

Rotor-router aggregation on the layered square lattice

Wouter Kager

VU University Amsterdam
Department of Mathematics
De Boelelaan 1081, 1081 HV
Amsterdam, The Netherlands

wkager@few.vu.nl

Lionel Levine*

Massachusetts Institute of Technology
Department of Mathematics
Cambridge, MA 02139, USA

levine@math.mit.edu

Submitted: Mar 31, 2010; Accepted: Oct 19, 2010; Published: Nov 5, 2010

Mathematics Subject Classification 2010: 82C24

Abstract

In rotor-router aggregation on the square lattice \mathbb{Z}^2 , particles starting at the origin perform deterministic analogues of random walks until reaching an unoccupied site. The limiting shape of the cluster of occupied sites is a disk. We consider a small change to the routing mechanism for sites on the x - and y -axes, resulting in a limiting shape which is a diamond instead of a disk. We show that for a certain choice of initial rotors, the occupied cluster grows as a perfect diamond.

1 Introduction

Recently there has been considerable interest in low-discrepancy deterministic analogues of random processes. An example is rotor-router walk [PDDK96], a deterministic analogue of random walk. Based at every vertex of the square grid \mathbb{Z}^2 is a *rotor* pointing to one of the four neighboring vertices. A chip starts at the origin and moves in discrete time steps according to the following rule. At each time step, the rotor based at the location of the chip turns clockwise 90 degrees, and the chip then moves to the neighbor to which that rotor points.

Holroyd and Propp [HP09] show that rotor-router walk captures the mean behavior of random walk in a variety of respects: stationary measure, hitting probabilities and hitting times. Cooper and Spencer [CS06] study rotor-router walks in which n chips starting at arbitrary even vertices each take a fixed number t of steps, showing that the final locations of the chips approximate the distribution of a random walk run for t steps to within constant error independent of n and t . Rotor-router walk and other low-discrepancy deterministic processes have algorithmic applications in areas such as

*The author was partly supported by a National Science Foundation Postdoctoral Fellowship.

broadcasting information in networks [DFS08] and iterative load-balancing [FGS10]. The common theme running through these results is that the deterministic process captures some aspect of the mean behavior of the random process, but with significantly smaller fluctuations than the random process.

Rotor-router aggregation is a growth model defined by repeatedly releasing chips from the origin $o \in \mathbb{Z}^2$, each of which performs a rotor-router walk until reaching an unoccupied site. Formally, we set $A_0 = \{o\}$ and recursively define

$$A_{m+1} = A_m \cup \{z_m\} \tag{1}$$

for $m \geq 0$, where z_m is the endpoint of a rotor-router walk started at the origin in \mathbb{Z}^2 and stopped on exiting A_m . We do not reset the rotors when a new chip is released.

It was shown in [LP08, LP09] that for any initial rotor configuration, the asymptotic shape of the set A_m is a Euclidean disk. It is in some sense remarkable that a growth model defined on the square grid, and without any reference to the Euclidean norm $|x| = (x_1^2 + x_2^2)^{1/2}$, nevertheless has a circular limiting shape. Here we investigate the dependence of this shape on changes to the rotor-router mechanism.

The *layered square lattice* $\widehat{\mathbb{Z}}^2$ (see Figure 2, left, below) is the directed multigraph obtained from the usual square grid \mathbb{Z}^2 by reflecting all directed edges on the x - and y -axes that point to a vertex closer to the origin. For example, for each positive integer n , the edge from $(n, 0)$ to $(n - 1, 0)$ is reflected so that it points from $(n, 0)$ to $(n + 1, 0)$. Thus the vertex $(n, 0)$ of $\widehat{\mathbb{Z}}^2$ has a pair of parallel directed edges to $(n + 1, 0)$, and one directed edge to each site $(n, \pm 1)$. All other edges of \mathbb{Z}^2 , in particular those that do not lie on the x - or y -axis, remain unchanged in $\widehat{\mathbb{Z}}^2$.

Rotor-router walk on $\widehat{\mathbb{Z}}^2$ is equivalent to rotor-router walk on \mathbb{Z}^2 with one modification: the reflection of the edges of the lattice carries over to the rotors. Thus, the rotor directions on the axes alternate between the directions of the two parallel edges of $\widehat{\mathbb{Z}}^2$ and the two perpendicular ones.

For $n \geq 0$, let

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

We call D_n the *diamond of radius n* . Our main result is the following.

Theorem 1. *There is a rotor configuration ρ_0 , such that rotor-router aggregation $(A_m)_{m \geq 0}$ on $\widehat{\mathbb{Z}}^2$ with rotors initially configured as ρ_0 satisfies*

$$A_{2n(n+1)} = D_n \quad \text{for all } n \geq 0.$$

A formal definition of rotor-router walk on $\widehat{\mathbb{Z}}^2$ and an explicit description of the rotor configuration ρ_0 are given below.

Let us remark on two features of Theorem 1. First, note that the rotor mechanism on $\widehat{\mathbb{Z}}^2$ is identical to that on \mathbb{Z}^2 except for sites on the x - and y -axes. Nevertheless, changing the mechanism on the axes completely changes the limiting shape, transforming it from a disk into a diamond. Second, not only is the aggregate close to a diamond, it is exactly equal to a diamond whenever it has the appropriate size (Figure 1).

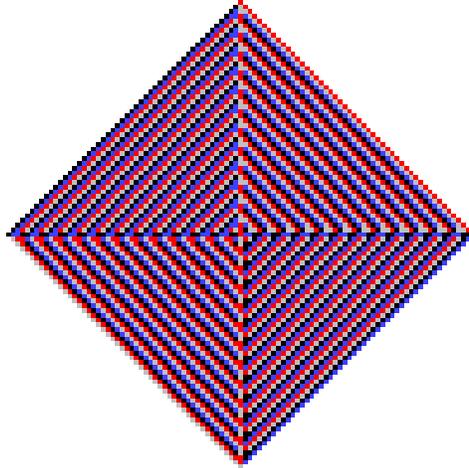


Figure 1: The rotor-router aggregate of 5101 chips in the layered square lattice $\widehat{\mathbb{Z}}^2$ is a perfect diamond of radius 50. The colors encode the directions of the final rotors at the occupied vertices: red = north, blue = east, gray = south and black = west.

Motivation and heuristic

In [KL10], we studied the analogous stochastic growth model, known as *internal DLA*, defined by the growth rule (1) using random walk on $\widehat{\mathbb{Z}}^2$. This random walk behaves like a simple random walk on \mathbb{Z}^2 except on the axes, where it takes steps with probability $1/2$ along the axis in the outward direction, and with probability $1/4$ in each of the two perpendicular directions. The walk has a *uniform layering property*: at any fixed time, its distribution is a mixture of uniform distributions on the diamond layers

$$L_m = \{(x, y) \in \mathbb{Z}^2 : |x| + |y| = m\}, \quad m \geq 1.$$

It is for this reason that we call $\widehat{\mathbb{Z}}^2$ the layered square lattice. The combinatorial feature of $\widehat{\mathbb{Z}}^2$ responsible for the uniform layering property is that each site in L_m has exactly two incoming edges from L_{m-1} and two incoming edges from L_{m+1} .

We have shown in [KL10] that, as a consequence of the uniform layering property, internal DLA on $\widehat{\mathbb{Z}}^2$ also grows as a diamond, but with random fluctuations at the boundary. Theorem 1 thus represents an extreme of discrepancy reduction: passing to the deterministic analogue removes *all* of the fluctuations from the random process, leaving only the mean behavior.

This work grew out of the uniformly layered walks in wedges studied in [Ka07]. The choice of transition probabilities on the axes — and hence the definition of the graph $\widehat{\mathbb{Z}}^2$ — was motivated by the idea that the uniform layering property of these walks could be extended to walks in the full plane.

Since the proof of Theorem 1 is a bit technical, we mention a heuristic that predicts the diamond shape without extensive calculation. The *uniform harmonic measure heuristic* says that a random walk started at the origin and stopped when it exits the cluster A_m

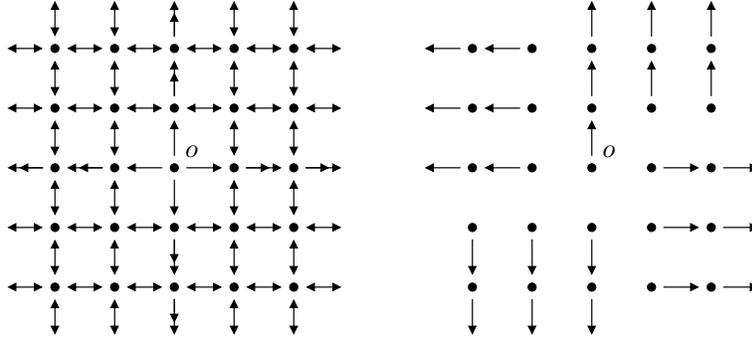


Figure 2: Left: The layered square lattice $\widehat{\mathbb{Z}}^2$. Each directed edge is represented by an arrow; parallel edges on the x - and y -axes are represented by double arrows. The origin o is in the center. Right: The initial rotor configuration ρ_0 .

should be roughly equally likely to stop at any boundary point. Intuitively, if this were not so, then those portions of the boundary more likely to be hit by the random walk would fill up faster as the cluster grows, changing the overall shape.

While it is usually not possible to convert this heuristic directly into a proof, note that it successfully predicts the limiting shape for growth models in both \mathbb{Z}^2 and $\widehat{\mathbb{Z}}^2$: simple random walk in \mathbb{Z}^2 has approximately uniform harmonic measure on a disk, while random walk in $\widehat{\mathbb{Z}}^2$ has *exactly* uniform harmonic measure on a diamond. This contrast helps explain why we could expect a “no discrepancy” result like Theorem 1 for $\widehat{\mathbb{Z}}^2$, as opposed to the “low discrepancy” results for \mathbb{Z}^2 .

Landau and Levine [LL09] prove a similar “no discrepancy” result to Theorem 1 when the underlying graph is a regular tree instead of $\widehat{\mathbb{Z}}^2$. The uniform harmonic measure heuristic predicts this behavior correctly as well. Still, more examples are needed: In other geometries, one expects that the shape may be controlled by a tradeoff between volume growth and harmonic measure rather than harmonic measure alone.

Formal definitions

To formally define rotor-router walk on $\widehat{\mathbb{Z}}^2$, write $\mathbf{e}_1 = (1, 0)$, $\mathbf{e}_2 = (0, 1)$ and let $R = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ be clockwise rotation by 90 degrees. The *layered square lattice* $\widehat{\mathbb{Z}}^2$ is the directed multigraph with vertex set $V = \mathbb{Z}^2$ and edge set E defined as follows. Every edge $e \in E$ is directed from its source vertex $\mathbf{s}(e)$ to its target vertex $\mathbf{t}(e)$. For every site $z \in \mathbb{Z}^2$ there are precisely 4 edges $e_z^0, e_z^1, e_z^2, e_z^3 \in E$ whose source vertex is z . For the origin o , the corresponding target vertices are $\mathbf{t}(e_o^i) = R^i \mathbf{e}_2$, meaning that $e_o^0, e_o^1, e_o^2, e_o^3$ are respectively directed north, east, south and west.

To specify the target vertices for $z \in \mathbb{Z}^2 \setminus \{o\}$, note that there is a unique choice of a number $j \in \{0, 1, 2, 3\}$ and a point w in the quadrant

$$Q = \{(x, y) \in \mathbb{Z}^2 : x \geq 0, y > 0\}$$

such that $z = R^j w$. Given j and $w = (x, y)$, we set

$$\mathfrak{t}(e_z^i) = \begin{cases} z + R^j \mathbf{e}_2 & \text{if } i = 2 \text{ and } x = 0; \\ z + R^{i+j} \mathbf{e}_2 & \text{otherwise.} \end{cases} \quad (2)$$

Thus, for $z \in Q$ (hence $j = 0$ and $w = z$) the edges $e_z^0, e_z^1, e_z^2, e_z^3$ are respectively directed north, east, north, west when z is on the y -axis; and north, east, south, west when z is off the y -axis. For z in another quadrant, the directions of $e_z^0, e_z^1, e_z^2, e_z^3$ are obtained by rotational symmetry.

Figure 2, left, gives a picture of $\widehat{\mathbb{Z}}^2$. Note that every vertex of $\widehat{\mathbb{Z}}^2$ has out-degree 4, and every vertex except for the origin and its immediate neighbors has in-degree 4. If $e = e_z^i \in E$, we will denote by e^+ the next edge $e_z^{i+1 \bmod 4}$ emanating from z , using the cyclic shift. Observe in particular that for $z \neq o$ on an axis, this sequence of consecutive edges alternates between the two parallel edges directed along the axis and the two perpendicular ones.

The initial rotor configuration ρ_0 appearing in Theorem 1 is given by

$$\rho_0(z) = e_z^0, \quad z \in \mathbb{Z}^2. \quad (3)$$

It has every rotor in the quadrant Q pointing north, and is chosen symmetric under R in accordance with the expected limiting shape (Figure 2, right).

We may now describe rotor-router walk on $\widehat{\mathbb{Z}}^2$ as follows. Given a rotor configuration ρ with a chip at vertex z , a single step of the walk consists of changing the rotor $\rho(z)$ to $\rho(z)^+$, and moving the chip to the vertex pointed to by the new rotor $\rho(z)^+$. This yields a new rotor configuration and a new chip location. Note that if the walk visits z infinitely many times, then it visits all out-neighbors of z infinitely many times, and hence visits every vertex of $\widehat{\mathbb{Z}}^2$ (except for o) infinitely many times. It follows that rotor-router walk exits any finite subset of $\widehat{\mathbb{Z}}^2$ in a finite number of steps; in particular, rotor-router aggregation terminates in a finite number of steps.

Outline

The rest of the paper proceeds as follows. In the next section we prove a ‘‘Strong Abelian Property’’ of the rotor-router model, Theorem 2. This theorem holds on any finite directed multigraph, and may be useful beyond its particular application to aggregation in $\widehat{\mathbb{Z}}^2$. Roughly speaking, the Strong Abelian Property allows us to reason about rotor-router moves without regard to whether particles are actually available to perform those moves. In Section 3, we prove Theorem 1 by applying the Strong Abelian Property to the induced subgraph D_n of $\widehat{\mathbb{Z}}^2$. Section 4 presents some open problems and discusses possible extensions of our methods.

2 Strong Abelian Property

Let $G = (V, E)$ be a finite directed multigraph (it may have loops and multiple edges). Each edge $e \in E$ is directed from its source vertex $\mathfrak{s}(e)$ to its target vertex $\mathfrak{t}(e)$. For a

vertex $v \in V$, write

$$E_v = \{e \in E : \mathfrak{s}(e) = v\}$$

for the set of edges emanating from v . The *outdegree* d_v of v is the cardinality of E_v .

Fix a nonempty subset $S \subset V$ of vertices called *sinks*. Let $V' = V \setminus S$, and for each vertex $v \in V'$, fix a numbering $e_v^0, \dots, e_v^{d_v-1}$ of the edges in E_v . If $e = e_v^i \in E_v$, we denote by e^+ the next element $e_v^{i+1 \bmod d_v}$ of E_v under the cyclic shift.

A *rotor configuration* on G is a function

$$\rho : V' \rightarrow E$$

such that $\rho(v) \in E_v$ for all $v \in V'$. A *chip configuration* on G is a function

$$\sigma : V \rightarrow \mathbb{Z}.$$

Note that we do not require $\sigma \geq 0$. If $\sigma(v) = m > 0$, we say there are m chips at vertex v ; if $\sigma(v) = -m < 0$, we say there is a hole of depth m at vertex v .

Fix a vertex $v \in V'$. Given a rotor configuration ρ and a chip configuration σ , the operation F_v of *firing* v yields a new pair

$$F_v(\rho, \sigma) = (\rho', \sigma')$$

where

$$\rho'(w) = \begin{cases} \rho(w)^+ & \text{if } w = v; \\ \rho(w) & \text{if } w \neq v; \end{cases}$$

and

$$\sigma'(w) = \begin{cases} \sigma(w) - 1 & \text{if } w = v; \\ \sigma(w) + 1 & \text{if } w = \mathfrak{t}(\rho(v)^+); \\ \sigma(w) & \text{otherwise.} \end{cases}$$

In words, F_v first rotates the rotor at v , then sends a single chip from v along the new rotor $\rho(v)^+$. We do not require $\sigma(v) > 0$ in order to fire v . Thus if $\sigma(v) = 0$, i.e., no chips are present at v , then firing v will create a hole of depth 1 at v ; if $\sigma(v) < 0$, so that there is already a hole at v , then firing v will increase the depth of the hole by 1.

Observe that the firing operators commute: $F_v F_w = F_w F_v$ for all $v, w \in V'$. Denote by \mathbb{N} the nonnegative integers. Given a function

$$u : V' \rightarrow \mathbb{N}$$

we write

$$F^u = \prod_{v \in V'} F_v^{u(v)}$$

where the product denotes composition. By commutativity, the order of the composition is immaterial.

A rotor configuration ρ is *acyclic on the set* $U \subset V'$ if the spanning subgraph $(V, \rho(U))$ has no directed cycles or, equivalently, if for every nonempty subset $A \subset U$ there is a vertex $v \in A$ such that $\mathfrak{t}(\rho(v)) \notin A$.

In the following theorem and lemmas, for functions f, g defined on a set of vertices $A \subset V$, we write “ $f = g$ on A ” to mean that $f(v) = g(v)$ for all $v \in A$, and “ $f \leq g$ on A ” to mean that $f(v) \leq g(v)$ for all $v \in A$.

Theorem 2 (Strong Abelian Property). *Let ρ be a rotor configuration and σ a chip configuration on G . Given two functions $u_1, u_2 : V' \rightarrow \mathbb{N}$, write*

$$F^{u_i}(\rho, \sigma) = (\rho_i, \sigma_i), \quad i = 1, 2.$$

If $\sigma_1 = \sigma_2$ on V' , and ρ_i is acyclic on the support of u_i for $i = 1, 2$, then $u_1 = u_2$.

Remark. If ρ_i is not acyclic on the support of u_i , one can always reduce u_i so that ρ_i becomes acyclic on its support without affecting σ_i , by a procedure called *reverse cycle-popping*, which is explained towards the end of the paper.

Note that the equality $u_1 = u_2$ in Theorem 2 implies that $\rho_1 = \rho_2$, and that $\sigma_1 = \sigma_2$ on all of V . For a similar idea with an algorithmic application, see [FL10, Theorem 1].

In a typical application of Theorem 2, we take $\sigma_1 = \sigma_2 = 0$ on V' , and u_1 to be the usual rotor-router odometer function

$$u_1(v) = \#\{1 \leq j \leq k : v_j = v\}$$

where v_1, v_2, \dots, v_k is a *complete legal firing sequence* for the initial configuration (ρ, σ) ; that is, a sequence of vertices which, when fired in order, causes all chips to be routed to the sinks without ever creating any holes. The resulting rotor configuration is necessarily acyclic on $A = \{v \in V' : u_1(v) > 0\}$: indeed, for any nonempty subset B of A , the rotor at the last vertex of B to fire points to a vertex not in B .

The usual abelian property of rotor-router walk [DF91, Theorem 4.1] says that any two complete legal firing sequences have the same odometer function. The Strong Abelian Property allows us to drop the hypothesis of legality: any two complete firing sequences whose final rotor configurations are acyclic on the set of vertices that have fired at all have the same odometer function, *even if one or both of these firing sequences temporarily creates holes*.

In our application to rotor-router aggregation on the layered square lattice, we take $V = D_n$ and $S = L_n$. We will take σ to be the chip configuration consisting of $2n(n+1)+1$ chips at the origin, and ρ to be the initial rotor configuration ρ_0 . Letting the chips at the origin in turn perform rotor-router walk until finding an unoccupied site defines a legal firing sequence (although not a complete one, since not all chips reach the sinks). In the next section, we give an explicit formula for the corresponding odometer function, and use Theorem 2 to prove its correctness. The proof of Theorem 1 is completed by showing that each nonzero vertex in D_n receives exactly one more chip from its neighbors than the number of times it fires.

To prove Theorem 2 we start with the following lemma.

Lemma 3. *Let $u : V' \rightarrow \mathbb{N}$, and write*

$$F^u(\rho, \sigma) = (\rho_1, \sigma_1).$$

If $\sigma = \sigma_1$, and u is not identically zero, then ρ_1 is not acyclic on the support $A = \{v \in V' : u(v) > 0\}$.

Proof. Since u is not identically zero, A is nonempty. Suppose for a contradiction that ρ_1 is acyclic on A . Then there is a vertex $v \in A$ whose rotor $\rho_1(v)$ points to a vertex not in A . The final time v is fired, it sends a chip along this rotor; thus, at least one chip exits A . Since the vertices not in A do not fire, no chips enter A , hence

$$\sum_{v \in A} \sigma_1(v) < \sum_{v \in A} \sigma(v)$$

contradicting $\sigma = \sigma_1$. □

Theorem 2 follows immediately from the next lemma.

Lemma 4. *Let $u_1, u_2 : V' \rightarrow \mathbb{N}$, and write*

$$F^{u_i}(\rho, \sigma) = (\rho_i, \sigma_i), \quad i = 1, 2.$$

If ρ_1 is acyclic on the support of u_1 , and $\sigma_2 \leq \sigma_1$ on V' , then $u_1 \leq u_2$ on V' .

Proof. Let

$$(\hat{\rho}, \hat{\sigma}) = F^{\min(u_1, u_2)}(\rho, \sigma).$$

Then (ρ_1, σ_1) is obtained from $(\hat{\rho}, \hat{\sigma})$ by firing only vertices in the set $A = \{v \in V' : u_1(v) > u_2(v)\}$, so

$$\hat{\sigma} \leq \sigma_1 \quad \text{on } V - A.$$

Likewise, (ρ_2, σ_2) is obtained from $(\hat{\rho}, \hat{\sigma})$ by firing only vertices in $V - A$, so

$$\hat{\sigma} \leq \sigma_2 \leq \sigma_1 \quad \text{on } A.$$

Thus $\hat{\sigma} \leq \sigma_1$ on V . Since $\sum_{v \in V} \hat{\sigma}(v) = \sum_{v \in V} \sigma_1(v)$ it follows that $\hat{\sigma} = \sigma_1$. Taking

$$u = u_1 - \min(u_1, u_2)$$

in Lemma 3, since $F^u(\hat{\rho}, \hat{\sigma}) = (\rho_1, \sigma_1)$ and the support of u is contained in the support of u_1 , we conclude that $u = 0$. □

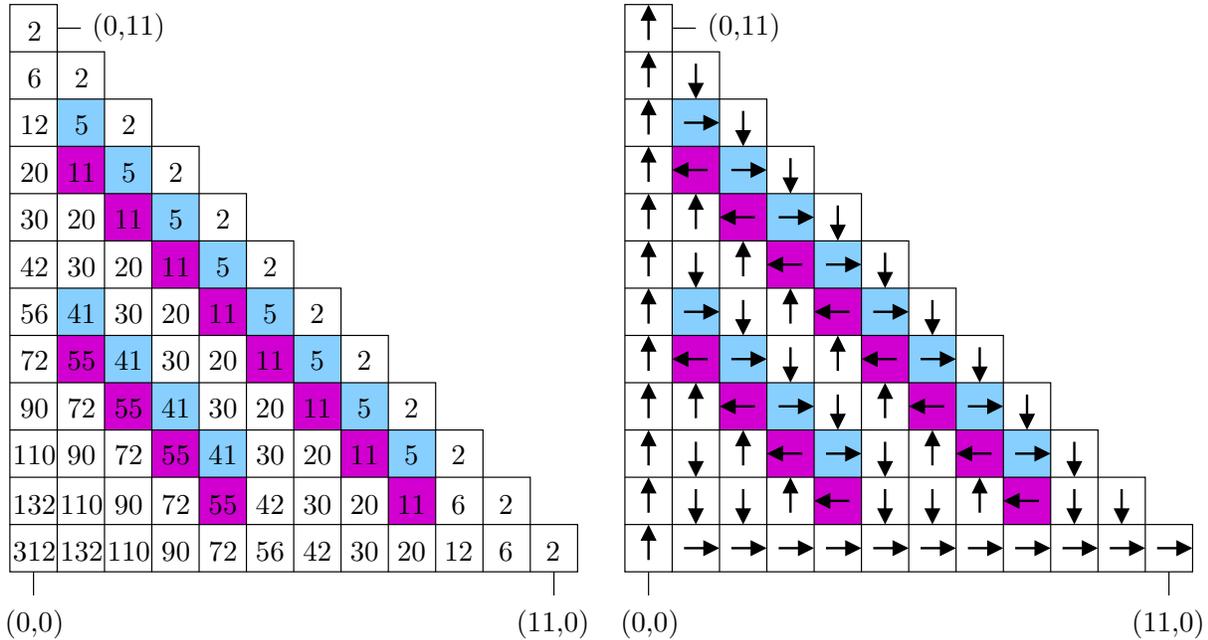


Figure 3: Left: the odometer u_{12} in the first quadrant and along the axes. Right: the corresponding rotor configuration ρ_{12} . The sets C_2 and C_3 are depicted in blue and purple, respectively. The rotor configuration is acyclic since following the rotors from any point of C_2 or the layer above it produces an alternating south-east path to the x -axis, while following the rotors from any point of C_3 or the layer below it produces an alternating north-west path to the y -axis.

3 Proof of Theorem 1

Consider again the rotor-router model on the layered square lattice $\widehat{\mathbb{Z}}^2$. We will work with the induced subgraph D_n of $\widehat{\mathbb{Z}}^2$, taking the sites in the outermost layer L_n as sinks.

Recall our notation

$$Q = \{(x, y) \in \mathbb{Z}^2 : x \geq 0, y > 0\}$$

for the first quadrant of \mathbb{Z}^2 . We have $\mathbb{Z}^2 = \{o\} \cup (\bigcup_{i=0}^3 R^i Q)$, where $R = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ is clockwise rotation by 90 degrees. Fix n , and for $z = (x, y) \in D_n$ write

$$\ell_z = n - |x| - |y|.$$

Then ℓ_z is the number of the diamond layer that z is on, where L_n is counted as layer 0, L_{n-1} as layer 1, and so on. Consider the sets

$$C_2 = \{(x, y) \in Q \cap D_{n-1} : x > 0, y \geq 2, \ell_{(x,y)} \equiv 2 \pmod{4}\}$$

$$C_3 = \{(x, y) \in Q \cap D_{n-1} : x > 0, y \geq 1, \ell_{(x,y)} \equiv 3 \pmod{4}\}$$

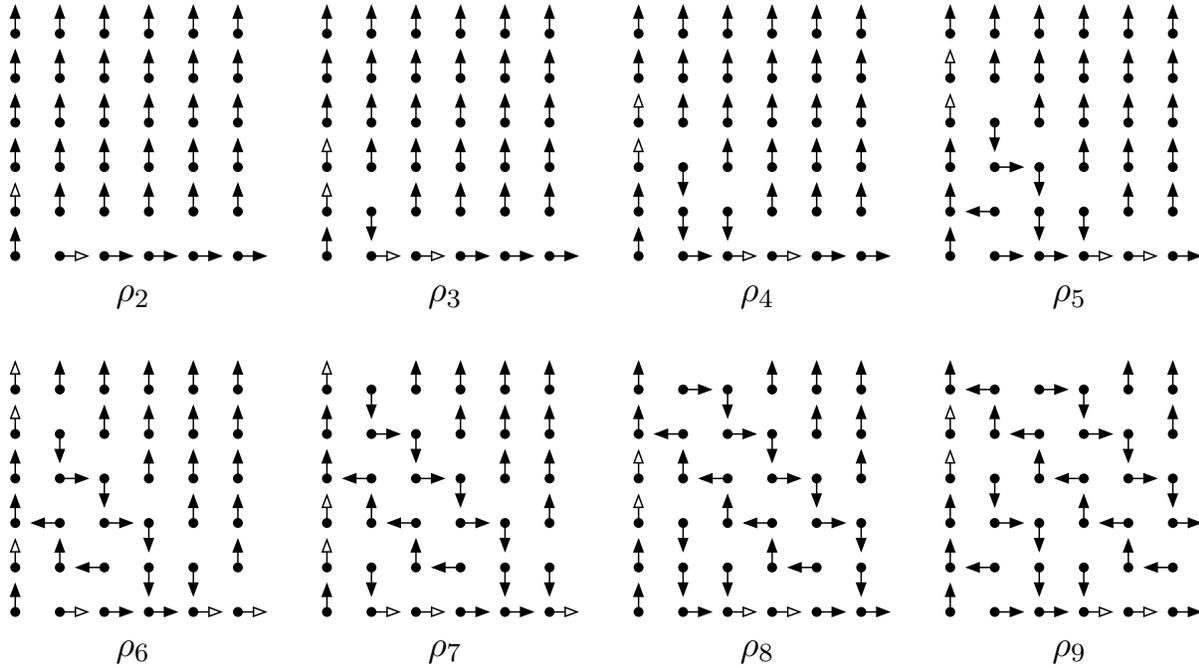


Figure 4: The rotor configurations $\rho_2, \rho_3, \dots, \rho_9$ on the set of vertices $\{(x, y) : 0 \leq x, y \leq 5\}$. On the axes, the black arrows correspond to the directed edge e_z^0 in (2), and open-headed arrows to e_z^2 .

and

$$C = \bigcup_{i=0}^3 R^i(C_2 \cup C_3).$$

Define $u_n : D_{n-1} \rightarrow \mathbb{N}$ by

$$u_n = u'_n - 1_C \tag{4}$$

where

$$u'_n(z) = \begin{cases} 2n(n+1) & \text{if } z = o; \\ \ell_z(\ell_z + 1) & \text{if } z \neq o; \end{cases} \tag{5}$$

and $1_C(z)$ is the indicator function which is 1 for $z \in C$ and 0 for $z \notin C$. See Figure 3, left, for a picture of the odometer u_{12} and the set C .

Let ρ_0 be the initial rotor configuration (3), and define the rotor configuration ρ_n on D_{n-1} and chip configuration σ_n on D_n by setting

$$F^{u_n}(\rho_0, (2n^2 + 2n + 1)\delta_o) = (\rho_n, \sigma_n).$$

From the formula (4) it follows that in the quadrant Q , all rotors of ρ_n point east on the set C_2 (since $\ell_z \equiv 2 \pmod{4}$ there), while the rotors on the diagonal above C_2 point south (see Figure 3, right). Thus, starting from any of these points, the rotors form

an alternating south-east path to the x -axis. Likewise, from any point in the set C_3 or the diagonal below it, the rotors form an alternating north-west path to the y -axis. Since the rotors on the axes all point in their initial directions ($\ell_z(\ell_z + 1)$ being even), it follows that the rotor configurations ρ_n are acyclic for all $n \geq 1$. Figure 4 shows how these rotor configurations develop; the periodicity mod 4 is apparent here by comparing a configuration in the upper row with the one below it.

Lemma 5. *For all $n \geq 1$, we have $\sigma_n = 1_{D_n}$.*

Proof. Recall that we work in the induced subgraph D_n of $\widehat{\mathbb{Z}}^2$, where we take the sites in the outermost layer L_n as sinks. The origin o has no incoming edges in $\widehat{\mathbb{Z}}^2$, so it receives no chips from its neighbors. Since $u_n(o) = 2n^2 + 2n$, the origin is left with exactly one chip after firing. The sink vertices L_n do not fire and only receive chips. Since $u_n(z) = 2$ for all $z \in L_{n-1}$, it follows from (2) and (3) that exactly one chip is sent to each sink vertex.

It remains to show that $\sigma_n(z) = 1$ for each vertex $z \in D_{n-1} \setminus \{o\}$, i.e., that the number of chips sent to z by its neighbors is one more than the number of times z is fired (that is, $1 + u_n(z)$). To show this, write

$$F^{u'_n}(\rho_0, (2n^2 + 2n + 1)\delta_o) = (\rho'_n, \sigma'_n)$$

where u'_n is given by (5). We will argue that $\sigma'_n(z) = 1$ and that $\sigma_n(z) = \sigma'_n(z)$. By symmetry, it suffices to consider points $z = (x, y)$ in $D_{n-1} \cap Q$. We argue separately in the two cases $x = 0$ and $x > 0$ (on the axis and off the axis).

Case 1: $x = 0$. Under $F^{u'_n}$, the site z fires $\ell_z(\ell_z + 1)$ times. If $y > 1$, its neighbor $z - \mathbf{e}_2$ fires $(\ell_z + 1)(\ell_z + 2)$ times, and from (2) and (3) we see that it sends a chip to z every even time it is fired. Since $(\ell_z + 1)(\ell_z + 2)$ is even, it follows that $z - \mathbf{e}_2$ sends $\frac{1}{2}(\ell_z + 1)(\ell_z + 2)$ chips to z . The same is true if $y = 1$, since then $\ell_z = n - 1$, and the origin $o = z - \mathbf{e}_2$ sends $\frac{1}{2}n(n + 1)$ chips to z .

The only other vertices that send chips to z under $F^{u'_n}$ are its left and right neighbors $z \pm \mathbf{e}_1$. Since $\ell_{z \pm \mathbf{e}_1} = \ell_z - 1$, these neighbors fire $\ell_z(\ell_z - 1)$ times. We claim that together they send $\frac{1}{2}\ell_z(\ell_z - 1)$ chips to z . To see this, note that if we fire these two vertices in parallel, they send one chip to z every two times we fire. We therefore conclude that

$$\sigma'_n(z) = \frac{1}{2}(\ell_z + 1)(\ell_z + 2) + \frac{1}{2}\ell_z(\ell_z - 1) - \ell_z(\ell_z + 1) = 1.$$

To show that $\sigma_n(z) = \sigma'_n(z)$, note first that neither z nor $z - \mathbf{e}_2$ is in C because $x = 0$. The right neighbor $z + \mathbf{e}_1$ might be in C , but since $\ell_z(\ell_z - 1)$ is even, the last chip sent from $z + \mathbf{e}_1$ by $F^{u'_n}$ does not move to z . The left neighbor $z - \mathbf{e}_1$ is in C only if $\ell_{z - \mathbf{e}_1} = \ell_z - 1 \equiv 3 \pmod{4}$, which implies $\ell_z(\ell_z - 1) \equiv 0 \pmod{4}$. Hence if $z - \mathbf{e}_1$ is in C , the last chip sent from $z - \mathbf{e}_1$ by $F^{u'_n}$ moves west. It follows that F^{u_n} and $F^{u'_n}$ fire z the same number of times and send the same number of chips to z , hence $\sigma_n(z) = \sigma'_n(z) = 1$.

Case 2: $x > 0$. To argue that $\sigma'_n(z) = 1$, as an initial step we *unfire* every vertex on the positive x -axis $B = \{(m, 0) \in \mathbb{Z}^2 : m > 0\}$ once. Since all initial rotors on B point

east, this turns all these rotors north without affecting the number of chips at z (nor at any other vertex of Q).

Now we apply $F^{u'_n}$. By firing the four neighbors of z in parallel, it is easy to see from (2) that they send one chip to z every firing round, since after every round exactly one of their rotors points to z . Hence, firing these neighbors $\ell_z(\ell_z - 1)$ times each sends $\ell_z(\ell_z - 1)$ chips to z . Since $\ell_{z-\mathbf{e}_1} = \ell_{z-\mathbf{e}_2} = \ell_z + 1$, the two neighbors $z - \mathbf{e}_1$ and $z - \mathbf{e}_2$ each fire

$$(\ell_z + 1)(\ell_z + 2) - \ell_z(\ell_z - 1) = 4\ell_z + 2$$

additional times under $F^{u'_n}$. Considering what happens when they are fired in parallel shows that they send one chip to z every two times they fire, meaning that $2\ell_z + 1$ additional chips are sent to z .

Finally, to obtain σ'_n we must fire every vertex in B once more. But since $F^{u'_n}$ fires each vertex in B an even number of times, their rotors are now pointing either north or south, so firing them once more does not affect the number of chips at z . Hence

$$\sigma'_n(z) = \ell_z(\ell_z - 1) + (2\ell_z + 1) - \ell_z(\ell_z + 1) = 1.$$

To finish the proof, we now argue that $\sigma_n(z) = \sigma'_n(z)$. First note that since $\ell_v(\ell_v + 1)$ is even for all $v \in D_{n-1}$, it follows from (2) that the last chips sent from $z \pm \mathbf{e}_1$ by $F^{u'_n}$ do not move to z . However, consider the neighbor $z + \mathbf{e}_2$. If $\ell_{z+\mathbf{e}_2} = \ell_z - 1 \equiv 3 \pmod{4}$, its final rotor points north after firing $\ell_z(\ell_z - 1)$ times, while if $\ell_z - 1 \equiv 2 \pmod{4}$, its final rotor points south. It therefore follows from the definition of C , that F^{u_n} sends one fewer chip from $z + \mathbf{e}_2$ to z than $F^{u'_n}$ in case $\ell_z \equiv 3 \pmod{4}$ and $y + 1 \geq 2$. Likewise, F^{u_n} sends one fewer chip from $z - \mathbf{e}_2$ to z than $F^{u'_n}$ in case $\ell_z \equiv 2 \pmod{4}$ and $y - 1 \geq 1$. But these are precisely the two cases when $z \in C$, hence F^{u_n} also fires z once fewer than $F^{u'_n}$. Therefore, $\sigma_n(z) = \sigma'_n(z) = 1$. \square

We remark that the rotor configuration ρ_n is obtained from ρ'_n by *reverse cycle-popping*: that is, for each directed cycle of rotors in ρ'_n , unfire each vertex in the cycle once. Reverse popping a cycle causes each vertex in the cycle to send one chip to the previous vertex, so there is no net movement of chips. Repeat the procedure if necessary until the rotor configuration is acyclic on the odometer's support. This is bound to happen after a finite number of steps, since reverse popping a cycle decreases the odometer on the cycle by 1. Let ρ''_n be the rotor configuration obtained from reverse cycle-popping, and let

$$u''_n = u'_n - c_n$$

where $c_n(z)$ is the number of times z is unfired during reverse cycle-popping. Then

$$F^{u''_n}(\rho_0, (2n^2 + 2n + 1)\delta_o) = (\rho''_n, 1_{D_n}).$$

By Lemma 5, we have

$$F^{u_n}(\rho_0, (2n^2 + 2n + 1)\delta_o) = (\rho_n, 1_{D_n}).$$

By the Strong Abelian Property (Theorem 2), it follows that $u''_n = u_n$. In particular, $c_n = 1_C$ and $\rho''_n = \rho_n$.

Proof of Theorem 1. For $0 \leq m \leq 3$, let $r_m = 1$, and for $m \geq 4$, let r_m be the unique integer n such that

$$2n(n-1) \leq m < 2n(n+1).$$

Consider a modified rotor-router aggregation defined by the growth rule

$$\tilde{A}_{m+1} = \tilde{A}_m \cup \{\tilde{z}_m\}$$

where \tilde{z}_m is the endpoint of a rotor-router walk started at the origin in $\widehat{\mathbb{Z}}^2$ and stopped on exiting $\tilde{A}_m \cap D_{r_m-1}$. For $z \in \mathbb{Z}^2$, let $v_m(z)$ be the number of times this walk visits z strictly prior to stopping at \tilde{z}_m .

Now fix $n \geq 1$, and define $\tilde{u}_n : D_{n-1} \rightarrow \mathbb{N}$ by setting

$$\tilde{u}_n(z) = \sum_{m=0}^{2n(n+1)-1} v_m(z).$$

In other words, $\tilde{u}_n(z)$ is the total number of times z fires during the formation of the cluster $\tilde{A}_{2n(n+1)}$. We will induct on n to show that $u_n = \tilde{u}_n$ for all $n \geq 1$. Since $u_n = \tilde{u}_n$ implies $A_{2n(n+1)} = \tilde{A}_{2n(n+1)} = D_n$ by Lemma 5, this proves the theorem.

The base case of the induction is immediate: $u_1 = \tilde{u}_1 = 4\delta_o$. For $n \geq 2$, in the induced subgraph D_n of $\widehat{\mathbb{Z}}^2$ with sink vertices L_n we have

$$F^{u_n}(\rho_0, (2n^2 + 2n + 1)\delta_o) = (\rho_n, 1_{D_n})$$

by Lemma 5. On the other hand,

$$F^{\tilde{u}_n}(\rho_0, (2n^2 + 2n + 1)\delta_o) = (\tilde{\rho}_n, \tilde{\sigma}_n)$$

for some rotor configuration $\tilde{\rho}_n$ on D_{n-1} and chip configuration $\tilde{\sigma}_n$ on D_n . By the inductive hypothesis, $\tilde{A}_{2n(n-1)} = D_{n-1}$. For all m such that $2n(n-1) \leq m < 2n(n+1)$, since $r_m = n$, it follows that \tilde{z}_m is the endpoint of a rotor-router walk in $\widehat{\mathbb{Z}}^2$ stopped on exiting D_{n-1} . This implies that $\tilde{\sigma}_n = 1$ on the set D_{n-1} . Moreover, since ρ_0 is acyclic, $\tilde{\rho}_n$ is acyclic (each rotor points in the direction a chip last exited). The Strong Abelian Property (Theorem 2) now gives $u_n = \tilde{u}_n$, which completes the inductive step. \square

4 Concluding Remarks

Theorems 1 and 2 raise several further questions. We treat these in order of increasing generality (and, we suspect, increasing difficulty!).

Intermediate cluster shapes

It is natural to ask about the shape of the cluster A_m when m is not of the form $2n(n+1)$. As m increases from $2n(n-1)$ to $2n(n+1)$, the sites in layer L_n appear to fill up in a predictable order.

General rotor configurations

We believe that for a general initial rotor configuration ρ on the layered square lattice $\widehat{\mathbb{Z}}^2$, the shape of the aggregate remains very close to a diamond. How close? Does there exist an absolute constant c such that for all ρ and n ,

$$D_{n-c} \subset A_{2n(n+1)} \subset D_{n+c} ?$$

As evidence that the initial rotor configuration should not change the shape by very much, consider the following modification of our aggregation model: stop each chip when it reaches either an empty site or a site containing just one other chip. That is, in the modified model it is legal to fire a vertex only if it contains at least 3 chips. The proof of Lemma 5 shows that starting with $4n^2 + 4n + 2$ chips at the origin, the odometer $2u'_n$ leads to the chip configuration $2 \cdot 1_{D_n}$. But observe that $2u'_n$ fires every vertex in D_{n-1} a multiple of 4 times. Therefore, *the final chip configuration does not depend on the initial rotor configuration, and the initial and final rotor configurations are equal.* Following the proof of Theorem 1, the Strong Abelian Property now implies that for any acyclic initial rotor configuration, the odometer for any complete legal firing sequence is $2u'_n$. Hence, the aggregate grows as a perfect diamond of height 2.

In fact, using reverse cycle-popping as in the remark following Lemma 5, one can show that the modified aggregation model yields a perfect diamond of height 2 for *any* initial rotor configuration.

Conceptual proofs

The Strong Abelian Property (Theorem 2) can be viewed as a tool for converting simulations into proofs. Specifically, if simulation of a rotor-router aggregation model reveals an odometer function with simple structure, yielding a conjectural explicit formula based on the behavior for small values of n , then the Strong Abelian Property provides a way of proving that this formula holds for all n . Unfortunately, the proofs produced in this way tend to be unenlightening. They provide formal verification, but not understanding. One would like to have a way of predicting the behavior of a growth model that does not rely on first explicitly simulating it; or an approach that provides a conceptual explanation of its behavior rather than a technical verification. In what other situations should we expect a “zero discrepancy” result like Theorem 1?

Local regularities

A final challenge concerns the (more typical) case when the odometer function reveals intriguing local regularities, but is beyond the reach of a global explicit formula. The odometer function of the usual rotor-router aggregation in \mathbb{Z}^2 has this character. Simulations indicate that near certain special points (the preimages of a square lattice in the complex plane under the conformal map $z \mapsto 1/z^2$) the odometer is exactly equal to a quadratic function of the coordinates. These points are visible in the image produced recently by a new large-scale simulation algorithm [FL10]. They lie in regions of the picture

where the final rotors all point in the same direction (or more generally, alternate in a simple periodic fashion).

The abelian sandpile model has a similar phenomenon, wherein the final state and odometer function appear to have a simple behavior near the boundary but increasingly complex and intractable behavior as one moves toward the origin [FLP10]. Simulations show the power of rotor-router and sandpile systems to form large-scale regular structures using simple local rules, but the underlying mechanism for this kind of pattern formation remains poorly understood. Extending the methods used here to prove local rather than global exact formulas for the odometer function would be a possible approach to understanding pattern formation in these systems.

Acknowledgements

We thank an anonymous referee, whose suggestions have significantly improved our presentation.

References

- [CS06] J. Cooper and J. Spencer. Simulating a random walk with constant error. *Combin. Probab. Comput.* **15**: 815–822 (2006). arXiv:math/0402323
- [DF91] P. Diaconis and W. Fulton. A growth model, a game, an algebra, Lagrange inversion, and characteristic classes. *Rend. Sem. Mat. Univ. Politec. Torino* **49**(1): 95–119 (1991).
- [DFS08] B. Doerr, T. Friedrich and T. Sauerwald. Quasirandom rumor spreading. *19th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2008), 773–781.
- [FLP10] A. Fey, L. Levine and Y. Peres. Growth rates and explosions in sandpiles. *J. Stat. Phys.* **38**: 143–159 (2010). arXiv:0901.3805
- [FGS10] T. Friedrich, M. Gairing and T. Sauerwald. Quasirandom load balancing. *21st Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2010), 1620–1629.
- [FL10] T. Friedrich and L. Levine. Fast simulation of large-scale growth models. arXiv:1006.1003 See also the rotor-router image of 10 billion pixels navigable by google maps: <http://rotor-router.mpi-inf.mpg.de/10Bio/>
- [HP09] A. E. Holroyd and J. Propp. Rotor walks and Markov chains. In *Algorithmic Probability and Combinatorics*, American Mathematical Society, 2010. arXiv:0904.4507.
- [Ka07] W. Kager. Reflected Brownian motion in generic triangles and wedges. *Stoch. Process. Appl.* **117**(5): 539–549 (2007). arXiv:math/0410007
- [KL10] W. Kager and L. Levine. Diamond aggregation. *Math. Proc. Cambridge Phil. Soc.* **149**(2): 351–372 (2010). arXiv:0905.1361
- [LBG92] G. F. Lawler, M. Bramson and D. Griffeath. Internal diffusion limited aggregation. *Ann. Probab.* **20**(4): 2117–2140 (1992).

- [LL09] I. Landau and L. Levine. The rotor-router model on regular trees. *J. Combin. Theory. A* (2009). arXiv:0705.1562.
- [LP08] L. Levine and Y. Peres. Spherical asymptotics for the rotor-router model in \mathbb{Z}^d . *Indiana Univ. Math. J.* **57**(1): 431–450 (2008). arXiv:math/0503251
- [LP09] L. Levine and Y. Peres. Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile. *Potential Anal.* **30**: 1–27 (2009). arXiv:0704.0688
- [PDDK96] 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).