

The range of a rotor walk

Laura Florescu^{*1}, Lionel Levine^{†2}, and Yuval Peres^{‡3}

¹New York University

²Cornell University

³Microsoft Research

August 20, 2014

Abstract

In a *rotor walk* the exits from each vertex follow a prescribed periodic sequence. On an infinite Eulerian graph embedded periodically in \mathbb{R}^d , we show that any simple rotor walk, regardless of rotor mechanism or initial rotor configuration, visits at least on the order of $t^{d/(d+1)}$ distinct sites in t steps. We prove a shape theorem for the rotor walk on the comb graph with i.i.d. uniform initial rotors, showing that the range is of order $t^{2/3}$ and the asymptotic shape of the range is a diamond. Using a connection to the mirror model and critical percolation, we show that rotor walk with i.i.d. uniform initial rotors is recurrent on two different directed graphs obtained by orienting the edges of the square grid, the Manhattan lattice and the F -lattice. We end with a short discussion of the time it takes for rotor walk to cover a finite Eulerian graph.

1 Introduction

In a *rotor walk* on a graph, the exits from each vertex follow a prescribed periodic sequence. Such walks were first studied in [18] as a model of mobile agents exploring a territory, and in [17] as a model of self-organized criticality. Propp proposed rotor walk as a deterministic analogue of random walk, a perspective explored in [6, 7, 10]. This paper is concerned with the following questions: How much territory does a rotor walk cover in t steps? Conversely, how many steps does it take for a rotor walk to completely explore a given finite graph?

Let $G = (V, E)$ be a finite or infinite directed graph. For $v \in V$ let $E_v \subset E$ be the set of outbound edges from v , and let \mathcal{C}_v be the set of all cyclic permutations of E_v . A *rotor configuration* on G is a choice of outbound edge $\rho(v) \in E_v$ for each $v \in V$. A *rotor mechanism* on G is a choice of cyclic permutation $m(v) \in \mathcal{C}_v$ for each $v \in V$. Given ρ and m , the *simple rotor walk* started at X_0 is a sequence of vertices $X_0, X_1, \dots \in \mathbb{Z}^d$ and rotor configurations $\rho = \rho_0, \rho_1, \dots$ such that for all integer times $t \geq 0$

$$\rho_{t+1}(v) = \begin{cases} m(v)(\rho_t(v)), & v = X_t \\ \rho_t(v), & v \neq X_t \end{cases}$$

and

$$X_{t+1} = \rho_{t+1}(X_t)^+$$

where e^+ denotes the target of the directed edge e . In words, the rotor at X_t “rotates” to point to a new neighbor of X_t and then the walker steps to that neighbor.

^{*}florescu@cims.nyu.edu

[†]<http://www.math.cornell.edu/~levine>. Partially supported by NSF grant DMS-1243606.

[‡]peres@microsoft.com

In a simple rotor walk the sequence of exits from v is periodic with period $\#E_v$. All rotor walks in this paper will be simple. (One can also study more general rotor walks in which the period is longer [8, 10].) We have chosen the retrospective rotor convention—each rotor at an already visited vertex indicates the direction of the most recent exit from that vertex—because it makes a few of our results such as Lemma 2.2 easier to state.




Figure 1: The range of a clockwise uniform rotor walk on \mathbb{Z}^2 after 80 returns to the origin: the mechanism m cycles through the four neighbors in clockwise order (North, East, South, West), and the initial rotors $\rho(v)$ were oriented independently North, East, South or West, each with probability $1/4$. Colors indicate the first twenty excursion sets A_1, \dots, A_{20} , defined in §2.

The *range* of rotor walk at time t is the set

$$R_t = \{X_1, \dots, X_t\}.$$

We investigate the growth rate of the number of distinct sites visited, $\#R_t$. A directed graph is called *Eulerian* if each vertex has as many incoming as outgoing edges: $\text{indeg}(v) = \text{outdeg}(v)$ for all $v \in V$. Any undirected graph can be made Eulerian by converting each undirected edge into a pair of oppositely oriented directed edges.

Theorem 1.1. *For any Eulerian graph G with a periodic embedding in \mathbb{R}^d , the number of distinct sites visited by a rotor walk in t steps satisfies*

$$\#R_t \geq ct^{d/(d+1)}.$$

for a constant $c > 0$ depending only on G (and not on ρ or m).

Priezzhev et al. [17] and Povolotsky et al. [16] gave a heuristic argument that $\#R_t$ has order $t^{2/3}$ for the clockwise rotor walk on \mathbb{Z}^2 with uniform random initial rotors ($\rho(x) = \pm e_1, \pm e_2$ each with probability $1/4$, independently for each site x). Theorem 1.1 gives a lower bound of this order, and our proof is directly inspired by their argument.

The upper bound promises to be more difficult because it depends on the initial rotor configuration ρ . Indeed, the next theorem shows that for certain ρ , the number of visited sites $\#R_t$ grows linearly in t . Rotor walk is called *recurrent* if $X_t = X_0$ for infinitely many t , and *transient* otherwise.

Theorem 1.2. *For any Eulerian graph G and any mechanism m , if the initial rotor configuration ρ has an infinite path of rotors directed toward X_0 , then rotor walk is transient and*

$$\#R_t \geq \frac{t}{\Delta},$$

where Δ is the maximal degree of a vertex in G .

One can also ask about the shape of the random set R_t , pictured in Figure 1. Each pixel in this figure corresponds to a vertex of \mathbb{Z}^2 , and R_t is the set of all colored pixels (the different colors correspond to *excursions* of the rotor walk, defined in §2); the mechanism m is clockwise, and the initial rotors ρ are i.i.d. uniform. Although the set R_t of Figure 1 looks far from round, Kapri and Dhar have conjectured that for very large t it becomes nearly a circular disk! From now on, by **uniform rotor walk** we will always mean that the initial rotors $\{\rho(v)\}_{v \in V}$ are independent and uniformly distributed on E_v .

Conjecture 1.3 (Kapri-Dhar [13]). The set of sites R_t visited by the clockwise uniform rotor walk in \mathbb{Z}^2 is asymptotically a disk: There exists a constant c such that for any $\epsilon > 0$,

$$P\{D_{(c-\epsilon)t^{1/3}} \subset R_t \subset D_{(c+\epsilon)t^{1/3}}\} \rightarrow 1$$

as $t \rightarrow \infty$, where $D_r = \{(x, y) \in \mathbb{Z}^2 : x^2 + y^2 < r^2\}$.




Figure 2: A piece of the comb graph (left) and the set of sites visited by a uniform rotor walk on the comb graph in 10000 steps.

We are a long way from proving anything like Conjecture 1.3, but we can show that an analogous shape theorem holds on a much simpler graph, the two dimensional comb (Figure 2).

Theorem 1.4. *For uniform rotor walk on the comb graph, $\#R_t$ has order $t^{2/3}$ and the asymptotic shape of R_t is a diamond.*

For the precise statement, see §4. This result contrasts with random walk on the comb, for which the expected number of sites visited is only on the order of $t^{1/2} \log t$ as shown by Pach and Tardos [15]. Thus the uniform rotor walk explores the comb more efficiently than random walk. (On the other hand, it is conjectured to explore \mathbb{Z}^2 *less* efficiently than random walk!)

The main difficulty in proving upper bounds for $\#R_t$ lies in showing that the uniform rotor walk is recurrent. This seems to be a difficult problem in \mathbb{Z}^2 , but we can show it for two different directed graphs obtained by orienting the edges of \mathbb{Z}^2 : the Manhattan lattice and the F -lattice, pictured in Figure 3.

Theorem 1.5. *Uniform rotor walk is recurrent on both the F -lattice and the Manhattan lattice.*

The proof uses a connection to the mirror model and critical bond percolation on \mathbb{Z}^2 ; see §5.

Theorems 1.1-1.5 bound the rate at which rotor walk explores various infinite graphs. In §6 we bound the time it takes a rotor walk to completely explore a given finite graph.

Related work

By comparing to a branching process, Angel and Holroyd [2] showed that uniform rotor walk on the infinite b -ary tree is transient for $b \geq 3$ and recurrent for $b = 2$. In the latter case the corresponding




Figure 3: Two different periodic orientations of the square grid with indegree and outdegree 2.

branching process is critical, and the distance traveled by rotor walk before returning n times to the root is doubly exponential in n . They also studied rotor walk on a singly infinite comb with the “most transient” initial rotor configuration ρ . They showed that if n particles start at the origin then order \sqrt{n} of them escape to infinity (more generally, order $n^{1-2^{1-d}}$ for a d -dimensional analogue of the comb).

In *rotor aggregation*, each of n particles starting at the origin performs rotor walk until reaching an unoccupied site, which it then occupies. For rotor aggregation in \mathbb{Z}^d , the asymptotic shape of the set of occupied sites is a Euclidean ball [14]. For the layered square lattice (\mathbb{Z}^2 with an outward bias along the x - and y -axes) the asymptotic shape becomes a diamond [12]. Huss and Sava [11] studied rotor aggregation on the 2-dimensional comb with the “most recurrent” initial rotor configuration. They showed that at certain times the boundary of the set of occupied sites is composed of four segments of exact parabolas. It is interesting to compare their result with Theorem 1.4: The asymptotic shape, and even the scaling required (elliptic for rotor walk, parabolic for rotor aggregation), is different.

2 Excursions

Let $G = (V, E)$ be a connected Eulerian graph. In this section G can be either finite or infinite, and the rotor mechanism m can be arbitrary. The main idea of the proof of Theorem 1.1 is to decompose rotor walk on G into a sequence of excursions. This idea was also used in [3] to construct recurrent rotor configurations on \mathbb{Z}^d for all d , and in [4, 5, 19] to bound the cover time of rotor walk on a finite graph (about which we say more in §6).

Definition. Fix a vertex $o \in V$. An *excursion* from o is a rotor walk started at o and run until it returns to o exactly $\deg(o)$ times.

More formally, let $(X_t)_{t \geq 0}$ be a rotor walk started at $X_0 = o$. For $t \geq 0$ let

$$u_t(x) = \#\{0 \leq s < t : X_s = x\}$$

and let $u_\infty(x) \in \mathbb{N} \cup \{\infty\}$ be the increasing limit of $u_t(x)$. For $n \geq 0$ let

$$T(n) = \inf\{t \geq 0 : u_t(o) \geq n \deg(o)\} \in \mathbb{N} \cup \{\infty\}$$

be the time taken for the rotor walk to complete n excursions from o . For all $n \geq 1$ such that $T(n-1) < \infty$, define

$$e_n = u_{T(n)} - u_{T(n-1)}, \quad n \geq 1.$$

Our first lemma says that each $x \in V$ is visited at most $\deg(x)$ times per excursion. The assumption that G is Eulerian is crucial here.

Lemma 2.1. [3, Lemma 8]; [5, §4.2] For any initial rotor configuration ρ ,

$$e_1(x) \leq \deg(x) \quad \forall x \in V.$$

Proof. If the rotor walk never traverses the same directed edge twice, then $T(1) = \infty$ and $u_\infty \leq \deg$, so we are done. Otherwise, consider the smallest t such that $(X_s, X_{s+1}) = (X_t, X_{t+1})$ for some $s < t$. Rotor walk reuses an outgoing edge from X_t only after it has used all of the outgoing edges from X_t . Therefore, at time t the vertex X_t has been visited $\deg(X_t) + 1$ times, but each incoming edge to X_t has been traversed at most once. Since G is Eulerian it follows that $X_t = o$ and $t = T(1)$. \square

Lemma 2.2. If $T(1) < \infty$ and there is a directed path of initial rotors from x to o , then

$$e_1(x) = \deg(x).$$

Proof. Let y be the first vertex on the path of initial rotors from x to o . By induction on the length of this path, y is visited exactly $\deg(y)$ times in an excursion from o . Each incoming edge to y is traversed at most once by Lemma 2.1, so in fact each incoming edge to y is traversed exactly once. In particular, the edge (x, y) is traversed. Since $\rho(x) = (x, y)$, the edge (x, y) is the last one traversed out of x , so x must be visited at least $\deg(x)$ times. \square

If G is finite, then $T(n) < \infty$ for all n by Lemma 2.1. If G is infinite, then depending on the rotor mechanism m and initial rotor configuration ρ , rotor walk may or may not complete an excursion from o . In particular, Lemma 2.2 implies the following.

Corollary 2.3. If ρ has an infinite path directed toward o , then $T(1) = \infty$.

Now let

$$A_n = \{x \in V : e_n(x) > 0\}$$

be the set of sites visited during the n th excursion. We also set $e_0 = \delta_o$ and $A_0 = \{o\}$. For a subset $A \subset V$, define

$$\partial A = \{y \in V : (x, y) \in E \text{ for some } x \in A\}.$$

Lemma 2.4. If $T(n+1) < \infty$, then

- (i) $e_{n+1}(x) \leq \deg(x)$ for all $x \in V$.
- (ii) $e_{n+1}(x) = \deg(x)$ for all $x \in A_n$.
- (iii) $A_{n+1} \supseteq A_n \cup \partial A_n$.

Proof. Part (i) is immediate from Lemma 2.1.

Part (ii) follows from Lemma 2.2 and the observation that in the rotor configuration $\rho_{T(n)}$, the rotor at each $x \in A_n$ points along the edge traversed most recently from x , so for each $x \in A_n$ there is a directed path of rotors in $\rho_{T(n)}$ leading to $X_{T(n)} = o$.

Part (iii) follows from (ii): the $(n+1)$ st excursion traverses each outgoing edge from each $x \in A_n$, so in particular it visits each vertex in $A_n \cup \partial A_n$. \square

For $x \in V$ and $r \in \mathbb{N}$ denote by $B(x, r)$ the set of vertices reachable from x by a directed path of length $\leq r$. Inducting on n using Lemma 2.4(ii), we obtain the following.

Corollary 2.5. If $T(n) < \infty$, then $B(o, n) \subseteq A_n$.

Rotor walk is called *recurrent* if $T(n) < \infty$ for all n . Consider the rotor configuration $\rho_{T(n)}$ at the end of the n th excursion. By Lemma 2.4, each vertex in $x \in A_n$ is visited exactly $\deg(x)$ times during the N th excursion for each $N \geq n+1$, so we obtain the following.

Corollary 2.6. For a recurrent rotor walk, $\rho_{T(N)}(x) = \rho_{T(n)}(x)$ for all $x \in A_n$ and all $N \geq n$.

The following proposition is a kind of converse to Lemma 2.4 in the case of undirected graphs.

Proposition 2.7. [4, Lemma 3]; [3, Prop. 11] Let $G = (V, E)$ be an undirected graph. For sequence of connected sets $S_1, S_2, \dots \subset V$ such that $S_{n+1} \supseteq S_n \cup \partial S_n$ for all $n \geq 1$, and any vertex $o \in S_1$, there exists a rotor mechanism m and initial rotors ρ such that the n th excursion for rotor walk started at o traverses each edge incident to S_n exactly once in each direction, and no other edges.

3 Lower bound on the range

In this section $G = (V, E)$ is an infinite connected Eulerian graph. For $x \in V$ and $r \in \mathbb{N}$ denote by $B(x, r)$ the set of vertices reachable from x by a directed path of length $\leq r$. Fix an origin $o \in V$ and let $v(r)$ be the number of directed edges incident to $B(o, r)$. Let $W(r) = \sum_{n=0}^{r-1} v(n)$. Write $W^{-1}(t) = \min\{r \in \mathbb{N} : W(r) > t\}$.

Fix a rotor mechanism m and an initial rotor configuration ρ on G . For $x \in V$ let $u_t(x)$ be the number of times x is visited by a rotor walk started at o and run for t steps. The *range* of rotor walk is the set $R_t = \{x \in V : u_t(x) > 0\}$.

Theorem 3.1. *For any rotor mechanism m , any initial rotor configuration ρ on G , and any time $t \geq 0$, the following bounds hold.*

- (i) $u_t(o) < \deg(o)W^{-1}(t)$.
- (ii) $u_t(x) \leq u_t(o) + \deg(x)$ for all $x \in V$.
- (iii) Let $\Delta_t = \max_{x \in B(o, t)} \deg(x)$. Then

$$\#R_t \geq \frac{t}{\deg(o)W^{-1}(t) + \Delta_t - 1} \quad (1)$$

Before proving this theorem we discuss a few examples. If $G = \mathbb{Z}^2$ then $B(o, r)$ is a diamond and $W(r) \sim \frac{8}{3}r^3$, which gives $\#R_t \geq ct^{2/3}$ with $c = (8/3)^{1/3} \approx 1.387$. More generally, if G is any graph with a periodic embedding in \mathbb{R}^d , then $W(r) = \Omega(r^{d+1})$, so by part (iii) the range of any rotor walk on G is at least $\#R_t = \Omega(t^{d/(d+1)})$, which shows that Theorem 3.1 implies Theorem 1.1. If G has exponential volume growth we get $\#R_t = \Omega(t/\log t)$.

Proof of Theorem 3.1. By Lemma 2.4 and Corollary 2.5, the n th excursion from o traverses each directed edge incident to $B(o, n-1)$, so the total length of the first r excursions is at least $W(r)$. Therefore if $t < W(r)$ then the rotor walk has not yet completed its r th excursion at time t , so $u_t(o) < r \deg(o)$. Taking $r = W^{-1}(t)$ yields part (i).

Part (ii) is immediate from Lemma 2.1.

Part (iii) follows from the fact that $t = \sum_{x \in B(o, t)} u_t(x)$: By parts (i) and (ii), each term is at most $\deg(o)W^{-1}(t) - 1 + \Delta_t$, so there are at least $t/(\deg(o)W^{-1}(t) - 1 + \Delta_t)$ nonzero terms. \square

Remark 3.2. Theorem 2.7 shows that if G is undirected, then (1) is the best possible lower bound on $\#R_t$ that does not depend on m or ρ . For example, taking $S_n = B(o, n)$ in \mathbb{Z}^2 yields a rotor walk with $\#R_t \sim (8/3)^{1/3}t^{2/3}$; the rotor mechanism is clockwise and the initial rotors are shown in Figure 4. More generally, by taking S_n to be a suitably growing sequence of sets, one can obtain any growth rate for $\#R_t$ intermediate between $t/W^{-1}(t)$ and t .

Part (i) of the next theorem gives a sufficient condition for rotor walk to be transient. Part (iii) shows that on a graph of bounded degree, the number of visited sites $\#R_t$ of a transient rotor walk grows linearly in t .

Theorem 3.3. *On any Eulerian graph, the following hold.*

- (i) If ρ has an infinite path of initial rotors directed toward the origin o , then $u_t(o) < \deg(o)$ for all $t \geq 1$.
- (ii) If $u_t(o) < \deg(o)$, then $\#R_t \geq t/\Delta_t$ where $\Delta_t = \max_{x \in B(o, t)} \deg(x)$.
- (iii) If rotor walk is transient, then there is a constant $C = C(m, \rho)$ such that

$$\#R_t \geq \frac{t}{\Delta_t} - C$$

for all $t \geq 1$.




Figure 4: Minimal range rotor configuration for \mathbb{Z}^2 : the excursion sets are diamonds.

Proof. (i) By Corollary 2.3, if ρ has an infinite path directed toward o , then rotor walk never completes its first excursion from o .

(ii) If rotor walk does not complete its first excursion, then it visits each vertex x at most $\deg(x)$ times by Lemma 2.1, so it must visit at least t/Δ_t distinct vertices.

(iii) If rotor walk is transient, then for some n it does not complete its n th excursion, so this follows from part (b) taking C to be the total length of the first $n - 1$ excursions. \square

4 Uniform rotor walk on the comb

The 2-dimensional *comb* is the subgraph of the square lattice \mathbb{Z}^2 obtained by removing all of its horizontal edges except for those on the x -axis (Figure 2). Vertices on the x -axis have degree 4, and all other vertices have degree 2.

Recall that the **uniform rotor walk** starts with independent random initial rotors $\rho(v)$ with the uniform distribution on outgoing edges from v . The following result shows that the range of the uniform rotor walk on the comb is close to the diamond

$$D_n := \{(x, y) \in \mathbb{Z}^2 : |x| + |y| < n\}.$$

Theorem 4.1. *Consider uniform rotor walk on the comb with any rotor mechanism. Let $n \geq 2$ and $t = \lfloor \frac{16}{3}n^3 \rfloor$. For any $a > 0$ there exist constants $c, C > 0$ such that*

$$\mathbb{P}\{D_{n-\sqrt{cn\log n}} \subset R_t \subset D_{n+\sqrt{cn\log n}}\} > 1 - Cn^{-a}.$$

Since the bounding diamonds have area $2n^2(1 + o(1))$, it follows that the size of the range is of order $t^{2/3}$: More precisely, by Borel-Cantelli,

$$\frac{\#R_t}{t^{2/3}} \rightarrow \left(\frac{3}{2}\right)^{2/3}$$

as $t \rightarrow \infty$, almost surely.

The proof of Theorem 4.1 is based on the observation that rotor walk on the comb, viewed at the times when it is on the x -axis, is a rotor walk on \mathbb{Z} . If $0 < x_1 < x_2 < \dots$ are the positions of rotors on the positive x -axis that will send the walker left before right, and $0 > x_{-1} > x_{-2} > \dots$ are the positions on the negative x -axis that will send the walker right before left, then the x -coordinate of the rotor walk on the comb follows a zigzag path: right from 0 to x_1 , then left to x_{-1} , right to x_2 , left to x_{-2} , and so on (Figure 5).

Likewise, rotor walk on the comb, viewed at the times when it is on the a fixed vertical line $x = k$, is also a rotor walk on \mathbb{Z} . Let $0 < y_{k,1} < y_{k,2} < \dots$ be the heights of the rotors on the line




Figure 5: An initial rotor configuration on \mathbb{Z} (top) and the corresponding rotor walk.

$x = k$ above the x-axis that initially send the walker down, and let $0 > y_{k,-1} > y_{k,-2} > \dots$ be the heights of the rotors on the line $x = k$ below the x-axis that initially send the walker up.

If the initial rotors are i.i.d. uniform then the random variables x_i and $y_{k,i}$ have mean $2|i|$. Consider the “bad event” that one of the random variables x_i or $y_{k,i}$ for $|i| \leq n$ is particularly far from its mean:

$$\mathcal{B}_1 = \mathcal{B}_1(n, c) = \bigcup_{i=-n}^n \{|x_i - 2|i|| > \sqrt{cn \log n}\} \cup \bigcup_{j,k=-n}^n \{|y_{k,j} - 2|j|| > \sqrt{cn \log n}\}.$$

The proof will be completed by the following three lemmas. The first shows that \mathcal{B}_1 is unlikely. The second shows that if \mathcal{B}_1 does not occur then the odometer function of the first m excursions is close to the function

$$f_m(x, y) = (m - \frac{|x|}{2} - \frac{|y|}{2})^+, \quad (2)$$

where we write $t^+ := \max(t, 0)$. Note that the contour lines of this function are diamonds! The third lemma says that this forces the range R_t to be close to a diamond.

Lemma 4.2. *For each $b > 0$ there exist $C, c > 0$ such that*

$$\mathbb{P}(\mathcal{B}_1(n, c)) < Cn^{-b}.$$

Lemma 4.3. *For rotor walk on the comb, $u_m(x, y)$ be the total number of full turns made by the rotor at position (x, y) during the first m excursions. Let $a = \sqrt{cn \log n}$ and*

$$\mathcal{B}_2 = \mathcal{B}_2(n, c) = \bigcup_{m=1}^n \bigcup_{(x,y) \in \mathbb{Z}^2} \{f_{m-2a}(x, y) \leq u_m(x, y) \leq f_{m+2a}(x, y)\}^c.$$

Then $\mathcal{B}_2 \subset \mathcal{B}_1$.

Lemma 4.4. *Let $t = \lfloor \frac{16}{3}n^3 \rfloor$ and $a = \sqrt{cn \log n}$ and*

$$\mathcal{B}_3 = \mathcal{B}_3(n, c) = \{D_{n-6a} \subset R_t \subset D_{n+6a}\}^c.$$

Then $\mathcal{B}_3(n, c) \subset \mathcal{B}_2(2n, c)$.

Proof of Lemma 4.2. Consider the random variable $B_N = \xi_1 + \xi_2 + \dots + \xi_N$, where ξ_i is 1 or 0 according to whether the rotor at $(i, 0)$ will send a particle left before right. Note that $\mathbb{P}(x_k \leq N) = \mathbb{P}(B_N \geq k) = \mathbb{P}(S_N \geq 2k - N)$, where $S_N = 2B_N - N$ is a sum of independent zero-mean ± 1 -valued random variables. By the usual Chernoff bound (see, for example, [1, A.1]),

$$\mathbb{P}(S_N > a) < e^{-a^2/2N}.$$

Taking $a = \sqrt{cn \log n}$ and $N = 2k - a$ we obtain

$$P(x_k < 2k - a) < e^{-cn \log n / 2N} < e^{-cn \log n / 4n} = n^{-c/4}$$

where we have used that $k \leq n$. Likewise, taking $N = 2k + a$ yields

$$P(x_k > 2k + a) < P(S_N < a) < Cn^{-c/5}$$

for sufficiently large constant C . Analogous bounds hold for $k < 0$ and for the $y_{k,j}$'s. Now the proof is completed with a union bound

$$\begin{aligned} \mathbb{P}(\mathcal{B}_1) &\leq \sum_{i=-n}^n \mathbb{P}(|x_i - 2i| > a) + \sum_{j,k=-n}^n \mathbb{P}(|y_{k,j} - 2j| > a) \\ &\leq ((2n) + (2n)^2)(2Cn^{-c/5}) \\ &\leq C'n^{-b} \end{aligned}$$

where $b = 2 - \frac{c}{5}$. □

Proof of Lemma 4.3. Fix $1 \leq m \leq n$ and a point $(x, y) \in \mathbb{Z}^2$. We must show that on the event \mathcal{B}_1^c we have

$$f_{m-2a}(x, y) \leq u_m(x, y) \leq f_{m+2a}(x, y). \quad (3)$$

By symmetry we may assume $x, y \geq 0$. In order to complete m excursions on the comb, the rotor walk viewed on the x -axis must complete m zigzags as in Figure 5,

$$0 \rightarrow x_1 \rightarrow x_{-1} \rightarrow \dots \rightarrow x_m \rightarrow x_{-m} \rightarrow 0.$$

If $x \in (x_K, x_{K+1}]$ then exactly $(m - K)^+$ of these zigzags cross x , so

$$u_m(x, 0) = (m - K)^+.$$

We have used a capital letter to remind you that the index K is random! On the event $\mathcal{B}_1^c \cap \{K < n\}$ we have the inequalities

$$\begin{aligned} 2K - a &\leq x_K \\ &< x \\ &\leq x_{K+1} \\ &\leq 2(K + 1) + a, \end{aligned}$$

which imply $|K - \frac{x}{2}| < a$. Hence

$$(m - \frac{x}{2} - a)^+ \leq u_m(x, 0) \leq (m - \frac{x}{2} + a)^+. \quad (4)$$

On the event $\mathcal{B}_1^c \cap \{K \geq n\}$ we have $u_m(x, 0) = 0$ and $x > 2n - a$, so (4) holds also in this case.

Having taken care of the x -axis, we now apply the same argument on each vertical tooth of the comb. The rotor walk viewed on the tooth passing through the point $(x, 0)$ completes M zigzags,

$$0 \rightarrow y_{x,1} \rightarrow y_{x,-1} \rightarrow \dots \rightarrow y_{x,M} \rightarrow y_{x,-M} \rightarrow 0$$

where $M = u_m(x, 0)$. So on the event \mathcal{B}_1^c we have

$$(u_m(x, 0) - \frac{y}{2} - a)^+ \leq u_m(x, y) \leq (u_m(x, 0) - \frac{y}{2} + a)^+.$$

This bound together with the x -axis bound (4) (and the fact that if $u_m(x, 0) = 0$ then $u_m(x, y) = 0$ for all $y \in \mathbb{Z}$) yields (3). Hence $\mathcal{B}_2 \subset \mathcal{B}_1$. □

Proof of Lemma 4.4. For a function f on the vertices of the comb, write

$$\|f\| := \sum_z \deg(z)f(z) = 2 \sum_{x \in \mathbb{Z}} f(x, 0) + 2 \sum_{(x,y) \in \mathbb{Z}^2} f(x, y).$$

Summing in diamond layers shows that for the function f_m of (2),

$$\|f_m\| = 2 \sum_{\ell=1}^{2m} (4\ell)(m - \frac{\ell}{2}) + O(m^2) = \frac{16}{3}m^3 + O(m^2).$$

Now let M be the (random) number of excursions completed by time t . Then

$$\|u_M\| \leq t \leq \|u_{M+2}\|.$$

We first argue that $\{M > \frac{3}{2}n\} \subset \mathcal{B}_2(2n)$ for sufficiently large n ; indeed, on the event $\mathcal{E} := \{M > \frac{3}{2}n\} \cap \mathcal{B}_2(2n)^c$ we have

$$t \geq \|u_M\| \geq \|f_{\frac{3}{2}n-2a}\|.$$

Recall that $t = \frac{16}{3}n^3$. The right side is $(1 - o(1))(\frac{3}{2})^3 \frac{16}{3}n^3$, so \mathcal{E} is empty for sufficiently large n .

Therefore on the event $\mathcal{B}_2(2n)^c$ we have

$$\|f_{M-2a}\| \leq \|u_M\| \leq t \leq \|u_{M+2}\| \leq \|f_{M+2+2a}\|$$

and taking cube roots we obtain

$$|M - n| \leq 3a.$$

Finally, writing $A_m = \{u_m > 0\}$, on $\mathcal{B}_2(2n)^c$ we have

$$D_{n-5a} \subset A_{n-3a} \subset R_t \subset A_{n+3a+1} \subset D_{n+5a+1}$$

which completes the proof. \square

5 Directed lattices and the mirror model

Figure 3 shows two different orientations of the square grid \mathbb{Z}^2 : The *F-lattice* has outgoing vertical arrows (N and S) at even sites, and outgoing horizontal arrows (E and W) at odd sites. The *Manhattan lattice* has every even row pointing *E*, every odd row pointing *W*, every even column pointing *S* and every odd column pointing *N*. In these two lattices every vertex has outdegree 2, so there is a unique rotor mechanism on each lattice (namely, exits from a given vertex alternate between the two outgoing edges) and a rotor walk is completely specified by its starting point and the initial rotor configuration ρ .

In this section we relate the uniform rotor walk on these lattices to percolation and the Lorenz mirror model [9, §13.3]. Consider the *half dual lattice* \mathbb{L} , a square grid whose vertices are the points $(x + \frac{1}{2}, y + \frac{1}{2})$ for $x, y \in \mathbb{Z}$ with $x + y$ even. We consider critical bond percolation on \mathbb{L} : each edge of \mathbb{L} is either open or closed, independently with probability $\frac{1}{2}$.

Note that each vertex v of \mathbb{Z}^2 lies on a unique edge e_v of \mathbb{L} . We consider two different rules for placing two-sided mirrors at the vertices of \mathbb{Z}^2 .

- Manhattan lattice: If e_v is closed then v has a mirror oriented parallel to e_v ; otherwise v has no mirror.
- F-lattice: Each vertex v has a mirror, which is oriented parallel to e_v if e_v is closed and perpendicular to e_v if e_v is open.

Consider now the *first glance mirror walk*: Starting at the origin o , it travels along a uniform random outgoing edge $\rho(o)$. On its first visit to each vertex $v \neq \mathbb{Z}^2 - \{o\}$, the walker behaves like a light ray: if there is a mirror at v then the walker reflects by a right angle, and if there is no mirror then the walker continues straight. At this point v is assigned the rotor $\rho(v) = (v, w)$ where w is the vertex of \mathbb{Z}^2 visited immediately after v . On all subsequent visits to v , the walker follows the usual rules of rotor walk.

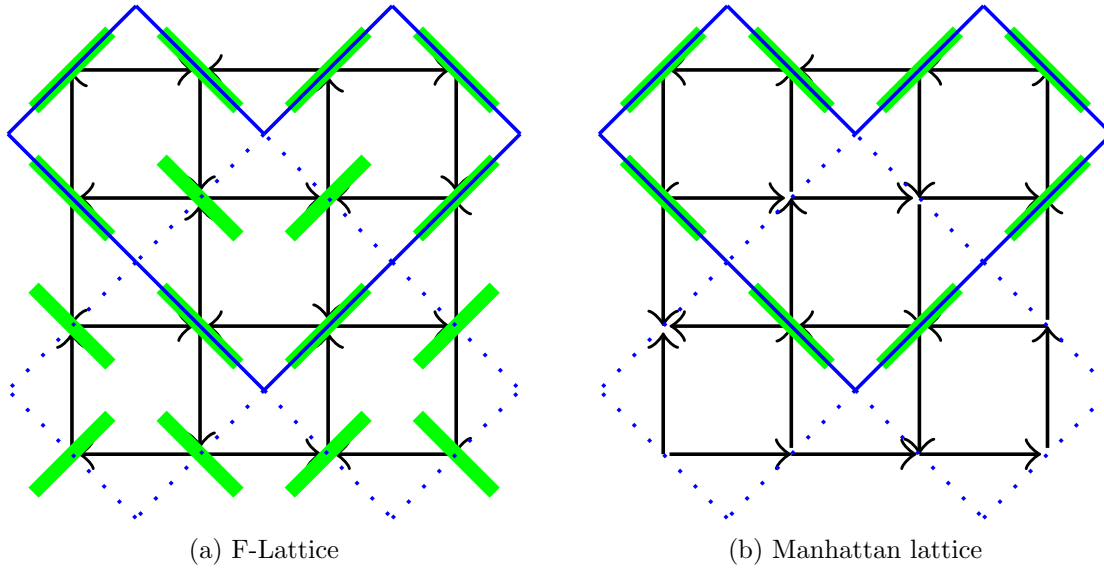


Figure 6: Percolation on \mathbb{L} : dotted blue edges are open, solid blue edges are closed. Shown in green are the corresponding mirrors on the F -lattice (left) and Manhattan lattice.

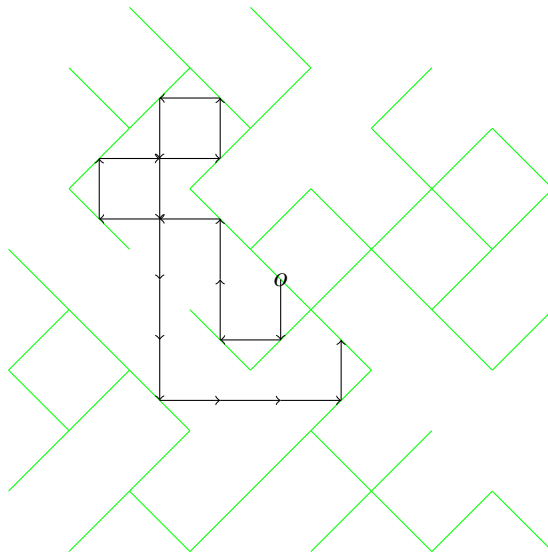


Figure 7: Mirror walk on the Manhattan lattice.

Lemma 5.1. *With the mirror assignments described above, uniform rotor walk on the Manhattan lattice or the F -lattice has the same law as the first glance mirror walk.*

Proof. The mirror placements are such that the first glance mirror walk must follow a directed edge of the corresponding lattice. The rotor $\rho(v)$ assigned by the first glance mirror walk when it first visits v is uniform on the outgoing edges from v ; this remains true even if we condition on the past, because all previously assigned rotors are independent of the status of the edge e_v (open or closed), and changing the status of e_v changes $\rho(v)$. \square

Write $\beta_e = 1\{e \text{ is open}\}$. Given the random variables $\beta_e \in \{0, 1\}$ indexed by the edges of \mathbb{L} , we have described how to set up mirrors and run a rotor walk, using the mirrors to reveal the initial rotors as needed. The next lemma holds pointwise in β .

Lemma 5.2. *If there is a cycle of closed edges in \mathbb{L} surrounding o , then rotor walk started at o returns to o at least twice before visiting any vertex outside the cycle.*

Proof. Denote by C the set of vertices v such that e_v lies on the cycle by C , and by A the set of vertices enclosed by the cycle. Let w be the first vertex not in $A \cup C$ visited by the rotor walk. Since the cycle surrounds o , the walker must arrive at w along an edge (v, w) where $v \in C$. Since e_v is closed, the walker reflects off the mirror e_v the first time it visits v , so only on the second visit to v does it use the outgoing edge (v, w) . Moreover, the two incoming edges to v are on opposite sides of the mirror. Therefore by minimality of w , the walker must use the same incoming edge (u, v) twice before visiting w . The first edge to be used twice is incident to the origin by Lemma 2.1, so the walk must return to the origin twice before visiting w . \square

Now we use a well-known theorem about critical bond percolation: there are infinitely many disjoint cycles of closed edges surrounding the origin. Together with Lemma 5.2 this completes the proof that the uniform rotor walk is recurrent both on the Manhattan lattice and the F -lattice.

To make a quantitative statement, consider the probability of finding a closed cycle within a given annulus. The following result is a consequence of the Russo-Seymour-Welsh estimate and FKG inequality.

Theorem 5.3. [9, 11.72] *Let $S_\ell = [-\ell, \ell] \times [-\ell, \ell]$. Then for all $\ell \geq 1$,*

$$P(\text{there exists a cycle of closed edges surrounding the origin in } S_{3\ell} - S_\ell) > p$$

for a constant p that does not depend on ℓ .

Let $u_t(o)$ be the number of visits to o by the first t steps of uniform rotor walk in the Manhattan or F -lattice.

Theorem 5.4. *For any $a > 0$ there exists $c > 0$ such that*

$$P(u_t(o) < c \log t) < t^{-a}.$$

Proof. By Lemma 5.2, the event $\{u_t(o) < k\}$ is contained in the event that at most $k/2$ of the annuli $S_{3^j} - S_{3^{j-1}}$ for $j = 1, \dots, \frac{1}{10} \log t$ contain a cycle of closed edges surrounding the origin. Taking $k = c \log t$ for sufficiently small c , this event has probability at most t^{-a} by Theorem 5.3. \square




Figure 8: Set of sites visited by uniform rotor walk after 250000 steps on the F -lattice and the Manhattan lattice (right). Green represents at least two visits to the vertex and red one visit.

Although we used the same technique to show that the uniform rotor walk on these two lattices is recurrent, experiments suggest that behavior of the two walks is rather different: the number of distinct sites visited in t steps appears to be of order $t^{2/3}$ on the Manhattan lattice but of order t for F -lattice. This difference is clearly visible in Figure 8.

6 Time for rotor walk to cover a finite Eulerian graph

Let $(X_t)_{t \geq 0}$ be a rotor walk on a finite connected Eulerian directed graph $G = (V, E)$. The *vertex cover time* is defined by

$$t_{\text{vertex}} = \min\{t : \{X_s\}_{s=1}^t = V\}.$$

The *edge cover time* is defined by

$$t_{\text{edge}} = \min\{t : \{(X_s, X_{s+1})\}_{s=0}^t = E\}.$$

Yanovski, Wagner and Bruckstein [19] show $t_{\text{edge}} \leq 2D\#E$ for any Eulerian directed graph. Our next result improves this bound slightly, replacing $2D$ by $D + 1$.

Theorem 6.1. *For rotor walk on a finite Eulerian graph G of diameter D , with any rotor mechanism m and any initial rotor configuration ρ ,*

$$t_{\text{vertex}} \leq D\#E$$

and

$$t_{\text{edge}} \leq (D + 1)\#E.$$

Proof. Consider the time $T(n)$ for rotor walk to complete n excursions from o . If G has diameter D then $A_D = V$ by Corollary 2.5, and $e_{D+1} = \deg$ by Lemma 2.4. It follows that $t_{\text{vertex}} \leq T(D)$ and $t_{\text{edge}} \leq T(D + 1)$. By Lemma 2.1, each edge is used at most once per excursion so $T(n) \leq n\#E$ for all $n \geq 0$. \square

Bampas et al. [4] prove a corresponding lower bound: on any finite undirected graph there exist a rotor mechanism m and initial rotor configuration ρ such that $t_{\text{vertex}} \geq \frac{1}{4}D\#E$.

6.1 Hitting times for random walk

The upper bounds for t_{vertex} and t_{edge} in Theorem 6.1 match (up to a constant factor) those found by Friedrich and Sauerwald [8] on an impressive variety of graphs: regular trees, stars, tori, hypercubes, complete graphs, lollipops and expanders. Intriguingly, the method of [8] is different: using a theorem of Holroyd and Propp [10] relating rotor walk to the expected time $H(u, v)$ for random walk started at u to hit v , they infer that $t_{\text{vertex}} \leq K + 1$ and $t_{\text{edge}} \leq 3K$, where

$$K := \max_{u, v \in V} H(u, v) + \frac{1}{2} \left(\#E + \sum_{(i, j) \in E} |H(i, v) - H(j, v) - 1| \right).$$

A curious consequence of the upper bound $t_{\text{vertex}} \leq K + 1$ of [8] and the lower bound $\max_{m, \rho} t_{\text{vertex}}(m, \rho) \geq \frac{1}{4}D\#E$ of [4] is the following inequality.

Corollary 6.2. *For any undirected graph G of diameter D we have*

$$K \geq \frac{1}{4}D\#E - 1.$$

Is K always within a constant factor of $D\#E$? It turns out the answer is no. To construct a counterexample we will build a graph $G = G_{\ell, N}$ of small diameter which has so few long-range edges that random walk effectively does not feel them (Figure 9). Let $\ell, N \geq 2$ be integers and set $V = \{1, \dots, \ell\} \times \{1, \dots, N\}$ with edges $(x, y) \sim (x', y')$ if either $x' \equiv x \pm 1 \pmod{\ell}$ or $y' = y$. The diameter of G is 2: any two vertices (x, y) and (x', y') are linked by the path $(x, y) \sim (x + 1, y) \sim (x', y')$. Each vertex (x, y) has $2N$ short-range edges to $(x \pm 1, y)$ and $\ell - 3$ long-range edges to (x', y) . We will argue that if ℓ is sufficiently large and $N = \ell^5$, then $K > \frac{1}{10}\ell\#E$, showing that K can exceed $D\#E$ by an arbitrarily large factor.

Write $Z_t^i = (X_t^i, Y_t^i)$ for random walk on G started at i . We couple the walks $(Z_t^i)_{t \geq 0}$ and $(Z_t^j)_{t \geq 0}$ as follows. At each time t either both walks will use a short-range edge or both walks will use a




Figure 9: The thick cycle $G_{\ell, N}$ with $\ell = 4$ and $N = 2$. Long range edges are dotted and short range edges are solid.

long-range edge. If they use a short-range edge, then we move the X coordinates independently and take $Y_t^i = Y_t^j$. If they use a long-range edge, then we take $X_t^i = X_t^j$.

Let T_{short} be the first time a short-range edge is used, and consider the hitting times

$$\begin{aligned}\tau(i) &= \min\{t \geq 1 : (X_t^i, Y_t^i) = (1, 1)\} \\ \sigma(i) &= \min\{t \geq T_{short} : X_t^i = 1\}.\end{aligned}$$

Now fix starting vertices i and j with $i_2, j_2 > 1$. Then neither walk can hit $v = (1, 1)$ before time T_{short} , so $\sigma(i) \leq \tau(i)$. Decompose the hitting time $\tau(i)$ into two pieces

$$\tau(i) = \sigma(i) + \sigma_2(i).$$

Both $Y_{\sigma(i)}^i$ and $Y_{\sigma(j)}^j$ are uniformly distributed on $\{1, \dots, L\}$, so $\mathbb{E}\sigma_2(i) = \mathbb{E}\sigma_2(j)$. Hence

$$H(i, v) - H(j, v) = \mathbb{E}\sigma(i) - \mathbb{E}\sigma(j). \quad (5)$$

To estimate the right side, we couple the random walks on G to random walks $\tilde{X}_t^i, \tilde{X}_t^j$ on the ℓ -cycle as follows. Let T_{long} be the first time a long-range edge is used. For $k \in \{i, j\}$ let $\tilde{X}_t^k = X_t^k$ for $0 \leq t < T_{long}$, and let the increments $\tilde{X}_t^k - \tilde{X}_{t-1}^k$ be independent of X for $t \geq T_{long}$. Let

$$\tilde{\sigma}(k) = \min\{t \geq 1 : \tilde{X}_t^k = 1\}.$$

On the event

$$\mathcal{A} := \{\max(\sigma(i), \sigma(j)) < T_{long}\}$$

we have $T_{short} = 1$ and $\tilde{\sigma}(k) = \sigma(k)$ for $k \in \{i, j\}$. Hence

$$\mathbb{E}|\sigma(i) - \sigma(j) - (\tilde{\sigma}(i) - \tilde{\sigma}(j))| \leq \mathbb{P}(\mathcal{A}^c) \left(\max_{i', j' \in V} \mathbb{E}|\sigma(i') - \sigma(j')| + \mathbb{E}(\tilde{\sigma}(i) + \tilde{\sigma}(j)) \right). \quad (6)$$

Since the probability of using a long-range edge at each fixed time t is $(\ell - 3)/(2N + \ell - 3)$, we have

$$\mathbb{P}(\mathcal{A}^c) \leq \mathbb{E} \sum_{t=0}^{\infty} 1\{t \leq \max(\sigma(i), \sigma(j))\} \frac{\ell - 3}{2N + \ell - 3} < \mathbb{E}(\sigma(i) + \sigma(j)) \frac{\ell}{2N}.$$

Now we use the explicit formula for random walk on the ℓ -cycle, $\mathbb{E}\tilde{\sigma}(i) = (i_1 - 1)(\ell - i_1 + 1)$. In particular, we have

$$|\mathbb{E}\tilde{\sigma}(i) - \mathbb{E}\tilde{\sigma}(j)| \geq \frac{\ell}{3} \quad \forall (i, j) \in E_1$$

where

$$E_1 := \left\{ (i, j) \in E : |i_i - j_1| = 1, i_2 > 1, j_2 > 1, \left| i_1 - \frac{\ell}{2} \right| > \frac{\ell}{5} \right\}.$$

Now take $N = \ell^5$ and ℓ sufficiently large. Since $\max_{i' \in V} \mathbb{E}\sigma(i') < \ell^2/2$, the right side of (6) is < 1 . Therefore by (5) we have

$$K > \sum_{(i, j) \in E_1} |H(i, v) - H(j, v) + 1| > \left(\frac{1}{2}\#E\right)\left(\frac{\ell}{3} - 2\right) > \frac{\ell}{10}\#E$$

as desired.

Note that Corollary 6.2 is a fact purely about random walk on a graph. Can it be proved without resorting to rotor walk?

Acknowledgements

This work was initiated while the first two authors were visiting Microsoft Research in Redmond, WA. We thank Sam Watson for help with some of the simulations, and Tobias Friedrich for bringing to our attention references [4] and [19].

References

- [1] Noga Alon and Joel Spencer. *The Probabilistic Method*. John Wiley & Sons, third edition, 2008.
- [2] Omer Angel and Alexander E. Holroyd. Rotor walks on general trees. *SIAM Journal on Discrete Mathematics*, 25:423–446, 2011. [arXiv:1009.4802](#).
- [3] Omer Angel and Alexander E. Holroyd. Recurrent rotor-router configurations. *Journal of Combinatorics*, 3(2):185–194, 2012. [arXiv:1101.2484](#).
- [4] Evangelos Bampas, Leszek Gasieniec, Nicolas Hanusse, David Ilcinkas, Ralf Klasing, and Adrian Kosowski. Euler tour lock-in problem in the rotor-router model. In *Distributed Computing*, pages 423–435. Springer, 2009.
- [5] Sandeep N. Bhatt, Shimon Even, David S. Greenberg, and Rafi Tayar. Traversing directed Eulerian mazes. *J. Graph Algorithms Appl.*, 6(2):157–173, 2002.
- [6] Joshua N. Cooper and Joel Spencer. Simulating a random walk with constant error. *Combinatorics, Probability and Computing*, 15(06):815–822, 2006. [arXiv:math/0402323](#).
- [7] Laura Florescu, Shirshendu Ganguly, Lionel Levine, and Yuval Peres. Escape rates for rotor walks in \mathbb{Z}^d . *SIAM Journal on Discrete Mathematics*, 28(1):323–334, 2014. [arXiv:1301.3521](#).
- [8] Tobias Friedrich and Thomas Sauerwald. The cover time of deterministic random walks. *The Electronic Journal of Combinatorics*, 17:R167, 2010. [arXiv:1006.3430](#).
- [9] Geoffrey Grimmett. *Percolation*. Springer, second edition, 1999.
- [10] Alexander E. Holroyd and James G. Propp. Rotor walks and Markov chains. *Algorithmic Probability and Combinatorics*, 520:105–126, 2010. [arXiv:0904.4507](#).
- [11] Wilfried Huss and Ecaterina Sava. Rotor-router aggregation on the comb. *Electron. J. Combin.*, 18:P224, 2011. [arXiv:1103.4797](#).
- [12] Wouter Kager and Lionel Levine. Rotor-router aggregation on the layered square lattice. *The Electronic Journal of Combinatorics*, 17(1):R152, 2010. [arXiv:1003.4017](#).
- [13] Rajeev Kapri and Deepak Dhar. Asymptotic shape of the region visited by an Eulerian walker. *Phys. Rev. E*, 80(5), November 2009. [arXiv:0906.5506](#).
- [14] Lionel Levine and Yuval Peres. Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile. *Potential Analysis*, 30:1–27, 2009. [arXiv:0704.0688](#).
- [15] János Pach and Gábor Tardos. The range of a random walk on a comb. *The Electronic Journal of Combinatorics*, 20(3):P59, 2013. [arXiv:1309.6360](#).
- [16] A. M. Povolotsky, V. B. Priezzhev, and R. R. Shcherbakov. Dynamics of Eulerian walkers. *Physical Review E*, 58(5):5449, 1998. [arXiv:cond-mat/9802070](#).
- [17] V. B. Priezzhev, Deepak Dhar, Abhishek Dhar, and Supriya Krishnamurthy. Eulerian walkers as a model of self-organised criticality. *Phys. Rev. Lett.*, 77:5079–5082, 1996. [arXiv:cond-mat/9611019](#).
- [18] Israel A. Wagner, Michael Lindenbaum, and Alfred M. Bruckstein. Smell as a computational resource – a lesson we can learn from the ant. In *4th Israel Symposium on Theory of Computing and Systems (ISTCS '96)*, pages 219–230, 1996.
- [19] Vladimir Yanovski, Israel A Wagner, and Alfred M Bruckstein. A distributed ant algorithm for efficiently patrolling a network. *Algorithmica*, 37(3):165–186, 2003.