These notes are for a talk in the UC Berkeley student probability seminar, and are based mainly on [3].

1 Self-organized criticality

In statistical mechanics, we often think of criticality as a special or rare phenomenon. For instance, if you consider the renormalization group procedure on the Ising model, the critical point is a repulsive fixed point, as starting even just below the critical temperature leads to extremely low temperature behavior, and starting just above leads to extremely high temperature behavior.

However, in nature we often see things which bear the hallmarks of criticality, namely self-similarity or some other critical behavior. Moreover, this critical behavior arises naturally. The classic example is a pile of sand: if you pour sand on a table, then it will keep building up until the slope reaches criticality, at which point the slope stays roughly constant as all the subsequent sand slides off.

This is known as self-organized criticality, and it has captivated physicists for a while and spawned pop-science books with evocative titles such as "How Nature Works" [1]. It even got the attention of former vice president Al Gore, who wrote the following in his book "Earth in the Balance" [4]:

> "The sand pile theory—self-organized criticality—is irresistible as a metaphor... The formation of identity is akin to the formation of the sand pile... One reason I am drawn to this theory is that it has helped me understand change in my own life."

2 The activated random walk

Here we discuss the first model to provably exhibit the self-organized criticality phenomenon.

2.1 Definition of the model

A configuration of the activated random walk on a graph G consists of a certain number of particles on each vertex, each particle being either active or sleeping. Active particles perform independent continuous-time random walks with jump rate 1, and when an active particle is alone at a vertex it falls asleep with rate $\lambda > 0$. A sleeping particle stops moving and is instantaneously reactivated when an active particle lands on the vertex where it is sleeping.

This can be done in any graph, finite or infinite. We'll only consider $\mathbb{Z}^d$ or finite lattice chunks though. For finite graphs which have a natural boundary, such as a finite piece of a lattice, we may define a *driven-dissipative* version of the activated random walk, which is a Markov chain on configurations of sleeping particles, at most one per vertex.

To get dissipation, we allow particles to "fall off the table" i.e. get killed at the boundary of the graph. They instantly leave the system, never to return. For the driven part, we just continuously add particles to the configuration. One step of the Markov chain consists of the following two sub-steps:

1. Add an active particle to a vertex of the graph uniformly at random.

2. Let the ARW dynamics run until all particles fall either off the table or asleep.

This is a finite, ergodic, aperiodic Markov chain, so it converges to a stationary distribution, which we denote by μ, on configurations of sleeping particles $\{0, \mathfrak{s}\}^{\Lambda}$.

2.2 Phase transition and self-organized criticality

The following theorem shows the existence of a phase transition for the activated random walk in terms of the density of particles in the starting configuration, from a phase where all particles go to sleep forever, to a phase where all particles are awake infinitely often.

Theorem 1 (lots of people)**.** *Consider the activated random walk in* $\mathbb{Z}^d$ *with sleep rate* $\lambda > 0$ *and initial configuration of active particles i.i.d. with density* $\rho \in (0, 1)$*. There exists* $\rho_c = \rho_c(\lambda) \in (0, 1)$ *such that, when* $\rho < \rho_c$ *the entire system sleeps, whereas when* $\rho > \rho_c$*, it stays awake.*

The density conjecture states that the density under the stationary distribution of the Markov chain in a finite chunk of the lattice converges to the critical density of the infinite lattice, as the chunk gets larger and larger. The main theorem we'll discuss is the resolution of this conjecture in $d = 1$:

Theorem 2 ([5])**.** *Let* S_n *denote the number of particles in the interval* $\{1, \ldots, n\}$ *under the stationary distribution* μ*. Then* $\frac{S_n}{n} \xrightarrow{\mathbb{P}} \rho_c$*.*

For the rest of the talk, we'll focus on the proof of this theorem, and so we'll only consider the ARW in a segment $\{1, \ldots, n\}$ with particles killed when they jump to the left from 1 or to the right from n. We'll follow a new simplification of a key portion of the proof due to [3].

2.3 Convergence via superadditivity

The main thing we'll talk about today is the fact that the density actually converges. The following two lemmas show that if the density converges, then it must converge to ρ_c. Let M_n denote the number of particles which "fall off the table" during the stabilization of a particular instance of the activated random walk in a segment.

Lemma 3 ([6])**.** *For each* $\rho < \rho_c$*, if the initial configuration is i.i.d. with density* ρ*, then* $\mathbb{E}M_n = o(n)$*.*

Lemma 4 ([2])**.** *For each* $\rho > \rho_c$*, there are* $\varepsilon > 0$ *and* $c > 0$ *such that for all initial configurations with at least* ρn *particles,* $\mathbb{P}[M_n > \varepsilon n] \geq c$*.*

In other words, if we start with fewer than the critical density of particles, then very few of them fall off the table during stabilization, and on the other hand if we start with more than the critical density, then many of them fall off the table.

To prove that the density actually converges, we'll use a stochastic version of Fekete's lemma.

Definition 5. A sequence of random variables (X_n) is stochastically superadditive if for every $n, m \geq 1$, the variable X_{n+m} stochastically dominates the sum of X_n and an *independent* copy of X_m.

Lemma 6. *Let* (X_n) *be a stochastically superadditive sequence of non-negative random variables, and define*

$$\rho_\star = \sup_{n \geq 1} \frac{\mathbb{E}X_n}{n} \in [0, \infty].$$

Then we have $\frac{X_n}{n} \xrightarrow{\mathbb{P}} \rho_\star$*.*

The proof is similar in spirit to Fekete's lemma, and I won't present it here. The key proposition is then the following

Proposition 7. *For every* $n, m \geq 1$*, the variable* S_{n+m+1} *stochastically dominates* S_n *plus an independent copy of* S_m*.*

The extra $+1$ may seem out of place, but we can apply Lemma 6 to $X_n = S_{2n-1}$ which is stochastically superadditive in the above sense, to conclude the convergence in probability of $\frac{S_n}{n}$ and thus, using Lemmas 3 and 4, conclude the proof of the density conjecture in $d = 1$.

By the way, the proof that we'll see for Proposition 7 is quite one-dimensional, and it would be interesting to see how, if at all, this generalizes to higher dimensions (this is still open).

3 Tools

Here we give some of the tools which will be used in the proof of Proposition 7.

3.1 Abelian property and monotonicity

The main tool is the Abelian property, which says that it doesn't matter which *order* we move things in. To make this precise, we have to change perspective a bit. Instead of each particle making its own decisions, suppose we have an infinite list of instructions, either $\leftarrow$, $\rightarrow$, or $\mathfrak{s}$, at each vertex.

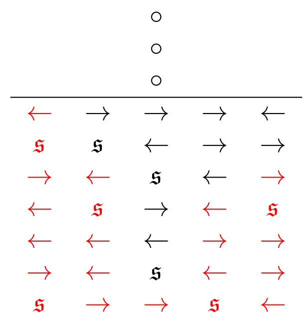

Now we can move a particle or put it to sleep by using up an instruction at the vertex where the particle is sitting. One disambiguation: if there are multiple particles at a vertex and a sleep instruction is attempted, it gets used up but nobody goes to sleep.

Let us keep track of the list of vertices where instructions are used, in order, and call it α. Clearly we can determine the final state based on the initial configuration, the instruction field, and α. We also keep track of how many instructions are used at each site via an *odometer* m_α, which is just a function $V \to \mathbb{N}$.

We'll call a sequence α *legal* for an initial condition η if all steps involve moving an *active* particle, and we'll say it's just *decriminalized* (the real term is *acceptable*) if some steps involve waking up a solitary particle. We'll say that a sequence α *stabilizes* an initial condition η if after all steps, all particles are sleeping or dead. These things satisfy a few nice properties, the first of which is the aforementioned Abelian property.

Lemma 8. *If α is acceptable and stabilizes η, and β is legal for η, then $m_\alpha \geq m_\beta$. In particular, if α and β are both legal for and stabilize η, then $m_\alpha = m_\beta$, and also the final configurations agree.*

I won't prove this, but it's just a bit of case checking. Now, to make this setup the same as λ-ARW, we need to take the sequence of instructions to be random. Each instruction will thus be chosen, i.i.d., such that the relative probabilities of $\leftarrow$, $\rightarrow$, and $\mathfrak{s}$ are $\frac{1}{2}$, $\frac{1}{2}$, and λ. Using Lemma 8, we can also prove a stochastic monotonicity property for the ARW. Let $S(\eta)$ denote the number of sleeping particles after η stabilizes under the ARW.

Lemma 9. *If $\eta \geq \xi$ are two initial configurations containing only active particles, then $S(\eta)$ stochastically dominates $S(\xi)$.*

Proof. It's enough to consider $\eta = \xi + \delta_x$. Then we can stabilize $\xi + \delta_x$ by first forcing the particle at x to walk off the table, waking it up if necessary (so this is only a decriminalized sequence of moves), and then stabilizing ξ. The resulting number of particles has the same distribution as $S(\xi)$, and by Lemma 8, we used more instructions than we would for a legal sequence, meaning that, if anything, we forced more particles off the table. ∎

3.2 Exact sampling

Here's a neat trick giving exact samples from the stationary distribution of the driven-dissipative chain.

Lemma 10. *Let the initial configuration have one active particle at each site, and then let it stabilize. The resulting distribution on configurations of sleeping particles is μ.*

Proof. Recall that in order to get from one state (configuration of sleeping particles) of the driven-dissipative chain to the next, we add an active particle somewhere uniformly at random and let the resulting configuration stabilize. So the *next* state after the state we sample using the procedure in the lemma statement can be sampled directly by starting with a configuration with one active particle at each vertex, and an extra one at some uniform vertex. By the Abelian property, we can force the extra particle to move first, and since it can't fall asleep, it will eventually fall off the table. So the state that we end up with is the same in distribution as if we hadn't added the particle at all: in other words, we're at the stationary distribution. ∎

4 Proof of superadditivity

Let's consider the interval $V = \{-n, \dots, +m\}$, with our array of instructions as before. Let $L = \{-n, \dots, -1\}$ and $R = \{1, \dots, m\}$. Let S_V, S_L, and S_R be the result of stabilizing $\mathbf{1}_V$, $\mathbf{1}_L$, and $\mathbf{1}_R$ in V, L, and R respectively, *using the same array of instructions* but where "stabilizing in L" means that particles fall off the table at 0, and similarly for R.

Note that $S_V \stackrel{d}{=} S_{n+m+1}$, $S_L \stackrel{d}{=} S_n$, $S_R \stackrel{d}{=} S_m$, and $S_L \perp\!\!\!\perp S_R$ since S_L and S_R depend on disjoint sets of instructions. So it suffices to prove that S_V stochastically dominates $S_L + S_R$ (note that it is not always true that $S_V \geq S_L + S_R$).

4.1 Interpolatively adding a hole to the table

For each $k \in \mathbb{N}$, we can see what happens when we replace all instructions at 0 after a certain point with kill instructions, effectively adding a hole in the middle of the table after k instructions are used there. Let N_k denote the number of sleeping particles after stabilizing $\mathbf{1}_V$ using this modified instruction field.

Note that $N_0 = S_L + S_R$, and $N_k \to S_V$ almost surely as $k \to \infty$, because stabilizing $\mathbf{1}_V$ in V requires only finitely many instructions, for almost every instance of the instruction field. So it suffices to prove that N_{k+1} stochastically dominates N_k for each $k \in \mathbb{N}$.

4.2 Replacing one instruction with a hole

There are a few cases to showing that N_{k+1} stochastically dominates N_k. First, if the kth instruction at 0 is $\mathfrak{s}$, then either we'll use this instruction as the last one at 0, and hence, end up with a sleeping particle on 0 (i.e. $N_{k+1} = N_k + 1$), or we'll use strictly fewer or more than k instructions at 0, in which case both stabilized configurations are the same (i.e. $N_{k+1} = N_k$).

Now suppose that the kth instruction is $\to$. In this case, we'll use the Abelian property and start stabilizing the configuration by only using the leftmost instruction, until k instructions have been used at 0. After that, once the hole at 0 has been revealed, we can use whatever legal instructions we want. By the way, if the configuration stabilizes *before* k instructions are used at 0, then both final configurations are the same, and so $N_{k+1} = N_k$.

Now suppose that we use the kth instruction at 0. In the instructions generating N_k, this kills a particle, and in the instructions generating N_{k+1}, this moves a particle from 0 to 1. By our leftwards priority, there are only sleeping particles in L at this point, and in R we have either some configuration η, or $\eta + \delta_1$, of *all active* particles. By the monotonicity property, Lemma 9, in this case we thus have N_{k+1} *stochastically dominating* N_k. Note that this is the only case where we don't necessarily have $N_{k+1} \geq N_k$.

Anyway, this finishes the proof since the case where the kth instruction is $\leftarrow$ is symmetric.

References

[1] Per Bak. *How nature works: the science of self-organized criticality.* Springer Science & Business Media, 2013.

[2] Nicolas Forien. Macroscopic flow out of a segment for activated random walks in dimension 1. *arXiv preprint arXiv:2405.04510*, 2024.

[3] Nicolas Forien. A new proof of superadditivity and of the density conjecture for activated random walks on the line. *arXiv preprint arXiv:2502.02579*, 2025.

[4] Albert Gore Jr. Earth in the balance. 1992.

[5] Christopher Hoffman, Tobias Johnson, and Matthew Junge. The density conjecture for activated random walk. *arXiv preprint arXiv:2406.01731*, 2024.

[6] LT Rolla and L Tournier. Non-fixation for biased activated random walks. *Annales de L'Institut Henri Poincare Section (B) Probability and Statistics*, 54(2):938–951, 2018.