Algebra, Dynamics and Computation in Abelian Networks

Abelian networks, invented by Deepak Dhar, tie together recent strands of research in probability, combinatorics, statistical physics and computer science. The PI proposes to analyze these networks from three complementary viewpoints. (1) From a dynamical point of view, abelian networks produce intricate patterns using only simple local rules. (2) From a computational point of view, an abelian network yields the same output regardless of the timing of events at its individual nodes. (3) From an algebraic point of view, an abelian network associates an invariant, in the form of an abelian group or commutative algebra, to its underlying graph.

Intellectual Merit

(1) *How do small-scale interactions combine to produce complex large-scale structure?* Abelian networks are among the simplest systems known that produce complex patterns from local rules. Most systems of this sort turn out to be mathematically intractable, but the additional algebraic structure of abelian networks makes them a rare exception. As such, they are ideal candidates for teasing out the mechanisms of pattern formation. In addition to its intrinsic importance, a better understanding of pattern formation would aid in engineering desired large-scale outcomes by modifying only small-scale features of a system. This problem arises, for example, in designing a road network to minimize traffic delays and in using nanotechnology to engineer stronger materials.

(2) *What kinds of computational tasks can be performed in a distributed system with no central control over timing?* From a fundamental viewpoint, this question bears on whether or not the large-scale structural complexity of the physical world is contingent on the noncommutativity of interactions at the microscopic scale. From a more practical standpoint, answering this question will help identify algorithms on graphs that can be implemented in a completely asynchronous distributed manner. The current ubiquity of data sets on large graphs (arising from biological networks, the brain and the internet) has produced a growing demand for such algorithms.

(3) *What kinds of algebraic invariants can help distinguish between nonisomorphic graphs?* The abelian group associated to an abelian network on a graph captures certain information missed by traditional spectral invariants like eigenvalues of the Laplacian and adjacency matrices. The sandpile group is essentially the only such group studied so far, but other abelian networks may yield distinct invariants.

Broader Impacts

Abelian networks are a new and rapidly developing research area teeming with accessible open problems. The PI will publicize these problems widely, and envisions supervising a focused research group of undergradute and graduate students working on a handful of closely related projects. The PI has begun work on a book *Discrete Laplacian Growth*, coauthored with Yuval Peres, with the aim of making new results in this field accessible to a wider audience of students and researchers.

The beautiful pictures of pattern formation in abelian networks appeal widely even to those without the expertise to appreciate all of the mathematics involved. Iconic images such as the "Propp circle" [FL10] can play a role in enhancing the popular appeal of mathematics. The PI will explore avenues to promote these images, such as an art exhibition, interactive exhibit at a science museum, or an online video channel, with the goal of rendering mathematics interesting and relevant to a wider segment of society.

Project Description:

Algebra, Dynamics and Computation in Abelian Networks

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. A familiar example is earth's moon: We never see the dark side of the moon due to the fact that 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. In addition to many examples of this type in astronomy, mode-locking phenomena are pervasive in biological systems ranging from fireflies that synchronize their flashes to pacemaker cells that time the beating of the heart [Str03].

Synchronization and pattern formation have fascinated generations of physicists and mathematicians. Despite some notable successes, however, "emergent" behaviors like these have largely defied rigorous mathematical understanding. Since the mathematics of most emergent systems is intractable, it is important to focus on the simplest examples. One of the simplest is the Bak-Tang-Wiesenfeld abelian sandpile model [BTW87], which forms complex patterns using simple local rules [DSC09] (Figure 1). In a variant known as parallel chip-firing, it also functions as a model of mode-locking and synchronization [BCFV03, Lev10a].

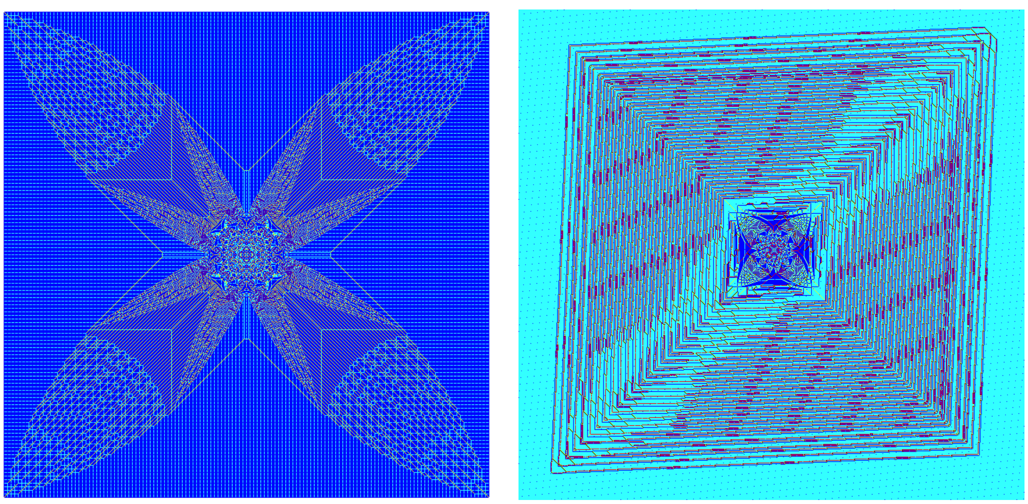

Figure 1: A stable sandpile (left) and an exploding sandpile (right) in $\mathbb{Z}^2$. Each site is colored according to the number of sand grains ("chips") present.

Over the past two 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. These new ideas have borne fruit on several fronts: they played a key role in our bound of sandpile growth rates in [FLP10], helped prove an exact formula for a rotor-router odometer function [KL10b], and inspired our new fast simulation algorithm for growth models [FL10]. The models which obey a least action principle appear to be quite diverse: for example, abelian sandpiles (deterministic), IDLA (a growth model based on random walk) and rotor-router (derandomized random walk).

What these models 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. I give a formal definition in §4.

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.) I argue that abelian networks are ideal candidates for studying synchronization and pattern formation: Their simple local dynamics give rise to interesting large-scale structure, and their extra algebraic structure makes them tractable. Moreover, 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.

After describing motivation, background and recent results in §1-3, I develop a research program for the computational point of view in §4. The dynamical point of view, including a discussion of stochastic abelian networks, follows in §5. The next section §6 presents the algebraic point of view, and the conclusion is devoted to broader impacts.

Results From Prior NSF Support

My current grant, an NSF Mathematical Sciences Postdoctoral Research Fellowship (DMS-0803064, $108,000) started September 1, 2008 with a break in tenure for the period January-June 2010 spent at Microsoft Research, and was eligible to be open through August 2012. It will be completed as of September 30, 2011. This award has resulted in 11 papers to date (10 research articles and 1 short survey [LP10b] exposing the results to a wider mathematical audience) on the topics of parallel chip-firing [Lev10a], abelian sandpiles [Lev10b, FLP10, FLW10a, FLW10b], random and derandomized aggregation [FL10, JLS10, KL10a, KL10b] and combinatorial group theory [HLR10]. The next three sections survey a selection of these results most relevant to the present proposal.

1 Mode-locking and synchronization in parallel chip-firing

A telltale sign of synchronization 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. The intervals of constancy correspond to preferred modes of the system: if the system is "locked in" to one of these modes, then even a relatively large perturbation may not change the observable behavior.

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 $\sigma : V \to \mathbb{Z}_{\geq 0}$, where V is the vertex set of a fixed finite undirected graph G. A vertex $v \in V$ is called *unstable* if $\sigma(v) \geq \deg(v)$, and an unstable vertex *fires* 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, $\lambda = \frac{1}{n} \sum_v \sigma(v)$, which is a conserved quantity. Here $n = \#V$ is the number of vertices. It also has a natural observable, the *activity*, defined by

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

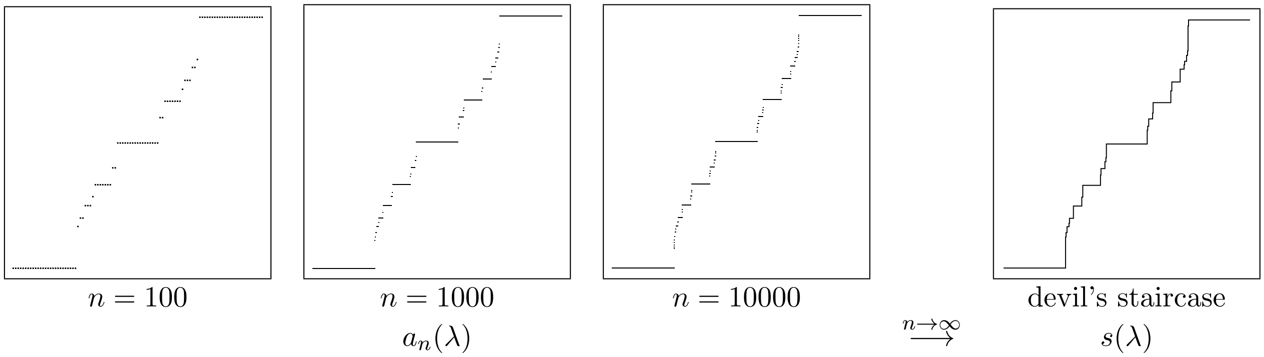

Figure 2: As the density of chips λ increases, the activity $a_n(\lambda)$ of a parallel chip-firing system on K_n has long intervals of constancy punctuated by sudden jumps. According to Theorem 1, as $n \to \infty$ the function $a_n(\lambda)$ tends to a limit $s(\lambda)$ whose graph has a flat horizontal 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.

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 λ varies, the activity a remains constant for long intervals punctuated by sudden jumps.

My motivation in [Lev10a] was to prove a theorem that would explain the striking findings of [BCFV03]. 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. (See in particular Problem 4 in §5.)

A second manifestation of mode-locking is the prevalence of *short period attractors.*

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 $k = 1, \ldots, n$ there exists a parallel chip-firing configuration on K_n of period k.*

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]. The proof of the following theorem was completed in [FLP10], building on earlier results of [LP09].

Theorem 3. *[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.*

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*

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

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.

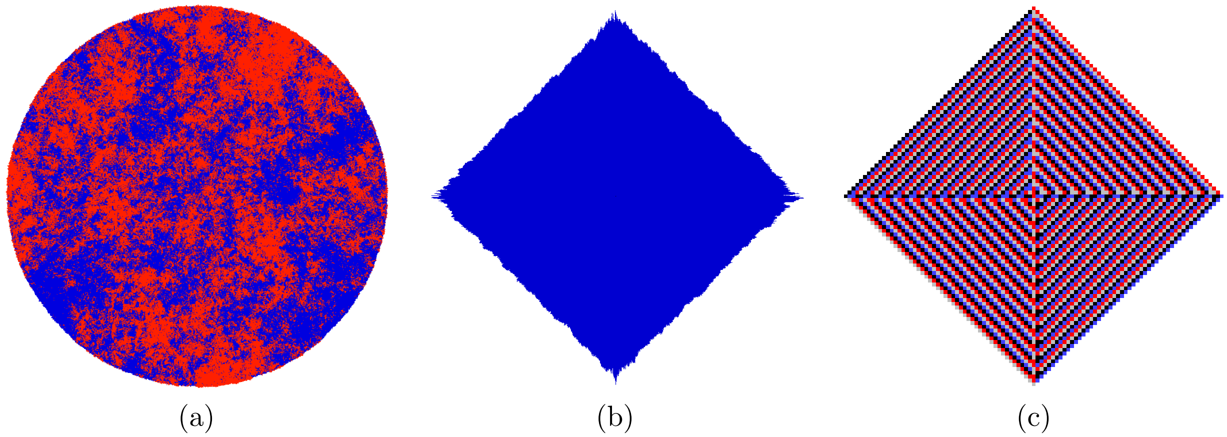

Figure 3: (a) IDLA cluster built from simple random walk in $\mathbb{Z}^2$ (b) IDLA cluster built from a uniformly layered walk (c) rotor-router cluster built from a uniformly layered walk.

Independently, Asselah and Gaudillière [AG10] have obtained a slightly weaker bound ($\log^2 r$ replaces the second $\log r$ in (1)) by a different method. 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 3a, 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 3b). 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. In joint work with Kager, I prove 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 3c).

Theorem 7. *[KL10b] There is an initial rotor configuration ρ_0, such that uniformly layered rotor-*

router aggregation with rotors initially configured as ρ_0 satisfies

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

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

Fast simulation

In joint work with Friedrich, I have developed fast algorithms for simulating IDLA and rotor-router aggregation. These algorithms are 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 $u : V \to \mathbb{Z}_{\geq 0}$ 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. We prove correctness of this algorithm using Theorem 8. The choice of initial approximation u_1 rests on earlier work joint with Peres [LP09].

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.

Future Directions in Abelian Networks

4 Computation: Exploring the limits of asynchronicity

This section begins with the formal definition of an abelian network, which can be skimmed on first reading. Let $G = (V, E)$ be a directed graph. Associated to each vertex $i \in V$ is a *processor* $\mathcal{P}_i$, which is a finite automaton with a single input feed and multiple output feeds, one for each edge $(i, j) \in E$. Each processor reads the letters in its input feed in first-in-first-out order.

The network is specified by *transition functions* associated to each vertex $i \in V$ and *message passing functions* M_{ij} associated to each edge $(i, j) \in E$. Formally, if the automaton at vertex i has state space Σ_i and input alphabet A_i, then these are maps

Transition function	$T_i : \Sigma_i \times A_i \to \Sigma_i$	(new internal state)
Message passing function	$M_{ij} : \Sigma_i \times A_i \to A_j^*$	(messages sent from i to j)

Here A_j^* denotes the free monoid on the alphabet A_j. We interpret these functions as follows. If the processor $\mathcal{P}_i$ is in state σ and processes input a, then two things happen:

1. The processor $\mathcal{P}_i$ transitions to state $\tau = T_i(\sigma, a)$; and
2. For each edge $(i, j) \in E$, the processor $\mathcal{P}_j$ receives input $M_{ij}(\tau, a)$.

Dhar writes that

> "In many applications, especially in computer science, one considers such networks where the speed of the individual processors is unknown, and where the final state and outputs generated should not depend on these speeds. Then it is essential to construct protocols for processing such that the final result does not depend on the order at which messages arrive at a processor." [Dha06]

To realize this vision, we ask that the following aspects of the computation *do not depend on the order in which individual processors act*:

- The **halting status**.
- The **final output** (final states of the processors).
- The **run time** (total number of letters processed by the entire network).
- The **local run times** (number of letters processed by a given $\mathcal{P}_i$).
- The **type-specific local run times** (number of times a given $\mathcal{P}_i$ processes a given letter $a \in A_i$).

A priori it is far from obvious that these goals are actually achievable by any nontrivial network. Remarkably, however, a simple local commutativity condition suffices to ensure that all five goals are achieved. For words w, w', write $w \sim w'$ if the letters of w' are a permutation of the letters of w. The processor $\mathcal{P}_i$ is called *abelian* if for any words $w \sim w' \in A_i^*$ we have for all $\sigma \in \Sigma_i$ and all edges (i, j)

$$T_i(\sigma, w) = T_i(\sigma, w') \qquad \text{and} \qquad M_{ij}(\sigma, w) \sim M_{ij}(\sigma, w').$$

That is, permuting the inputs to $\mathcal{P}_i$ does not change the final state of the processor $\mathcal{P}_i$, and may change the output words sent to the neighboring processors $\mathcal{P}_j$ only by a permutation.

Definition. An *abelian network* on a directed graph $G = (V, E)$ is a collection of automata $\mathcal{N} = (\mathcal{P}_i)_{i \in V}$ indexed by the vertices of G, such that each $\mathcal{P}_i$ is abelian.

It turns out that the purely local requirement that each $\mathcal{P}_i$ is abelian implies that the entire network is abelian, in the sense that it satisfies all five of the above properties. 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 type-specific local run times depend 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, which asks whether a given network halts on a *given* input.

Question 1. *Is the halting problem for abelian networks decidable in polynomial time?*

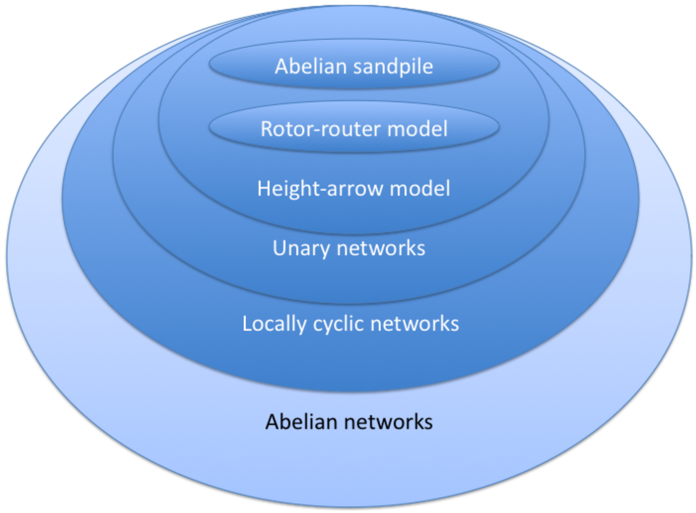

Figure 4: Venn diagram illustrating several classes of abelian networks, ranging from most specific at the top to most general at the bottom.

Figure 4 shows increasingly general classes of abelian networks. The two most widely studied examples are the Bak-Tang-Wiesenfeld abelian sandpile model [BTW87] and the rotor-router or "Eulerian walkers" model [PDDK96]. Dartois and Rossin [DR04] study a common generalization called the height-arrow model. These are all are examples of *unary networks*, that is, networks in which each alphabet A_i has cardinality 1. Another example of a unary network is the "periodically mutating game" of Eriksson [Eri96]. 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?*

Problem 1. *Define a precise notion of when one abelian network can simulate another. For each level of the hierarchy of Figure 4, either prove that it collapses to the previous level, or find an example that cannot be simulated by any member of the previous level.*

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 2. *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.

Certain nonabelian networks display intriguing remnants of an abelian property. These networks are nondeterministic: the final outcome depends on the order in which processors act. Yet in some cases this dependence is only of a very restricted type. For instance, in a model with particles and antiparticles each obeying sandpile dynamics and annihilating on contact, the initial configuration determines a one-parameter family to which the final configuration must belong.

Problem 3. *Develop a theory of "nonabelian" networks. How can the condition that each processor $\mathcal{P}_i$ is abelian be weakened without throwing it away altogether?*

Comparison with cellular automata

Abelian networks may be viewed as special types of cellular automata. They have several desirable features that cellular automata in general lack.

Abelian networks can update asynchronously. Traditionally, cellular automata use a *parallel update rule*: each cell updates its own state based on the states of its neighbors, and all cells perform their updates simultaneously. Implicit in this simultaneity is absolute centralized control over timing. Since perfect simultaneity is hard to achieve in practice, the physical significance of parallel updating cellular automata is open to debate. Abelian networks avoid this problem completely, because they reach the same final state no matter in what order the updates occur. Accordingly, no central control over timing is required. While it is still of interest in some cases to study them under parallel update (see §5), it is no longer a requirement.

Abelian networks do not rely on shared memory. Implicit in the update rule of cellular automata is some kind of action at a distance, through which each cell is constantly and instantaneously kept informed of the states of its neighbors. The lower-level interactions needed to facilitate this exchange of information in a physical implementation are missing from the model. Abelian networks remove this deficiency by operating in a "message passing" framework instead of the "shared memory" framework of cellular automata: An individual processor in an abelian network cannot access the states of neighboring processors, it can only read the messages they send.

Abelian networks can have any underlying geometry. Although cellular automata may be defined on any graph G, they are traditionally studied on the grid $\mathbb{Z}^d$ or on other lattices. I suggest that the study of cellular automata on an arbitrary graph G could be a fruitful means of revealing interesting graph-theoretic properties of G, and this is the perspective I take in developing the theory of abelian networks.

5 Dynamics: Synchronization and pattern formation

Bitar made a striking conjecture about parallel chip-firing (defined in §1) 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 parallel chip-firing must *eventually* settle into a periodic state; 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 3. *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. 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 $\mathbb{Z}_a \times \mathbb{Z}_b$.

The main insight that has enabled recent progress is a correspondence between parallel chip-firing dynamics on the complete graph and iteration of a circle map $S^1 \to S^1$. I believe that the right way to extend this correspondence to general graphs is by replacing the circle with a higher-dimensional torus.

Problem 4. *Generalize the circle map construction of [Lev10a] to model parallel chip-firing on a general graph G by iteration of a torus map $T^n \to T^n$.*

Parallel chip-firing systems arise from the abelian sandpile model by imposing the *parallel update rule* that all vertices fire simultaneously.

Problem 5. *Investigate whether other abelian networks with the parallel update rule display the same sorts of mode-locking phenomena and short period lengths characteristic of parallel chip-firing.*

The example of rotor-routing provides initial evidence of short periods: after an initial transient period, rotor-router walk with a single chip on a finite graph G repeatedly traces out an Eulerian tour of G [PDDK96, HLMPPW08]. This tour traverses every edge once in each direction, so it has length $2m$. A natural extension is to systems with more than one chip: at each time step, all chips simultaneously take a single rotor-router step. The analogue of Bitar's conjecture for these systems is the following.

Question 4. *Must any parallel rotor-routing configuration on a connected graph with m edges have eventual period at most $2m$?*

5.1 Pattern formation from local rules

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. A general theory of pattern formation with real predictive power remains a distant goal, and understanding how patterns form in specific instances is a necessary first step. In this section I outline approaches to several particular instances of pattern formation.

Dimensional Reduction. 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 5). In [FLP10] we formalized Dhar's observation as a *dimensional reduction conjecture*: Amazingly, large portions of the two pictures in Figure 5 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.

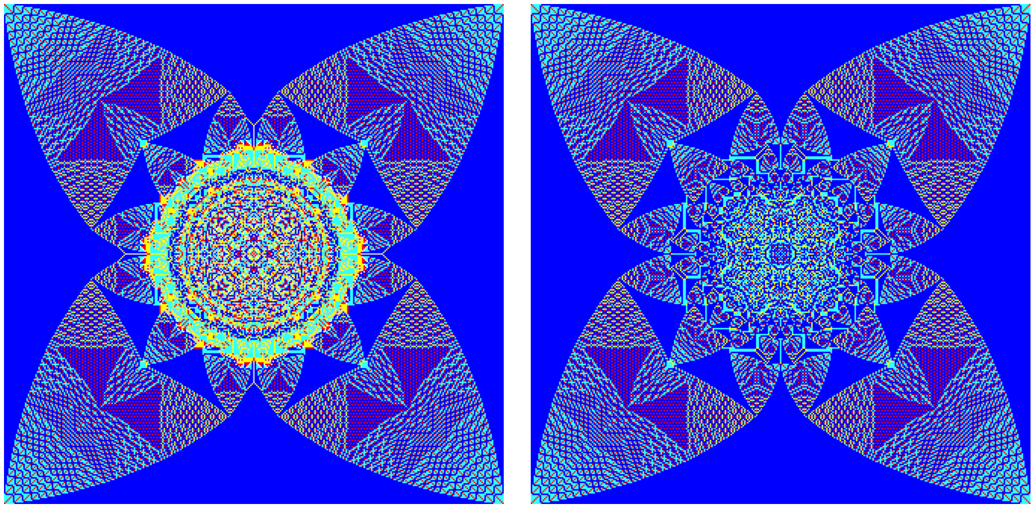

Figure 5: Illustration of the dimensional reduction conjecture. Left: A 2-dimensional slice through the origin of the sandpile of $n = 5 \cdot 10^6$ chips in $\mathbb{Z}^3$ on background height $h = 4$. Right: The sandpile of $m = 47465$ chips in $\mathbb{Z}^2$ on background height $h = 2$. Each site is colored according to the number of chips present. Remarkably, the two pictures agree pixel-for-pixel except in a region near the center.

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.

The Ostojic heuristic. The final configurations in sandpile aggregation contain many large "patches," visible in Figure 5, 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 (defined in §3). 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 6. *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 6 bring us any closer to a proof? Possibly yes. 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. Often the odometer function reveals intriguing local regularities but is beyond the reach of a global exact formula. The odometer function for rotor-router aggregation in $\mathbb{Z}^2$ has this character.

Problem 7. *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.*

5.2 Stochastic Abelian Networks

In a stochastic abelian network, the transition functions T_i (but not the message-passing functions!) depend on a probability space Ω.

$$\begin{aligned} T_i : \Sigma_i \times A_i \times \Omega &\longrightarrow \Sigma_i && \text{(new internal state)} \\ M_{ij} : \Sigma_i \times A_i \quad &\longrightarrow A_j^* && \text{(messages passed from } i \text{ to } j\text{)} \end{aligned}$$

We then require that pointwise for all $\omega \in \Omega$, the processor $\mathcal{P}_i(\omega)$ with transition function $T_i(\cdot, \cdot, \omega)$ is abelian. Examples of stochastic abelian networks include 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. The traditional view of random walk on a graph $G = (V, E)$ is that it is a single Markov chain X_t with state space V, giving the location of the walk at time t. We instead adopt the "stack-based" view of random walk developed by Diaconis and Fulton [DF91] and Wilson [Wil96], in which the steps of the random walk are indexed by location as well as time: each vertex i has a countably infinite stack of cards each labeled by a random outgoing edge (i, j). For each $k \geq 1$, the label of the k-th card in the stack is the edge traveled by the walk on its k-th exit from i. More generally, we can label each card by a state $\sigma \in \Sigma_i$ of the processor $\mathcal{P}_i$. A *locally Markov walk* is a stochastic unary network such that for each $i \in V$ and $\sigma \in \Sigma_i$,

$$\sum_j |M_{ij}(\sigma, i)| = 1.$$

This condition ensures that exactly one message is passed on for every message read. (Relaxing it would yield branching random walks, if the right side is allowed to be > 1, or walks with killing if the right side is allowed to be 0.)

Locally Markov walks obey a weak version of the Markov property: if the walker's current location is i, then the distribution of its next step depends only on the history of previous steps from i. These walks provide a way to interpolate between random walks (which correspond to the case of i.i.d. stacks) and their deterministic analogues, rotor walks (which correspond to the case of periodic stacks, [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 5. *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: among many other variations, one can consider a p-rotor walk on $\mathbb{Z}^2$ in which the rotor turns 90° clockwise with probability p and 90° counterclockwise with probability $1-p$. Is this walk recurrent for i.i.d. initial rotors? What happens in higher dimensions? Angel and Holroyd [AH10] prove several striking results about recurrence and transience of rotor walk on trees; which of their results extend to p-rotor walks?

Given a locally Markov walk $\mathcal{N}$ on a finite graph G, we can construct its *total chain*, which is a Markov chain on the state space $V \times \prod_{i \in V} \Sigma_i$, indicating the location of the walker and the current states of all the processors. The *Markov chain tree theorem* (see e.g. [Bro89, Ald90] both of which cite Diaconis) gives the stationary distribution of the total chain in the case that $\mathcal{N}$ itself is Markov chain.

Problem 8. *Generalize the Markov chain tree theorem to characterize the stationary distribution of the total chain of a locally Markov walk.*

6 Algebra: Transformation Monoids and Critical Groups

Determining whether two graphs are isomorphic is one of a relatively small number of decision problems not known to be computable in polynomial time and also not known to be NP-complete. It is therefore of interest to find new efficiently computable isomorphism invariants of graphs. Traditional invariants such as eigenvalues of the adjacency matrix, Laplacian, and other matrices associated to the graph throw away some useful information: by treating these integer matrices as linear transformations $\mathbb{Q}^n \to \mathbb{Q}^n$, they ignore the additional structure coming from the corresponding maps abelian groups $\mathbb{Z}^n \to \mathbb{Z}^n$.

The *sandpile group* [Dha90, Lor91, Big99] is an isomorphism invariant that captures some of this extra "integral" information: Figure 6 shows two graphs Laplacian-cospectral graphs that have nonisomorphic sandpile groups.

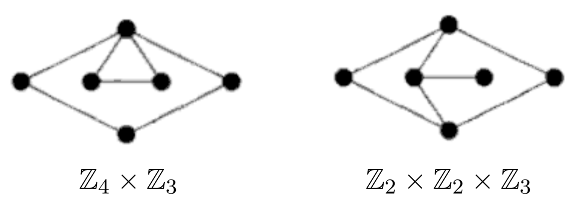

Figure 6: Two graphs with the same Laplacian spectrum but different sandpile groups.

Babai and Toumpakari [BT05] construct the sandpile group in a way that generalizes to abelian networks. Let $\mathcal{N}$ be an abelian network on a finite graph G that halts on all of its inputs, and write $A := \coprod A_i$ and $\Sigma := \prod \Sigma_i$ for the total alphabet and total state space of $\mathcal{N}$. Let $L_a(\sigma)$ be the final state of the network if we input letter a starting from initial state σ. This defines maps

$$L_a : \Sigma \to \Sigma$$

such that $L_a \circ L_b = L_b \circ L_a$ for all $a, b \in A$. These maps generate a finite commutative monoid $\mathcal{M}$ under composition, whose minimial ideal is an abelian group.

Definition. The *critical group* $K(\mathcal{N})$ is the minimal ideal of the transformation monoid $\mathcal{M}$.

The terms *critical group* and *sandpile group* are used more or less interchangeably in the mathematical literature. In the more general theory of abelian networks, we can make a distinction between them: any finite abelian network $\mathcal{N}$ has an associated critical group $K(\mathcal{N})$, and the sandpile group is simply the critical group of the particular network $\texttt{Sand}_G$.

Question 6. *Do abelian networks yield genuinely new algebraic invariants of graphs, or are the groups $K(\mathcal{N})$ always closely related to sandpile groups?*

Of course, finding genuinely new invariants would be extremely interesting. On the other hand, the alternative finding, that $K(\mathcal{N})$ is always closely related to $K(\texttt{Sand}_G)$, would greatly enlarge the class of models to which existing tools apply: Sandpile groups have been computed exactly for a wide variety of graphs, and the sandpile group is efficiently computable by putting the Laplacian in Smith normal form.

Given a graph G and specified "sink" vertex, the critical groups of the rotor-router model $\texttt{Rotor}_G$ and abelian sandpile model $\texttt{Sand}_G$ are isomorphic [PDDK96, LL09], a fact that has enriched our understanding of both models. In joint work with Landau [LL09], we use this isomorphism to prove that rotor-router aggregation grows as a perfect ball when the underlying graph is a regular tree. In [HLMPPW08] we show that via this isomorphism, $K(\texttt{Sand}_G)$ acts freely and transitively on the set of spanning trees of G. The latter result provides a high-level conceptual explanation of a fact that has fascinated combinatorialists for many years: the number of recurrent sandpiles on G equals the number of spanning trees of G.

A first step toward understanding the critical group $K(\mathcal{N})$ of a general network $\mathcal{N}$ is to classify the *recurrent states* of $\mathcal{N}$, defined as the states σ such that $\sigma = L_{a_1} \circ \ldots \circ L_{a_k}\sigma$ for some $k \geq 1$ and $a_1, \ldots, a_k \in A$. Dhar [Dha90] solved this problem for sandpiles by giving a procedure called the *burning algorithm* which decides in linear time whether a sandpile state is recurrent.

Problem 9. *Generalize Dhar's burning algorithm to abelian networks.*

The burning algorithm applies to undirected graphs (Speer[Spe93] gave a more involved algorithm for directed graphs). Thus the first challenge in approaching Problem 9 is to define an adequate notion of "undirectedness" (or what might be called "reversibility," extending the notion of a reversible Markov chain) for abelian networks. A second challenge is how to eliminate the special role of the "sink" vertex in the burning algorithm; a general abelian network may halt on all inputs even if it has no completely inert vertex serving as a sink.

Using a linear-algebraic version of the transformation monoid, Dhar [Dha06] associates to a finite stochastic abelian network $\mathcal{N}$ a finite-dimensional commutative $\mathbb{R}$-algebra, $\mathcal{A}(\mathcal{N})$. One of the simplest stochastic abelian networks is the network $\texttt{Markov}(P)$ associated to a stochastic transition matrix P with an absorbing state (see §5.2). Viewing this network as a random walk on the directed graph G whose edges correspond to the nonzero entries of P, the state of each processor represents the edge along which the random walker last exited, and the recurrent states of $\texttt{Markov}(P)$ correspond to spanning trees of G rooted at the absorbing state.

Question 7. *What is the structure of the algebra $\mathcal{A}(\texttt{Markov}(P))$?*

Broader Impacts:

Student Research and Popular Appreciation of Mathematics

Abelian networks are concrete objects that lend themselves to experimentation, which makes them an ideal topic for student research projects at all levels. I have made an effort to popularize widely some accessible open problems in the area and will continue to do so. When a student asks me for a research problem, I try hard find one to match her skills, interests and mathematical style. Two recent student-authored papers in combinatorics [BK10] and [Jia10] are based on problems I proposed. In both cases the students surpassed my expectations with truly beautiful original work.

Last year, Jim Propp and I started the Chips and Rotors Research Initiative to involve MIT undergraduates in our research. The group has two papers in the works, both with undergraduate coauthors. Just as important as the research itself is the effect of the experience on the students. One of the undergraduates I am currently working with wrote in a recent review,

> "I wanted to see if I would enjoy research. I was considering graduate school, but I knew I needed to test the waters first. I found that it is very satisfying to be able to shape my ideas from vague thoughts into publishable material, and that I enjoy collaborating with other people. I will be applying to graduate school for next fall. I am invested in the research topic and in the paper and I am excited to continue working. I have a goal to publish a paper before I graduate, and I will achieve that goal this term."

Over the next few years, I envision enlarging our student research group to 3-5 undergraduates and adding one or more graduate students. In the present proposal, Question 4 and Problem 5 are projects I have earmarked as future research problems for students.

Abelian networks lie at a nexus of algebraic combinatorics, probability, statistical mechanics and theoretical computer science. In addition to involving students, I intend to draw in researchers in all of these fields, with a special effort to reach out to departments other than mathematics and to researchers working in industry, including my ongoing collaborations researchers at Microsoft Research. Yuval Peres and I have recently made arrangements with Princeton University Press to write a book *Discrete Laplacian Growth*, with the aim of making new results in this field accessible to a wider audience of graduate students and researchers.

The beautiful pictures of pattern formation in abelian networks appeal widely even to those without a technical background. Iconic images such as the "Propp circle" [FL10] have the potential to enter mainstream culture just as the Mandelbrot set and other fractals did in the 1980's and 1990's. I will explore avenues such as an art show, interactive exhibit at a science museum, or a public online video channel to promote these images with the goal of exposing a wider segment of society to a taste of mathematics. If these initial forays into popular culture prove successful, the next step is to get the public more actively involved. Modeled on the success of free software like Fractint — which in the 1990's enabled anyone with a PC to experiment with fractals — I will work with students to develop software that allows users to tinker with abelian networks, create their own mathematical images and videos, and share them with the community.

References Cited

[Ald90] D. J. Aldous, The random walk construction of uniform spanning trees and uniform labelled trees, *SIAM J. Discrete Math.* **3**(4):450–465, 1990.

[AH10] O. Angel and A. E. Holroyd, Rotor walks on general trees, 2010. http://arxiv.org/abs/1009.4802

[AG10] A. Asselah and A. Gaudillière, From logarithmic to subdiffusive polynomial fluctuations for internal DLA and related growth models, 2010. http://arxiv.org/abs/1009.2838

[BT05] L. Babai and E. Toumpakari, A structure theory of the sandpile monoid for digraphs, *Formal Power Series in Algebraic Combinatorics*, 2005.

[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.

[BK10] H. Bidkhori and S. Kishore, A bijective proof of a theorem of Knuth, *Combin. Probab. Comput.*, to appear.

[Big99] N. L. Biggs, Chip-firing and the critical group of a graph, *J. Algebraic Combin.* **9**(1):25–45, 1999.

[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.

[Bro89] A. Broder, Generating random spanning trees, In *30th Annual Symposium on Foundations of Computer Science* (FOCS89), 442447, 1989.

[CDVV06] M. Casartelli, L. Dall'Asta, A. Vezzani and P. Vivo, Dynamical invariants in the deterministic fixed-energy sandpile, *European Phys. J. B* **52**(1):91–105, 2006.

[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. Manuscript 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.*, to appear. 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.

[Lor91] D. J. Lorenzini, A finite group attached to the Laplacian of a graph, *Discrete Math.* **91**(3):277–282, 1991.

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

[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.

[Spe93] E. R. Speer, Asymmetric abelian sandpile models. *J. Stat. Phys.* **71**:61–74, 1993.

[Str03] S. Strogatz, *Sync: The emerging science of spontaneous order*, Hyperion, 2003.

[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.

Biographical Sketch: Lionel Levine

Professional Preparation

1998 – 2002 **A.B. in Mathematics**, Harvard University, *magna cum laude.*
2002 – 2007 **Ph. D. in Mathematics**, University of California, Berkeley.
Advisor: Yuval Peres
Thesis: "Limit Theorems for Internal Aggregation Models"

Appointments

2008 – 2012 **C. L. E. Moore Instructor**, Department of Mathematics, MIT.

Relevant Publications

2010
1. Tobias Friedrich and Lionel Levine, *Fast simulation of large-scale growth models*, submitted for publication July 2010. http://arxiv.org/abs/1006.1003
2. Lionel Levine, *Sandpile groups and spanning trees of directed line graphs.* Journal of Combinatorial Theory A, accepted April 2010. http://arxiv.org/abs/0906.2809
3. Lionel Levine, *Parallel chip-firing on the complete graph: devil's staircase and Poincaré rotation number.* Ergodic Theory and Dynamical Systems, accepted January 2010. http://arxiv.org/abs/0811.2800
4. Anne Fey, Lionel Levine and David B. Wilson, *Driving sandpiles to criticality and beyond.* Physical Review Letters **104**: 145703, 2010. http://arxiv.org/abs/0912.3206
5. Anne Fey, Lionel Levine and Yuval Peres, *Growth rates and explosions in sandpiles.* Journal of Statistical Physics **138**: 143–159, 2010. http://arxiv.org/abs/0901.3805

Other Significant Publications

2010
1. David Jerision, Lionel Levine and Scott Sheffield, *Logarithmic fluctuations for internal DLA*, submitted for publication October 2010. http://arxiv.org/abs/1010.2483
2. Lionel Levine and Yuval Peres, *Scaling limits for internal aggregation models with multiple sources.* Journal d'Analyse Mathematique, accepted February 2009. http://arxiv.org/abs/0712.3378
3. Wouter Kager and Lionel Levine, *Diamond aggregation.* Mathematical Proceedings of the Cambridge Philosophical Society **149**(2): 351–372, 2010. http://arxiv.org/abs/0905.1361

2009
4. Lionel Levine and Yuval Peres, *Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile.* Potential Analysis **30**: 1–27, 2009. http://arxiv.org/abs/0704.0688

2007
5. Christopher J. Hillar and Lionel Levine, *Polynomial recurrences and cyclic resultants.* Proceedings of the American Mathematical Society **135**: 1607–1618, 2007. http://arxiv.org/abs/math.AG/0411414

Synergistic Activities

2011	Co-organizer of the special session on "Visual Mathematics: Laplacian Growth" at the 2011 Joint Meetings of the American Mathematical Society and Mathematical Association of America.
2008 – 2010	Co-organizer of the MIT Probability Seminar.
2008 – 2010	Supervised MIT undergraduates Giuliano Giacaglia, Damien Jiang and Linda Zayas-Palmer as part of MIT's Undergraduate Research Opportunities Program.
2006 – 2007	Mentored Berkeley undergraduates Itamar Landau and Parran Vanniasegaram on a research project in probability and combinatorics. In addition to the research component, this project involved developing software for aggregation models.
2000 – 2001	Award Peer Tutor Program. Tutored Harvard undergraduates in calculus, linear algebra, group theory, mechanics, basic probability and combinatorics.

Collaborators and Other Affiliations

Anne Fey (TU Delft, The Netherlands)
Christopher J. Hillar (Mathematical Sciences Research Institute)
Alexander E. Holroyd (Microsoft Research)
David Jerison (MIT)
Wouter Kager (VU Amsterdam, The Netherlands)
Itamar Landau (University of California, Berkeley)
Karl Mahlburg (Princeton)
Karola Meszaros (MIT)
Yuval Peres (Microsoft Research and University of California, Berkeley)
James Propp (University of Massachusetts, Lowell)
Darren Rhea (University of California, Berkeley)
Scott Sheffield (MIT)
Katherine E. Stange (Simon Fraser University)
David B. Wilson (Microsoft Research)

Graduate Advisor and Postdoctoral Sponsor

Yuval Peres (Microsoft Research and University of California, Berkeley): *Ph.D. thesis advisor.*
Richard Stanley (MIT): *Sponsoring scientist for NSF postdoctoral fellowship.*

Advisees

Total number of graduate students advised: 0
Total number of postdoctoral scholars sponsored: 0