The range of a rotor walk

Laura Florescu[*], **Lionel Levine**[†], **Yuval Peres**[‡]

Abstract. In *rotor walk* on a graph, the exits from each vertex follow a prescribed periodic sequence. We show that any rotor walk on the d-dimensional lattice $\mathbb{Z}^d$ visits at least on the order of $t^{d/(d+1)}$ distinct sites in t steps. This result extends to Eulerian graphs with a volume growth condition. In a *uniform* rotor walk the first exit from each vertex is to a neighbor chosen uniformly at random. We prove a shape theorem for the uniform rotor walk on the comb graph, showing that the size of 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, we show that the uniform rotor walk 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

Imagine walking your dog on an infinite square grid of city streets. At each intersection, your dogs tugs you one block further North, East, South or West. After you've been dragged in this fashion down t blocks, how many distinct intersections have you seen?

The answer depends of course on your dog's algorithm. If she makes a beeline for the North then every block brings you to a new intersection, so you see $t+1$ distinct intersections. At the opposite extreme, she could pull you back and forth repeatedly along her favorite block so that you see only ever see 2 distinct intersections.

In the *clockwise rotor walk* each intersection has a signpost pointing the way when you first arrive there. But your dog likes variety, and she has a capacious memory. If you come back to an intersection you have already visited, your dog chooses the direction 90° clockwise from the direction you went the last time you were there. We can formalize the city grid as the infinite graph $\mathbb{Z}^2$. The intersections are all the points (x, y) in the plane with integer coordinates, and the city blocks are the line segments from (x, y) to $(x \pm 1, y)$ and $(x, y \pm 1)$. More generally, we can consider a d-dimensional city $\mathbb{Z}^d$ or even an arbitrary graph, but the 90° clockwise rule will have to be replaced by something more abstract (a rotor mechanism, defined below).

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 [5, 6, 10]. This paper is concerned with the following questions. How much territory does a rotor walk cover in a fixed number of 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 an 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 $X_0, X_1, \ldots \in \mathbb{Z}^d$ and rotor configurations $\rho = \rho_0, \rho_1, \ldots$ such that for all integer times $t \geq 0$

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

[*]New York University, florescu@cims.nyu.edu

[†]Cornell University, http://www.math.cornell.edu/~levine. Partially supported by NSF grant DMS-1243606 and a Sloan Fellowship.

[‡]Microsoft Research, peres@microsoft.com

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.

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.

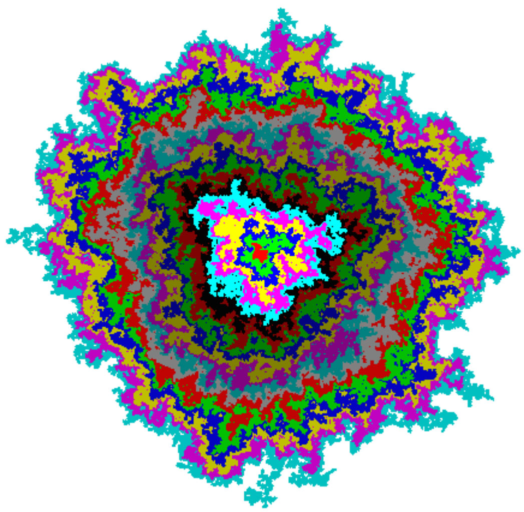

Figure 1. The range of a clockwise uniform rotor walk on $\mathbb{Z}^2$ after 80 returns to the origin. The mechanism m cycles through the four neighbors in clockwise order (North, East, South, West), and the initial rotors $\rho(v)$ were oriented independently North, East, South or West, each with probability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

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

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

We investigate the size of the range, $\#R_t$, in terms of the growth rate of balls in the underlying graph G. Fix an origin $o \in V$ (the starting point of our rotor walk). For $r \in \mathbb{N}$ the *ball* of radius r centered at o, denoted $B(o, r)$, is the set of vertices reachable from o by a directed path of length $\leq r$. Suppose that there are constants $d, k > 0$ such that

$$\#B(o, r) \geq kr^d \tag{1}$$

for all $r \geq 1$. Intuitively, this condition says that G is "at least d-dimensional."

A directed graph is called *Eulerian* if each vertex has as many incoming as outgoing edges. In particular, any undirected graph can be made Eulerian by converting each edge into a pair of oppositely oriented directed edges.

Theorem 1.1. *For any Eulerian graph G of bounded degree satisfying (1), the number of distinct sites visited by a rotor walk started at o 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. 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 (which we need not point out is much faster than $t^{2/3}$!). 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 directed toward o, then rotor walk started at o is transient and*

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

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

Theorems 1.1 and 1.2 are proved in §3. But enough about the size of the range; what about its shape? Each pixel in Figure 1 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 independently point North, East, South, or West with probability $1/4$ each. 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 $\epsilon > 0$,

$$\mathbb{P}\{\mathcal{D}_{(c-\epsilon)t^{1/3}} \subset R_t \subset \mathcal{D}_{(c+\epsilon)t^{1/3}}\} \to 1$$

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

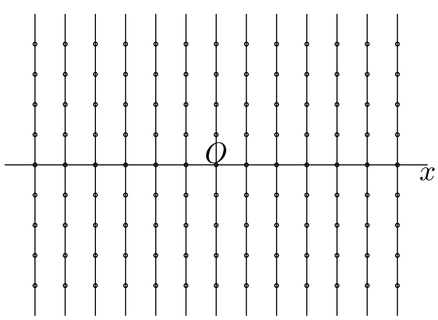

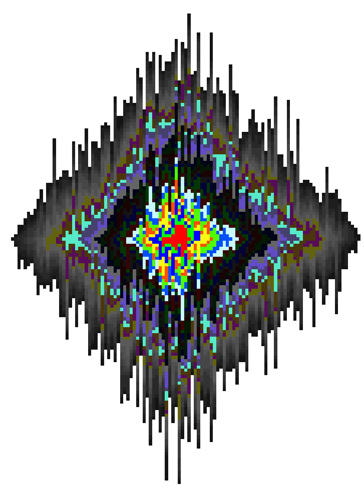

Figure 2. A piece of the comb graph (left) and the set of sites visited by a uniform rotor walk on the comb graph in 10000 steps.

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 *comb* obtained from $\mathbb{Z}^2$ by deleting all horizontal edges except those along the x-axis (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. The F-lattice has two outgoing horizontal edges at every odd node and two outgoing vertical edges at every even node (we call (x, y) odd or even according to whether $x + y$ is odd or even). The Manhattan lattice is full of one-way streets: rows alternate pointing left and right, while columns alternate pointing up and down.

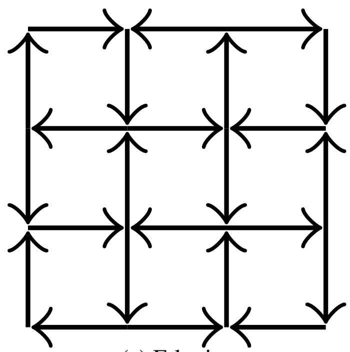

(a) F-lattice

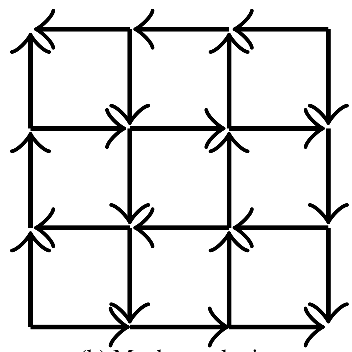

(b) Manhattan lattice

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

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 [1] 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 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, 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 [2] to construct recurrent rotor configurations on $\mathbb{Z}^d$ for all d, and in [3, 4, 19] to bound the cover time of rotor walk on a finite graph (about which we say more in §6). For a vertex $o \in V$ we write $\deg(o)$ for the number of outgoing edges from o, which equals the number of incoming edges since G is Eulerian.

Definition. 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

$$u_t(x) = \#\{1 \leq s \leq t \,:\, X_s = x\}.$$

For $n \geq 0$ let

$$T(n) = \min\{t \geq 0 \,:\, u_t(o) \geq n \deg(o)\},$$

be the time taken for the rotor walk to complete n excursions from o (with the convention that $\min$ of the empty set is ∞). For all $n \geq 1$ such that $T(n-1) < \infty$, define

$$e_n \equiv u_{T(n)} - u_{T(n-1)}$$

so that $e_n(x)$ counts the number of visits to x during the nth excursion. To make sense of this expression when $T(n) = \infty$, we write $u_\infty(x) \in \mathbb{N} \cup \{\infty\}$ for the increasing limit of the sequence $u_t(x)$.

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. *[2, Lemma 8]; [4, §4.2] For any initial rotor configuration ρ,*

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

Proof. If the rotor walk never traverses the same directed edge twice, then $u_t(x) \leq \deg(x)$ for all t and x, so we are done. Otherwise, consider the smallest t such that $(X_s, X_{s+1}) = (X_t, X_{t+1})$ for some $s < t$. By definition, 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 by the minimality of t each incoming edge to X_t has been traversed at most once. Since G is Eulerian it follows that $X_t = X_0 = o$ and $t = T(1)$.

Therefore every directed edge is used at most once during the first excursion, so each $x \in V$ is visited at most $\deg(x)$ times during the first excursion. ∎

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

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

Proof. Let y be the first vertex after x 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. ∎

If G is finite, then $T(n) < \infty$ for all n, since by Lemma 2.1 the number of visits to a vertex is at most or equal to the degree of that vertex. 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 nth excursion. We also set $e_0 = \delta_o$ (where, as usual, $\delta_o(x) = 1$ if $x = o$ and 0 otherwise) and $A_0 = \{o\}$. For a subset $A \subset V$, define its outer boundary ∂A as the set

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

Lemma 2.4. *For each $n \geq 0$, 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 $A_n \cup \partial A_n$. ∎

Note that the balls $B(o, n)$ can be defined inductively by $B(o, 0) = \{o\}$ and

$$B(o, n+1) = B(o, n) \cup \partial B(o, n)$$

for each $n \geq 0$. Inducting on n using Lemma 2.4(iii), we obtain the following.

Corollary 2.5. *For each $n \geq 1$, 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 nth excursion. By Lemma 2.4, each vertex in $x \in A_n$ is visited exactly $\deg(x)$ times during the Nth 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. *[3, Lemma 3]; [2, Prop. 11] Let $G = (V, E)$ be an undirected graph. For a sequence $S_1, S_2, \ldots \subset V$ of sets inducing connected subgraphs 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 nth 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. Fix an origin $o \in V$ and let $v(n)$ be the number of directed edges incident to the ball $B(o, n)$. Let $W(m) = \sum_{n=0}^{m-1} v(n)$. Write $W^{-1}(t) = \min\{m \in \mathbb{N} : W(m) > t\}$.

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. In the proof of the next theorem, our strategy for lower bounding the size of the range

$$R_t = \{x \in V \,:\, u_t(x) > 0\}$$

will be to (i) upper bound the number of excursions completed by time t, in order to (ii) upper bound the number of times each vertex is visited, so that (iii) many distinct vertices must be visited.

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) $\frac{u_t(o)}{\deg(o)} < W^{-1}(t)$.

(ii) $\frac{u_t(x)}{\deg(x)} \leq \frac{u_t(o)}{\deg(o)} + 1$ *for all* $x \in V$.

(iii) *Let* $\Delta_t = \max_{x \in B(o,t)} \deg(x)$. *Then*

$$\#R_t \geq \frac{t}{\Delta_t(W^{-1}(t)+1)}. \tag{2}$$

Before proving this theorem, let us see how it implies Theorem 1.1. The volume growth condition (1) implies $v(r) \geq kr^d$, so $W(r) \geq k'r^{d+1}$ for a constant k', so $W^{-1}(t) \leq (t/k')^{1/(d+1)}$. Now if G has bounded degree, then the right side of (2) is at least $ct^{d/(d+1)}$ for a constant c (which depends only on k and the maximal degree).

Proof of Theorem 3.1. We first argue that the total length $T(m)$ of the first m excursions is at least $W(m)$. By Corollary 2.5, the nth excursion visits every site in the ball $B(o,n)$. Therefore, by Lemma 2.4(ii), the $(n+1)$st excursion visits every site $x \in B(o,n)$ exactly $\deg(x)$ times, so the $(n+1)$st excursion traverses each directed edge incident to $B(o,n)$. The length $T(n+1) - T(n)$ of the $(n+1)$st excursion is therefore at least $v(n)$. Summing over $n < m$ yields the desired inequality $T(m) \geq W(m)$. Now let $m = W^{-1}(t)$. Since $t < W(m)$, the rotor walk has not yet completed its mth excursion at time t, so $u_t(o) < m \deg(o)$, which proves (i).

Part (ii) now follows from Lemma 2.1, since $e_1(x) = u_{T(1)}(x) \leq \deg(x)$. During each completed excursion, the origin o is visited $\deg(o)$ times while x is visited at most $\deg(x)$ times. The $+1$ accounts for the possibility that time t falls in the middle of an excursion.

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

Pausing to reflect on the proof, we see that an essential step was the inclusion $B(o,n) \subseteq A_n$ of Corollary 2.5. Can this inclusion ever be an equality? Yes! By Proposition 2.7, if G is undirected then there exists a rotor walk (that is, a particular m and ρ) for which

$$A_n = B(o,n) \qquad \text{for all } n \geq 1.$$

If $G = \mathbb{Z}^d$ (or any undirected graph satisfying (1) along with its upper bound counterpart, $\#B(o,n) \leq Kn^d$ for a constant K) then the range of this particular rotor walk satisfies $R_{W(n)} = B(o,n)$ and hence

$$\#R_t \leq \#B(o, W^{-1}(t)) \leq Ct^{d/(d+1)}$$

for a constant C. So in this case the exponent in Theorem 1.1 is best possible. We derived this upper bound just for a *particular* rotor walk, by choosing a rotor mechanism m and initial rotors ρ. For example, when $G = \mathbb{Z}^2$ the rotor mechanism is clockwise and the initial rotors are shown in Figure 4. Next we are going to see that by varying ρ we can make $\#R_t$ a lot larger.

Part (i) of the next theorem gives a sufficient condition for rotor walk to be transient. Parts (i) and (ii) together prove Theorem 1.2. 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.

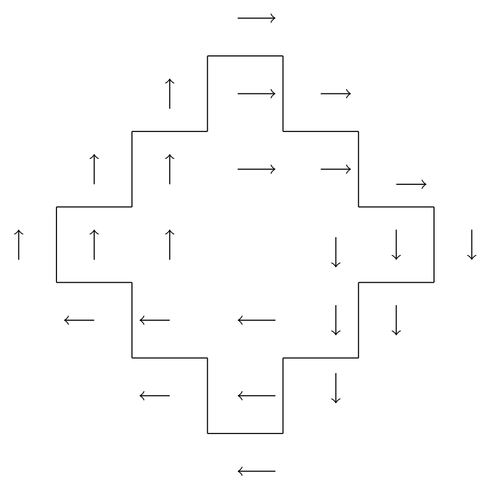

Figure 4. Minimal range rotor configuration for $\mathbb{Z}^2$. The excursion sets are diamonds.

Theorem 3.2. *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 $t \geq 1$.

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 nth excursion, so this follows from part (ii) taking C to be the total length of the first $n-1$ excursions. ∎

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 $2n^2(1+o(1))$ while t has order n^3, it follows that the size of the range is of order $t^{2/3}$. More precisely, by the first Borel-Cantelli lemma,

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

as $t \to \infty$, almost surely. See [7] for more details.

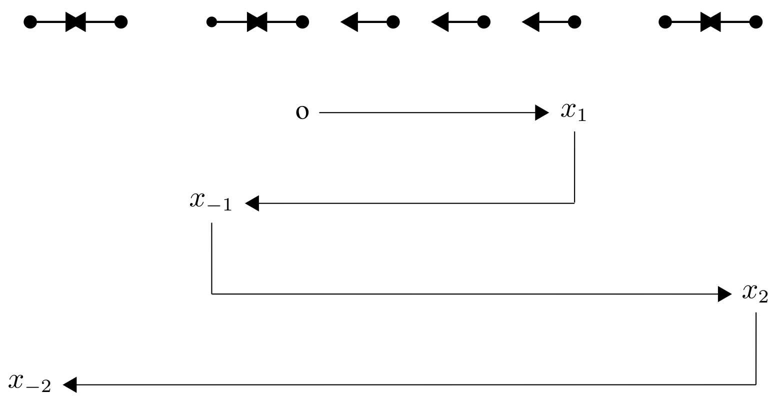

Figure 5. An initial rotor configuration on $\mathbb{Z}$ (top) and the corresponding rotor walk.

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 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 $x = k$ 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.

We only sketch the remainder of the proof; the full details are in [7]. For uniform initial rotors, the quantities x_i and $y_{k,i}$ are sums of independent geometric random variables of mean 2. We have $\mathbb{E}x_i = 2|i|$ and $\mathbb{E}y_{k,j} = 2|j|$. Standard concentration inequalities ensure that these quantities are close to their expectations, so that a rotor walk on the comb run for $n/2$ excursions visits each site $(x, 0) \in D_n$ about $(n - |x|)/2$ times, and hence visits each site $(x, y) \in D_n$ about $(n - |x| - |y|)/2$ times. Summing over $(x, y) \in D_n$ shows that the total time to complete these $n/2$ excursions is about $\frac{16}{3}n^3$. With high probability, every site in the smaller diamond $D_{n-\sqrt{cn\log n}}$ is visited at least once during these $n/2$ excursions, whereas no site outside the larger diamond $D_{n+\sqrt{cn\log n}}$ is visited.

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, and the usual lattice edges: $(x+\frac{1}{2}, y+\frac{1}{2}) - (x+\frac{1}{2}, y-\frac{1}{2})$, $(x+\frac{1}{2}, y+\frac{1}{2}) - (x-\frac{1}{2}, y+\frac{1}{2})$, $(x+\frac{1}{2}, y+\frac{1}{2}) - (x+\frac{3}{2}, y+\frac{1}{2})$, $(x+\frac{1}{2}, y+\frac{1}{2}) - (x+\frac{1}{2}, y+\frac{3}{2})$. We consider critical bond percolation on $\mathbb{L}$. Each possible lattice 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 $\mathbb{Z}^2$.

- 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.
- Manhattan lattice: If e_v is closed then v has a mirror oriented parallel to e_v; otherwise v has no mirror.

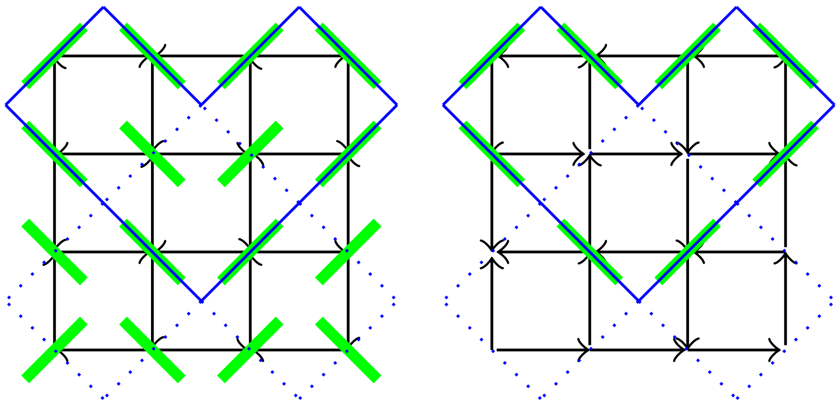

(a) F-Lattice (b) Manhattan lattice

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.

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.

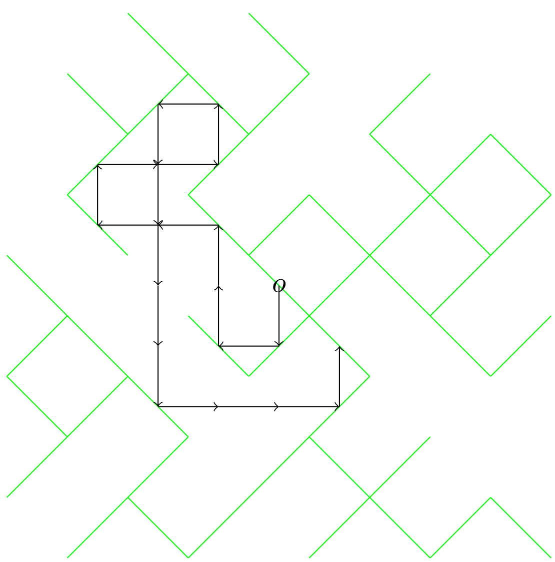

Figure 7. Mirror walk on the Manhattan lattice.

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 $\rho(v)$. ■

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, 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. ■

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; see [9, 11.72].

Theorem 5.3. *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 $j = 1, \ldots, \frac{1}{10} \log t$ 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. ■

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.

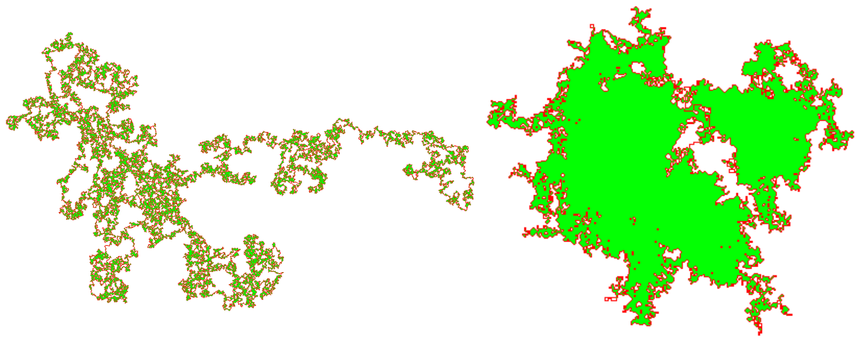

Figure 8. Set of sites visited by uniform rotor walk after 250000 steps on the F-lattice and the Manhattan lattice (right). Green represents at least two visits to the vertex and red one visit.

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 $G = (V, E)$ with diameter D. 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

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

where E is the set of directed edges. 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,*

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

and

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

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 $e_{D+1} \equiv \deg$ by Lemma 2.4(ii). It follows that $t_{\text{vertex}} \leq T(D)$ and $t_{\text{edge}} \leq T(D+1)$. By Lemma 2.1, each directed edge is used at most once per excursion so $T(n) \leq n\#E$ for all $n \geq 0$. ∎

Bampas et al. [3] prove a corresponding lower bound: on any finite undirected graph there exist a rotor mechanism m and initial rotor configuration ρ such that $t_{\text{vertex}} \geq \frac{1}{4}D\#E$.

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

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

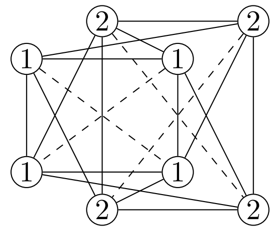

Figure 9. The thick cycle $G_{\ell,N}$ with $\ell = 4$ and $N = 2$. Long-range edges are dotted and short-range edges are solid.

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 [3] 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 $V = \{1, \ldots, \ell\} \times \{1, \ldots, N\}$ 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). It turns out that if ℓ is sufficiently large and N is much larger still ($N = \ell^5$), then $K > \frac{1}{10}\ell\#E$, showing that K can exceed $D\#E$ by an arbitrarily large factor. The details can be found in [7].

We conclude with a curious observation and a question. 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, Tobias Friedrich for bringing to our attention references [3] and [19], and the referees for their careful reading.

REFERENCES

1. O. Angel and A. E. Holroyd, Rotor walks on general trees, *SIAM Journal on Discrete Mathematics* **25** (2011) 423–446, arXiv:1009.4802.
2. ———, Recurrent rotor-router configurations, *Journal of Combinatorics* **3** (2012) 185–194, arXiv:1101.2484.
3. E. Bampas, L. Gasieniec, N. Hanusse, D. Ilcinkas, R. Klasing, and A. Kosowski, Euler tour lock-in problem in the rotor-router model, in *Distributed Computing*, Springer, 2009, 423–435.
4. S. N. Bhatt, S. Even, D. S. Greenberg, and R. Tayar, Traversing directed Eulerian mazes, *J. Graph Algorithms Appl.* **6** (2002) 157–173.
5. J. N. Cooper and J. Spencer, Simulating a random walk with constant error, *Combinatorics, Probability and Computing* **15** (2006) 815–822, arXiv:math/0402323.
6. L. Florescu, S. Ganguly, L. Levine, and Y. Peres, Escape rates for rotor walks in $\mathbb{Z}^d$, *SIAM Journal on Discrete Mathematics* **28** (2014) 323–334, arXiv:1301.3521.
7. L. Florescu, L. Levine, and Y. Peres, The range of a rotor walk (long version), arXiv:1408.5533.
8. T. Friedrich and T. Sauerwald, The cover time of deterministic random walks, *The Electronic Journal of Combinatorics* **17** (2010) R167, arXiv:1006.3430.
9. G. Grimmett, *Percolation*, second ed., Springer, 1999.
10. A. E. Holroyd and J. G. Propp, Rotor walks and Markov chains, *Algorithmic Probability and Combinatorics* **520** (2010) 105–126, arXiv:0904.4507.
11. W. Huss and E. Sava, Rotor-router aggregation on the comb, *Electronic Journal of Combinatorics* **18** (2011) P224, arXiv:1103.4797.

12. W. Kager and L. Levine, Rotor-router aggregation on the layered square lattice, *The Electronic Journal of Combinatorics* **17** (2010) R152, arXiv:1003.4017.
13. R. Kapri and D. Dhar, Asymptotic shape of the region visited by an Eulerian walker, *Phys. Rev. E* **80** (2009), arXiv:0906.5506.
14. L. Levine and Y. Peres, Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile, *Potential Analysis* **30** (2009) 1–27, arXiv:0704.0688.
15. J. Pach and G. Tardos, The range of a random walk on a comb, *The Electronic Journal of Combinatorics* **20** (2013) P59, arXiv:1309.6360.
16. A. M. Povolotsky, V. B. Priezzhev, and R. R. Shcherbakov, Dynamics of Eulerian walkers, *Physical Review E* **58** (1998) 5449, arXiv:cond-mat/9802070.
17. V. B. Priezzhev, D. Dhar, A. Dhar, and S. Krishnamurthy, Eulerian walkers as a model of self-organised criticality, *Phys. Rev. Lett.* **77** (1996) 5079–5082, arXiv:cond-mat/9611019.
18. I. A. Wagner, M. Lindenbaum, and A. 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)*, 1996, 219–230.
19. V. Yanovski, I. A. Wagner, and A. M. Bruckstein, A distributed ant algorithm for efficiently patrolling a network, *Algorithmica* **37** (2003) 165–186.

LAURA FLORESCU (florescu@cims.nyu.edu)
Courant Institute of Mathematical Sciences, NYU, New York, NY 10012
florescu@cims.nyu.edu

Laura Florescu is a Ph.D. student at Courant Institute, New York University.

LIONEL LEVINE (levine@math.cornell.edu)
Cornell University, Ithaca, NY 14853
levine@math.cornell.edu

Lionel Levine is an assistant professor at Cornell University. His Ph.D. is from Berkeley, where his advisor was Yuval Peres. You can usually find him thinking about why things are the way they are, or else about why they aren't the way they aren't.

YUVAL PERES (peres@microsoft.com)
Microsoft Research, Redmond, WA 98052
peres@microsoft.com

Yuval Peres is a Principal Researcher at Microsoft Research in Redmond, WA. He has written more than 250 research papers in probability theory, ergodic theory, analysis and theoretical computer science, in particular on fractals, random walks, Brownian motion, percolation, online learning and Markov chain mixing times. He obtained his Ph.D. at the Hebrew University of Jerusalem in 1990, and was later a faculty member there and at the University of California, Berkeley. Yuval was awarded the Rollo Davidson Prize in 1995, the Love Prize in 2001, and was a co-recipient of the David P. Robbins Prize in 2011. He was an invited speaker at the International Congress of Mathematicians in 2002. He is most proud of the 21 PhD students he has mentored, and hopes to keep learning from them.