RESEARCH STATEMENT: DYNAMICS AND COMPUTATION IN ABELIAN NETWORKS

LIONEL LEVINE

The broad goal of my research is to understand how large-scale forms and complex patterns emerge from simple local rules. My approach is to analyze specific mathematical models which isolate just one or a few features of pattern formation. A good model is one that captures some aspect of "scaling up" from local to global, yet is tractable enough to prove theorems about! Some of the models I'm working on include

- Abelian sandpiles [BTW87], a model of self-organization and pattern formation.
- Parallel chip-firing [BG92], a model of mode-locking and synchronization.
- Internal DLA [LBG92], a model of fluid flow and random interfaces.
- Rotor-router aggregation [LP09], a model based on derandomizing random walks.

Abelian networks, invented by Deepak Dhar [Dha06], tie these models and many others together in a common mathematical framework. The mathematics involved is a mixture of probability and combinatorics, draws on techniques that originated in the study of partial differential equations, and has close connections with statistical physics and computer science.

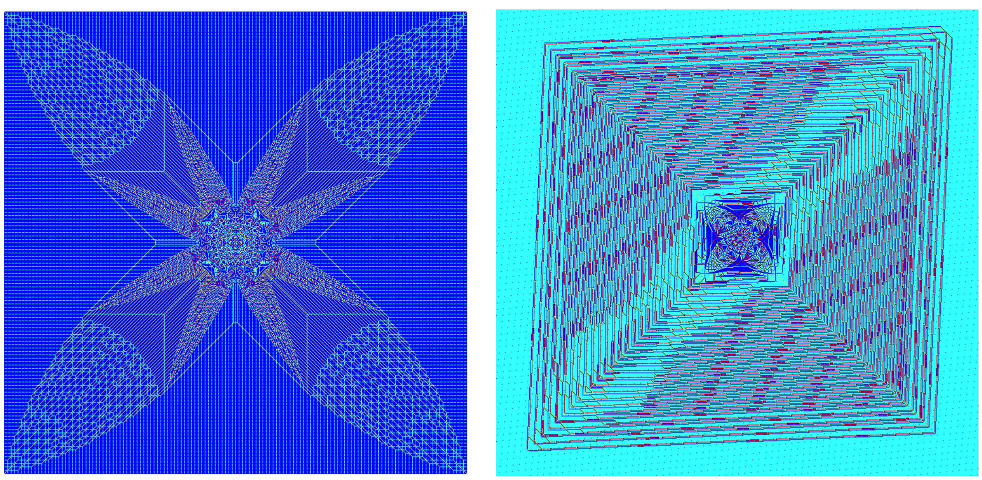

FIGURE 1. Pattern formation in a stable sandpile (left) and an exploding sandpile (right) in $\mathbb{Z}^2$. Each site is colored according to the number of sand grains present.

Synchronization and pattern formation have fascinated generations of physicists and mathematicians. One of the simplest systems known that forms complex patterns from local rules is the Bak-Tang-Wiesenfeld abelian sandpile model [BTW87], pictured in Figure 1. Over the past few years, beginning with the "least action principle" of [FLP10], my coauthors Fey, Friedrich, Kager, Peres and I have developed new variational tools for the abelian sandpile and related models. I survey a selection of results and open problems in §1-§3 and discuss some new directions in §4.

1. Mode-locking and synchronization in parallel chip-firing

Mode-locking, defined as the "tendency of weakly coupled oscillators to synchronize their motion" [Lag92], is a widespread phenomenon in dynamical systems in the physical and biological sciences.

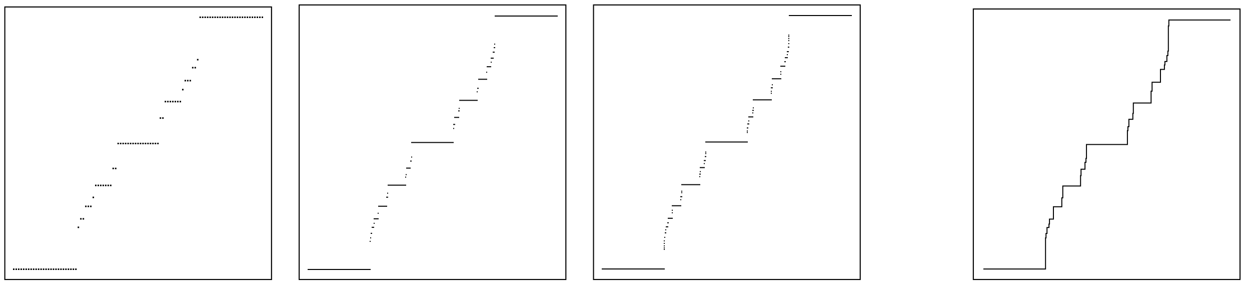

$n = 100$ | $n = 1000$ | $n = 10000$ | devil's staircase

$a_n(\lambda)$ $\xrightarrow{n\to\infty}$ $s(\lambda)$

FIGURE 2. As the density of chips λ increases, the activity $a_n(\lambda)$ of a parallel chip-firing system on K_n has long intervals of constancy punctuated by sudden jumps. According to Theorem 1, as $n \to \infty$ the function $a_n(\lambda)$ tends to a limit $s(\lambda)$ whose graph has a flat horizontal segment at height y for each rational number y in the interval $[0, 1]$. Within each box above, λ runs from 0 to 1 on the horizontal axis, and the activity runs from 0 to 1 on the vertical axis.

A familiar example is the dark side of the moon, which is never visible from earth because the moon's two modes of oscillation — its revolution around the earth and rotation about its own axis — are synchronized. This synchronization arose over millions of years from a weak coupling caused by action of the earth's tidal forces on the moon. A telltale sign of mode-locking in a dynamical system is the appearance of a *devil's staircase* dependence of an observable on a parameter (Figure 2): as the parameter varies, the observable has long intervals of constancy punctuated by sudden jumps.

In [Lev10a] I studied *parallel chip-firing* as a combinatorial model of mode-locking. This is a discrete dynamical system whose state space consists of functions $\sigma : V \to \mathbb{N}$, where V is the vertex set of a fixed finite undirected graph G. We interpret $\sigma(v)$ as the number of "chips" at vertex v. If $\sigma(v)$ is at least the number of edges incident to v, then v is called *unstable* and can *fire* by sending one chip along each incident edge. A single time step of the system consists of firing all unstable vertices in parallel. For background, see [BG92, Pri94].

Parallel chip-firing has a natural control parameter, the average number of chips per vertex $\lambda = \frac{1}{n}\sum_v \sigma(v)$. Here $n = |V|$ is the number of vertices. It also has a natural observable, the *activity*, defined by

$$a = \lim_{t\to\infty} \frac{\alpha_t}{nt}$$

where α_t is the total number of firings before time t. In numerical experiments, Bagnolli et al. [BCFV03] discovered the telltale devil's staircase: as we vary λ by adding additional chips to the system, the activity a remains constant for long intervals punctuated by sudden jumps. My motivation in [Lev10a] was to prove a theorem that would explain these experimental findings. For each $n \geq 2$ let σ_n be a stable chip configuration on the complete graph K_n, and let $a_n(\lambda)$ be the activity of the configuration $\sigma_n + \lambda n$.

Theorem 1. [Lev10a] *Under mild hypotheses on σ_n there exists a function $s : [0, 1] \to [0, 1]$ such that for each $\lambda \in [0, 1]$*

$$a_n(\lambda) \to s(\lambda) \qquad \text{as } n \to \infty.$$

The function s is continuous and nondecreasing. Moreover,

- *If $y \in [0, 1]$ is irrational, then $s^{-1}(y)$ is a single point.*
- *If $y \in [0, 1]$ is rational, then $s^{-1}(y)$ is an interval of positive length.*

I prove this theorem in [Lev10a] by constructing, for any chip configuration σ on K_n, a circle map $f_\sigma : S^1 \to S^1$ whose Poincaré rotation number equals the activity of σ. This construction reveals a connection between certain discrete and continuous dynamical systems ripe for further exploration. I believe that the right way to extend this correspondence to general graphs is by replacing the circle with a higher-dimensional torus.

Problem 1. *Generalize the circle map construction of* [Lev10a] *to model parallel chip-firing on a general graph G by iteration of a torus map $T^n \to T^n$.*

A second manifestation of mode-locking is the prevalence of *short period attractors*. Bitar made a striking conjecture about parallel chip-firing in 1989.

Conjecture 1. [Bit89] *Any parallel chip-firing configuration on a connected undirected n-vertex graph G has eventual period at most n.*

To see what is striking about this conjecture, note that as a deterministic system with a finite state space, parallel chip-firing must *eventually* settle into a periodic cycle; but the number of distinct states is superexponential in n, so in principle the period could be extremely long. Bitar's conjecture posits that such long periods do not occur.

Kiwi et al. [KNTG94] constructed a counterexample to Bitar's conjecture by making the graph extremely thinly connected: G is a disjoint union of cycles each connected by a single edge to a central hub vertex. To date the only known counterexamples are of this type.

Question 1. *How "well connected" must G be in order for Bitar's conjecture to hold?*

Theorem 2 shows that Bitar's conjecture holds when G is a complete graph.

Theorem 2. [Lev10a]

- *Every parallel chip-firing configuration on K_n has eventual period at most n.*
- *Every parallel chip-firing configuration on K_n with strictly between $n^2 - n$ and n^2 chips has eventual period 2.*
- *For each integer $k = 1, \dots, n$ there exists a parallel chip-firing configuration on K_n of period k.*

Jiang [Jia10] has extended the methods of [Lev10a] to prove Bitar's conjecture for complete bipartite graphs, and recently he proposed to me some novel ideas for attacking the case of grid graphs $\mathbb{Z}_a \times \mathbb{Z}_b$.

2. Abelian Sandpile Model

In joint work with Fey and Peres, I prove bounds on the growth rate of the abelian sandpile model in the integer lattice $\mathbb{Z}^d$. The model starts from a stable *background* configuration in which each site $x \in \mathbb{Z}^d$ ($d \geq 1$) has a pile of $\sigma(x) \leq 2d - 1$ chips. To this background, n chips are added at the origin. Typically, n is large. We *stabilize* this configuration by *toppling* every unstable site; that is, every site with at least $2d$ chips gives one chip to each of its neighbors, until there are no more unstable sites. For more background, see [BTW87, BLS91, Dha90].

Theorem 3. [LP09, FLP10] *For any $h \leq 2d-2$, if h chips start at every site of $\mathbb{Z}^d$ and n additional chips start at the origin, then the number of sites that topple is of order n, and the diameter of the set of sites that topple is of order $n^{1/d}$.*

The main tool used in the proof is the *least action principle* characterizing the odometer function $u(x)$, which is defined as the number of times x topples. We say that a function $f : \mathbb{Z}^d \to \mathbb{Z}$ is *stabilizing* for a sandpile σ if the inequality

$$\sigma + \Delta f \leq 2d - 1$$

holds pointwise on $\mathbb{Z}^d$. Here Δ is the discrete Laplacian on functions on $\mathbb{Z}^d$, defined by $\Delta f(x) = \sum_{y \sim x} f(y) - 2df(x)$, where the sum is over the $2d$ neighbors of x.

Lemma 4. [FLP10] *(Least Action Principle) The odometer function is the pointwise minimum of all nonnegative stabilizing functions $f : \mathbb{Z}^d \to \mathbb{Z}$.*

The least action principle states that sandpiles are "efficient" in a certain strong sense: each vertex performs exactly the minimum amount of work required of it in order to produce a globally stable configuration.

In [FLP10] we extended Theorem 3 to certain backgrounds in which an arbitrarily high proportion of sites start with $2d-1$ chips. On the other hand, the next result shows that the density of chips in the background is not the only controlling parameter: if just a few extra chips are placed at random on the background $2d-2$, then it has a dramatically different behavior.

Theorem 5. [FLP10] *Fix $\epsilon > 0$, and let $(\beta(x))_{x\in\mathbb{Z}^d}$ be independent Bernoulli random variables with $\mathbb{P}(\beta(x) = 1) = \epsilon$. Then with probability 1, there exists $n < \infty$ such that adding n chips at the origin to the background $2d-2+\beta$ causes every site in $\mathbb{Z}^2$ to topple infinitely often.*

Certain notable features of the patterns formed by sandpiles call out for a mathematical explanation.

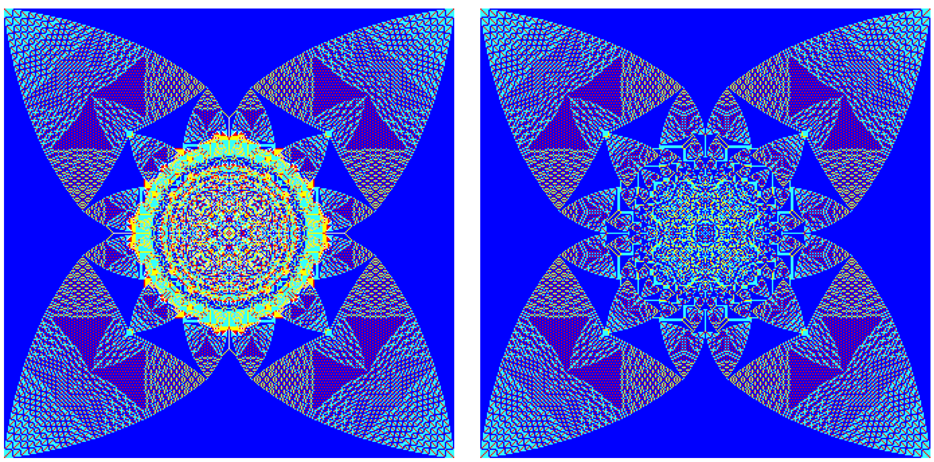

slice of a 3-dimensional sandpile ($n = 5\,000\,000$, $h = 4$) | 2-dimensional sandpile ($n = 47\,465$, $h = 2$)

FIGURE 3. Illustration of the dimensional reduction conjecture. Left: The $z = 0$ slice of a sandpile in $\mathbb{Z}^3$. Right: a sandpile in $\mathbb{Z}^2$. Each site is colored according to the number of chips present. Remarkably, the two pictures agree pixel-for-pixel except in a region near the center.

Dimensional reduction conjecture. Dhar observed that despite the difference between sandpile dynamics in different dimensions, sandpiles in $\mathbb{Z}^d$ intersected with a coordinate hyperplane look remarkably similar to sandpiles in $\mathbb{Z}^{d-1}$ (Figure 3). In [FLP10] we formalized Dhar's observation as a *dimensional reduction conjecture*: Amazingly, large portions of the two pictures in Figure 3 match up exactly pixel-for-pixel. At the same time, our conjecture was cumbersome to state due to the need to carve out exceptions for regions where the two pictures do not agree. Karl Mahlburg recently discovered a variant of the initial configuration which causes the two pictures match up exactly, with only a single exceptional pixel at the origin. For this variant, I am working with Holroyd and Mahlburg to prove the dimensional reduction conjecture.

Unexpected symmetries. Mahlburg's configuration has another remarkable property: extra symmetries appear in the final configuration that were not present in the initial configuration. Using the least action principle, we have succeeded in proving one of these extra symmetries in the two-dimensional case and are investigating whether they also appear in higher dimensions.

3. Aggregation: Limit shapes, fluctuations, and fast simulation

Starting with n particles at the origin in $\mathbb{Z}^2$, let each particle in turn perform simple random walk until reaching a site where no other particles are present. The resulting random set $A(n)$, consisting of n occupied sites, is called *internal diffusion-limited aggregation* (IDLA). Lawler, Bramson and Griffeath [LBG92] showed that with high probability $A(n)$ is close to a disk. In joint work with Jerison and Sheffield, we address the question "how close"?

Theorem 6. [JLS10] *Let $\mathbf{B}_r$ denote the set of lattice points in the disk of radius r centered at 0. There is an absolute constant C such that*

$$\mathbb{P}\left\{\mathbf{B}_{r-C\log r} \subset A(\pi r^2) \subset \mathbf{B}_{r+C\log r} \text{ for all sufficiently large } r\right\} = 1. \tag{1}$$

This answers a question posed by Lawler [Law95] and confirms numerical predictions of Meakin and Deutch [MD86] from over 20 years ago, who wrote that it is "of some fundamental significance to know just how smooth a surface formed by diffusion limited processes may be." Theorem 6 shows that such surfaces are extremely smooth: their fluctuations are visible only at the logarithmic scale.

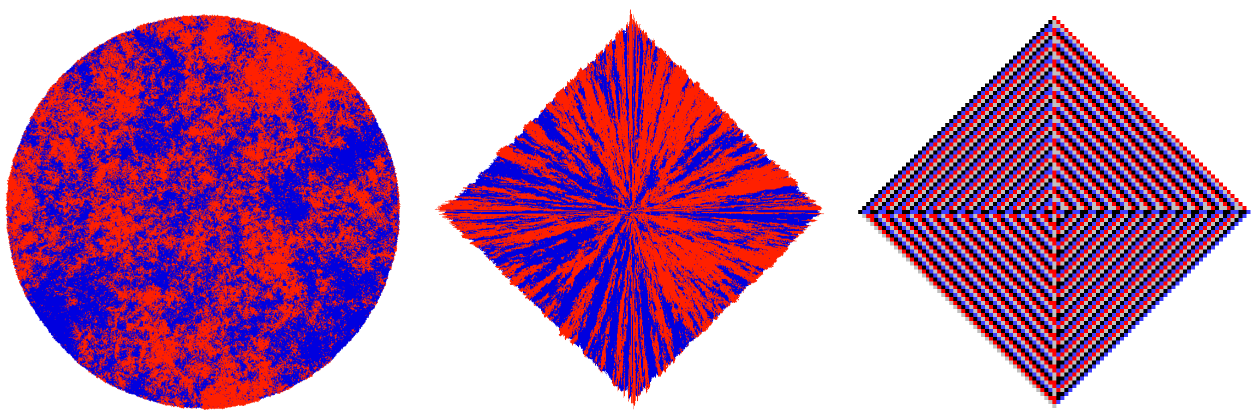

(a) (b) (c)

Figure 4. (a) The IDLA cluster built from simple random walk in $\mathbb{Z}^2$ has only logarithmic error from a disk, by Theorem 6. (b) The IDLA cluster built from a uniformly layered walk grows as a diamond instead of a disk. (c) Rotor-router cluster built from a uniformly layered walk yields a perfect diamond by Theorem 7.

Our paper [JLS10] is the first in a planned series of three. The second in the series proves weak convergence of *averaged* IDLA fluctuations to a variant of the Gaussian free field [Sh07]; these fluctuations are shown in Figure 4a, where points that arrive early to the cluster are colored red and those that arrive late are colored blue. The final paper in the series will prove an analogue of Theorem 6 in dimensions $d \geq 3$.

Diamond aggregation. In joint work with Kager, I defined the class of *uniformly layered walks* in $\mathbb{Z}^2$ and studied IDLA based on these walks. This work was motivated by the surprising observation that a suitable change to the transition probabilities of simple random walk only at sites on the x- and y-axes transforms the asymptotic shape from a disk into a diamond (Figure 4b). In [KL10a] we prove a result like Theorem 6 for these walks, with the disk $\mathbf{B}_r$ replaced by the set $\mathbf{D}_r$ of points $(x, y) \in \mathbb{Z}^2$ such that $|x| + |y| \leq r$. We show that IDLA based on certain uniformly layered walks – those with a drift toward the origin – has fluctuations at most of order $\log r$, while other such walks with drift away from the origin produce much larger fluctuations, of order $\sqrt{r}$ up to logarithmic factors.

James Propp proposed a derandomization of IDLA known as *rotor-router aggregation*, in which the sequence of successive exits from each vertex is periodic instead of random. It turns out that certain instances of rotor-router aggregation corresponding to uniformly layered walks yield not just a shape close to the diamond $\mathbf{D}_r$ as might be expected, but a perfect replica of $\mathbf{D}_r$ (Figure 4c).

Theorem 7. [KL10b] *There is a rotor configuration ρ_0, such that uniformly layered rotor-router aggregation with rotors initially configured as ρ_0 satisfies*

$$A(2r^2 + 2r + 1) = \mathbf{D}_r \qquad \text{for all } r \geq 0.$$

Theorem 7 represents an extreme of discrepancy reduction of the sort studied in [CS06, CDST07, DF09, HP10]: derandomization removes not just most but *all* of the fluctuations from the random process.

Fast simulation. Building on the work of Moore and Machta [MM00], in joint work with Friedrich I develop a new algorithm for simulating IDLA and rotor-router aggregation without simulating every step of every walk. The algorithm is based on the following theorem of [FL10] characterizing the odometer function, which measures how many particles are emitted from each vertex.

Theorem 8. [FL10] *The odometer is the unique function $u : V \to \mathbb{N}$ with the following properties: u has finite support, and the chip configuration σ_* and rotor configuration ρ_* obtained by forcing each vertex x to fire $u(x)$ times obey the "*Three No's*:"*

- *No hills: $\sigma_* \leq 1$ everywhere;*
- *No holes: $\sigma_* \equiv 1$ on the support of u;*
- *No cycles: ρ_* is acyclic on the support of u.*

In [FL10], we give an algorithm that takes an approximation u_1 to the odometer function as input, cancels out hills and holes a multiscale annihilation process, and eliminates cycles using Wilson cycle-popping [Wil96] to arrive at the exact odometer function u. The choice of initial approximation u_1 rests on earlier work joint with Peres [LP09, LP10a].

Propp circle. Friedrich and I have implemented our algorithm to generate a rotor-router aggregate ("Propp circle") of size 10^{10}, exceeding previous simulations by over three orders of magnitude. The resulting 10 gigapixel image, which we have posted online [FL10], is so large that it requires a google maps interface to navigate. This image reveals the intricate patterns formed by the final rotors in unprecedented detail, and promises to be a rich source of new conjectures.

The Ostojic heuristic. The patterns formed by sandpile aggregation contain many large "patches," visible in Figure 3, where the number of chips is constant or periodic. Ostojic [Ost03] gave a heuristic, involving the conformal map $\zeta \mapsto 1/\zeta^2$, for the locations and certain features of these patches. Dhar et al. [DSC09] carried this idea further for sandpiles on certain directed graphs obtained from orientations of the square grid $\mathbb{Z}^2$. Converting these ideas into fully rigorous proofs remains a substantial challenge.

Intriguingly, although it was developed for the sandpile model, the Ostojic heuristic appears to apply even more precisely to rotor-router aggregation. For the aggregate of size $n = \pi r^2$, which approximates a disk of radius r, simulations indicate that near certain special points — points $\zeta \in \mathbb{C}$ such that the real and imaginary parts of r^2/ζ^2 are rational numbers with small denominators — the odometer function has the form

$$u(x, y) = Q_{r,\zeta}(x, y) + P_{r,\zeta}(x, y)$$

where $Q_{r,\zeta}$ is a quadratic function of the coordinates and $P_{r,\zeta}$ is a periodic function. These points are visible in the image produced by our new large-scale simulation algorithm [FL10]. They lie in regions of the picture ("studs") where the final rotors alternate in a simple periodic fashion.

Problem 2. *Use the Ostojic heuristic to make a mathematically precise prediction about the studs in rotor-router aggregation: where are they located, how large are they, and what are the functions $Q_{r,\zeta}$ and $P_{r,\zeta}$ in terms of r and ζ?*

Proving local regularities. Does a prediction coming from Problem 2 bring us any closer to a proof? The "strong abelian property" used in [KL10b] to prove Theorem 7 can be viewed as a tool for converting an exact prediction for the odometer function into a proof. In the most interesting cases, the odometer function reveals intriguing local regularities but is beyond the reach of a global exact formula.

Problem 3. *Adapt the methods of* [FLP10] *(least action principle) and* [KL10b] *(strong abelian property) to prove local rather than global exact formulas for the odometer function.*

4. Future Directions: Abelian Networks

What the diverse models discussed in §1-3 have in common is that they are all examples of *abelian networks*. An abelian network is a system of communicating finite automata satisfying a certain local commutativity condition. Each finite automaton, or *processor*, lives at a vertex of a directed graph G and communicates with the processors at neighboring vertices.

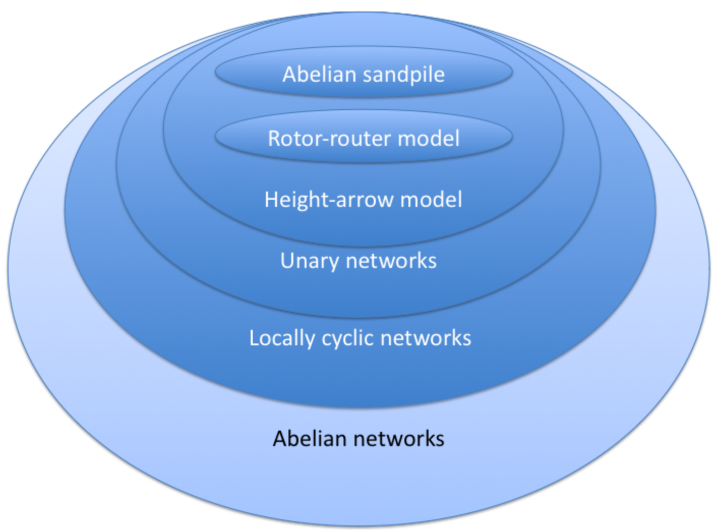

Figure 5. Venn diagram illustrating several classes of deterministic abelian networks, ranging from most specific at the top to most general at the bottom.

Deepak Dhar, the inventor of abelian networks, called them "abelian distributed processors" [Dha06]. (The shorter term "abelian network," which I coined, is useful when one wants to refer to several communicating processors as a single entity.) In a series of papers beginning with [Lev10c], I plan to develop a theory of abelian networks from three complementary viewpoints.

- From a dynamical point of view, abelian networks produce intricate large-scale patterns using only simple local rules.
- From a computational point of view, an abelian network yields the same output in the same number of steps regardless of the order of events at its individual nodes.
- From an algebraic point of view, an abelian network associates an isomorphism invariant, in the form of an abelian group or commutative algebra, to its underlying graph.

Abelian networks represent a vast unexplored territory: the class of such networks is fairly large, and only a very few particular examples have been studied in any detail (Figure 5). In particular, the abelian sandpile and rotor-router models, as well as their common generalizations studied by Eriksson [Eri96] and Dartois and Rossin [DR04], all belong to the subclass of *unary networks*: each automaton in the network reads input from a one-letter alphabet. Informally, a unary network on a graph G is a system of local rules by which *indistinguishable* particles move around on the vertices of G.

Question 2. *How much more general are abelian networks than unary networks? Can an arbitrary abelian network be modeled by a unary network on a larger graph?*

In work in progress [Lev10c] I prove a least action principle for abelian networks, which in particular implies that if a network halts on a given input, then neither the final output nor the run time depends on the order in which processors act. I am also working on a polynomial-time algorithm to decide whether a given abelian network halts on all inputs. A refinement of this question is the halting problem for abelian networks, which asks whether a given network halts on a *given* input.

Question 3. *Is the halting problem for abelian networks decidable in polynomial time?*

The requirement that a distributed network produce the same output regardless of the order in which its processors act would seem to place a severe restriction on the kinds of algorithmic tasks it can perform. Yet abelian networks can perform some highly nontrivial tasks, such as solving certain integer programs. One form of the least action principle of [FLP10] asserts that the odometer function u for a sandpile σ on a directed graph $G = (V, E)$ solves the integer program

$$\text{Minimize } \sum_{i \in V} u(i) \qquad \text{subject to } u \geq 0 \text{ and } \Delta^* u \leq \delta - 1 - \sigma \tag{2}$$

where δ is the degree vector and Δ^* is the adjoint Laplacian matrix of G. (Lemma 4 is the case $G = \mathbb{Z}^d$.) Tseng [Tse90] analyzed a class of diagonally dominant linear programs that includes (2). These programs involve multiple Laplacian-type matrices $\Delta^1, \ldots, \Delta^k$.

Problem 4. *Design k-ary abelian networks that solve integer programs analogous to the linear programs of* [Tse90].

Informally, these k-ary networks may be viewed as sandpile models in which the chips come in k distinct "colors." Chips of color j obey sandpile dynamics corresponding to the Laplacian Δ^j, with firing restrictions based on the presence of particles of other colors.

Stochastic Abelian Networks. By allowing randomness in the transition function of an abelian network, one obtains a class of models that includes branching random walks, stochastic sandpiles [RS09] and the activated random walkers model [DRS10] as well as classical Markov chains. Viewing a Markov chain as a stochastic abelian network leads naturally to a larger class of "locally Markov walks," which include excited walks [BW03] and directed edge-reinforced walks such as those studied in [KR02].

Locally Markov walks obey a weak version of the Markov property: if the walker's current location is x, then the distribution of its next step depends only on the history of previous exits from x. These walks provide a way to interpolate between random walks (in which successive exits from a vertex are i.i.d.) and their deterministic analogues, rotor walks (in which successive exits are periodic, [HP10]). For concreteness we describe one example, the *p-rotor walk* on $\mathbb{Z}$. Each integer n has a *rotor* pointing right or left. The walker repeatedly takes steps according to the following rule: if the walker's current at location is n, then the rotor at n flips direction with probability p, and the walker then moves to $n-1$ or $n+1$ as indicated by the rotor at n.

Question 4. *Is the p-rotor walk on $\mathbb{Z}$ recurrent for all $0 < p < 1$ and any initial configuration of the rotors?*

This question sits at the tip of a rather large iceberg. In addition to questions of recurrence and transience, one can ask for martingales, stationary distributions, mixing times and scaling limits for locally Markov walks.

References

[BCFV03] F. Bagnoli, F. Cecconi, A. Flammini, and A. Vespignani, Short-period attractors and non-ergodic behavior in the deterministic fixed-energy sandpile model, *Europhys. Lett.* **63**:512–518, 2003.

[BTW87] P. Bak, C. Tang and K. Wiesenfeld, Self-organized criticality: an explanation of the $1/f$ noise, *Phys. Rev. Lett.* **59**(4):381–384, 1987.

[BW03] I. Benjamini and D. B. Wilson, Excited random walk, *Elect. Comm. Probab.* **8**:86–92, 2003.

[Bit89] J. Bitar, Juegos Combinatoriales en Redes de Autómatas, Tesis Ingeniero Matemático, Dept. Ingeniería Matemática, U. de Chile, 1989.

[BG92] J. Bitar and E. Goles, Parallel chip firing games on graphs, *Theoretical Comp. Sci.* **92**(2):291–300, 1992.

[BLS91] A. Björner, L. Lovász and P. Shor, Chip-firing games on graphs, *European J. Combin.* **12**(4):283–291, 1991.

[CS06] J. Cooper and J. Spencer, Simulating a random walk with constant error, *Combin. Probab. Comput.* **15**:815–822, 2006.

[CDST07] J. Cooper, B. Doerr, J. Spencer and G. Tardos, Deterministic random walks on the integers, *European J. Combin.* **28**:2072–2090, 2007.

[DR04] A. Dartois and D. Rossin, Height-arrow model, *Formal Power Series and Algebraic Combinatorics*, 2004.

[Dha90] D. Dhar, Self-organized critical state of sandpile automaton models, *Phys. Rev. Lett.* **64**:1613–1616, 1990.

[Dha06] D. Dhar, Theoretical studies of self-organized criticality, *Physica A* **369**:29–70, 2006.

[DSC09] D. Dhar, T. Sadhu and S. Chandra, Pattern formation in growing sandpiles, *Europhysics Lett.* **85**:48002, 2009.

[DF91] P. Diaconis and W. Fulton, A growth model, a game, an algebra, Lagrange inversion, and characteristic classes, *Rend. Sem. Mat. Univ. Pol. Torino* **49**(1):95–119, 1991.

[DRS10] R. Dickman, L. Rolla and V. Sidoravicius, Activated random walkers: facts, conjectures andchallenges, **138**(1-3):126–142, 2010.

[DF09] B. Doerr and T. Friedrich, Deterministic random walks on the two-dimensional grid, *Combin. Probab. Comput.* **18**(1-2):123–144, 2009.

[Eri96] K. Eriksson, Chip-firing games on mutating graphs, *SIAM J. Discrete Math.* **9**(1):118–128, 1996.

[FLP10] A. Fey, L. Levine and Y. Peres, Growth rates and explosions in sandpiles, *J. Stat. Phys.* 138:143–159, 2010. http://arxiv.org/abs/0901.3805

[FLW10a] A. Fey, L. Levine and D. B. Wilson, Driving sandpiles to criticality and beyond, *Phys. Rev. Lett.* **104**:145703, 2010. http://arxiv.org/abs/0912.3206

[FLW10b] A. Fey, L. Levine and D. B. Wilson, The approach to criticality in sandpiles, *Phys. Rev. E*, **82**:031121, 2010. http://arxiv.org/abs/1001.3401

[FL10] T. Friedrich and L. Levine, Fast simulation of large-scale growth models, 2010. http://arxiv.org/abs/1006.1003. See the resulting "Propp circle" of 10^{10} pixels navigable by Google Maps at http://rotor-router.mpi-inf.mpg.de.

[HLR10] C. J. Hillar, L. Levine and D. Rhea, Equations solvable by radicals in a uniquely divisible group, 2010. http://arxiv.org/abs/1004.5239

[HLMPPW08] A. E. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp and D. B. Wilson, Chip-firing and rotor-routing on directed graphs, *In and out of equilibrium 2*, 331–364, Progr. Probab. **60**, Birkhäuser, 2008. http://arxiv.org/abs/0801.3306

[HP10] A. E. Holroyd and J. G. Propp, Rotor walks and Markov chains, in *Algorithmic Probability and Combinatorics*, American Mathematical Society, 2010. http://arxiv.org/abs/0904.4507

[JLS10] D. Jerison, L. Levine and S. Sheffield, Logarithmic fluctuations for internal DLA, 2010. http://arxiv.org/abs/1010.2483

[Jia10] T.-Y. Jiang, On the period lengths of the parallel chip-firing game, 2010. http://arxiv.org/abs/1003.0943

[KL10a] W. Kager and L. Levine, Diamond Aggregation, *Math. Proc. Cambridge Phil. Soc.* **149**(2): 351–372, 2010. http://arxiv.org/abs/0905.1361

[KL10b] W. Kager and L. Levine, Rotor-router aggregation on the layered square lattice, 2010. http://arxiv.org/abs/1003.4017

[KR02] M. S. Keane and S. W. W. Rolles, Tubular recurrence, *Acta Math. Hungarica* **97**(3):207–221, 2002.

[KNTG94] M. A. Kiwi, R. Ndoundam, M. Tchuente and E. Goles, No polynomial bound for the period of the parallel chip firing game on graphs, *Theoretical Comp. Sci.* **136**(2):527–532, 1994.

[Lag92] J. C. Lagarias, Number theory and dynamical systems, *Proc. Symp. Applied Math.* **46**:35–72, 1992.

[LL09] I. Landau and L. Levine, The rotor-router model on regular trees, *J. Combin. Theory A* **116**: 421–433, 2009. http://arxiv.org/abs/0705.1562

[LBG92] G. F. Lawler, M. Bramson and D. Griffeath, Internal diffusion limited aggregation, *Ann. Probab.* **20**(4):2117–2140, 1992.

[Law95] G. F. Lawler, Subdiffusive fluctuations for internal diffusion limited aggregation, *Ann. Probab.* **23**(1):71–86, 1995.

[Lev10a] L. Levine, Parallel chip-firing on the complete graph: devil's staircase and Poincaré rotation number, *Ergodic Theory Dynam. Systems*, to appear. http://arxiv.org/abs/0811.2800

[Lev10b] L. Levine, Sandpile groups and spanning trees of directed line graphs, *J. Comb. Theory A*, to appear. http://arxiv.org/abs/0906.2809

[Lev10c] L. Levine, Abelian networks I: foundations and examples. In preparation, 2010.

[LP09] L. Levine and Y. Peres, Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile, *Potential Anal.* **30**:1–27, 2009. http://arxiv.org/abs/0704.0688

[LP10a] L. Levine and Y. Peres, Scaling limits for internal aggregation models with multiple sources, *J. d'Analyse Math.* **111**:151–219, 2010. http://arxiv.org/abs/0712.3378

[LP10b] L. Levine and J. Propp, What is a sandpile?, *Notices Amer. Math. Soc.* 57(8):976–979, 2010.

[MD86] P. Meakin and J. M. Deutch, The formation of surfaces by diffusion-limited annihilation, *J. Chem. Phys.* **85**:2320, 1986.

[MM00] C. Moore and J. Machta, Internal diffusion-limited aggregation: parallel algorithms and complexity, *J. Stat. Phys.*, **99**(3-4): 661–690, 2000. http://arxiv.org/abs/cond-mat/9909233

[Ost03] S. Ostojic, Patterns formed by addition of grains to only one site of an abelian sandpile, *Physica A* **318**:187–199, 2003.

[PDDK96] V. B. Priezzhev, D. Dhar, A. Dhar and S. Krishnamurthy, Eulerian walkers as a model of self-organised criticality, *Phys. Rev. Lett.* **77**:5079–5082, 1996.

[Pri94] E. Prisner, Parallel chip-firing on digraphs, *Complex Systems* **8**:367–383, 1994.

[RS09] L. Rolla and V. Sidoravicius, Absorbing-state phase transition for stochastic sandpiles and activated random walks. http://arxiv.org/abs/0908.1152

[Sh07] S. Sheffield., Gaussian free fields for mathematicians. *Probab. Theory Related Fields*, 139(3-4):521–541, 2007.

[Tse90] P. Tseng, Distributed computation for linear programming problems satisfying a certain diagonal dominance condition, *Mathematics of Operations Research* **15**(1):33–48, 1990.

[Wil96] D. B. Wilson, Generating random spanning trees more quickly than the cover time, In *28th Annual ACM Symposium on the Theory of Computing (STOC '96)*, 296–303, 1996.

MIT Department of Mathematics, 77 Massachusetts Ave. 2-335, Cambridge, MA 02139. http://math.mit.edu/~levine

E-mail address: levine@math.mit.edu