references). In this paper, we explore several properties of the rotor mechanism that imply corresponding properties of the *hitting sequence* when G comes with a set of designated *target vertices*:

- Given a (periodic) rotor mechanism at each vertex, the hitting sequence of the associated rotor walk is periodic (Theorem 1).
- If every rotor mechanism is palindromic, then the hitting sequence is palindromic (Theorem 3).
- If every rotor mechanism is m-repetitive, then the hitting sequence is m-repetitive (Theorem 4).

See below for precise definitions. Since the rotor mechanisms are local features of the walk — each one depends only on the exits from a particular vertex — while the hitting sequence is a global feature, we regard these theorems as *local-global principles*.

Let $G = (V, E)$ be a finite directed graph, with self-loops and multiple edges permitted. For vertices $v, w \in V$, let $d(v, w)$ denote the number of arcs from v to w, and let $d(v) = \sum_{w \in V} d(v, w)$ denote the outdegree of v. A *rotor mechanism* at v is an ordering of the directed edges (or "arcs") emanating from v, say as e_v^i for $1 \leq i \leq d(v)$. Let v^i denote the endpoint of the arc e_v^i. We extend the definition of e_v^i and v^i to all $i \in \mathbb{Z}$ by taking e_v^i and v^i to be periodic in i with period $d(v)$. We often indicate the rotor mechanism at v using the notation

$$v \to v^1, v^2, \ldots, v^{d(v)}, \ldots \text{ (period } d(v)).$$

Given a rotor mechanism at each vertex v, a *rotor walk* on G is a finite or infinite sequence of vertices $x_0, x_1, x_2, \ldots$ in which the i-th occurrence of v is followed immediately by an occurrence of v^i. For example, if the vertex set of G is $\{1, 2, 3, 4, 5\}$ and the rotor mechanisms are

$$\begin{aligned}
1 &\to 3, 4, 5, \ldots \text{ (period 3)} \\
2 &\to 3, \ldots \text{ (period 1)} \\
3 &\to 4, 2, \ldots \text{ (period 2)} \\
4 &\to 1, \ldots \text{ (period 1)} \\
5 &\to 1, \ldots \text{ (period 1)}
\end{aligned}$$

then the rotor walk starting from 1 is

$$1, 3, \mathbf{4}, 1, \mathbf{4}, 1, \mathbf{5}, 1, 3, 2, 3, \mathbf{4}, 1, \mathbf{4}, 1, \mathbf{5}, 1, 3, 2, 3, \mathbf{4}, 1, \ldots$$

which is eventually periodic with period 9. Note that this sequence is not itself periodic (the initial 1 does not repeat) but if we isolate the terms equal to 4 or 5 we obtain the sequence

$$4, 4, 5, 4, 4, 5, \ldots$$

which is periodic with period 3.

Vertex v is *reachable* from vertex w if there is a directed path of edges in G from v to w. We assume that our graph G comes with a designated *source vertex* s and a non-empty set T of designated *target vertices*, such that

$$\begin{gathered}\text{For every vertex } v \text{ of } G \text{ that is reachable from } s,\\ \text{at least one vertex of } T \text{ is reachable from } v.\end{gathered} \tag{1}$$

The role of this condition is to ensure that a rotor walk started at s eventually reaches T. In the example above, $s = 1$ and $T = \{4, 5\}$. All of the rotor walks we consider in this paper will start at vertices reachable from s. Such a walk can never visit a vertex not reachable from s, so there is no loss of generality in assuming

$$\text{All vertices of } G \text{ are reachable from } s. \tag{2}$$

Now consider the following experiment. Start a rotor walk at s and stop it when it first visits the target set T, say at target t_1. Then (without resetting the rotors) start a new rotor walk at s and stop it when it first visits T, say at t_2. Continuing in this way, the sequence $t_1, t_2, \ldots$ of successively hit targets is called the *hitting sequence*. In the above example, the hitting sequence is $4, 4, 5, 4, 4, 5, \ldots$.

In the experiment just described, all rotor walks stop as soon as they hit T, so outgoing edges from vertices in T are never used. An equivalent way to define the hitting sequence uses a modified graph G' in which each target vertex $t \in T$ has just one outgoing edge, which points to s. The effect is that rotor walk in G' always returns to s immediately after visiting a target vertex. Hence, the sequence of rotor walks in G in the experiment can be realized by a single rotor walk in G'. Conditions (1) and (2) together on G are equivalent to requiring that the modified graph G' is *strongly connected*, that is, for any two vertices v and w, each is reachable from the other by a directed path of edges in G'.

It is easy to show that rotor walk in G' visits T infinitely often, so that the hitting sequence is infinite; see Lemma 6, below. As explained in §2, it is also easy to show that the hitting sequence is eventually periodic. Our first main result goes further:

Theorem 1. *The hitting sequence determined by a (periodic) rotor mechanism is periodic.*

As we have already seen, the rotor walk itself is typically *not* periodic. Let γ_n be the portion of the walk strictly between the n-th and $(n+1)$-st visits to T. In the preceding example, the sequence $\{\gamma_n\}_{n\geq 1}$ is

$$13, 1, 1, 1323, 1, 1, 1323, 1, 1, \ldots.$$

In general, this sequence is eventually periodic but is not periodic.

A natural question is how to determine the period of the hitting sequence. We will see that this period divides the order of a certain element of the *sandpile group* $S(G/T)$ of the graph G with the target set T collapsed to a single vertex (Lemma 20).

Our second main result states that if we *reverse* the rotor mechanism at each vertex by replacing

$$v \to v^1, v^2, \ldots, v^{d(v)}, \ldots \text{ (period } d(v))$$

by

$$v \to v^{d(v)}, v^{d(v)-1}, \ldots, v^1, \ldots \text{ (period } d(v))$$

for each vertex v, then the hitting sequence undergoes an analogous reversal; specifically, if the original hitting sequence has period D, then the new hitting sequence will also have period D, and for all $1 \leq i \leq D$ the i-th term of the new hitting sequence will equal the $(D+1-i)$-th term of the original hitting sequence. That is:

Theorem 2. *Reversing the periodic pattern of all rotor mechanisms results in reversing the periodic pattern of the hitting sequence.*

E.g., for the above example, the reversed rotor mechanism

$$\begin{aligned} 1 &\to 5, 4, 3, \ldots \text{ (period 3)} \\ 2 &\to 3, \ldots \text{ (period 1)} \\ 3 &\to 2, 4, \ldots \text{ (period 2)} \\ 4 &\to 1, \ldots \text{ (period 1)} \\ 5 &\to 1, \ldots \text{ (period 1)} \end{aligned}$$

gives the reversed hitting sequence $5, 4, 4, 5, 4, 4, \ldots$.

An immediate corollary of Theorem 2 is that if the rotors are all *palindromic* (that is, if each fundamental period of each rotor reads the same backwards and forwards) then the same is true of the hitting sequence.

Theorem 3. *If all rotor mechanisms are palindromic, then the hitting sequence is palindromic.*

One can think of the entire collection of rotor mechanisms on G as a single rotor — perhaps embedded as a component of a larger system — whose rotor mechanism is the hitting sequence. From this perspective, Theorems 1 and 3 are local-to-global principles asserting that if the sequence of exits from each vertex possesses a certain property (periodicity, palindromicity), then the hitting sequence has the same property. We now state one more result of this type, Theorem 4. Further examples of local-global principles include Lemma 6, below, and [12, Theorem 1].

Call a sequence $\{u_i\}_{i \geq 1}$ *m-repetitive* if it consists of blocks of m consecutive equal terms; that is,

$$u_{am+1} = u_{am+2} = \ldots = u_{am+m}$$

for all $a \geq 0$.

Theorem 4. *If all rotor mechanisms are m-repetitive, then the hitting sequence is m-repetitive.*

For example, consider the 2-repetitive rotor mechanism

$$\begin{aligned} 1 &\to 3, 3, 2, 2, \ldots \text{ (period 4)} \\ 2 &\to 1, 1, 4, 4, \ldots \text{ (period 4)} \\ 3 &\to 1, \ldots \text{ (period 1)} \\ 4 &\to 1, \ldots \text{ (period 1)} \end{aligned}$$

with source 1 and targets 3 and 4. The sequence of paths γ_n taken by the walker until it hits a target

$$13, 13, 121213, 13, 124, 124, \ldots \text{ (period 6)}$$

is not 2-repetitive, but the hitting sequence

$$3, 3, 3, 3, 4, 4, \ldots \text{ (period 6)}$$

is 2-repetitive.

The proof of Theorem 4 is not difficult (see §2) and uses only the abelian property of rotor walk (Lemma 7). The proofs of Theorems 1-3 make essential use of a new notion of equivalence of rotor configurations. We summarize the highlights here, referring the reader to §3.1 for the full definitions.

Let $V_0 = V - T$. A *rotor configuration* is a map $\rho : V_0 \to E$ such that $\rho(v)$ is an arc emanating from v; the arc $\rho(v)$ represents the arc by way of which a particle most recently exited vertex v. A *particle configuration* is a map $\sigma : V_0 \to \mathbb{N}$; we interpret $\sigma(v)$ as the number of particles present at vertex v. Following [11] we define an action $(\sigma, \rho) \mapsto \sigma\rho$ of particle configurations on rotor configurations. We then define rotor configurations ρ_1 and ρ_2 to be equivalent, written $\rho_1 \equiv \rho_2$, if there exists a particle configuration σ such that $\sigma\rho_1 = \sigma\rho_2$ (Lemma 10 will show that this is an equivalence relation. In fact, $\rho_1 \equiv \rho_2$ if and only if $\sigma\rho_1 = \sigma\rho_2$ for all "sufficiently large" σ, in a sense made precise by part (d) of Lemma 10.) We define an operation called *complete cycle pushing* which takes an arbitrary rotor configuration ρ as input and produces an acyclic rotor configuration $\rho^\dagger$ as output.

Theorem 5.

(i) Each equivalence class of rotor configurations contains a unique acyclic configuration.

(ii) The unique acyclic configuration equivalent to ρ is $e\rho$, where e is the recurrent identity element of the sandpile group.

(iii) $e\rho = \rho^\dagger$ is the result of performing complete cycle pushing on ρ.

To see the relevance of this notion of equivalence to Theorem 1, let ρ_n be the rotor configuration immediately after the rotor walk hits the target set T for the n-th time. The sequence $\{\rho_n\}_{n\geq 0}$ is not periodic, but we will show that the sequence of equivalence classes $[\rho_0], [\rho_1], [\rho_2], \ldots$ is periodic. We then show that which target is hit by rotor walk starting at s with rotor configuration ρ depends only on the equivalence class $[\rho]$.

In the proof of Theorem 2, a helpful trick is the use of *antiparticles* that behave like the "holes" considered in [10]: while a particle at vertex v first increments (progresses) the rotor at v and then moves to a neighbor according to the updated rotor, an antiparticle at v first moves to a neighbor according to the *current* rotor at v and then decrements (regresses) the rotor at v. Reversing the rotor mechanism at each vertex is equivalent to replacing all particles by antiparticles and vice versa.

Related Work

Rotor walk was first studied in computer science from the point of view of autonomous agents patrolling a territory [17], and in statistical physics as a model of self-organized criticality [16]. It is an example of a "convergent game" of the type studied by Eriksson [9] and more generally of an *abelian network* of communicating finite automata. Abelian networks were proposed by Dhar [7], and their theory is developed in [3].

Rotor walk on G reflects certain features of random walk on G [5]. Recalling that $d(v,w)$ denotes the number of arcs from v to w, so that $\sum_{w \in V} d(v,w)/d(v) = 1$, consider the Markov chain on state space V in which the transition probability from v to w equals $d(v,w)/d(v)$. The frequency p_i with which a particular target vertex t occurs in the hitting sequence for rotor walk equals the probability that the Markov chain when started from the source s reaches t before it reaches any other target vertex. A main theme of [12] is that the "global" discrepancy between np_i and the number of times the rotor walk hits the target t in the first n runs is bounded — independently of n — by a sum of "local discrepancies" associated with the rotors. Note also that the connectivity condition (1) has a natural interpretation in terms of random walk: it says that with probability 1, random walk in G started from s will eventually hit the target set T.

A special case of the periodicity phenomenon was noted by Angel and Holroyd. If G is the b-regular tree of height h and T is the set of leaves, it follows from the proof of Theorem 1.1 of [14] that the hitting sequence from the root is eventually periodic with period $\#T$, and its fundamental period is a permutation of T. In Proposition 22 of [1] Angel and Holroyd prove that for any initial setting of the rotors, the first $\#T$ terms of the hitting sequence are in fact a permutation of T.

2 Abelian property, monoid action and group action

This section collects the results from the literature that we will use. Most of these can be found in the survey [11], and many date from considerably earlier; where we know of an earlier reference, we indicate that as well. The idea behind Lemma 8 is well known, but we have not seen it stated anywhere in exactly this form, so we include the proof.

Let G be a finite directed graph with source vertex s and target set T satisfying (1) and (2). In §1, we defined a rotor walk in G as an infinite sequence of vertices $x_0, x_1, x_2, \ldots$ in which the i-th occurrence of v is followed immediately by an occurrence of v^i. The proofs make use of an alternative, "stack-based" picture of rotor walk, which we now describe. This viewpoint goes back at least to [8, 18].

At each vertex v is a bi-infinite stack of cards, in which each card is labeled by an arc of G emanating from v. The i-th card is labeled by the arc $e_i = (v, v^i)$. For $i > 0$, the i-th card in the stack represents an instruction for where the particle should step upon visiting vertex v for the i-th time. (When $i < 0$, the i-th card in the stack never gets used, but it is helpful to pretend that it was used in the past before the rotor walk began; this point of view will play an important role in the proof of Theorem 3.) We also have a pointer at v that keeps track of how many departures from v have already occurred; this

pointer moves as time passes. When i departures from v have occurred during the rotor walk thus far, we represent the state of the stack and pointer as

$$[\ldots, e_{i-2}, e_{i-1}, e_i | e_{i+1}, e_{i+2}, e_{i+3}, \ldots]$$

(i is 0 at the start of the rotor walk). Arcs e_j with $j \leq i$ to the left of the pointer constitute the "past" of the rotor (arcs previously traversed), while the e_j's with $j > i$ constitute the "future" of the rotor (arcs to be traversed on future visits to v). The arc e_i is called the *retrospective state* of the rotor. It represents the most recent arc traversed from v. Arc e_{i+1} is called the *prospective state* of the rotor. It represents the next arc to be traversed from v. When the particle next exits v (along arc e_{i+1}) the pointer moves to the right and the stack at v becomes

$$[\ldots, e_{i-2}, e_{i-1}, e_i, e_{i+1} | e_{i+2}, e_{i+3}, \ldots].$$

The defining property of rotor walk is that for each vertex v, the sequence of labels in its stack is periodic. Initially, however, we will not need this assumption. We use the term *stack walk* to describe the more general situation when the stack at each vertex v may be an arbitrary sequence of arcs emanating from v.

The connectivity conditions (1) and (2) are global conditions on G. In fact, they are the only non-local ingredient needed for our local-global principles. The next lemma provides a simple example of how connectivity parlays a local property — one that can be checked for each stack individually — into a corresponding global property of the hitting sequence.

A sequence $a_1, a_2, \ldots$ whose terms belong to an alphabet A is called *infinitive* if for every $a \in A$ there are infinitely many indices i such that $a_i = a$ [13]. Thus, we say that the stack at vertex v is infinitive if every outgoing arc from v appears infinitely often as a label. Likewise, the hitting sequence is infinitive if the walk hits every target t infinitely often.

Lemma 6. *If all stacks are infinitive, then the hitting sequence of the stack walk is infinitive.*

Proof. Since the graph is finite, the stack walk visits at least one vertex infinitely often. If the walk visits v infinitely often, then since the stack at v is infinitive, the walk traverses every outgoing arc from v infinitely often, so it visits all of the out-neighbors of v infinitely often. By the strong connectivity of G', every vertex is reachable by a directed path of arcs from v, so every vertex is visited infinitely often. In particular, the walk hits every target infinitely often. □

2.1 Abelian property

Suppose that several indistinguishable particles are present on vertices of G. At each moment, one has a choice of which particle to move; one chooses a particle, shifts the

pointer in the stack at the corresponding vertex, and advances that particle to a neighboring vertex according to the instruction on the card that the pointer just passed. We call this procedure a *firing*.

For example, if we begin with m particles at the source vertex s, we can repeatedly advance one of them until it hits a target, then repeatedly advance another particle until it too hits a target, and so on, until all the particles have hit (and remain at) targets.

The following lemma is known as the *abelian property* of rotor-routing (another name for it is the "strong convergence property," following Eriksson [9]; yet another term for it is "confluence"). For a proof, see [8, Theorem 4.1] or [11, Lemma 3.9].

Lemma 7. *Starting from particle configuration σ and rotor configuration ρ, let $v_1, \ldots, v_m$ be a sequence of firings that results in all particles reaching the target set. Let $N(t)$ be the number of particles that hit target t. The numbers $N(t)$ ($t \in T$) and the final rotor configuration depend only on σ and ρ; in particular, they do not depend on the sequence $v_1, \ldots, v_m$.*

The abelian property is all that is needed to prove Theorem 4, which says that if every rotor mechanism is m-repetitive, then the hitting sequence is m-repetitive.

Proof of Theorem 4. It suffices to show for all n that if we feed mn particles through the system in succession (starting them at s and stopping them when they hit T), then the number of particles that hit each target is a multiple of m; for, if we know this fact for both mn and $m(n+1)$, then it follows that the $(mn+1)$-st through $(mn+m)$-th particles must all hit the same target. By the abelian property (Lemma 7), if we let the mn particles walk in tandem, letting each particle take its i-th step before any particle takes its $(i+1)$-st step, then since each stack is m-repetitive, the particles travel in groups of size m, such that the particles in each group travel the same path and hit the same target. □

Note that Theorem 4 did not require the stacks to be periodic. Theorems 1–3 certainly do require periodic stacks. For the rest of the paper, we assume that the stack at vertex v has period equal to its outdegree $d(v)$, and that its fundamental period is a permutation of the outgoing edges from v. (Recall that multiple edges are permitted in G, so any periodic rotor mechanism can be realized in this way.)

2.2 Action of particle configurations on rotor configurations

Let $V_0 = V - T$ be the set of non-target vertices. Denote by Q the set of particle configurations

$$Q = \{\sigma : V_0 \to \mathbb{N}\}$$

and by R the set of rotor configurations

$$R = \{\rho : V_0 \to E \mid \mathtt{s}(\rho(v)) = v \text{ for all } v \in V_0\}$$

where $\mathtt{s}(e)$ denotes the source of the arc e. We give Q the structure of a commutative monoid under pointwise addition.

Next we recall from [11, Definition 3.11] the construction of the action

$$Q \times R \to R.$$

Associated to each vertex $v \in V_0$ is a *particle addition operator* E_v acting on the set of rotor configurations: given a rotor configuration ρ, we define $E_v(\rho)$ as the rotor configuration obtained from ρ by adding a particle at v and letting it perform rotor walk until it arrives at a target. Lemma 7 implies that the operators E_v commute: $E_v E_w = E_w E_v$ for all $v, w \in V_0$.

Now given a particle configuration σ on G, we define

$$E_\sigma = \prod_{v \in V_0} (E_v)^{\sigma(v)}$$

where the product denotes composition. Since the operators E_v commute, the order of composition is immaterial. The action of particle configuration σ on rotor configuration ρ is defined by $\sigma\rho := E_\sigma(\rho)$. In words, $\sigma\rho$ is the rotor configuration obtained from ρ by placing $\sigma(v)$ particles at each vertex v and letting all particles perform rotor walk until they hit the target set T. By Lemma 7, the order in which the walks are performed has no effect on the outcome. The fact that the operators E_v commute ensures that we have a well-defined action, that is, $(\sigma_1 + \sigma_2)\rho = \sigma_1(\sigma_2\rho)$.

2.3 The sandpile monoid and its action on rotor configurations

Next we show that certain ways of rearranging the particles (called *topplings* in the sandpile model) do not change the action on rotor configurations; this observation goes back to [16]. As a consequence, the action of the infinite monoid Q factors through a finite monoid Q° called the *sandpile monoid* of G (Lemma 8).

Recall that $d(v, w)$ is the number of directed arcs from v to w in G, and that $d(v) = \sum_{w \in V} d(v, w)$. A particle configuration σ is called *stable* if

$$\sigma(v) \leq d(v) - 1 \qquad \text{for all } \ v \in V_0.$$

If σ is not stable, we can *stabilize* it by repeatedly *toppling* unstable vertices: Set $\sigma_0 = \sigma$, choose a vertex $v_0 \in V_0$ such that $\sigma_0(v_0) \geq d(v_0)$ and topple it by sending one particle along each outgoing arc from v_0. The resulting configuration σ_1 is given by

$$\sigma_1(w) = \begin{cases} \sigma_0(w) + d(v_0, w) & \text{if } w \neq v_0, \\ \sigma_0(w) + d(v_0, w) - d(v_0) & \text{if } w = v_0. \end{cases}$$

If σ_1 is not stable, choose a vertex v_1 such that $\sigma_1(v_1) \geq d(v_1)$ and topple it in the same way to arrive at a new configuration σ_2. Conditions (1) and (2) ensure that after finitely many topplings we reach a stable configuration, which is called the *stabilization* of σ and denoted σ°. The stabilization σ° does not depend on the order of topplings [6].

Let Q° be the set of stable particle configurations. We give Q° the structure of a commutative monoid with the operation

$$(\sigma_1, \sigma_2) \mapsto (\sigma_1 + \sigma_2)^\circ.$$

That is, we sum the configurations pointwise, and then stabilize. By comparing two different toppling orders to stabilize $\sigma_1 + \sigma_2 + \sigma_3$, we see that $((\sigma_1 + \sigma_2)^\circ + \sigma_3)^\circ = (\sigma_1 + (\sigma_2 + \sigma_3)^\circ)^\circ$, which shows that this operation is associative. The monoid Q° is called the *sandpile monoid* of G; its structure has been investigated in [2, 4].

Lemma 8. *For any particle configuration σ and any rotor configuration ρ we have*

$$\sigma^\circ \rho = \sigma \rho.$$

Hence, the action of particle configurations on rotor configurations descends to an action of the sandpile monoid

$$Q^\circ \times R \to R.$$

Proof. (cf. the proof of [11, Lemma 3.12], which uses a similar trick) We compute $\sigma\rho$ by grouping the initial rotor moves into "batches" each consisting of $d(v)$ moves from a vertex v. The net effect of a batch of rotor moves is the same as that of a toppling at v: the rotor at v makes a full turn, so the rotor configuration is unchanged, and one particle is sent along each arc emanating from v. After finitely many batches, we arrive at particle configuration σ° with rotors still configured as ρ. Now letting each remaining particle perform rotor walk until reaching the target set yields the rotor configuration $\sigma^\circ\rho$. □

We may express Lemma 8 as a commutative diagram

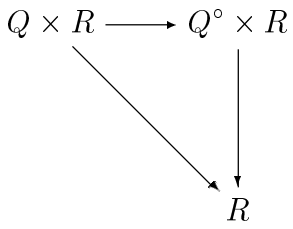

where the top arrow is $(\sigma, \rho) \mapsto (\sigma^\circ, \rho)$.

2.4 The sandpile group and its action on spanning forests

We say that a stable particle configuration $\tau \in Q^\circ$ is *reachable* from a particle configuration σ if there exists a particle configuration τ' such that $\tau = (\tau' + \sigma)^\circ$. We say that τ is *recurrent* if it is reachable from any $\sigma \in Q$.

Note that if τ is recurrent, then for any $\sigma \in Q$ there exists τ_1 such that $\tau_1(v) \geq \sigma(v)$ for all v and $\tau_1^\circ = \tau$; indeed, since τ is reachable from σ, there exists $\tau' \in Q$ with $\tau = (\tau' + \sigma)^\circ$, and we can take $\tau_1 = \tau' + \sigma$.

Denote by $S(G/T)$ the set of recurrent particle configurations. If τ is recurrent and σ is any particle configuration, then $(\sigma + \tau)^\circ$ is recurrent. That is, the set $S(G/T)$ is an

ideal of the monoid Q°. In fact, $S(G/T)$ is the minimal ideal of Q°, which shows that it is an abelian group [2]. This group is called the *sandpile group* of G relative to the target set T. The set T plays the role of the sink vertex in [11]. In the terminology of that paper, $S(G/T)$ is the sandpile group of the graph G/T obtained by collapsing T to a single vertex.

A rotor configuration ρ is *acyclic* if the graph $(V, \rho(V_0))$ contains no oriented cycles (where $\rho(V_0) = \{\rho(v) : v \in V_0\}$). Equivalently, the rotors $\{\rho(v)\}_{v \in V_0}$ form an oriented spanning forest of G rooted at T.

Lemma 9. *[11, Lemmas 3.10 and 3.12] Each addition operator E_σ acts as a permutation on the set R_0 of acyclic rotor configurations. Thus the action of the sandpile monoid Q° on rotor configurations restricts to an action*

$$S(G/T) \times R_0 \to R_0$$

of the sandpile group $S(G/T)$ on acyclic rotor configurations.

A further result proved in [11, Lemmas 3.13 and 3.17] is that this group action is free and transitive. In other words, if ρ and ρ' are two spanning forests of G rooted at T, then there is a unique element of the sandpile group $\sigma \in S(G/T)$ such that $\sigma\rho = \rho'$. In particular, the order $|S(G/T)|$ of the sandpile group equals the number of acyclic rotor configurations $|R_0|$, which is the number of spanning forests of G rooted at T. We will not use these facts, however, except for a brief aside (Lemma 20) where we identify the period of a sequence that arises in the proof of Theorem 1; there we use the freeness of the action.

The identity element $e \in S(G/T)$ is a highly nontrivial object (see for instance [11, Figures 4–6]) and plays a role in several of our lemmas below. If G has an oriented cycle, then e is distinct from the identity element $\mathbf{0}$ of Q° because the latter is not recurrent.

3 Equivalence and cycle pushing

In this section we develop a new notion of equivalence of rotor configurations and use it to prove Theorem 1. Here and throughout the rest of the article, we assume that the stack at each vertex v is periodic with period $d(v)$. For $e = e_v^i$ (the ith arc in the rotor mechanism at v) we define $e^+ = e_v^{i+1}$ and $e^- = e_v^{i-1}$, where $i + 1$ and $i - 1$ are to be interpreted modulo $d(v)$.

3.1 Equivalence of rotor configurations

Lemma 10. *Let ρ and ρ' be rotor configurations on G. The following are equivalent:*

(a) *$\sigma\rho = \sigma\rho'$ for some particle configuration σ.*

(b) *$\sigma\rho = \sigma\rho'$ for all recurrent configurations $\sigma \in S(G/T)$.*

(c) $e\rho = e\rho'$, *where* e *is the identity element of* $S(G/T)$.

(d) $\sigma\rho = \sigma\rho'$ *for all configurations* $\sigma \geq e$.

Proof. It suffices to show that (a) $\Rightarrow$ (b) and (c) $\Rightarrow$ (d), since (b) $\Rightarrow$ (c) and (d) $\Rightarrow$ (a) trivially.
(a) $\Rightarrow$ (b): Suppose that $\sigma\rho = \sigma\rho'$ for some particle configuration σ, and let $\tau \in S(G/T)$ be a recurrent configuration. Since τ is recurrent, there exists a particle configuration τ_1 such that $\tau_1^\circ = \tau$ and $\tau_1 \geq \sigma$. Writing $\tau_1 = \sigma + \sigma_1$ for some $\sigma_1 \geq 0$, we obtain

$$\begin{aligned}\tau\rho = \tau_1^\circ\rho = \tau_1\rho = \sigma_1(\sigma\rho) \\ = \sigma_1(\sigma\rho') = \tau_1\rho' = \tau_1^\circ\rho' = \tau\rho'.\end{aligned}$$

(c) $\Rightarrow$ (d): If $\sigma \geq e$, then writing $\sigma = e + \tau$ we have $\sigma\rho = \tau(e\rho) = \tau(e\rho') = \sigma\rho'$. □

Definition. Rotor configurations ρ and ρ' are *equivalent*, denoted $\rho \equiv \rho$, if the four equivalent conditions of Lemma 10 hold.

From condition (c) of Lemma 10 it is immediate that $\equiv$ is an equivalence relation. We write the equivalence class of ρ as $[\rho]$.

Lemma 11. *If* $\rho \equiv \rho'$ *then* $\tau\rho \equiv \tau\rho'$ *for all particle configurations* τ.

Proof. If $\rho \equiv \rho'$, then there exists σ such that $\sigma\rho = \sigma\rho'$, which implies $\sigma(\tau\rho) = \tau(\sigma\rho) = \tau(\sigma\rho') = \sigma(\tau\rho')$, which implies $\tau\rho \equiv \tau\rho'$. □

We say that a rotor configuration ρ is *reachable* from ρ' if there exists a particle configuration $\sigma \neq \mathbf{0}$ such that $\rho = \sigma\rho'$. We say that a rotor configuration ρ is *recurrent* if it is reachable from itself. The following lemma encapsulates the remaining results of [11] that we will need.

Lemma 12. *The following properties of a rotor configuration* ρ *are equivalent:*

(a) ρ *is recurrent.*

(b) ρ *is acyclic.*

(c) $\rho = e\rho$.

Proof. The equivalence of (a) and (b) is Lemma 3.15 of [11]. The implication (b)$\Rightarrow$(c) follows from the well-definedness of the action of $S(G/T)$ on acyclic rotor configurations (Lemma 9). To see that (c)$\Rightarrow$(a), note that $e\rho = (ee)\rho = e(e\rho)$ is reachable from itself, hence recurrent. □

See [11, Lemma 3.16] for several other conditions equivalent to being recurrent.

Lemma 13. *Each equivalence class of rotor configurations contains exactly one that is recurrent. The unique recurrent configuration equivalent to* ρ *is* $e\rho$, *where* e *is the identity element of* $S(G/T)$.

Proof. Let ρ be any rotor configuration. Then

$$e(e\rho) = (e^2)\rho = e\rho.$$

Hence $e\rho$ is recurrent by Lemma 12 and $e\rho \equiv \rho$. So each equivalence class contains at least one recurrent configuration.

For the reverse direction, suppose that ρ and ρ' are both recurrent and that $\rho \equiv \rho'$. By Lemma 12 we have $\rho = e\rho = e\rho' = \rho'$. □

As a consequence of Lemmas 9, 11, and 13, we have

Corollary 14. *Each addition operator E_σ acts as a permutation on the set $R/\equiv$ of equivalence classes of rotor configurations. Thus the action of the sandpile monoid Q° on rotor configurations projects to an action*

$$S(G/T) \times R/\equiv \;\to\; R/\equiv$$

of the sandpile group $S(G/T)$ on equivalence classes of rotor configurations.

3.2 Cycle pushing

Lemma 13 gives one way to compute the unique recurrent rotor configuration equivalent to ρ: first compute the identity element e of the sandpile group of G/T, then add $e(v)$ particles at each vertex $v \in V_0$ and stabilize. Note however that this is rather inefficient; for instance, in the case where ρ is already acyclic, a smart algorithm would recognize this fact and simply output ρ directly. We now describe a more efficient way to compute $e\rho$. The idea is to convert ρ into an acyclic configuration by successively removing cycles in the rotors. We call this process *complete cycle pushing*.

The stack picture underlying cycle popping and cycle pushing is integral to Wilson's work on random stacks [18], and cycle pushing in particular is a key idea in recent work on fast simulation of rotor-routing [10]. Suppose the rotor configuration ρ contains a cycle $\mathcal{C}$ with vertices $v_0, v_1, v_2, \ldots, v_r = v_0$; that is, for all $0 \le j \le r-1$, the arc $\rho(v_j)$ points from v_j to v_{j+1}.

The rotor configuration $\mathcal{C}\rho$ obtained by *pushing* $\mathcal{C}$ is given by

$$\mathcal{C}\rho(v) = \begin{cases} \rho(v)^- & \text{if } v \in \mathcal{C}, \\ \rho(v) & \text{otherwise.} \end{cases}$$

In other words, for each $j = 0, 1, \ldots, r-1$ the rotor $\rho(v_j)$ is regressed, and the other rotors remain unchanged. (Compare this with popping a cycle, wherein rotors are progressed, not regressed.)

Suppose we have a sequence of rotor configurations $\rho_0, \rho_1, \ldots, \rho_m$ where for each $i < m$ the configuration ρ_{i+1} is obtained from ρ_i by pushing a cycle $\mathcal{C}_i$ in ρ_i, and suppose moreover that ρ_m is acyclic. We say that ρ_m is obtained from ρ_0 by *complete cycle pushing*.

Complete cycle pushing involves a choice of ordering in which to push the cycles $\mathcal{C}_i$. Wilson [18] showed that these choices do not affect the outcome: if ρ, ρ' are acyclic

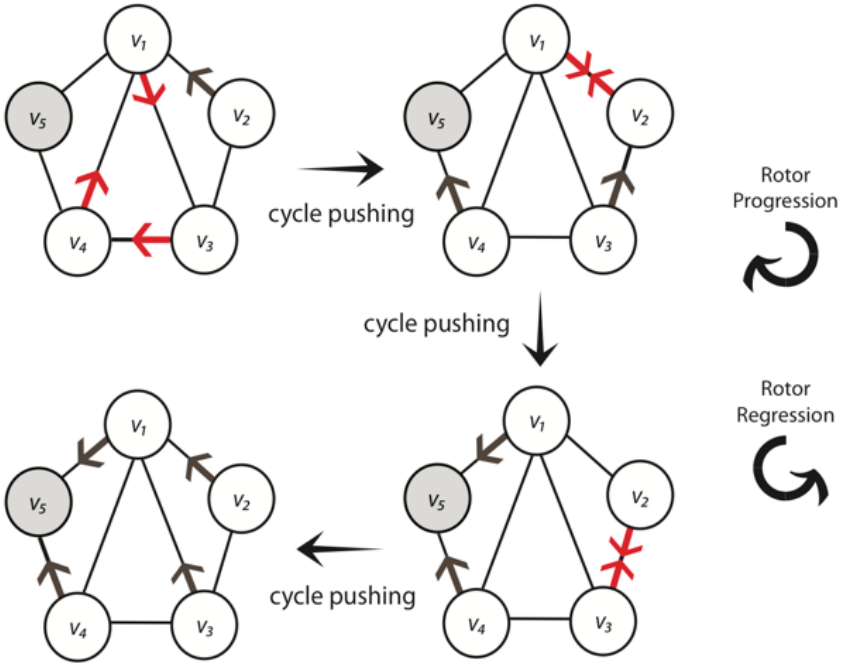

Figure 1: Example of complete cycle pushing, starting from an arbitrary rotor configuration (upper left) to obtain an acyclic rotor configuration (lower left). Here the target set is $T = \{v_5\}$. At each step, each rotor participating in a cycle (drawn in red) is regressed counterclockwise, until there are no more cycles.

configurations that can be obtained from ρ_0 by complete cycle pushing, then $\rho = \rho'$. We will not use the uniqueness in our proofs: in fact, Lemma 17 below gives another proof of Wilson's result.

Figure 1 shows how a rotor configuration is affected by cycle pushing. The shaded vertex v_5 is the target vertex, and the cycles that are pushed (first the 3-cycle $v_1 \to v_3 \to v_4 \to v_1$, then the 2-cycle $v_1 \to v_2 \to v_1$, and then the 2-cycle $v_2 \to v_3 \to v_2$) are shown in red.

Lemma 15. *If ρ' is obtained from ρ by cycle pushing, then $\rho' \equiv \rho$.*

Proof. Let $\rho' = \mathcal{C}\rho$, and let $\sigma = 1_{\mathcal{C}}$ be the particle configuration consisting of one particle at each vertex v_j of the cycle. We claim that $\sigma\rho' = \sigma\rho$. Starting from ρ', let each particle take a single step of rotor walk: for each j, the particle at v_j moves to v_{j+1} (taking indices mod r), and the rotor at v_j progresses to $\rho'(v_j)^+ = \rho(v_j)$. Since each vertex v_j on the cycle sends one particle to v_{j+1} and receives one particle from v_{j-1}, the resulting particle configuration is still σ; on the other hand, the rotor configuration has changed from ρ' to ρ. By the abelian property we conclude that $\sigma\rho' = \sigma\rho$, and hence $\rho' \equiv \rho$. □

Lemma 16. *For any initial rotor configuration, any sequence of cycle pushing moves yields an acyclic configuration in finitely many steps.*

Proof. Recall that target vertices do not have rotors. Hence if a vertex w has an arc to a target vertex t, then w can participate in only a finite number of cycle pushing moves, because at some point the rotor at w would point to t, and thereafter w cannot belong to a pushable cycle. Thereafter, each vertex v that has an arc to w can participate in only a finite number of cycle pushing moves, because at some point the rotor at v would point to w, and thereafter v cannot belong to a pushable cycle. Continuing in this fashion, using the fact (1)–(2) that every vertex v has a directed path to some target vertex, we see that every vertex can participate in only finitely many cycle pushing moves. □

Lemma 17. *Let ρ be a rotor configuration. Any sequence of cycle pushing moves that starts from ρ must terminate with $e\rho$, the unique acyclic rotor configuration equivalent to ρ.*

Proof. By Lemma 16, any sequence of cycle pushing moves starting from ρ must terminate in an acyclic configuration ρ'. By Lemma 15 we have $\rho' \equiv \rho$. Since ρ' is acyclic, ρ' is recurrent by Lemma 12, and hence $\rho' = e\rho$ by Lemma 13. □

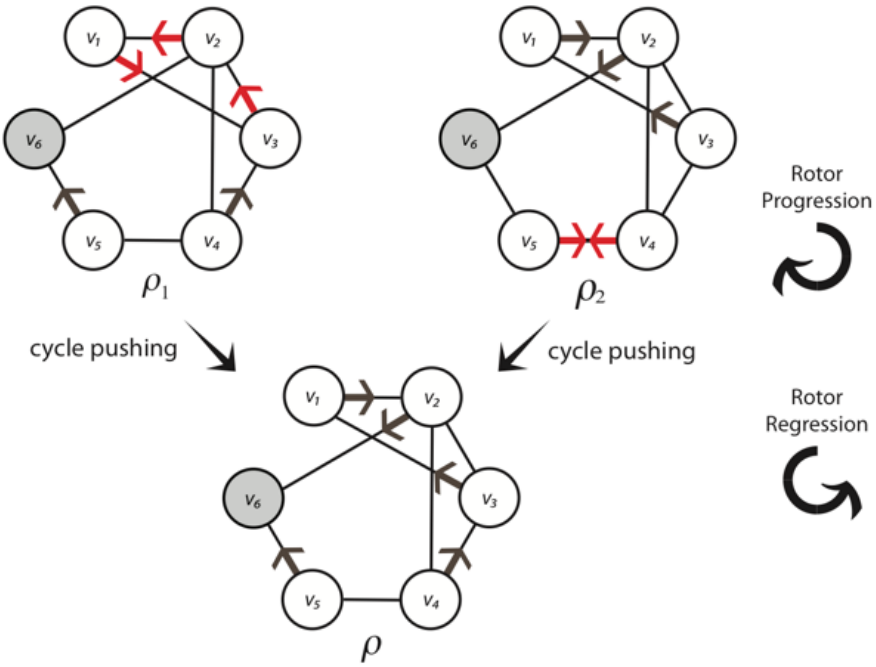

Figure 2: The rotor configurations ρ_1 at top left and ρ_2 at top right yield the same acyclic configuration ρ after complete cycle pushing, so they are equivalent by Lemma 18.

The next lemma shows that equivalence between rotor configurations is the reflexive-symmetric-transitive closure of the relation $\rho \sim \mathcal{C}\rho$ given by cycle pushing.

Lemma 18. *$\rho_1 \equiv \rho_2$ if and only if there exists a rotor configuration that is accessible from both ρ_1 and ρ_2 by a sequence of cycle pushing moves.*

Proof. If $\rho_1 \equiv \rho_2$ then $e\rho_1 = e\rho_2$ is accessible from both ρ_1 and ρ_2 by Lemma 17. Conversely, if ρ' is a configuration accessible from both ρ_1 and ρ_2, then $\rho_1 \equiv \rho' \equiv \rho_2$ by Lemma 15. □

For a pictorial example, see Figure 2. The rotor configuration ρ at the bottom is acyclic, and the other two rotor configurations lead to ρ after a single cycle pushing move (in one case, the 3-cycle $v_1 \to v_3 \to v_2 \to v_1$ is pushed, and in the other case, the 2-cycle $v_4 \to v_5 \to v_4$ is pushed). As in Figure 1, rotors progress by turning clockwise and regress by turning counterclockwise. Lemma 18 tells us that the two non-acyclic rotor configurations ρ_1 and ρ_2 must be equivalent, and indeed the reader can check that condition (a) of Lemma 10 is satisfied if one takes σ to be the particle configuration with a single particle at v_4; that is, if we add a single particle at v_4 and let it perform rotor walk until reaching the target vertex v_6, then the two rotor configurations become the same.

Denote by $t_v(\rho)$ the target vertex reached by a particle started at v if the initial rotor configuration is ρ.

Lemma 19. *If $\rho_1 \equiv \rho_2$, then $t_v(\rho_1) = t_v(\rho_2)$ for all $v \in V_0$.*

Proof. By Lemma 18 it suffices to consider the case where ρ_2 is obtained from ρ_1 by pushing a cycle $v_0, v_1, \ldots, v_r = v_0$. If the particle added to ρ_1 at v never hits the cycle, then the particle added to ρ_2 at v will traverse the exact same path, arriving at the same target. On the other hand, suppose the particle added to ρ_1 at v hits the cycle, say at v_0. Then the particle added to ρ_2 at v will take the same walk to v_0 and will then traverse the cycle, arriving back at v_0. At this point the rotor configuration will be the same as the rotor configuration for the first situation (i.e., starting from ρ_1) when the particle first hits v_0. Thereafter, the two processes evolve identically, since in both situations the particle is at v_0 and the rotor configurations at this stage are the same in both evolutions. In particular, the particle will end up at the same target vertex. □

3.3 Proof of the periodicity theorem

We can now prove our first main result, that the hitting sequence associated with a (periodic) rotor mechanism is periodic.

Proof of Theorem 1. Let $t_1, t_2, \ldots$ be the hitting sequence for initial rotor configuration ρ_0, and for $n \geq 1$ let ρ_n be the rotor configuration after n particles released from the source vertex s have hit the targets $t_1, \ldots, t_n$ (staying put after each hit). Then $\rho_n = E_s\rho_{n-1}$ for $n \geq 1$. Let $[\rho_n]$ denote the equivalence class of ρ_n. Recall that E_s acts as a permutation on equivalence classes (Corollary 14), so the sequence $[\rho_0], [\rho_1], [\rho_2], \ldots$ is periodic, say with period D. Then by Lemma 19, since $[\rho_{n+D}] = [\rho_n]$ for all $n \geq 0$, we conclude that $t_{n+D} = t_n$ for all $n \geq 1$, which shows that the hitting sequence is periodic with period D. □

Next we identify the period D of the sequence $[\rho_0], [\rho_1], \ldots$ arising in the proof of Theorem 1. This in turn gives an upper bound on the period of the hitting sequence $t_1, t_2, \ldots$, namely, the latter period is a divisor of D. Denote by δ_s the particle configuration consisting of 1 particle at the source vertex, and let $g_s = (\delta_s + e)^\circ$ be the corresponding recurrent configuration.

Lemma 20. *Let D be the order of g_s in the sandpile group $S(G/T)$. The sequence of equivalence classes of rotor configurations $\{[\rho_n]\}_{n\geq 0}$ has period D. Moreover, the hitting sequence satisfies $t_{n+D} = t_n$ for all $n \geq 1$.*

Proof. For any rotor configuration ρ, since $\rho \equiv e\rho$ we have by Lemma 11

$$g_s\rho = (\delta_s + e)^\circ\rho = (\delta_s + e)\rho = E_s(e\rho) \equiv E_s\rho.$$

Since $g_s^D = e$, we obtain

$$E_s^D\rho \equiv g_s^D\rho = e\rho \equiv \rho$$

which shows that $\rho_{n+D} \equiv \rho_n$ for all $n \geq 0$. Conversely, if $\rho_{n+k} \equiv \rho_n$ for some $n \geq 0$ and $k \geq 1$, then $g_s^k\rho_n \equiv \rho_n$, which implies $g_s^k = e$ since the action of $S(G/T)$ on equivalence classes of rotor configurations is free; hence k must be divisible by D.

The fact that $t_{n+D} = t_n$ for all $n \geq 1$ follows from Lemma 19. □

4 Time reversal and antiparticles

4.1 Stack flipping

Recall the stacks picture introduced in §2. Each vertex $v \in V_0$ has a stack ρ_v, which is a bi-infinite sequence of arcs

$$\rho_v = [\ldots, e_{-2}, e_{-1}, e_0|e_1, e_2, e_3, \ldots].$$

(We abuse notation slightly by using the same letter (ρ) to denote a stack configuration and its corresponding rotor configuration.) The e_i with $i \leq 0$ constitute the "past" of the stack, the e_i with $i > 0$ constitute the "future" of the stack, e_0 is the retrospective state of the stack, and e_1 is the prospective state of the stack; the pointer "|" marks the divide between past and future. When a particle at v takes a step, the pointer shifts to the right, so that the stack at v becomes

$$[\ldots, e_{-2}, e_{-1}, e_0, e_1|e_2, e_3, \ldots]$$

and the particle travels along arc e_1.

Shifting the pointer at v to the right corresponds to progressing the rotor at v, or in stack language, popping the stack at v; correspondingly, shifting the pointer at v to the left will be called regressing the rotor or *pushing* the stack at v. When we perform cycle pushing, the pointer for the vertex v moves one place to the left for all vertices v belonging to the cycle.

We define *stack flipping* as the operation on a bi-infinite stack that exchanges past and future, turning

$$[\ldots, e_{-2}, e_{-1}, e_0 | e_1, e_2, e_3, \ldots]$$

into

$$[\ldots, e_3, e_2, e_1 | e_0, e_{-1}, e_{-2}, \ldots].$$

Given a stack configuration $\rho = (\rho_v)_{v \in V_0}$, let $\Phi(\rho)$ denote the stack configuration obtained by flipping all its stacks. Note that $\Phi(\Phi(\rho)) = \rho$.

Lemma 21. *Let ρ be a rotor configuration that has a cycle $\mathcal{C}$. Then $\mathcal{C}$ is also a cycle of $\Phi(\mathcal{C}\rho)$, and*

$$\Phi(\mathcal{C}(\Phi(\mathcal{C}\rho))) = \rho.$$

Proof. Let v be a vertex of $\mathcal{C}$. Let $\rho' = \mathcal{C}\rho$, and write the rotor stack at v as

$$\rho_v = [\ldots, e_{-2}, e_{-1}, e_0 | e_1, e_2, e_3, \ldots].$$

If we push the cycle, the stack at v becomes

$$(\mathcal{C}\rho)_v = [\ldots, e_{-2}, e_{-1} | e_0, e_1, e_2, e_3, \ldots].$$

If we then flip all the stacks, we obtain

$$(\Phi(\mathcal{C}\rho))_v = [\ldots, e_3, e_2, e_1, e_0 | e_{-1}, e_{-2}, \ldots].$$

The retrospective rotors at the vertices $v \in \mathcal{C}$ are now as they were initially in ρ, so they form the same cycle $\mathcal{C}$. Pushing that cycle yields

$$(\mathcal{C}(\Phi(\mathcal{C}\rho)))_v = [\ldots, e_3, e_2, e_1 | e_0, e_{-1}, e_{-2}, \ldots].$$

Finally, flipping the stacks once more brings us to

$$(\Phi(\mathcal{C}(\Phi(\mathcal{C}\rho))))_v = [\ldots, e_{-2}, e_{-1}, e_0 | e_1, e_2, e_3, \ldots]$$

which equals ρ_v.

Meanwhile, for those vertices v that are not part of the cycle $\mathcal{C}$, the stack at v is simply reversed twice (with no intervening cycle pushing moves to complicate things), so this stack ends up in exactly the same configuration as it was in ρ. □

Diagrammatically, writing $\rho' = \mathcal{C}\rho$, we have

$$\begin{array}{ccc} \rho & \xrightarrow{\mathcal{C}} & \rho' \\ \Phi \uparrow & & \uparrow \Phi \\ \Phi(\rho) & \xleftarrow{\mathcal{C}} & \Phi(\rho') \end{array}$$

Note the reversal of the direction of the $\mathcal{C}$ arrow.

Lemma 22. *If $\rho \equiv \rho'$, then $\Phi(\rho) \equiv \Phi(\rho')$.*

Proof. By Lemma 18, it suffices to show that if two stack configurations are related by a cycle pushing move, then their flips are related by a cycle pushing move. But that is precisely what Lemma 21 tells us. □

4.2 Antiparticles

Next we introduce antiparticles. Like particles, they move from vertex to vertex in the graph, but they interact with the stacks in a different way. Suppose that the current stack configuration at v is

$$[\ldots, e_{-2}, e_{-1}, e_0 | e_1, e_2, e_3, \ldots]$$

and that there is an antiparticle at v. An antiparticle step consists of first moving the particle along the arc e_0 and then pushing the stack at v to obtain

$$[\ldots, e_{-2}, e_{-1} | e_0, e_1, e_2, e_3, \ldots].$$

(Compare: a particle step consists of first popping the stack at v to obtain

$$[\ldots, e_{-2}, e_{-1}, e_0, e_1 | e_2, e_3, \ldots]$$

and then moving the particle along the arc e_1.)

Lemma 23. *If ρ' is obtained from ρ by moving a particle from v along arc e, then $\Phi(\rho')$ is obtained from $\Phi(\rho)$ by moving an antiparticle from v along arc e.*

Proof. Write the stack at v for the rotor configuration ρ as

$$\rho_v = [\ldots, e_{-2}, e_{-1}, e_0 | e_1, e_2, e_3, \ldots].$$

When a particle at v advances by one step, the particle moves along the arc e_1 and the stack at v becomes

$$\rho'_v = [\ldots, e_{-2}, e_{-1}, e_0, e_1 | e_2, e_3, \ldots].$$

On the other hand, the stack at v for the flipped rotor configuration $\Phi(\rho)$ is

$$\Phi(\rho)_v = [\ldots, e_3, e_2, e_1 | e_0, e_{-1}, e_{-2}, \ldots].$$

When an antiparticle at v advances by one step, the antiparticle moves along the arc e_1 and the stack at v becomes

$$[\ldots, e_3, e_2 | e_1, e_0, e_{-1}, e_{-2}, \ldots]$$

which equals $\Phi(\rho')_v$. □

Just as one defines particle addition operators E_v, one can define antiparticle addition operators E_v^- on rotor configurations: to apply E_v^-, add an antiparticle at v and let it perform rotor walk on G (using the antiparticle dynamics described above) until it arrives at a vertex in the target set T. To highlight the symmetry between particles and antiparticles we will sometimes write E_v^+ instead of E_v for particle addition operators. Note that in general, E_v^+ and E_w^- do not commute.

Write $t_v^+(\rho)$ (resp. $t_v^-(\rho)$) for the target vertex hit by a particle (resp. antiparticle) started at v if the initial rotor configuration is ρ.

Lemma 24. *For any rotor configuration ρ and any $v \in V_0$ we have $\Phi(E_v^+(\rho)) = E_v^-(\Phi(\rho))$, and $t_v^+(\rho) = t_v^-(\Phi(\rho))$.*

Proof. This follows by repeated application of Lemma 23: the sequence of vertices traveled by the particle added to ρ at v is the same as the sequence of vertices traveled by the antiparticle added to $\Phi(\rho)$ at v. □

Diagrammatically, writing $\rho' = E_v^+\rho$, we have:

$$\begin{array}{ccc} \rho & \xrightarrow{E_v^+} & \rho' \\ \Phi \updownarrow & & \updownarrow \Phi \\ \Phi(\rho) & \xrightarrow{E_v^-} & \Phi(\rho') \end{array}$$

Lemma 25. *If $\rho \equiv \rho'$, then $E_v^-\rho \equiv E_v^-\rho'$ and $t_v^-(\rho) = t_v^-(\rho')$ for all $v \in V_0$.*

Proof. We have $E_v^- = \Phi \circ E_v^+ \circ \Phi$ by Lemma 24. Moreover, Φ preserves equivalence by Lemma 22 and E_v^+ preserves equivalence by Lemma 11, so E_v^- must preserve equivalence. This proves the first statement. For the second, since $\Phi(\rho) \equiv \Phi(\rho')$, we have by Lemmas 24 and 19

$$t_v^-(\rho) = t_v^+(\Phi(\rho)) = t_v^+(\Phi(\rho')) = t_v^-(\rho'). \quad \square$$

4.3 Loop-erasure

If a path $(x_0, \ldots, x_r)$ in the directed graph G contains a cycle, i.e., a sub-path $(x_p, x_{p+1}, \ldots, x_q)$ with $x_q = x_p$, define the *first cycle* as the unique cycle with q as small as possible; we may replace the path by the shorter path $(x_0, \ldots, x_{p-1}, x_p, x_{q+1}, \ldots, x_r)$ from which the $q - p$ vertices of the first cycle have been removed. If this new path contains a cycle, we may erase the first cycle of the new path, obtaining an even shorter path. If we continue in this fashion, we eventually obtain a simple path from x_0 to x_r, called the *loop-erasure* of the original path.

The notion of loop-erasure is due to Lawler [15], who studied the loop-erasure of random walk. As is mentioned at the end of §5 of [11], there is also a connection between loop-erasure and rotor walk. Given a rotor configuration ρ and a set $S \subset V_0$, define *popping* S as the operation of popping the stack at each vertex in S to obtain the new rotor configuration

$$S^+\rho(v) = \begin{cases} \rho(v)^+ & \text{if } v \in S, \\ \rho(v) & \text{otherwise.} \end{cases}$$

(Compare to cycle pushing §3.2, in which the rotors are regressed instead of progressed.) For a rotor configuration ρ and a vertex $v \in V_0$, let Γ be the loop-erasure of the path

$x_0, \ldots, x_r$ traveled by a particle performing rotor walk starting from $x_0 = v$ until it hits the target set. Let $C_1, \ldots, C_m$ be the cycles erased to obtain Γ. For any vertex w, the number of i $(0 \le i \le r-1)$ with $x_i = w$ is equal to the number of j $(1 \le j \le m)$ for which $w \in C_j$, plus either 1 or 0 according to whether or not $w \in \Gamma$. Hence the final rotor configuration $E_v^+\rho$ can be obtained from ρ by popping the cycles $C_1, \ldots, C_m$ and the path $\gamma := \Gamma - \{x_r\}$; that is,

$$E_v^+\rho = \gamma^+ C_1^+ \ldots C_m^+ \rho. \tag{3}$$

Lemma 26. *For every rotor configuration ρ and every $v \in V_0$ we have $E_v^- E_v^+ \rho \equiv \rho$, and the path traversed by the antiparticle is the loop-erasure of the path traversed by the particle. In particular, the antiparticle hits the same target as the particle:*

$$t_v^-(E_v^+\rho) = t_v^+(\rho).$$

Proof. After the particle has been added to ρ at v, changing the rotor configuration to $E_v^+\rho$ and arriving at target $t = t_v^+(\rho)$, the retrospective rotor at each vertex v is the arc that the particle traversed the last time it left v. Hence the rotors of $E_v^+\rho$ give a simple (cycle-free) path γ from v to t, and the antiparticle will travel this path, arriving at the same target t. By (3), the rotor configuration $E_v^+\rho$ is obtained from ρ by a sequence of cycle popping moves followed by a "path-popping move" along γ. The motion of the antiparticle from v to t undoes the path-popping move, so all that survives in $E_v^- E_v^+ \rho$ are the cycle popping moves. Since cycle popping doesn't change the equivalence class of a rotor configuration (by Lemma 18), we conclude that $E_v^- E_v^+ \rho \equiv \rho$. □

Likewise, for every ρ we have $E_v^+ E_v^- \rho \equiv \rho$. Lemma 26 thus says that the products $E_v^+ E_v^-$ and $E_v^- E_v^+$ act as the identity operation on equivalence classes of rotor configurations. That is, if we view E_v^+ and E_v^- as elements of the sandpile group (acting on equivalence classes of rotor configurations), they are inverses.

4.4 Proof of the rotor-reversal theorem

Now we turn to the proof of our second main result, that reversal of the periodic pattern of the rotor mechanism at all vertices causes reversal of the periodic pattern of the hitting sequence. To save unnecessary notation in the proof, we write $E^\pm := E_s^\pm$ and $t^\pm := t_s^\pm$.

Proof of Theorem 2. As in the proof of Theorem 1, the sequence of equivalence classes, $[\rho_0], [\rho_1], [\rho_2], \ldots$ is periodic, say with period D. Now consider the hitting sequence for *antiparticles* released from the source vertex s from initial configuration ρ_0. Define $\eta_0 = \rho_0$ and $\eta_i = E^-(\eta_{i-1})$ for $i \ge 1$. We first show by induction on i that $\eta_i \equiv \rho_{D-i}$ for all $i = 0, \ldots, D$. The base case $i = 0$ is the fact that $\rho_0 \equiv \rho_D$; and for $1 \le i \le D$, if $\eta_{i-1} \equiv \rho_{D-i+1}$ then by Lemmas 25 and 26,

$$\eta_i = E^-(\eta_{i-1}) \equiv E^-(\rho_{D-i+1}) = E^- E^+(\rho_{D-i}) \equiv \rho_{D-i}$$

which completes the inductive step.

Now for $i \geq 1$, let $u_i^+ = t^+(\rho_{i-1})$ and $u_i^- = t^-(\eta_{i-1})$ be the hitting sequences for a particle (resp. antiparticle) started at s with initial rotor configuration ρ_0. Using the second statements of Lemmas 25 and 26, we have for $i = 0, \ldots, D-1$

$$u_{i+1}^- = t^-(\eta_i) = t^-(\rho_{D-i}) = t^-(E^+\rho_{D-i-1}) = t^+(\rho_{D-i-1}) = u_{D-i}^+.$$

By Lemma 24, the hitting sequence for a particle starting at s with rotor configuration $\Phi(\rho_0)$ equals the hitting sequence for an *antiparticle* starting at s with rotor configuration ρ_0, that is, the sequence $\{u_i^-\}_{i\geq 1}$. Moreover, since $\eta_D = \rho_0 \equiv \rho_D = \eta_0$, the sequence $\{u_i^-\}_{i\geq 1}$ satisfies $u_{i+D}^- = u_i^-$ for all $i \geq 1$ by Lemma 25. Hence the particle hitting sequences for $\Phi(\rho_0)$ and ρ_0 are both periodic modulo D, and reversing the first D terms of the latter hitting sequence yields the first D terms of the former. □

Acknowledgments

We thank Peter Winkler for launching this investigation with his suggestion that the palindromic period-4 rotor 1,2,2,1,... would have special properties worthy of study. This work was supported by MIT's Undergraduate Research Opportunities Program, NSF grants 0644877 and 1001905, an NSF postdoctoral research fellowship, and NSF-REU funding provided through the courtesy of Prof. Richard Stanley.

References

[1] O. Angel and A. E. Holroyd, Rotor walks on general trees, *SIAM J. Discrete Math.* **25**:423–446, 2011. arXiv:1009.4802.

[2] L. Babai and E. Toumpakari, A structure theory of the sandpile monoid for directed graphs, 2010. http://people.cs.uchicago.edu/~laci/REU10/evelin.pdf

[3] B. Bond and L. Levine, Abelian networks I. Foundations and examples, 2011. http://www.math.cornell.edu/~levine/abelian-networks-I.pdf

[4] S. Chapman, L. Garcia-Puente, R. Garcia, M. E. Malandro and K. W. Smith, Algebraic and combinatorial aspects of sandpile monoids on directed graphs. arXiv:1105.2357.

[5] J. N. Cooper and J. Spencer, Simulating a random walk with constant error, *Combin. Probab. Comput.* **15**:815–822, 2006.

[6] D. Dhar, Self-organized critical state of sandpile automaton models, *Phys. Rev. Lett.* **64**, 1613–1616, 1990.

[7] D. Dhar, Theoretical studies of self-organized criticality, *Physica A* **369**, 29–70, 2006.

[8] P. Diaconis and W. Fulton, A growth model, a game, an algebra, Lagrange inversion, and characteristic classes, *Rend. Sem. Mat. Univ. Pol. Torino* **49**(1):95–119, 1991.

[9] K. Eriksson, Strong convergence and a game of numbers, *European J. Combin.* **17**(4): 379–390, 1996.

[10] T. Friedrich and L. Levine, Fast simulation of large-scale growth models, *Random Structures & Algorithms*, to appear. arXiv:1006.1003

[11] A. E. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp and D. B. Wilson, Chip-firing and rotor-routing on directed graphs, *In and out of equilibrium 2*, 331–364, Progr. Probab. **60**, Birkhäuser, 2008. arXiv:0801.3306

[12] A. E. Holroyd and J. Propp, Rotor walks and Markov chains, in: *Algorithmic Probability and Combinatorics*, volume 520 of *Contemporary Mathematics*, 105–126, American Mathematical Society, 2010. arXiv:0904.4507.

[13] C. Kimberling, Fractal sequences and interspersions, *Ars Combin.* **45**:157–168, 1997.

[14] I. Landau and L. Levine, The rotor-router model on regular trees, J. Comb. Theory A (2009) 116: 421–433. arXiv:0705.1562

[15] G. F. Lawler, A self-avoiding random walk, *Duke Math. J.* **47**(3): 655–693, 1980.

[16] V. B. Priezzhev, D. Dhar, A. Dhar and S. Krishnamurthy. Eulerian walkers as a model of self-organized criticality. *Phys. Rev. Lett.* **77**: 5079–5082 (1996).

[17] I. A. Wagner, M. Lindenbaum and A. M. Bruckstein, Smell as a computational resource — a lesson we can learn from the ant, *Proc. ISTCS96* pp. 219-230.

[18] D. B. Wilson, Generating random spanning trees more quickly than the cover time, In *28th Annual ACM Symposium on the Theory of Computing (STOC '96)*, 296–303, 1996.