Rotor Walks and Markov Chains

Alexander E. Holroyd and James Propp

28 April 2009 (revised 6 December 2009)

Dedicated to Oded Schramm, 1961–2008

Abstract

The rotor walk is a derandomized version of the random walk on a graph. On successive visits to any given vertex, the walker is routed to each of the neighboring vertices in some fixed cyclic order, rather than to a random sequence of neighbors. The concept generalizes naturally to countable Markov chains. Subject to general conditions, we prove that many natural quantities associated with the rotor walk (including normalized hitting frequencies, hitting times and occupation frequencies) concentrate around their expected values for the random walk. Furthermore, the concentration is stronger than that associated with repeated runs of the random walk; the discrepancy is at most C/n after n runs (for an explicit constant C), rather than $c/\sqrt{n}$.

1 Introduction

Let $X_0, X_1, \ldots$ be a Markov chain on a countable set V with transition probabilities $p : V \times V \to [0,1]$ (see e.g. [20] for background). We call the elements of V **vertices**. We write $\mathbb{P}_u$ for the law of the Markov chain started at vertex u (so $\mathbb{P}_u$-a.s. we have $X_0 = u$).

Key words: rotor router; quasirandom; Markov chain; discrepancy
2000 Mathematics Subject Classifications: 82C20; 20K01; 05C25
AEH funded in part by Microsoft Research and an NSERC Grant.
JP funded in part by an NSF grant.

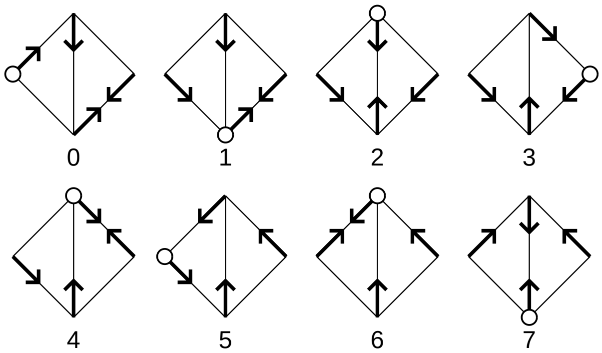

Figure 1: Steps $0, \ldots, 7$ of a rotor walk associated with the simple random walk on a graph with 4 vertices. The thin lines represent the edges of graph, the circle is the particle location, and the thick arrows are the rotors. The rotor mechanism in this case is such that each rotor successively points to the vertex's neighbors in anticlockwise order.

The **rotor-router walk** or **rotor walk** is a deterministic cellular automaton associated with the Markov chain, defined as follows. Assume that all transition probabilities $p(u, v)$ are rational (later we will address relaxation of this assumption) and that for each u there are only finitely many v such that $p(u, v) > 0$. To each vertex u we associate a positive integer $d(u)$ and a finite sequence of (not necessarily distinct) vertices $u^{(1)}, \ldots, u^{(d(u))}$, called the **successors** of u, in such a way that

$$p(u, v) = \frac{\#\{i : u^{(i)} = v\}}{d(u)} \quad \text{for all } u, v \in V. \tag{1}$$

(This is clearly possible under the given assumptions; $d(u)$ may be taken to be the lowest common denominator of the transition probabilities from u.) The set V together with the quantities $d(u)$ and the assignments of successors will sometimes be called the **rotor mechanism**.

A **rotor configuration** is a map r that assigns to each vertex v an integer $r(v) \in \{1, \ldots, d(v)\}$. (We think of an arrow or **rotor** located at each vertex, with the rotor at v pointing to vertex $v^{(r(v))}$). We let a rotor configuration evolve in time, in conjunction with the position of a particle moving from vertex to vertex: the rotor at the current location v of the particle is incremented, and the particle then moves in the new rotor direction. More formally, given a rotor mechanism, an initial particle location $x_0 \in V$ and an initial rotor configuration r_0, the **rotor walk** is a sequence of vertices $x_0, x_1, \ldots \in V$ (called **particle locations**) together with rotor configurations $r_0, r_1, \ldots$, constructed inductively as follows. Given x_t and r_t at time t we set:

(i) $r_{t+1}(v) := \begin{cases} (r_t(v)+1) \mod d(v), & v = x_t; \\ r_t(v), & v \neq x_t \end{cases}$
(increment the rotor at the current particle location); and

(ii) $x_{t+1} := (x_t)^{(r_{t+1}(x_t))}$
(move the particle in the new rotor direction).

See Figure 1 for a simple illustration of the mechanism.

Given a rotor walk, write

$$n_t(v) := \#\{s \in [0, t-1] : x_s = v\}$$

for the number of times the particle visits vertex v before (but not including) time t.

We next state general results, Theorems 1–4, relating basic Markov chain objects to their rotor walk analogues (under suitable conditions). We then state a more refined result (Theorem 5) for the important special case of simple random walk on $\mathbb{Z}^2$, followed by extensions to infinite times (Theorem 8) and irrational transition probabilities (Theorem 12). We postpone discussion of history and background to the end of the introduction, and proofs to the later sections.

1.1 Hitting probabilities

Let $T_v := \min\{t \geq 0 : X_t = v\}$ be the first hitting time of vertex v by the Markov chain (where $\min \emptyset := \infty$). Fix two distinct vertices b, c and consider the hitting probability

$$h(v) = h_{b,c}(v) := \mathbb{P}_v(T_b < T_c). \tag{2}$$

Note that $h(b) = 1$ and $h(c) = 0$. In order to connect hitting probabilities with rotor walks, fix a starting vertex $a \notin \{b, c\}$, and modify the transition probabilities from b and c so that $p(b, a) = p(c, a) = 1$. (Thus, after hitting b or c, the particle is returned to a.) Note that this modification does not change the function h. Modify the rotor mechanism accordingly by setting $d(b) = d(c) = 1$ and $b^{(1)} = c^{(1)} = a$. Let $x_0, x_1, \ldots$ be a rotor walk associated with the modified chain. The following is our most basic result.

Theorem 1 (Hitting probabilities). *Under the above assumptions, suppose that the quantity*

$$K_1 := 1 + \frac{1}{2} \sum_{\substack{u \in V \setminus \{b,c\}, \\ v \in V}} d(u)p(u,v)|h(u) - h(v)|$$

is finite. Then for any rotor walk and all t,

$$\left| h(a) - \frac{n_t(b)}{n_t(b) + n_t(c)} \right| \le \frac{K_1}{n_t(b) + n_t(c)}.$$

Theorem 1 implies that the proportion of times that the rotor walk hits b as opposed to c converges to the Markov chain hitting probability $h(a)$, provided the rotor walk hits $\{b, c\}$ infinitely often (we will consider cases where this does not hold in the later discussion on transfinite rotor walks). Furthermore, after n visits to $\{b, c\}$, the discrepancy in this convergence is at most K/n for a fixed constant K. In contrast, for the proportion of visits by the Markov chain itself, the discrepancy is asymptotically a random multiple of $1/\sqrt{n}$ (by the central limit theorem).

The condition $K_1 < \infty$ holds in particular whenever V is finite, as well as in many cases when it is infinite; for examples see [22].

In the case when the Markov chain (before modification) is a simple random walk on an undirected graph $G = (V, E)$ (thus, $p(u, v)$ equals $1/d(u)$ if (u, v) is an edge, and 0 otherwise, with $d(u)$ being the degree of u), we obtain the particularly simple bound $K_1 \le 1 + \sum_{(u,v)\in E} |h(u) - h(v)|$.

Theorem 1 can be easily adapted to give similar results for the probability of *returning* to b before hitting c when started at $a = b$, and for the probability of hitting one *set* of vertices before another. This can be done either by adapting the proof or by adding appropriate extra vertices and then appealing to Theorem 1. For brevity we omit such variations.

We next discuss extensions of Theorem 1 in the following directions: hitting times and stationary distributions, an example where $K_1 = \infty$, cases where the particle can escape to infinity, and irrational transition probabilities.

1.2 Hitting times

Fix a vertex b and let

$$k(v) = k_b(v) := \mathbb{E}_v T_b \tag{3}$$

be its expected hitting time. Fix also an initial vertex $a \neq b$ and modify the transition probabilities from b so that $p(b,a) = 1$. (Then $k(a)$ is also the expected return time from b to b in the reduced chain in which the vertices a and b are conflated.) Let $x_0, x_1, \dots$ be a rotor walk associated with the modified chain.

Theorem 2 (Hitting times)**.** *Under the above assumptions, suppose that V is finite, and let*

$$K_2 := \max_{v \in V} k(v) + \frac{1}{2} \sum_{\substack{u \in V \setminus \{b\}, \\ v \in V}} d(u)p(u,v)|k(u) - k(v) - 1|.$$

Then for any rotor walk and all t,

$$\left| (k(a)+1) - \frac{t}{n_t(b)} \right| \leq \frac{K_2}{n_t(b)}.$$

Thus the average time for the rotor walk to get from a to b concentrates around the expected hitting time. The "+1" term corresponds to the time step to move from b to a.

Note that, in contrast with Theorem 1, in the above result we require V to be finite. Leaving aside some degenerate cases, such a bound cannot hold when V is infinite. Indeed, if V is infinite and the Markov chain is irreducible, then $|(k(a)+1)n_t(b) - t|$ is unbounded in t, since the rotor walk has arbitrarily long excursions between successive visits to b; hence the conclusion of Theorem 2 cannot hold (for any constant K_2) in this case. In contrast, in the next result we again allow V to be infinite.

1.3 Stationary vectors

Suppose that the Markov chain is irreducible and recurrent, and let $\pi : V \to (0, \infty)$ be a stationary vector (so that $\pi p = \pi$ as a matrix product). Let $x_0, x_1, \dots$ be an associated rotor walk. Fix two vertices $b \neq c$ and let $h = h_{b,c}$ be as in (2) above. Also let $T_u^+ := \min\{t \geq 1 : X_t = u\}$ denote the first return time to u, and define the escape probability $e_{u,v} := \mathbb{P}_u(T_v < T_u^+)$.

Theorem 3 (Occupation frequencies)**.** *For any irreducible, recurrent Markov chain, with the above notation, suppose that the quantity*

$$K_3 := 1 + \frac{1}{2}\Big(d(b) + d(c) + \sum_{u,v\in V} d(u)p(u,v)|h(u) - h(v)|\Big)$$

is finite. Then for all t,

$$\Big|\frac{n_t(b)}{\pi(b)} - \frac{n_t(c)}{\pi(c)}\Big| \le \frac{K_3}{\pi(b)e_{b,c}}.$$

Thus, the ratio of times spent at different vertices by the rotor walk concentrates around the ratio of corresponding components of the stationary vector.

Now suppose that the Markov chain is irreducible and positive recurrent, and let π be the stationary *distribution* (so that $\sum_{v\in V}\pi(v) = 1$). Fix a vertex b and let $k = k_b$ be as in (3). The following result states that the proportion of time spent by the rotor walk at b concentrates around $\pi(b)$.

Theorem 4 (Stationary distribution)**.** *For an irreducible positive recurrent Markov chain with V finite, with the above notation, let*

$$K_4 := \max_{v\in V} k(v) + \frac{1}{2}\Big(\frac{d(b)}{\pi(b)} + \sum_{u,v\in V} d(u)p(u,v)|k(u) - k(v) - 1|\Big).$$

Then for all t,

$$\Big|\pi(b) - \frac{n_t(b)}{t}\Big| \le \frac{K_4\pi(b)}{t}.$$

1.4 Logarithmic discrepancy for walks on $\mathbb{Z}^2$

While Theorem 1 requires the quantity K_1 to be finite, experiments suggest that similar conclusions hold in many cases where it is infinite. We next treat one interesting example in which such a conclusion provably holds, but with an additional logarithmic factor in the bound on the discrepancy. (Additional such examples will appear in [22]).

Consider simple symmetric random walk on the square lattice $\mathbb{Z}^2$. That is, let $V = \mathbb{Z}^2$, and let $p(u,v) := 1/4$ for all $u,v \in V$ with $\|u-v\|_1 = 1$ and $p(u,v) := 0$ otherwise. Let each rotor rotate anticlockwise; that is for each $u \in V$, we set $d(u) := 4$ and

$$u^{(i)} := u + (\cos\tfrac{i\pi}{2}, \sin\tfrac{i\pi}{2}), \quad i = 1,\dots,4. \tag{4}$$

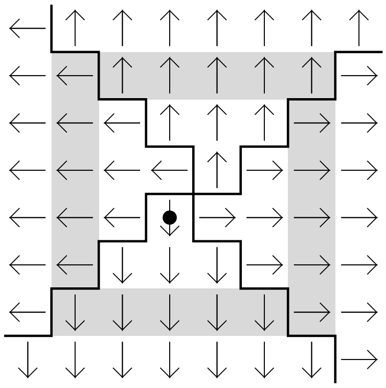

Figure 2: The initial rotor configuration in Theorem 5. The dot shows the location of $(0,0)$. The third layer is shaded (see the later proofs).

Consider the particular initial rotor configuration r given by

$$r\big((x,y)\big) := \left\lfloor \tfrac{1}{2} + \tfrac{2}{\pi}\arg(x - \tfrac{1}{2}, y - \tfrac{1}{2}) \right\rfloor \mod 4 \tag{5}$$

(where $\arg(x,y)$ denotes the angle $\theta \in [0,2\pi)$ such that $(x,y) = r(\cos\theta, \sin\theta)$ with $r > 0$). See Figure 2.

Fix vertices a, b, c of $\mathbb{Z}^2$ with $b \neq c$ and modify p by setting $p(b,a) = p(c,a) = 1$. If $a = b$ then also split this vertex into two vertices a and b, let b inherit all the incoming transition probabilities of the original random walk, and let a inherit the outgoing probabilities; similarly if $a = c$. Also modify the rotor mechanism and the rotor configuration r accordingly.

Theorem 5 (Hitting probabilities for walk on $\mathbb{Z}^2$)**.** *Let a, b, c be vertices of $\mathbb{Z}^2$ with $b \neq c$, and consider the rotor walk associated with the random walk, rotor mechanism and initial rotor configuration described above, started at vertex a. Then for any t, with $h(a) = h_{b,c}(a)$ and $n = n_t(b) + n_t(c)$,*

$$\left| h(a) - \frac{n_t(b)}{n} \right| \leq \frac{C \ln n}{n}.$$

Furthermore, $t \leq C' n^3$. Here C, C' are finite constants depending on a, b, c.

In contrast to the above result for the rotor walk, for the Markov chain itself, after n visits to $\{b, c\}$ the proportion of visits to b differs from its limit $h(a)$ by $K/\sqrt{n}$ in expected absolute value (by the central limit theorem),

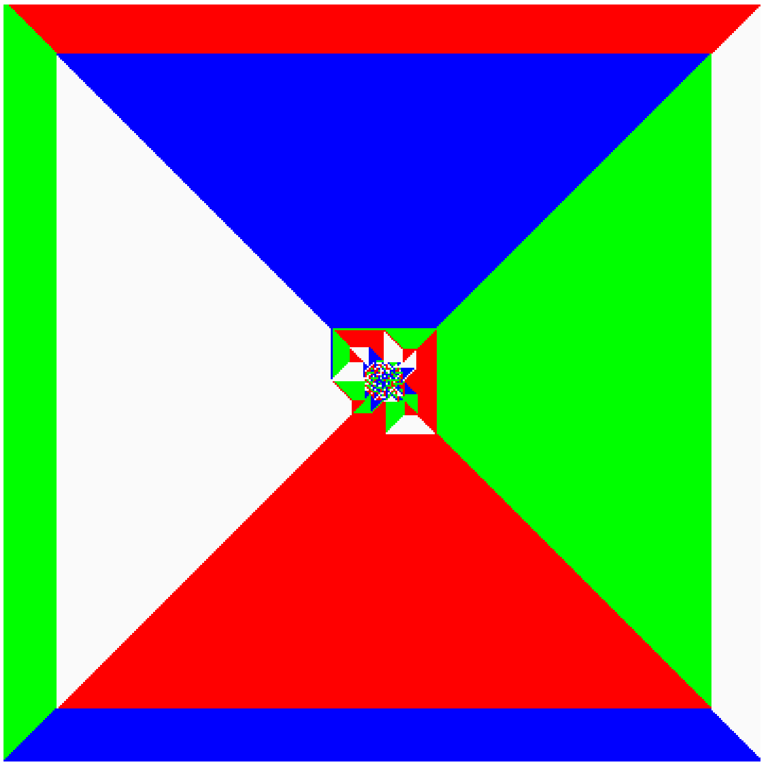

Figure 3: The rotor configuration after 500 visits to b, starting in the configuration of Figure 2 with $a = c = (0, 0)$ and $b = (1, 1)$. The rotor directions are: East=white, North=red, West=green, South=blue.

while the median number of time steps needed to achieve n visits is at least $(K')^n$ where $K > 0$ and $K' > 1$ are constants depending on a, b, c. (The latter fact is an easy consequence of the standard fact [24] that the expected number of visits to the origin of $\mathbb{Z}^2$ after t steps of random walk is $O(\ln t)$ as $t \to \infty$.)

Simulations suggest that a much tighter bound on the discrepancy should actually hold in the situation of Theorem 5, and in fact the results seem consistent with a bound of the form const$/n$. The rotor configurations at large times are very interesting; see Figure 3. (Also compare with Figure 4 below). Further discussion of these issues will appear in [22].

1.5 Transfinite walks

As mentioned above, Theorem 1 implies convergence of $n_t(b)/(n_t(b)+n_t(c))$ to $h(a)$ only if $n_t(b)+n_t(c) \to \infty$ as $t \to \infty$; we now investigate when this holds, and what can be done if it does not. We say that a rotor walk is **recurrent** if it visits every vertex infinitely often, and **transient** if it visits every vertex only finitely often.

Lemma 6 (Recurrence and transience)**.** *Any rotor walk associated with an irreducible Markov chain is either recurrent or transient.*

Note in particular that if V is finite and p is irreducible then every rotor walk is recurrent.

Fix an initial rotor configuration r_0 and an initial vertex $x_0 = a$. Suppose that the rotor walk $x_0, x_1, \ldots$ is transient. Then we can define a rotor configuration r_ω by $r_\omega(v) := \lim_{t\to\infty} r_t(v)$ (the limit exists since the sequence $r_t(v)$ is eventually constant). Now **restart** the particle at a by setting $x_\omega := a$, and define a rotor walk $x_\omega, x_{\omega+1}, x_{\omega+2}, \ldots$ according to the usual rules. If this is again transient we can set $r_{2\omega} := \lim_{t\to\infty} r_{\omega+t}$ and restart at $x_{2\omega} := a$ and so on. Continue in this way up to the first m for which the walk $x_{m\omega}, x_{m\omega+1}, \ldots$ is recurrent, or indefinitely if it is transient for all m. Call this sequence of walks a **transfinite rotor walk** started at a.

A **transfinite time** is a quantity of the form $\tau = \omega^2$, or $\tau = m\omega + t$ where m, t are non-negative integers. There is a natural order on transfinite times given by $m\omega+t < m'\omega+t'$ if and only if either $m < m'$ or both $m = m'$ and $t < t'$, while $m\omega + t < \omega^2$ for all m and t. For a transfinite walk and a transfinite time τ we write $n_\tau(v) = \#\{\alpha < \tau : x_\alpha = v\}$ for the number of visits to v before time τ. We sometimes say that the walk **goes to infinity** just before each of the times $\omega, 2\omega, \ldots$ at which it is defined.

Lemma 7 (Transfinite recurrence and transience)**.** *For an irreducible Markov chain and a transfinite rotor walk started at a, for any transfinite time τ, either $n_\tau(v)$ is finite for all v or $n_\tau(v)$ is infinite for all v. Also there exists $M \in \{1, 2, \ldots, \omega\}$ such that $n_{M\omega}(v)$ is infinite for all v and the rotor walk is defined at all $\tau < M\omega$.*

Note that while it is not obvious how to algorithmically compute transfinite rotor walks in general, it is at least possible in certain settings, such as a random walk on the integers with a periodic initial rotor configuration.

Theorem 8 (Transfinite walks and hitting probabilities)**.** *Under the assumptions of Theorem 1, suppose further that p is irreducible, and that*

$$\limsup_{v\in V} h(v) = 0.$$

Then for any transfinite time $\tau = m\omega + t$ at which all vertices have been visited only finitely often,

$$\left| h(a) - \frac{n_\tau(b)}{n_\tau(b)+n_\tau(c)+m} \right| \le \frac{K_1}{n_\tau(b)+n_\tau(c)+m}.$$

Thus the proportion of times the particle hits b as opposed to hitting c or going to infinity concentrates around $h(a)$. Furthermore, Lemma 7 ensures that $n_\tau(b)+n_\tau(c)+m \to \infty$ as $\tau \to M\omega$, so that the proportion converges to $h(a)$. The proof of Theorem 8 may be easily adapted to cover the probability of hitting a single vertex b as opposed to escaping to infinity.

Next, for a vertex b, write $g(v) = g_b(v) := \mathbb{E}_v \sum_{t=0}^\infty \mathbf{1}[X_t = b]$ for the expected total number of visits to b. Note that this is finite for an irreducible, transient Markov chain.

Theorem 9 (Transfinite walks and number of visits)**.** *Consider an irreducible, transient Markov chain and fix vertices a, b. Suppose that*

$$\limsup_{v\in V} g(v) = 0.$$

Suppose moreover that the quantity

$$K_5 := \sup_{v\in V} g(v) + \frac{1}{2}\Big(d(b) + \sum_{u,v\in V} d(u)p(u,v)\big|g(u)-g(v)\big| \Big)$$

is finite. Then for any transfinite walk started at a, and for any transfinite time $\tau = m\omega + t$ at which all vertices have been visited only finitely often,

$$\left| g(a) - \frac{n_\tau(b)}{m} \right| \le \frac{K_5}{m}.$$

It is natural to ask how recurrence and transience of rotor walks are related to recurrence and transience of the associated Markov chain. The following variant of an unpublished result of Oded Schramm provides an answer in one direction: in a certain asymptotic sense, the rotor walk is no more transient than the Markov chain. For a transfinite rotor walk started at vertex a, let I_n be the number of times the walk goes to infinity before the nth return to a (i.e. $I_n := \max\{m \geq 0 : n_{m\omega}(a) < n\}$ — this is well defined by Lemma 7; recall that the walk is restarted at a after each escape to infinity).

Theorem 10 (Transience density; Oded Schramm)**.** *Consider an irreducible Markov chain, and an associated transfinite rotor walk started at vertex a. With I_n as defined above we have*

$$\limsup_{n\to\infty} \frac{I_n}{n} \leq \mathbb{P}_a(T_a^+ = \infty).$$

In particular we note that for a recurrent Markov chain the right side in Theorem 10 is zero, so the sequence of escapes to infinity has density zero in the sequence of returns to a. On the other hand, for a recurrent Markov chain it is possible for a rotor walk to go to infinity, for example in the case of simple symmetric random walk on $\mathbb{Z}$, with all rotors initially pointing in the same direction.

Moreover, for simple random walk on $\mathbb{Z}^2$ with all rotors initially in the same direction, the rotor walk goes to infinity infinitely many times. (To check this, suppose the rotors rotate anticlockwise and initially point East. Whenever the particle's horizontal coordinate achieves a new maximum, it is immediately sent directly Northwards to infinity. This happens infinitely often by Lemma 7.) See Figure 4 for a simulation of this remarkable process, and see [22] for further discussion.

On the other hand, it should be noted that the rotor walk on $\mathbb{Z}^2$ is recurrent for the initial configuration in Theorem 5 (see Figure 2). It is also possible for the rotor walk to be recurrent for a transient Markov chain, for example in the case of simple random walk on an infinite binary tree, with all rotors arranged so as to next send the particle towards the root. Landau and Levine [14] studied the rotor walk on regular trees in great detail, in particular identifying exactly which sequences $(I_n)_{n\geq 0}$ are possible on the binary tree. Further work on rotor walks on trees will appear in [1].

1.6 Stack walks

To generalize rotor walks to Markov chains with irrational transition probabilities, we must allow the particle to be routed to a non-periodic sequence of vertices on its successive visits to a given vertex.

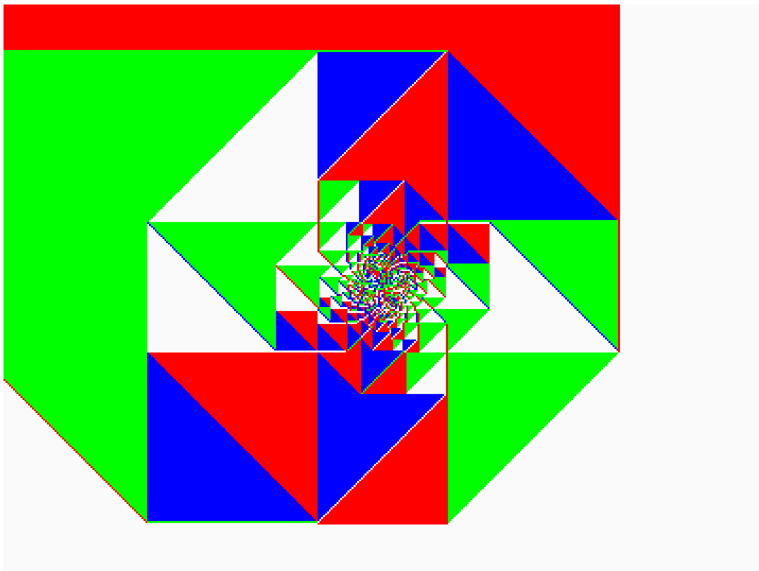

Figure 4: The rotor configuration after 500 restarts from $a = (0,0)$, for the transfinite rotor walk on $\mathbb{Z}^2$ with all rotors initially pointing East. The rotor directions are: East=white, North=red, West=green, South=blue. The red region extends infinitely far to the North.

Given a set V, a **stack mechanism** is an assignment of an infinite sequence of successors $u^{(1)}, u^{(2)}, \ldots$ to each vertex $u \in V$. The **stack walk** started at x_0 is a sequence of vertices $x_0, x_1, \ldots$ defined inductively by

$$x_{t+1} := x_t^{(n_t(x_t)+1)}$$

where

$$n_t(v) := \#\{s \in [0, t-1] : x_s = v\}.$$

(Note that, in the case of rational transition probabilities considered previously, the rotor walk can be regarded as a special case of a stack walk, with the periodic stacks given by $u^{(kd(u)+j)} = u^{(j)}$ for $1 \le j \le d(u)$ and $k \ge 0$.)

We illustrate the use of stacks with Theorem 12 below on hitting probabilities. The following will enable us to choose a suitable stack mechanism.

Proposition 11 (Low-discrepancy sequence)**.** *Let $p_1, \ldots, p_n \in (0,1]$ satisfy $\sum_i p_i = 1$. There exists a sequence $z_1, z_2, \ldots \in \{1, \ldots, n\}$ such that for all i and t,*

$$\Big| p_i t - \#\{s \le t : z_s = i\} \Big| \le 1. \tag{6}$$

Let p be a Markov transition kernel on V, and suppose that for each vertex u there are only finitely many vertices v such that $p(u,v) > 0$. We may then choose a stack mechanism according to Proposition 11. More precisely, for each vertex u, enumerate the vertices v such that $p(u,v) > 0$ as $v_1, \ldots, v_n$, and set $p_i = p(u, v_i)$. Then let $u^{(j)} := v_{z_j}$ where z is the sequence given by Proposition 11. Now let a, b, c be distinct vertices and assume that $p(b,a) = p(c,a) = 1$ and $b^{(i)} = c^{(i)} = a$ for all i. Write $h = h_{b,c}$.

Theorem 12 (Stack walks)**.** *Under the above assumptions, suppose that*

$$K_6 := 1 + \sum_{\substack{u \in V \setminus \{b,c\}, v \in V: \\ p(u,v) > 0}} |h(u) - h(v)|$$

is finite. For the stack mechanism described above, and any t,

$$\left| h(a) - \frac{n_t(b)}{n_t(b) + n_t(c)} \right| \leq \frac{K_6}{n_t(b) + n_t(c)}.$$

Proposition 11 can in fact be extended to the case of infinite probability vectors [2], and Theorem 12 carries over straightforwardly to this case. However, the result appears to have few applications in this broader context, since $\sum_v |h(u) - h(v)|$ is typically infinite when u has infinitely many successors.

1.7 Further Remarks.

History The rotor-router model was introduced by Priezzhev, Dhar, Dhar and Krishnamurthy [21] (under the name "Eulerian walkers model") in connection with self-organized criticality. It has been rediscovered several times, including in [23] as a tool in load balancing and in [9] in the analysis of combinatorial games. The present article reports the first work on the close connection between rotor walks and Markov chains, originating in discussions between the two authors at a meeting in 2003. (Such a connection was however anticipated in the "whirling tours" theorem of [9], which shows that for random walk on a tree, the expected hitting time from one vertex to another can be computed by means of a special case of rotor walk; see also [22].) A special case of results presented here was reported in [13], and earlier drafts of the current work provided partial inspiration for some of the recent progress in [3, 4, 5, 6, 8, 12, 17, 19, 18], which we discuss below.

The idea of stack walks has its roots in Wilson's approach to random walks via random stacks; see [25].

Time-dependent bounds. We have chosen to focus on upper bounds of the form K_i/n, where K_i is a fixed constant not depending on time t. If this latter requirement is relaxed, our proofs may be adapted to give bounds that are stronger in some specific cases (at the expense of less clean formulations). Specifically, in each of Theorems 1–4 and 12, the claimed bound still holds if the relevant constant K_i is replaced with a modified quantity $K_i(t)$ obtained from K_i by:

(i) multiplying the initial additive term "1" or "$\max k(v)$" or "$\sup g(v)$" by the indicator $\mathbf{1}[x_t \neq x_0]$ (so that the term vanishes when the particle returns to its starting point); and

(ii) multiplying the summand in the sum $\sum_{u,v}$ by $\mathbf{1}[r_t(u) \neq r_0(u)]$ (so in particular terms corresponding to vertices u that have not been visited by time t vanish).

The same holds for Theorems 8 and 9 in the transfinite case, but replacing t with τ.

The above claims follow by straightforward modifications to our proofs. Indeed, our proof of Theorem 5 employs a special case of this argument. These and other refinements will be discussed more fully in the forthcoming article [22].

Abelian property. The rotor-router model has a number of interesting properties that will not be used directly in most of our proofs but which are nonetheless relevant. In particular, it enjoys an "Abelian property" which allows rotor walks to be parallelized. Specifically, consider a Markov chain on a finite set V with one or more **sinks**, i.e. vertices s with $p(s,s) = 1$, and suppose that from every vertex, some sink is accessible (so that the Markov chain eventually enters a sink almost surely). Then we may run several rotor walks simultaneously as follows. Start with an initial rotor configuration, and some non-negative number of particles at each vertex. At each step, choose a particle and route it according to the usual rotor mechanism; i.e. increment the rotor at its current vertex and move the particle in the new rotor direction. Continue until all particles are at sinks. It turns out that the resulting configuration of particles and rotors is independent of the order in which we chose to route the particles. This is the Abelian property; see e.g. [12, Lemma 3.9] for a proof (and generalizations).

In the situation of Theorem 1, for example, assume that V is finite and the Markov chain is irreducible, and then modify it to make vertices b and c sinks. Start n particles at vertex a and perform simultaneous rotor walks. The Abelian property implies that the number of particles eventually at b is the same as the number $n_t(b)$ of times that b is visited when $n_t(b)+n_t(c)=n$ in the original set-up of Theorem 1, and the bound of Theorem 1 therefore applies.

A similar Abelian property holds for the "chip-firing" model introduced by Engel [10, 11] (later re-invented by Dhar [7] under the name "abelian sandpile model" as another model for self-organized criticality). The two models have other close connections, and in particular there is a natural group action involving sandpile configurations acting on rotor configurations. More details may be found in [12] and references therein. Engel's work was motivated by an analogy between Markov chains and chip-firing (indeed, he viewed chip-firing as an "abacus" for Markov chain calculations).

Periodicity. In the case when V is finite, we note the following very simple argument which gives bounds similar to Theorems 1–4 but with (typically) much worse constants. Since there are only finitely many rotor configurations, the sequence of vertices $((x_t, r_t))_{t\geq 0}$ is eventually periodic (with explicit upper bounds on the period and the time taken to become periodic which are exponentially large in the number of vertices). Therefore the proportion of time $n_t(v)/t$ spent at vertex v converges as $t \to \infty$ to some quantity $\mu(v)$, say, with a discrepancy bounded by const$/t$. Furthermore, as a consequence of the rotor mechanism, we have $\mu(u) = \sum_{v\in V} p(u,v)\mu(v)$ for all vertices u (because after many visits to u, the particle will have been routed to each successor approximately equal numbers of times). Thus μ is a stationary distribution for the Markov chain. This implies the bound in Theorem 4, except with a different (and typically much larger) constant in place of $K_4\pi(b)$. Similar arguments yield analogues of Theorems 1–3, but only in the case where V is finite.

Related work. As remarked earlier, rotor walks on trees were studied in detail by Landau and Levine [14]. Further results on rotor walks on trees will appear in a forthcoming work of Angel and Holroyd [1], and further refinements and discussions of the results presented here will appear in Propp [22].

Cooper and Spencer [6] studied the following closely related problem. For the rotor walk associated with simple symmetric random walk on $\mathbb{Z}^d$, start with n particles at the origin, or more generally distributed in any fashion on vertices $(i_1, \ldots, i_d)$ with $i_1+\cdots+i_d$ even, and apply one step of the rotor walk to each particle; repeat this t times. (It should be noted that the Abelian property does *not* apply here — the result is not the same as applying t rotor steps to each particle in an arbitrary order; see [12].) It is proved in [6] that the number of particles at a given vertex differs from the expected number of particles for n random walks by at most a constant (depending only on d). Further more precise estimates are proved in dimension $d = 1$ in [5] and in dimension $d = 2$ in [8].

The following rotor-based model for internal diffusion-limited aggregation (IDLA) was proposed by the second author, James Propp, and studied by Levine and Peres in [16, 17, 19, 18]. Starting with a rotor configuration on $\mathbb{Z}^d$, perform a sequence of rotor walks starting at the origin, stopping each walk as soon as it reaches a vertex not occupied by a previously stopped particle. It is proved in [18] that, as the number of particles n increases, the shape of the set of occupied vertices converges to a d-dimensional Euclidean ball; generalizations and more accurate bounds are proved in [17, 19].

2 Proofs of basic results

Theorems 1–4 will all follow as special cases of Proposition 13 below, and the remaining results will also follow by adapting the same proof. For any Markov transition kernel p and any function $f : V \to \mathbb{R}$ we define the **Laplacian** $\Delta f : V \to \mathbb{R}$ by

$$\Delta f(u) := \sum_{v\in V} p(u,v)f(v) - f(u). \tag{7}$$

Proposition 13 (Key bound)**.** *For any rotor walk $x_0, x_1, \ldots$ associated with p, any function f and any t we have*

$$\Big|\sum_{s=0}^{t-1} \Delta f(x_s)\Big| \le |f(x_t) - f(x_0)| + \frac{1}{2}\sum_{u,v\in V} d(u)p(u,v)\big|f(u) - f(v) + \Delta f(u)\big|.$$

The proofs of Theorems 1–4 will proceed by applying Proposition 13 to a suitable f. The proof of Proposition 13 will use the following simple fact.

Lemma 14. *If $\sum_{i=1}^n a_i = 0$ then $\big|\sum_{i=1}^j a_i - \sum_{i=1}^k a_i\big| \le \frac{1}{2}\sum_{i=1}^n |a_i|$ for all $j, k \in [1, n]$.*

Proof. We prove the stronger statement that $|\sum_{i\in S} a_i| \leq \frac{1}{2}\sum_{i=1}^n |a_i|$ for any subset S of $\{1,\ldots,n\}$: assuming without loss of generality that $\sum_{i\in S} a_i$ is positive, it is at most $\sum_{a_i:a_i>0} a_i = \frac{1}{2}(\sum_{i:a_i>0} a_i - \sum_{i:a_i<0} a_i) = \frac{1}{2}\sum_{i=1}^n |a_i|$. $\square$

Proof of Proposition 13. Recall that r_0 denotes the initial rotor configuration. For a vertex x and a rotor configuration r, consider the quantity

$$\Phi(x,r) := f(x) + \sum_{u\in V} \big[\phi(u,r(u)) - \phi(u,r_0(u))\big]$$

where

$$\phi(u,j) := \sum_{i=1}^{j} \big[f(u) - f(u^{(i)}) + \Delta f(u)\big].$$

Note that $\Phi(x,r_t)$ is finite if r_t is any rotor configuration encountered by the rotor walk, since the only non-zero terms in the sum over u are those corresponding to vertices that the walk has visited (this is the reason for including the term "$-\phi(u,r_0(u))$" in the above definition). Note also that the definition of the Laplacian (7) and the rotor property (1) imply for all $u \in V$ that

$$\phi(u,d(u)) = 0. \tag{8}$$

Let us compute the change in Φ produced by a step of the rotor walk from (x_t,r_t) to (x_{t+1},r_{t+1}). The only term in the sum over u that changes is the one corresponding to $u = x_t$, and thus

$$\begin{aligned}\Phi(x_{t+1},r_{t+1}) - \Phi(x_t,r_t) &= f(x_{t+1}) - f(x_t) + [\phi(x_t,r_{t+1}(x_t)) - \phi(x_t,r_t(x_t))]\\ &= f(x_{t+1}) - f(x_t) + [f(x_t) - f(x_t^{(r_{t+1}(x_t))}) + \Delta f(x_t)]\\ &= \Delta f(x_t),\end{aligned}$$

where we have used (8) in the case when $r_{t+1}(x_t) = 1$. Therefore $\Phi(x_t,r_t) - \Phi(x_0,r_0) = \sum_{s=0}^{t-1} \Delta f(x_s)$. Also $\Phi(x_0,r_0) = f(x_0)$, so we obtain

$$\sum_{s=0}^{t-1} \Delta f(x_s) = f(x_t) - f(x_0) + \sum_{u\in V} \big[\phi(u,r_t(u)) - \phi(u,r_0(u))\big]. \tag{9}$$

In order to bound the last sum in (9), we use (8) together with Lemma 14 and the definition of ϕ to deduce

$$\begin{aligned}\bigl|\phi(u,r_t(u))-\phi(u,r_0(u))\bigr| &\le \frac{1}{2}\sum_{i=1}^{d(u)}\bigl|f(u)-f(u^{(i)})+\Delta f(u)\bigr| \\ &= \frac{1}{2}\sum_{v\in V} d(u)p(u,v)\bigl|f(u)-f(v)+\Delta f(u)\bigr| \qquad (10)\end{aligned}$$

(since $d(u)p(u,v)$ is the number of i such that $u^{(i)}=v$). We conclude by applying the triangle inequality to (9). $\square$

Proof of Theorem 1. We will apply Proposition 13 with $f=h_{b,c}$. Note that $h(b)=1$ and $h(c)=0$, while conditioning on the first step of the Markov chain gives $h(u)=\sum_{v\in V}p(u,v)h(v)$ for $u\neq b,c$. Hence, using $p(b,a)=p(c,a)=1$,

$$\Delta h(u)=\begin{cases}0, & u\neq b,c;\\ h(a)-1, & u=b;\\ h(a), & u=c,\end{cases}$$

and thus $\sum_{s=0}^{t-1}\Delta h(x_s)=h(a)[n_t(b)+n_t(c)]-n_t(b)$.

Turning to the other terms in Proposition 13, note that $|h(x_t)-h(x_0)|\le 1$, and $h(u)-h(v)+\Delta h(u)=0$ when $u\in\{b,c\}$ and $v=a$. Substituting into Proposition 13 gives $\bigl|h(a)[n_t(b)+n_t(c)]-n_t(b)\bigr|\le K_1$ as required. $\square$

Proof of Theorem 2. We will apply Proposition 13 with $f=k_b$. In this case

$$\Delta k(u)=\begin{cases}-1, & u\neq b;\\ k(a) & u=b,\end{cases}$$

and thus $\sum_{s=0}^{t-1}\Delta k(x_s)=(n_t(b))(k(a))+(t-n_t(b))(-1)=(k(a)+1)n_t(b)-t$. Substituting into Proposition 13 and using $|k(x_t)-k(x_0)|\le\max_{v\in V}k(v)$ and $k(b)-k(a)+\Delta k(b)=0$ completes the proof. $\square$

To prove Theorem 3 we note some elementary facts about Markov chains.

Lemma 15. *Let b,c be two distinct vertices of an irreducible recurrent Markov chain, and let π be a stationary vector. Then $\pi(b)e_{b,c}=\pi(c)e_{c,b}$. Also the hitting probabilities $h=h_{b,c}$ satisfy $\Delta h(b)=-e_{b,c}$ and $\Delta h(c)=e_{c,b}$.*

Proof. Let N denote the number of visits to c before the first return to b when started from b. It is a standard fact (see e.g. Theorem 1.7.6 in [20]) that $\mathbb{E}N = \pi(c)/\pi(b)$. On the other hand $\mathbb{P}(N = n) = e_{b,c}(1 - e_{c,b})^{n-1}e_{c,b}$ for $n \geq 1$, so $\mathbb{E}N = e_{b,c}/e_{c,b}$, and the first claim follows. For the remaining claims we compute Δh by conditioning on the first step: $\Delta h(b) = (1 - e_{b,c}) - h(b) = -e_{b,c}$ and $\Delta h(c) = e_{c,b} - h(c) = e_{c,b}$. ☐

Proof of Theorem 3. We will again apply Proposition 13 with $f = h = h_{b,c}$ (now without the restriction $p(b, a) = p(c, a) = 1$). Lemma 15 gives

$$\Delta h(u) = \begin{cases} 0, & u \neq b, c; \\ -e_{b,c}, & u = b; \\ e_{c,b}, & u = c, \end{cases}$$

and so $\sum_{s=0}^{t-1} \Delta h(x_s) = -n_t(b)e_{b,c} + n_t(c)e_{c,b}$. In order to bound the terms in the last sum in Proposition 13 in the cases $u = b, c$ note that $|\Delta h(u)| \leq 1$ in these cases, and so, for $u = b, c$,

$$\sum_{v \in V} p(u, v)|h(u) - h(v) + \Delta h(u)| \leq 1 + \sum_{v \in V} p(u, v)|h(u) - h(v)|.$$

Hence Proposition 13 gives

$$\left|n_t(b)e_{b,c} - n_t(c)e_{c,b}\right| \leq K_3.$$

Now divide through by $\pi(b)e_{b,c}$ (which equals $\pi(c)e_{c,b}$ by Lemma 15). ☐

Proof of Theorem 4. We will apply Proposition 13 with $f = k = k_b$. Note that $k(b) = 0$, while $\mathbb{E}_b T_b^+ = 1 + \sum_{v \in V} p(b, v)k(v)$. Also we have $\mathbb{E}_b T_b^+ = 1/\pi(b)$ (see [20]), hence

$$\Delta k(u) = \begin{cases} -1, & u \neq b; \\ 1/\pi(b) - 1 & u = b. \end{cases}$$

We bound the term for $u = b$ in Proposition 13 thus: $\sum_{v \in V} p(b, v)|k(b) - k(v) + \Delta k(b)| \leq 1/\pi(b) + \sum_{v \in V} p(b, v)|k(b) - k(v) - 1|$. We obtain

$$\left|\frac{n_t(b)}{\pi(b)} - t\right| \leq K_4,$$

and multiply by $\pi(b)/t$ to conclude. ☐

3 Proofs for walks on $\mathbb{Z}^2$

Our proof of Theorem 5 is based on the two lemmas below. For $k \geq 1$ we define the kth **box** $B(k) := (-k, k]^2 \cap \mathbb{Z}^2$ and the kth **layer** $\partial B(k) := B(k) \setminus B(k-1)$. See Figure 2.

Lemma 16. *Fix two distinct vertices b, c of $\mathbb{Z}^2$, and let $h = h_{b,c}$ be the hitting probability for the simple random walk on the square lattice. There exists $C = C(b,c) \in (0, \infty)$ such that for all positive integers k,*

$$\sum_{\substack{u,v \in B(k): \\ \|u-v\|_1 = 1}} |h(u) - h(v)| \leq C \ln k.$$

Proof. Fix b, c and write $C_1, C_2, \ldots$ for constants depending on b, c. We claim first that for all $v \in \mathbb{Z}^2$,

$$h(v) = C_1 + e_{b,c}[a(v-c) - a(v-b)], \tag{11}$$

where $a : \mathbb{Z}^2 \to \mathbb{R}$ is the *potential kernel* of $\mathbb{Z}^2$. (The function a may be expressed as $a(v) := \lim_{n\to\infty} \sum_{t=0}^{n} [\mathbb{P}(X_t = 0) - \mathbb{P}(X_t = v)]$, where (X_t) is the simple random walk on $\mathbb{Z}^2$ — for more information see e.g. [24, Ch. 3] or [15, Sect. 1.6].) To check (11), we note the following facts about a. Firstly,

$$a(v) = A + \tfrac{2}{\pi} \ln |v| + O(|v|^{-2}) \quad \text{as } |v| \to \infty, \tag{12}$$

where $|v| := \|v\|_2$ and A is an absolute constant (see [15, p. 39]). Since $\frac{d}{dx}(A + \frac{2}{\pi} \ln x) = \frac{2}{\pi} x^{-1}$ we deduce

$$|a(v-c) - a(v-b)| \leq C_2 |v|^{-1}.$$

Secondly, writing Δ for the Laplacian of the random walk on $\mathbb{Z}^2$, i.e. $\Delta f(u) := \frac{1}{4} \sum_{v:\|u-v\|_1=1} f(v) - f(u)$, we have

$$\Delta a(v) = \mathbf{1}[v = 0].$$

Hence, using Lemma 15 and the fact that $\pi \equiv 1$ is a stationary vector for the random walk, the function $v \mapsto h(v) - e_{b,c}[a(v-c) - a(v-b)]$ is bounded and harmonic, therefore constant, establishing (11).

We now claim that for all u, v with $\|u - v\|_1 = 1$,

$$|h(u) - h(v)| \leq C_3 |v|^{-2}. \tag{13}$$

Once this is established we obtain

$$\sum_{\substack{u,v\in B(k):\\ \|u-v\|_1=1}} |h(u)-h(v)| \le \sum_{j=1}^{k} C_4 j (C_3 j^{-2}) \le C \ln k.$$

as required.

Finally, turning to the proof of (13), combining (11) and (12) gives

$$h(u)-h(v) = C_5\big(\ln|u-c| - \ln|u-b| - \ln|v-c| + \ln|v-b|\big) + O(|v|^{-2}).$$

In order to bound the above expression, fix $u-v$ to be one of the 4 possible integer unit vectors, and write $v-c=z$ and $c-b=\alpha$ and $u-v=\beta$. For convenience identify the vector (x,y) with the complex number $x+iy$ and let $|\cdot|$ denote the modulus. We have

$$\begin{aligned}
&\ln|u-c| - \ln|u-b| - \ln|v-c| + \ln|v-b| \\
=&\ln\left|\frac{(z+\alpha)(z+\beta)}{(z+\alpha+\beta)z}\right| = \ln\left|1+\frac{\alpha}{z}+\frac{\beta}{z}-\frac{\alpha}{z}-\frac{\beta}{z}+O(|z|^{-2})\right| \\
=&O(|z|^{-2}) \qquad \text{as } z\to\infty.
\end{aligned}$$

□

Fix $a,b,c\in\mathbb{Z}^2$, and consider the rotor walk $x_0,x_1,\ldots$ started at a with rotor mechanism (4) and rotor configuration (5) modified so that $p(b,a)=p(c,a)=1$ as discussed in the paragraph preceding the statement of Theorem 5. We say that the walk **enters a new layer** at time t if for some k we have $x_0,\ldots,x_{t-1}\in B(k)$ but $x_t\notin B(k)$.

Lemma 17. *Under the above assumptions, between any two times at which the rotor walk enters a new layer, it must visit vertex a at least once. Also, between any two consecutive visits to vertex a, no vertex is visited more than 4 times.*

Proof. We start by proving the first assertion. The reader may find it helpful to consult Figure 2 throughout. Suppose for a contradiction that $a\in B(k-1)$, and that the rotor walk enters both the layers $\partial B(k)$ and $\partial B(k+1)$ for the first time without visiting a in between. Let s be the time of the last visit to a prior to entering $\partial B(k)$, and let t be the first time at which $\partial B(k+1)$ is entered.

We claim that some vertex v emitted the particle at least 5 times during $[s,t]$. To prove this, note first that $x_{t-1}\in\partial B(k)$, and consider the following two cases. If x_{t-1} is not one of the four "corner vertices" of $\partial B(k)$, then immediately after the particle moves from x_{t-1} to $x_t \in \partial B(k+1)$, the rotor at x_{t-1} is pointing in the same direction as in the initial rotor configuration r. Since this rotor did not move before time s, vertex x_{t-1} must have emitted the particle at least 4 times during $[s,t]$. Therefore, x_{t-1} must have received the particle at least 4 times from among its 4 neighbors in $[s,t]$ — but it has not received the particle from x_t, therefore by the pigeonhole principle it received it at least twice from some other neighbor v. And $v \notin \{b,c\}$ since $x_{t-1} \neq a$. By considering the rotor at v, we see that this implies that v emitted the particle at least 5 times during $[s,t]$. On the other hand, if x_{t-1} is a corner vertex of $\partial B(k)$, then on comparing with the initial rotor configuration r we see that x_{t-1} has emitted (and hence received) the particle 3 or 4 times, but two of its neighbors lie in $\partial B(k+1)$, so it did not receive the particle from them, and the same argument now applies. Thus we have proved the above claim.

Now let u be the first vertex to emit the particle 5 times during $[s,t]$. Then $u \notin \{a,b,c\}$, otherwise we would have a contradiction to our assumption that a is visited only once. But now repeating the argument above, u must have received the particle 5 times, so it must have received it at least twice from some neighbor, not in $\{a,b,c\}$, so this neighbor must have emitted the particle 5 times by some earlier time in $[s,t]$, a contradiction. Thus the first assertion is established.

The second assertion follows by an almost identical argument: if some vertex is visited at least 5 times between visits to a, then considering the first vertex to be so visited leads to a contradiction. $\square$

Proof of Theorem 5. We write $C_1, C_2, \ldots$ for constants which may depend on a, b, c. We use the proof of Proposition 13 in the case $f = h$. As in the proof of Theorem 1, equation (9) becomes

$$h(a)n - n_t(b) = h(x_t) - h(a) + \sum_{u \in V \setminus \{b,c\}} \big[\phi(u, r_t(u)) - \phi(u, r_0(u))\big],$$

where $n = n_t := n_t(b) + n_t(c)$. However, the term $\phi(u, r_t(u)) - \phi(u, r_0(u))$ is non-zero only for those vertices which have been visited by time t. Now the first assertion of Lemma 17 implies that at most one new layer is entered for each visit to a, and thus for each visit to $\{b,c\}$. Hence for some C_1, all the vertices visited by time t lie in $B(n + C_1)$ (where the constant C_1 depends on the layer of the initial vertex a).

Now proceeding as in the proof of Proposition 13 and using Lemma 16,

$$\begin{aligned}\Big|\sum_{u\in B(n+C_1)\setminus\{b,c\}} \big[\phi(u,r_t(u))-\phi(u,r_0(u))\big]\Big| &\le \tfrac{1}{2}\sum_{\substack{u,v\in B(n+C_1+1):\\ \|u-v\|_1=1}} |h(u)-h(v)| \\ &\le C\ln n.\end{aligned}$$

Combining this with the above facts gives

$$\big|h(a)n - n_t(b)\big| \le 1 + C\ln n,$$

as required.

Finally to prove the bound $t \le C'n^3$, we note by the second assertion of Lemma 17 that after n visits to vertex a, each of the at most $C_2 n^2$ vertices in $B(n+C_1)$ has been visited at most $4n$ times, so the total number of time steps is at most $4C_2 n^3$. □

4 Proofs for transfinite walks

Proof of Lemma 6. By irreducibility it is enough to show that if u is visited infinitely often and $p(u,v)>0$ then v is visited infinitely often. But this is immediate since $v = u^{(i)}$ for some i, so the rotor at u will be incremented to point to v infinitely often. □

Proof of Lemma 7. As in the preceding proof, if u is visited infinitely often and $p(u,v)>0$ then v is visited infinitely often, proving the first assertion. For the second assertion, let M be one greater than the first m for which the walk $x_{m\omega}, x_{m\omega+1}, \ldots$ is recurrent, or $M=\omega$ if all are transient. Then a is visited infinitely often before time $M\omega$, and we apply the first assertion. □

Proof of Theorem 8. We consider the quantity Φ defined in the proof of Proposition 13, with $f = h = h_{b,c}$ (as in the proof of Theorem 1). Suppose $x_0, x_1, \ldots$ is a transient rotor walk. We claim that

$$\Phi(x_\omega, r_\omega) - \lim_{t\to\infty} \Phi(x_t, r_t) = h(a). \tag{14}$$

The claim is proved as follows. The assumption of the theorem and the fact that the walk is transient imply that $\lim_{t\to\infty} h(x_t) = 0$. We clearly have $\lim_{t\to\infty} \phi(u, r_t(u)) = \phi(u, r_\omega(u))$ for each u, and by (10) and the definition of K_1 in Theorem 1 we have for all u and t that $|\phi(u, r_t(u)) - \phi(u, r_0(u))| \leq F(u)$ where $\sum_{u \in V} F(u) \leq 2(K_1 - 1)$. Hence by the dominated convergence theorem,

$$\lim_{t \to \infty} \Phi(x_t, r_t) = 0 + \sum_{u \in V} \big[\phi(u, r_\omega(u)) - \phi(u, r_0(u))\big] = \Phi(x_\omega, r_\omega) - h(a).$$

We have proved claim (14); thus whenever we "restart from infinity to a", the quantity Φ increases by $h(a)$. Combining this with the argument from the proof of Theorem 1, we get

$$\big[n_\tau(b) + n_\tau(c) + m\big] h(a) - n_\tau(b) = \Phi(x_\tau, r_\tau) - \Phi(x_0, r_0)$$

for $\tau = m\omega + t$, and the right side is bounded in absolute value by K_1 exactly as in the proof of Theorem 1. □

Proof of Theorem 9. We consider the quantity Φ defined in the proof of Proposition 13, with $f = g = g_b$. Note that

$$\Delta g(u) = \begin{cases} 0, & u \neq b; \\ -1, & u = b. \end{cases}$$

Mimicking the proof of Theorem 8, we obtain

$$g(a)m - n_\tau(b) = \Phi(x_\tau, r_\tau) - \Phi(x_0, r_0),$$

and we bound the right side as in the previous proofs, noting that when $u = b$ we have $|g(u) - g(v) + \Delta g(u)| \leq |g(u) - g(v)| + 1$. □

Our proof of Theorem 10 is based on an unpublished argument of Oded Schramm (although we present the details in a somewhat different way). We will need some preparation. It will be convenient to work with $R_n := n - I_n$, i.e. the number of times the transfinite rotor walk returns to a *without* going to infinity up to the time of the n return a. We also introduce some modified Markov chains and rotor mechanisms as follows.

Firstly, replace the vertex a with two vertices a_0 and a_1. Let $\widehat{V} = (V \setminus \{a\}) \cup \{a_0, a_1\}$ denote this modified vertex set. Introduce a modified transition kernel $\widehat{p}$ by letting a_0 inherit all the outgoing transition probabilities from a, and letting a_1 inherit all the incoming transition probabilities to a (i.e. let $\widehat{p}(a_0, v) = p(a, v)$ and $\widehat{p}(v, a_1) = p(v, a)$ for all $v \in V \setminus \{a\}$); also let $\widehat{p}(a_1, a_0) = 1$ and $\widehat{p}(a_0, a_1) = 0$, and let $\widehat{p}$ otherwise agree with p.

Secondly, for a positive integer d, let $B(d)$ denote the set of vertices that can be reached in at most d steps of the original Markov chain starting from a, and let $\partial B(d) := B(d) \setminus B(d-1)$. Let $\widehat{p}^d$ be $\widehat{p}$ modified so that $\widehat{p}^d(b, a_0) = 1$ for all $b \in \partial B(d)$. (Thus, on reaching distance d from a, the particle is immediately returned to a_0).

Fix a rotor mechanism and initial rotor configuration for the original Markov chain, and modify them accordingly to obtain a rotor walk associated with $\widehat{p}^d$, started at a_0. Let R_n^d be the number of times this rotor walk hits a_1 before the nth return to a_0 (i.e. before the $(n+1)$st visit to a_0). Also note that R_n is the number of times the transfinite rotor walk associated with $\widehat{p}$ and started at a_0 hits a_1 before the nth return to a_0.

Lemma 18. *For a fixed initial rotor configuration, and any non-negative integer n, we have $R_n^d \to R_n$ as $d \to \infty$ (i.e., $R_n^d = R_n$ for d sufficiently large).*

Proof. For $v \in \widehat{V}$, let $N_n^d(v)$ (respectively $N_n(v)$) be the number of visits to vertex v before the nth return to a_0 for the (transfinite) rotor walk associated with $\widehat{p}^d$ (respectively $\widehat{p}$). We claim that

$$N_n^d \to N_n \qquad \text{as } d \to \infty, \tag{15}$$

where the convergence is in the product topology on $\mathbb{N}^{\widehat{V}}$; in other words, for any finite set $F \subset \widehat{V}$, if d is sufficiently large then $N_n^d(v) = N_n(v)$ for all $v \in F$. The required result follows immediately from this, because $R_n^d = N_n^d(a_1)$ and $R_n = N_n(a_1)$.

We prove (15) by induction on n. It holds trivially for $n = 0$ because N_0^d and N_0 equal zero everywhere. Assume it holds for $n-1$. This implies in particular that the rotor configuration at the time of the $(n-1)$st return to a_0 similarly converges as $d \to \infty$ to the corresponding rotor configuration in the transfinite case. Now consider the portion of the transfinite rotor walk corresponding to $\widehat{p}$, starting just after the $(n-1)$st return to a_0, up until the nth return to a_0. Consider the following two possibilities. If this walk is recurrent (so that it returns to a_0 via a_1) then it visits only finitely many vertices, so if d is sufficiently large that N_{n-1}^d and N_{n-1} agree on all the vertices it visits, then N_n^d and N_n agree also agree on the same set of vertices, establishing (15) in this case. On the other hand, suppose the aforementioned walk is transient (so that it goes to infinity before being restarted at a_0). Given a finite set $F \subset \widehat{V}$, let d be such that that when this walk leaves F for the last time, it has never been outside $B(d)$. Now let d' be such that $N_{n-1}^{d'}$ and N_{n-1} agree on $B(d)$. Then N_n^d and N_n will agree on F. So (15) holds in this case also, and the induction is complete. $\square$

Lemma 19. *For all positive integers n and d we have $R_n^{d+1} \geq R_n^d$.*

Proof. This will follow by a special case of the Abelian property for rotor walks on finite graphs with a sink (see e.g. [12, Lemma 3.9]). First we slightly modify the mechanism yet again. Consider the rotor mechanism and initial rotor configuration corresponding to $\widehat{p}^{d+1}$. Remove all the vertices in $V \setminus B(d+1)$ (these cannot be visited by the rotor walk started at a_0 anyway). Introduce an additional *absorbing* vertex s (called the sink), and modify the transition probabilities so that on hitting a_1 or $\partial B(d+1)$, particles are sent immediately to s instead of to a_0. Modify the rotor mechanism accordingly, but do not otherwise modify the initial rotor configuration.

We now consider the following multi-particle rotor walk (see e.g. [12] or the discussion in the introduction for more information). Start with n particles at a_0, and perform a sequence of rotor steps. That is, at each step, choose any non-sink vertex which has a positive number of particles (if such exists), and **fire** the vertex; i.e. increment its rotor, and move one particle in the new rotor direction. Continue in this way until all particles are at the sink. [12, Lemma 3.9] states that the total number of times any given vertex fires during this procedure is independent of our choices of which vertex to fire.

In particular, consider the firing order in which we first move one particle repeatedly (so it performs an ordinary rotor walk) until it reaches s, then move the second particle in the same way, and so on. Thus the number of times a_1 fires is R_n^{d+1}. Alternatively, we may move one particle until the first time it reaches $\partial B(d) \cup \{s\}$, then "freeze" it, and move the second particle until it reaches $\partial B(d) \cup \{s\}$, and so on. At this stage, the number of times a_1 has fired is R_n^d. Now we can continue firing until the frozen particles reach s. Comparing the two procedures shows $R_n^{d+1} \geq R_n^d$. $\square$

Corollary 20. *For all positive integers n and d we have $R_n \geq R_n^d$.*

Proof. Immediate from Lemmas 18 and 19. $\square$

Proof of Theorem 10. Since $R_n = n - I_n$, the required result is clearly equivalent to $\liminf_{n\to\infty} R_n/n \geq \mathbb{P}_a(T_a^+ < \infty)$. Fix any $\epsilon > 0$. Then there exists d such that $\mathbb{P}_a(T_a^+ < T_{\partial B(d)}) \geq \mathbb{P}_a(T_a^+ < \infty) - \epsilon$. Now consider the modified rotor walk corresponding to $\widehat{p}^d$ as defined above. Since the set of vertices that can be reached from a_0 is finite (so in effect the vertex set is finite), Theorem 1 implies that $R^d_n/n \to \mathbb{P}_a(T_a^+ < T_{\partial B(d)})$ as $n \to \infty$. Putting these facts together with Corollary 20 we obtain

$$\liminf_{n\to\infty} \frac{R_n}{n} \geq \lim_{n\to\infty} \frac{R^d_n}{n} = \mathbb{P}_a(T_a^+ < T_{\partial B(d)}) \geq \mathbb{P}_a(T_a^+ < \infty) - \epsilon. \qquad \square$$

5 Proofs for stack walks

In this section we will prove Proposition 11, and use it together with Proposition 21 below to prove Theorem 12. Given a Markov chain and a stack mechanism, we define the discrepancy functions

$$D_n(u,v) := \#\{i \leq n : u^{(i)} = v\} - np(u,v).$$

Proposition 21. *For any Markov chain, any stack walk, any function f and any t,*

$$\sum_{s=0}^{t-1} \Delta f(x_s) = f(x_t) - f(x_0) + \sum_{u,v\in V} D_{n_t(u)}(u,v)\big[f(u) - f(v) + \Delta f(u)\big].$$

Proof. Consider the function

$$\Psi(t) := f(x_t) + \sum_{u\in V} \psi(u, n_t(u))$$

where

$$\psi(u,n) := \sum_{i=1}^{n} \big[f(u) - f(u^{(i)}) + \Delta f(u)\big].$$

As in the proof of Proposition 13 we have $\sum_{s=0}^{t-1} \Delta f(x_s) = \Psi(t) - \Psi(0)$.

From the definition of D we have

$$\psi(u,n) = \sum_{v\in V} \big[D_n(u,v) + np(u,v)\big]\big[f(u) - f(v) + \Delta f(u)\big].$$

But by the definition of the Laplacian, $\sum_{v\in V} p(u,v)[f(u)-f(v)+\Delta f(u)] = 0$; therefore

$$\psi(u,n) = \sum_{v\in V} D_n(u,v)\big[f(u) - f(v) + \Delta f(u)\big],$$

and the result follows on substituting. $\square$

Proof of Proposition 11. First note that it suffices to prove the case in which $p_1, \dots, p_n$ are all rational. The irrational case then follows by a limiting argument. Specifically, let $p_1^k, \dots, p_n^k$ be rational with $p_i^k \to p_i$ as $k \to \infty$, and let $z_1^k, z_2^k, \dots$ be a sequence satisfying (6) for the p_i^k's; then by a compactness argument (since $\{1, \dots, n\}$ is finite) there is a subsequence (k_j) and a sequence $z_1, z_2 \dots$ such that $z_t^{k_j} \to z_t$ for each t, and (z_t) then satisfies (6) for the p_i's.

Now suppose that $p_1, \dots, p_n$ are rational, and let d be their least common denominator. Consider the finite bipartite graph G with vertex classes $L := \{1, \dots, d\}$ and $R := \bigcup_{i=1}^n R_i$ where $R_i := \{(i, m) : m \in \{1, \dots, p_i d\}\}$, and with an edge from $t \in L$ to $(i, m) \in R$ if and only if

$$\left\lceil \frac{m-1}{p_i} \right\rceil \le t \le \left\lceil \frac{m}{p_i} \right\rceil. \tag{16}$$

We will prove that G has a perfect matching between L and R. Note first that $\#L = d = \sum_i p_i d = \#R$. We claim that any set $T \subseteq L$ has at least $p_i \#T$ neighbors in R_i; from this it follows that it has at least $\#T$ neighbors in R, and the existence of a perfect matching then follows from Hall's marriage theorem. To prove the claim, fix $i \in \{1, \dots, n\}$ and note that (16) is equivalent to $p_i t - p_i < m \le p_i t + 1$. Therefore in the case when T is an interval $[s, t]$, it is adjacent to all those pairs $(i, m) \in R_i$ for which m is an integer in $(p_i s - p_i, p_i t + 1] \cap [1, d]$; this includes all integers in $(p_i s - p_i, p_i t + 1)$ (since $p_i s - p_i \ge 0$ and $p_i t + 1 \le d + 1$). The latter interval has length $p_i(t - s + 1) + 1$, therefore it contains at least $p_i(t - s + 1) = p_i \#T$ integers as required. Now consider the case $T = [s, t] \cup [u, v]$ (where $u > t+1$). If the two intervals have disjoint neighborhoods in R_i, the claim follows by applying the single-interval case to each and summing. On the other hand if the neighborhoods of the two intervals intersect, we see from (16) that the neighborhood in R_i of T is the same as the neighborhood in R_i of the larger set $[s, v]$, so the claim again follows from the single-interval case. Finally, the case when T is a union of three or more intervals is handled by applying the same reasoning to each adjacent pair, proving the claim and hence the existence of a matching.

Fix a perfect matching of G, and for $t = 1, \dots, d$, let $z_t := i$ where R_i contains the partner of t. It follows from (16) that if (i, m) and $(i, m+1)$ have respective partners t and t' then $t < t'$. Therefore t and (i, m) are partners if and only if z_t is the mth occurrence of i in the sequence z; from (16) this

mth occurrence appears between positions $\lceil \frac{m-1}{p_i} \rceil$ and $\lceil \frac{m}{p_i} \rceil$. Thus,

$$\lfloor p_i t \rfloor \leq \#\{s \leq t : z_s = i\} \leq \lfloor p_i t \rfloor + 1,$$

and it follows that (6) holds for all $t \leq d$. Note also that the left side of (6) is zero for $t = d$, therefore continuing the sequence z so as to be periodic with period d completes the proof. $\square$

(For an alternative proof of Proposition 11 that applies also to infinite probability vectors, see [2].)

Proof of Theorem 12. Choosing the stack mechanism according to Proposition 11 ensures that $|D_n(u,v)| \leq 1$ for all u, v and n. Now apply Proposition 21 to $f = h$ to obtain

$$\begin{aligned} &\big|h(a)[n_t(b) + n_t(c)] - n_t(b)\big| \\ &\qquad \leq |h(x_t) - h(x_0)| + \sum_{\substack{u \in V \setminus \{b,c\}, \\ v \in V}} |D_{n_t(u)}(u,v)| \cdot |h(u) - h(v)|, \end{aligned}$$

and conclude by noting that $D_{n_t(u)}(u,v) = 0$ unless $p(u,v) > 0$. $\square$

6 Open questions

As the burgeoning literature on Eulerian walk and rotor-routing attests, there are numerous interesting open problems. Here we focus on a few that are related to rotor walk on Euclidean lattices.

(i) Can the bound $C \log n / n$ in Theorem 5 for the discrepancy in hitting probabilities for simple random walk on $\mathbb{Z}^2$ be improved to C/n? Do similar results hold in $\mathbb{Z}^d$ and for other initial rotor configurations?

(ii) For simple random walk on $\mathbb{Z}^2$, let all rotors initially point East, and consider the transfinite rotor walk restarted at the origin after each escape to infinity. What is the asymptotic behavior of I_n, the number of escapes to infinity before the nth return to the origin? Theorem 10 implies that $I_n/n \to 0$ as $n \to \infty$, but simulations suggest that the convergence is rather slow.

(iii) For simple random walk on $\mathbb{Z}^d$ with $d \geq 3$, does there exist an initial rotor configuration for which the rotor walk is recurrent?

Acknowledgments

We thank Omer Angel, David desJardins, Lionel Levine, Russell Lyons, Karola Mészáros, Yuval Peres, Oded Schramm and David Wilson for valuable discussions.

References

[1] O. Angel and A. E. Holroyd. Rotor walks on general trees. In preparation.

[2] O. Angel, A. E. Holroyd, J. Martin, and J. Propp. Discrete low-discrepancy sequences. arXiv:0910.1077. Preprint.

[3] J. Cooper, B. Doerr, T. Friedrich, and J. Spencer. Deterministic random walks on regular trees. In *Proceedings of SODA 2008*, pages 766–772, 2008.

[4] J. Cooper, B. Doerr, J. Spencer, and G. Tardos. Deterministic random walks. In *Proceedings of the Workshop on Analytic Algorithmics and Combinatorics*, pages 185–197, 2006.

[5] J. Cooper, B. Doerr, J. Spencer, and G. Tardos. Deterministic random walks on the integers. *European J. Combin.*, 28(8):2072–2090, 2007.

[6] J. N. Cooper and J. Spencer. Simulating a random walk with constant error. *Combin. Probab. Comput.*, 15(6):815–822, 2006.

[7] D. Dhar. Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.*, 64(14):1613–1616, 1990.

[8] B. Doerr and T. Friedrich. Deterministic random walks on the two-dimensional grid. In *Combinatorics, Probability and Computing*, volume 18, pages 123–144. Cambridge University Press, 2009.

[9] I. Dumitriu, P. Tetali, and P. Winkler. On playing golf with two balls. *SIAM J. Discrete Math.*, 16(4):604–615 (electronic), 2003.

[10] A. Engel. The probabilistic abacus. *Ed. Stud. Math.*, 6(1):1–22, 1975.

[11] A. Engel. Why does the probabilistic abacus work? *Ed. Stud. Math.*, 7(1–2):59–69, 1976.

[12] A. E. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp, and D. B. Wilson. Chip-firing and rotor-routing on directed graphs. In V. Sidoravicius and M. E. Vares, editors, *In and Out of Equilibrium 2*, volume 60 of *Progress in Probability*, pages 331–364. Birkhäuser, 2008.

[13] M. Kleber. Goldbug variations. *Math. Intelligencer*, 27(1):55–63, 2005.

[14] I. Landau and L. Levine. The rotor-router model on regular trees. *J. Combin. Theory Ser. A*, 116(2):421–433, 2009.

[15] G. F. Lawler. *Intersections of random walks*. Probability and its Applications. Birkhäuser Boston Inc., Boston, MA, 1991.

[16] L. Levine. The rotor-router model. arXiv:math/0409407.

[17] L. Levine and Y. Peres. Scaling limits for internal aggregation models with multiple sources. *J. d'Analyse Math.*, to appear, arXiv:0712.3378.

[18] L. Levine and Y. Peres. Spherical asymptotics for the rotor-router model in Z^d. *Indiana Univ. Math. J.*, 57(1):431–449, 2008.

[19] L. Levine and Y. Peres. Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile. *Potential Analysis*, 30(1):1–27, 2009, arXiv:0704.0688.

[20] J. R. Norris. *Markov chains*, volume 2 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998. Reprint of 1997 original.

[21] 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.

[22] J. Propp. Rotor walk and Markov chains: Refinements and examples. 2009. In preparation.

[23] Y. Rabani, A. Sinclair, and R. Wanka. Local divergence of markov chains and the analysis of iterative load-balancing schemes. In *IEEE Symp. on Foundations of Comp. Sci.*, pages 694–704, 1998.

[24] F. Spitzer. *Principles of random walk*. Springer-Verlag, New York, second edition, 1976. Graduate Texts in Mathematics, Vol. 34.

[25] D. Wilson. Generating spanning trees more quickly than the cover time. In *Proceedings of the Twenty-Eighth ACM Symposium on the Theory of Computing*, pages 296–303, 1996.

Alexander E. Holroyd: `holroyd at math dot ubc dot ca`
Microsoft Research, 1 Microsoft Way, Redmond WA, USA; and
University of British Columbia, 121-1984 Mathematics Rd., Vancouver, BC, Canada.

James Propp: `jpropp at cs dot uml dot edu`
University of Massachusetts Lowell, 1 University Ave., Olney Hall 428, Lowell, MA, USA.