Growth, Erosion, and Competition Driven by Random and Non-Random Walk

James Propp
(UMass Lowell)

August 4, 2012

Slides for this talk are on-line at

http://jamespropp.org/mathfest12b.pdf

Acknowledgments

This talk describes past and on-going work with David Einstein, Tobias Friedrich, Lionel Levine, and Yuval Peres.

Diffusion Limited Aggregation (DLA)

Diffusion-Limited Aggregation, or DLA (Witten and Sander, 1981), is a process in which randomly-walking particles cluster together.

One version:

- The 1-particle cluster $C_1 = \{s_1\}$ contains just $s_1 = (0, 0)$.
- For all $n > 1$, a particle does "random walk from infinity" in $\mathbb{Z}^2$, until it hits a site s_n adjacent to the $(n-1)$-particle cluster C_{n-1}; then $C_n = C_{n-1} \cup \{s_n\} = \{s_1, \ldots, s_n\}$ is the n-particle cluster.

A typical DLA aggregate

http://classes.yale.edu/fractals/panorama/physics/dla/BigDLA2.gif shows what a typical DLA aggregate looks like.

Although this model has been studied for thirty years, virtually nothing has been proved about it.

A related erosion model

Lionel Levine wrote a simulation of Diffusion-Limited Erosion (DLE) in which the initial aggregate is the **complement** of a large disk centered at $(0, 0)$.

Particles take a random 2-D walk starting from $(0, 0)$ until they join the aggregate.

One sees from
`http://www.math.cornell.edu/~levine/gallery/idledisc500.png`
that in DLE, as in DLA, random-looking dentritic structures form.

(Do the dendritic structures in two pictures look the same? Would the differences go away if we made our simulations larger?)

Internal DLA

On the other hand, consider the model in which the initial aggregate is $\{(0,0)\}$, and the way the aggregate grows is that a particle takes a random walk from $(0,0)$ until it reaches the complement of the aggregate, at which point the new site joins the aggregate.

Now instead of a random dentritic structure we see a circular cluster with small random fluctuations at its boundary; see Fig. 1 on page 4 of `http://arxiv.org/pdf/1010.2483.pdf` (and ignore the red/blue coloring, which is not relevant for today's purposes).

Stability vs. instability

It's easy to see intuitively why we get this qualitative difference between DLA and DLE on the one hand and IDLA on the other, by considering what happens to small fluctuations from circularity.

In the case of DLA and DLE, these fluctuations tend to be magnified by subsequent evolution of the interface between the aggregate and its complement; in the case of IDLA, the fluctuations tend to be dampened.

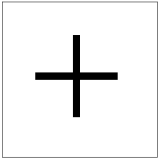

That's because a random walk started at point x tends to visit points close to x sooner than points that are farther away.

Deviation from circularity

So we know that the IDLA cluster “wants to be round”; but how badly does it want it?

Deviation from circularity

So we know that the IDLA cluster "wants to be round"; but how badly does it want it?

Theorem (Jerison-Levine-Sheffield, 2010; Asselah, Gaudillière, 2010): The deviations from circularity have typical magnitude on the order of $\log r$ or smaller, where r is the radius of the growing cluster.

The continuum limit

One could start with a lopsided cluster and see how it gets rounder.

Under appropriate time-dependent rescaling, this discrete stochastic interface process becomes a continuous deterministic interface process called **Laplacian growth**.

In Laplacian growth, a time-dependent region Ω_t in $\mathbb{R}^2$ grows so that the boundary $\partial\Omega_t$ moves outward with velocity proportional to harmonic measure on $\partial\Omega_t$.

Probabilistically, harmonic measure is the probability density for Brownian motion started at $(0,0)$ and stopped at $\partial\Omega_t$.

In other words: the velocity is proportional to the normal derivative of the Green's function for Ω_t, which is the function on Ω_t that vanishes on $\partial\Omega_t$ whose Laplacian is the Dirac delta at $(0,0)$.
See the upcoming book "Laplacian Growth" by Levine and Peres.

Competitive erosion

Another model exhibiting stable interface dynamics is competitive erosion (aka mutual diffusion-limited aggregation).

A large patch in $\mathbb{Z}^2$ is colored red and blue, with a vertices colored red and the remaining b vertices colored blue.

Two vertices are designated the "red source" and the "blue source".

1. A particle does a random walk from the red source until it arrives at a vertex colored blue; then that vertex is recolored red.

2. A particle does a random walk from the blue source until it arrives at a vertex colored red; then that vertex is recolored blue.

Steps 1 and 2 get repeated in alternation, over and over.

Interfaces for competitive erosion

See pictures at the top of
`http://www.math.cornell.edu/~levine/gallery/erosion.html`.

It is predicted on theoretical grounds, and “confirmed” by simulations, that if

- the patch is a large disk (that is, the intersection of a large disk with $\mathbb{Z}^2$, with edges joining vertices at Euclidean distance 1), and
- the two sources are on the boundary of the disk,

then

- the system evolves into a segregated coloring, and
- the red/blue interface is (a discrete approximation to) a circular arc perpendicular to the bounding circle.

A sketch of a physics argument

The stability heuristic for IDLA applies here too.

In the case where the two sources are diametrically opposite and the initial coloring has equal numbers of red and blue vertices, symmetry dictates that if there is a stable interface, it must be the diameter equidistant from the sources.

In the general case, we invoke the fact that Brownian motion (the continuum limit of random walk) exhibits conformal invariance, and use Möbius transformations to reduce the general case to the diametric case discussed above.

Is randomness necessary?

These interface processes are driven by random walks.
What if we replace the random walks by non-random walks with some of the same essential properties?

Specifically: a random walk on $\mathbb{Z}^2$ has the property that if, up to some time t, it has visited some site N times, then it has exited that site in each of the four directions approximately $N/4$ times.

What if we devise a walk that satisfies this law-of-large-numbers property but uses "less randomness" than ordinary random walk, or isn't random at all?

Might it still give rise to similar interface dynamics?

Partially derandomized DLA

One way for a particle to do "less random" random walk on $\mathbb{Z}^2$ is to obey the following protocol:

1. If the particle wants to take a step from a site that has not been exited before, the particle takes a random step (independent of everything it has already done).

2. If a particle wants to take a step from a site that HAS been exited before, the particle goes in the "successor direction" relative to the previous exit from that site. (That is, if the last exit from the site was in the East/North/West/South direction, the new exit must be in the North/West/South/East direction, respectively.)

Partially derandomized DLA on a torus

Here's what the 2000-particle cluster looks like for partially derandomized DLA on a 128-by-128 torus with initial 1-site cluster at $(64, 64)$ and a point source at $(0, 0)$:

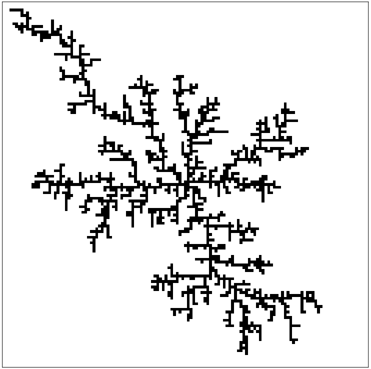

Partially derandomized DLA on a torus

Oded Schramm came up with this variant of DLA and did the first simulations of it in 2005 (unpublished work).

No theoretical work on this model has been done.

Getting rid of randomness

Going a step further, we could decree that our particle must obey the following protocol:

$1'$. If the particle wants to take a step from a site that has not been exited before, the particle goes East.

2. If a particle wants to take a step from a site that HAS been exited before, it goes in the "successor direction" relative to the previous exit from that site.

Note that **nothing** is random.

Fully derandomized DLA on a torus

Here's what the 2000-particle cluster looks like for fully derandomized DLA on a 128-by-128 torus:

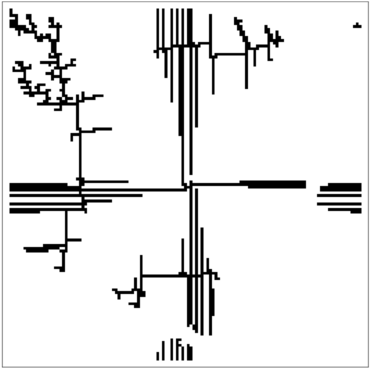

As Schramm observed, "the patterns look rather orderly".

Rotor-routers

But in fact, many **stable** random interface processes have fully derandomized analogues that display (or appear to display) the **same** interface dynamics in the continuum limit.

Walks that obey property 2 are said to follow the "rotor-router rule".

(See `http://jamespropp.org/csps12.pdf` for slides for a talk on rotor-routing in $\mathbb{Z}^2$ that I gave at Berkeley in Spring 2012, see `http://jamespropp.org/pims10.pdf` for slides for a talk on rotor-routing in more general graphs, and Google `"rotor-router"` to find lots of recent work in this general area.)

Rotor-router aggregation

To see rotor-router aggregation in action, visit
`http://rotor-router.mpi-inf.mpg.de/growing.mpg`.

A whole afternoon could be devoted to talks about what we know, and what we guess, about rotor-IDLA in $\mathbb{Z}^2$.

Theorem (Levine and Peres, 2010): The deviations from circularity have magnitude on the order of $\sqrt{r}\log r$ or smaller, where r is the of the radius of the growing cluster.

Even though this $\sqrt{r}\log r$ is a worse bound than the $\log r$ bound obtained for IDLA, experimental evidence suggests that deviations from circularity for rotor-router aggregation are smaller than typical deviations for IDLA.

Indeed, it is possible that for rotor-router aggregation, the radial fluctuations from circularity remain bounded as r goes to infinity!

Rotor-router competitive erosion

David Einstein did a simulation of derandomized competitive erosion in a disk (using one set of rotors for the particle emitted by the blue source at the top of the disk and another set of rotors for the particle emitted by the red source at the bottom of the disk):
`http://jamespropp.org/qmdle.gif`

The initial coloring of the disk is as far from equilibrium as possible.

The left panel shows the coloring, the middle panel shows the states of the rotors used by particles emitted by the blue source and the right panel shows the states of the rotors used by particles emitted by the red source.

The equilibrium state for the interface seems be the same circular-arc that is conjectured for the fully random case, and (just as in the case of IDLA), the derandomized version appears to have smaller deviations.

Rotor-router erosion

Recall that (internal) diffusion-limited erosion is an unstable interface process.

What happens when we derandomize it with rotor-routers?

Check out the pictures of rotor-DLE at
`http://www.math.cornell.edu/~levine/gallery/rotordledisc1000.png`

As in Schramm's (fully derandomized) rotor-DLA, we see very orderly spikes interspersed with dendritic structures displaying both orderly and random features.

Also note the very regular behavior of the rotors in some, but not all, of the region.

Random ballistic deposition

Tobias Friedrich did a simulation in which particles, starting at a source in the North, randomly step Southwest, South, or Southeast, and join an aggregate that is initially a horizontal line segment in the South.

As in DLA, random-looking dentritic structures form; see `http://jamespropp.org/snow_fullyrnd.mpg`.

Nonrandom ballistic deposition

Friedrich also did a version using 3-way rotor-routers: `http://jamespropp.org/snow.mpg`.

A striking feature of rotor-router deposition is that initially one sees very orderly structures, as in rotor-router aggregation, but then at some point a transition occurs, and one sees dendritic structures similar to the ones we see in rotor-DLE.

There should be a deterministic continuum interface model associated with Friedrich's rotor-router deposition model; it should exhibit a finite-time singularity, forming a cusp similar to what we see in the discrete model.

Slides for this talk are on-line at

`http://jamespropp.org/mathfest12b.pdf`