Rotor-Router Walk in Two Dimensions

Speaker/Performer James Propp
(UMass Lowell; visiting UCB and MSRI)

April 3, 2012

Slides for this talk are on-line at

http://jamespropp.org/csps12.pdf

Acknowledgments

Thanks to Doug Rizzolo for inviting me to give this talk.

This talk describes past and on-going work with Tobias Friedrich, Ander Holroyd, Lionel Levine, and Yuval Peres; with thanks also to David desJardins, Michael Kleber, and Oded Schramm.

Overview

In many number of fields of mathematics (e.g. number theory), many extremely deep and delicate theorems may be understood to assert that "thus-and-such behaves as if it were random".

Sometimes such assertions are proved on an ad hoc basis, but ideally one would want probabilistic methods that work in these non-probabilistic contexts.

I'm going to show you a very simple example in which we can get this kind of method.

It's an artificial example, because it's a deterministic system that doesn't occur in the wild; it was rigged to mimic a particular random system.

But it's of interest in its own right, and it may contain clues about the link between random and non-random processes.

What I won't talk about

A related example is rotor-router aggregation, which I won't have time to discuss, but you can read Lionel Levine and Yuval Peres' articles on the arXiv,

http://arxiv.org/abs/0704.0688

http://arxiv.org/abs/0712.3378

and you can go to Tobias Friedrich's stunning website

http://rotor-router.mpi-inf.mpg.de/

which shows images created with an algorithm Friedich developed with Levine, described in

http://arxiv.org/abs/1006.1003

The $\pi/8$ theorem for random walk in $\mathbf{Z}^2$

A particle that leaves $(0,0)$ and performs unbiased random walk in $\mathbf{Z}^2$ will hit the site $\{(0,0),(1,1)\}$ again with probability 1.

The probability that the particle will hit $(1,1)$ before it returns to $(0,0)$ (the "escape probability") is $p_{\mathrm{esc}} = \pi/8$.

(I'm not sure who first proved this, or the more general result with arbitrary targets, which involves more complicated rational functions of π.)

A much easier result is that a random walk from $(0,0)$ will hit $(1,0)$ before it returns to $(0,0)$ with probability $1/2$.

Rotor walk in $\mathbf{Z}^2$

To each site (i, j) in $\mathbf{Z}^2$ associate a 4-state "rotor" whose four states correspond to the four nearest neighbors of (i, j).

A particle executes a deterministic walk determined by two rules:

1. After the particle arrives at a target ($(0, 0)$ or $(1, 1)$), it gets replaced at the source $(0, 0)$. (It's helpful to have two copies of $(0, 0)$: one a source and one a target.)

Alternatively, one can imagine that when a particle arrives at a target, a new particle is placed at the source.

2. After the particle arrives at a non-target vertex (i, j), it advances the rotor at (i, j) to the next state in counterclockwise order, and moves to the neighbor of (i, j) indicated by the current state of the rotor.

Fylfot initial conditions

We also assume that the initial configuration of the rotors is as shown below (the black dot indicates where the origin is):

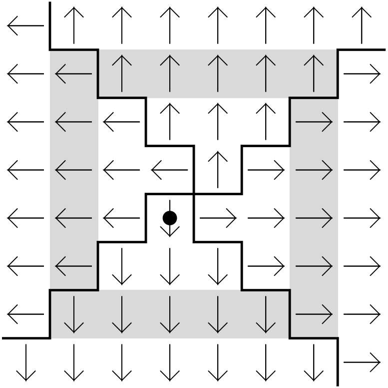

What it looks like

Go to http://www.cs.uml.edu/∼jpropp/rotor-router-model/

Click on "The Applet".

Change Graph/Mode to "2-D Walk".

(Note: This applet uses clockwise, rather than counterclockwise, progression of rotors.)

The weak $\pi/8$ theorem for rotor walk in $\mathbf{Z}^2$

Take the many-particles point of view (when a particle gets absorbed at a target, a new particle is released from the source).

Let $k(n)$ denote the number of particles absorbed at $(1,1)$ among the first n particles released from $(0,0)$.

Theorem: $k(n)/n \to p_{\mathrm{esc}} = \pi/8$.

The confluence property of rotor-routing

If n indistinguishable particles are placed at the source, the order in which you advance them (obeying the rotor-router rules) doesn't affect the number that get absorbed at each target.

(This property is often called the "abelian property".)

In particular, instead of letting one particle walk until it hits a target, and then another, and then another, etc. you could advance each particle one step, then another, etc.

Why the weak $\pi/8$ theorem is true

When n is really large compared to m, the first m steps of rotor-routing looks like a diffusion process on $\mathbf{Z}^2$.

In particular, the proportion of particles that have been absorbed at $(1,1)$ is close to $p_{\mathrm{esc},m} :=$ the probability that a particle released from $(0,0)$ gets absorbed at $(1,1)$ within m steps.

But $p_{\mathrm{esc},m}$ is close to p_{esc} when m is large.

This argument can be used to show that

$$k(n)/n - \pi/8 = O(1/(\log n)).$$

But with more work we can do much better.

The strong $\pi/8$ theorem for rotor walk in $\mathbf{Z}^2$

Theorem:

$$|k(n) - n\pi/8| = O(\log n).$$

(That is, $|k(n)/n - \pi/8| = O((\log n)/n) \ll O(1/\log n)$.)

For a full proof, see http://arxiv.org/abs/0904.4507: "Rotor Walks and Markov Chains" by Alexander Holroyd and James Propp.

Ingredients of the proof

1. Harmonic functions
2. Rotor values
3. Sum-rearrangement and telescoping
4. A recurrence lemma
5. A Green's function estimate
6. The truncated harmonic series estimate

1. Harmonic functions

Let $h(v)$ be the probability that a particle that executes a random walk starting from v will escape (i.e., arrive at $(1, 1)$ before arriving at $(0, 0)$).

(Technicality: We need two copies of $(0, 0)$.
At $(0, 0)_{\text{source}}$, $h()$ takes the value p_{esc};
at $(0, 0)_{\text{target}}$, $h()$ takes the value 0.)

Then $h()$ is harmonic away from the target set;
that is, for all v other than $(1, 1)$ and $(0, 0)_{\text{target}}$,
$h(v)$ equals the average of $h(w)$
as w varies over the neighbors of v.

2. Rotor values

Say that a particle at v has value $h(v)$.

We can assign values $h(e)$ to the rotor-states e so that when a particle updates the rotor at v from e to e' and moves to the associated neighbor w of v, the sum of the values is preserved:

$$h(v') + h(e') = h(v) + h(e).$$

(Call this the "conjugacy relation".)

Why can rotors be assigned values consistently?

For each vertex v, we can choose an arc e emanating from v and assign any value we like to $h(e)$. The conjugacy relation forces us to take $h(e') = h(v) + h(e) - h(v')$, $h(e'') = h(v) + h(e') - h(v'')$, etc., where e' is the cyclic successor of e at v, e'' is the cyclic successor of e' at v, etc., and where e' points from v to v', e'' points from v to v'', etc.

The only constraint we have to worry about is that when we've cycled around v, returning to e, our rule must assign $h(e)$ the same value as before.

That is, the quantities $h(e') - h(e)$, $h(e'') - h(e')$, $\ldots$ summed over a full period must equal 0.

But this sum equals
$(h(v) - h(v')) + (h(v) - h(v'')) + \cdots = \deg(v)h(v) - \sum_w h(w)$, where w ranges over the neighbors of v, and the harmonicity property of $h()$ says that this vanishes.

3. Rotor values (concluded)

When the particle travels from the source to the target set, the value of the particle plus the values of all the rotors that it visits along the way ("the value of the whole system") doesn't change.

When we move the particle from a target vertex t back to the source, the value of the rotors doesn't change and the value of the particle increases by $p - h(t)$, which is $p - 1 < 0$ if $t = (1, 1)$ and $p - 0 = p > 0$ if $t = (0, 0)_{\text{target}}$.

(Here $p = p_{\text{esc}}$ for short.)

3. Sum-rearrangement and telescoping

How does the value of the system change when we send n particles through it, all of which start at $(0,0)_{\mathrm{source}}$ and end at $(1,1)$ or $(0,0)_{\mathrm{target}}$?

On the one hand, the change equals

$$(k)(p-1)+(n-k)(p)=np-k;$$

on the other hand, the change equals the change in the total value of the rotors.

So if we can bound how much the total value of the rotors has changed, we can bound $k-np$.

4. A recurrence lemma

We use an ad hoc lemma that says that the first n particles all stay within distance $O(n)$ of the origin, so that there aren't that many rotors contributing to the change in the total value of the rotors.

(This is where we make use of the very special initial conditions.)

5. A Green's function estimate

We express $h()$ in terms of the discrete Green's function for $\mathbf{Z}^2$ and show that the total change in the value of the rotors in the nth square shell surrounding the origin can't contribute more than $O(1/n)$.

6. The truncated harmonic series estimate

The sum of the first $O(n)$ terms of the harmonic series is $O(\log n)$.

What's missing from the proof

1. The conclusion is too weak.
2. The hypotheses are too strong.

1. The conclusion is too weak

Empirically, it appears that the upper bound $\log n$ is far from tight.

Indeed, for more than half of the values of n between 1 and 10^4,

$$|k(n) - (\pi/8)n| < 1/2.$$

Note that for a third of these values of n, the fractional part of $(\pi/8)n$ lies between $1/3$ and $2/3$, so for these n there doesn't exist *any* integer k such that $|k - (\pi/8)n| < 1/3$!

Judging from the data, it's conceivable that $k(n) - (\pi/8)n$ is *bounded* (so that $k/n - \pi/8 = O(1/n)$).

This is clearly the most one could hope for, since $|k(n) - (\pi/8)| > 1/3$ for a set of n's of density $1/3$.

A digression about regression

There's a sense in which rotor-routing may "tune itself to $\pi/8$" with error $o(1/n)$ (even though I just explained to you why you can't expect to have $|k(n)/n - \pi/8| = o(1/n)$).

Specifically, if you draw the regression line through the points $(i, k(i))$ for all i between 1 and n, its slope seems to differ from $\pi/8$ by something like $1/n^{3/2}$.
I'll touch upon this in my upcoming probability seminar talk.
The method of Holroyd-Propp seems to fail completely in this context.
What can take its place?

2. The hypotheses are too strong

The Holroyd-Propp proof is not general enough: the initial conditions to which the method of proof applies are very special.

The proof does go through if the initial conditions are modified in only finitely many places. But that's not saying much.

To get a more general result, we would need a more flexible way to estimate the amount by which the rotor values change when n particles pass through the system.

To infinity and beyond

Lionel Levine and I have empirically studied what happens when all the rotors are initially aligned.

Particles can wander off to infinity, but that's okay: it can be shown that each rotor gets updated only finitely often, so it still makes sense to ask what happens when all the particles have either hit a target or wandered off to infinity.

A theorem of Holroyd and Propp (originally proved by Schramm) shows that the number of escapes to infinity in the first n trials (call it $e(n)$) is $o(n)$, and that $k(n)/n$ goes to $\pi/8$.

We don't know how big $e(n)$ is; we know that $e(n)/n$ goes to 0, but it appears to go to 0 quite slowly (empirically, one might even think it was converging to a non-zero constant, even though we've proved that it doesn't).

Getting random

What about random initial conditions?

I believe that there's a probability measure on initial conditions that's more natural than IID initial conditions, namely "maximally random acyclic initial conditions" where the rotors can't exhibit any cycles.

(This makes sense, since cycles can't spontaneously appear: one acyclic configuration can only yield another one. See the next slide for more on this.)

I'll touch on this in my probability seminar talk too. There I'll mostly focus on the case of finite-state Markov chains, since that's a case that I understand much better. I haven't thought through the infinite-state case.

Why cycles go away

Suppose that the rotor at v points to w, *and* that the particle has visited v (but is no longer at v). Then the last time the particle left v, it went to w.

Now suppose that the rotor-configuration has a cycle; that is, there are vertices $v_0, v_1, v_2, \ldots, v_m = v_0$ such that the rotor at v_i points to v_{i+1} ($0 \leq i \leq m-1$), and the particle is no longer at any of the v_i's.

This gives a contradiction (consider the last time the particle was at one of the v_i's)!

So cycles are transient (aside from a possible cycle involving the rotor at the vertex that the particle currently occupies).

One can use this to show that the IID measure on initial conditions is not preserved by rotor-routing.

What I suspect

Start with IID initial conditions for the rotors on $\mathbf{Z}^2$.

Perform "cycle popping" a la David Wilson until all cycles are gone.

I believe that the resulting measure is preserved by rotor-routing.

Are there other probability measures that are preserved by rotor-routing?

Maybe the one I've described is the only one.

Random versus non-random routing

One thing I'd like is a way to put random walk on $\mathbf{Z}^2$, and in particular the $|k(n) - pn| = O(\sqrt{n})$ estimate, into the same framework as rotor-router walk.

I know how to do this for finite-state Markov chains, but there are technical problems for random walk on $\mathbf{Z}^2$.

I'd like to show that when n successive particles go through the system, the value of the rotors typically changes by $O(\sqrt{n})$.

Can anyone see a way to prove this?

Getting rid of the targets

Suppose we put down fylfot initial conditions for the rotors but *don't* move particles from $(1, 1)$ back to $(0, 0)$; we just treat $(1, 1)$ as a site like any other, with a rotor of its own that determines where particles that leave $(1, 1)$ go next.

Then the particle that leaves $(0, 0)$ will execute an orderly sequence of ever-growing inbound and outbound spirals.

Let $n_t(v)$ be the number of times site v gets visited up to time t.

A theorem of Holroyd and Propp says that as $t \to \infty$, the ratio $n_t(v)/n_t(w)$ converges to the ratio $\mu(v)/\mu(w)$, where μ is the measure on $\mathbf{Z}^2$ that's stationary for the random walk.

In this case the stationary measure is (non-finite) uniform measure, so $\mu(v)/\mu(w) = 1$.

And indeed one has $h(v) - h(w)$ bounded, so the convergence to 1 is fairly fast: $n_t(v)/n_t(w) - 1 = O(1/t^{1/3})$.

A beautiful process in search of a beautiful theorem

If one does target-less rotor-router walk on $\mathbf{Z}^2$ with all rotors initially pointing in the same direction, Lionel Levine noticed that something beautiful happens:

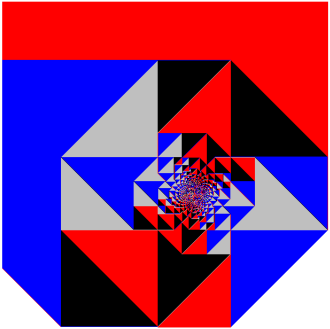

It would be great if we could prove something about this!

Directed walk in a quadrant

Suppose we only allow steps to the North and the East.

Even this rotor-walk has interesting (purely empirical) patterns.

See http://jamespropp.org/quincunx.gif (or http://jamespropp.org/galton.swf for the "guided tour").

Discrete analytic function theory?

I called $h(v') + h(e') = h(v) + h(e)$ a "conjugacy relation".

I suspect that there's a way to see h as a discrete analytic function of some kind, where its values on the edges are conjugate to its values on the vertices.

Can this point of view be pushed through? (It's not the standard flavor of discrete analytic function theory.)

Is complex dynamics already lurking here?

Such a point of view would probably improve our understanding of rotor-router aggregation.

Tobias Friedrich's pictures of rotor-router aggregation (see especially the "Irdu blob") bear an amazing resemblance to pull-backs of Apollonian gaskets in $\mathbf{C}$ via the map $z \mapsto 1/z^2$ (see http://jamespropp.org/RRcircles.pdf).

What kind of math could explain this?

The last slide of this talk

I'm happy to talk about this stuff further with anyone who's interested;

My office hour is at 11 am on Tuesdays and Thursdays in 1063 Evans.

I divide my out-of-the-house time between Evans, MSRI, and nearby cafes.

Slides for this talk are on-line at

http://jamespropp.org/csps12.pdf