Rotor walks and Markov chains

Jim Propp
(U. Mass. Lowell)

April 7, 2009

(joint work with Ander Holroyd)

(with thanks also to Hal Canary, Matt Cook, Dan Hoey, Michael Kleber, Lionel Levine, Yuval Peres, and Oded Schramm)

Slides for this talk are on-line at
`jamespropp.org/probsem09a.pdf`
and the article on which the talk is based is now available at
`http://www.math.ubc.ca/~holroyd/papers/rr.pdf`

I. Context

Non-probabilistic probability theory

Probabilistic phenomena keep turning up in non-probabilistic contexts.

E.g., the binary digits of π are not random in the usual measure-theoretic sense, or in the information-content sense, or in the algorithmic sense; yet if you pretended they were random, you could experimentally "verify" the Central Limit Theorem.

To explain this sort of phenomenon, we may need to go back to pre-Kolmogorov, frequentist definitions of randomness.

Probability vs. frequency

Let P be a stochastic matrix with eigenvalue 1 having multiplicity 1, and let π be the unique probability vector with $\pi P = \pi$.

Fundamental theorem of recurrent finite Markov chains (probability version):

Suppose $X_0, X_1, \ldots$ are random variables taking values in $\{1, 2, \ldots, m\}$ such that X_{n+1} is conditionally independent of $X_0, \ldots, X_{n-1}$ given X_n, with

$$\mathrm{Prob}(X_{n+1} = j \mid X_n = i) = p_{i,j}$$

for all i, j, n.

Then for each i, i almost surely occurs in $X_0, X_1, \ldots$ with asymptopic frequency $\pi(i)$; that is, a.s.

$$N(i; n)/n \to \pi(i)$$

as $n \to \infty$ where

$$N(i; n) = |\{k < n : \; X_k = i\}|.$$

Fundamental theorem of recurrent finite Markov chains (frequency version):

Suppose $x_0, x_1, ...$ are in $\{1, 2, ..., m\}$ such that for all i, j, i is followed by j with asymptotic frequency $p(i, j)$.

Then for all i, i occurs in $x_0, x_1, \ldots$ with asymptotic frequency $\pi(i)$.

Local discrepancy and global discrepancy: an example

Let P be a stochastic matrix with two sink vertices b, c with $p(b, b) = p(c, c) = 1$ such that the chain a.s. hits b or c from any starting state i.

For all i let $h(i)$ be the probability that a particle released from i will be absorbed at b (rather than c), so that $h(b) = 1$ and $h(c) = 0$.

Note that

$$\sum_j p(i, j)h(j) = h(i)$$

for all i; i.e., h is harmonic. (More compactly: $Ph = h$.)

Suppose a particle repeatedly walks from source vertex a until it hits a sink (b or c), restarting at a after each absorption.

Consider a finite path

$$x_0, x_1, \ldots, x_{T-1}, x_T$$

in which the particle starts and ends at a, arriving K times at b and $N - K$ times at c.

If the path were a typical sample path of the walk, we would expect $K \approx Np$ where $p = h(a)$.

Let's find a formula for $K - Np$.

Write

$$\begin{aligned} K - Np &= (K - Kp) - (Np - Kp) \\ &= K(1-p) + (N-K)(0-p) \\ &= K(h(b) - h(a)) \\ &\quad +(N-K)(h(c) - h(a)) \\ &= \sum (h(x_t) - h(x_{t-1})) \end{aligned}$$

where the sum is over all $1 \leq t \leq T$ with x_{t-1} not a sink.

We can gather together terms of the sum for which $x_{t-1} = i$ and $x_t = j$, obtaining the double sum

$$\sum_i \sum_j N(i,j)(h(j) - h(i))$$

where $N(i,j)$ is the number of times the particle moved from i to j up to time T, and where i is not a sink.

So the discrepancy $D := K - Np$ satisfies

$$D = \sum_i \sum_j N(i,j)(h(j) - h(i))$$

while the harmonicity of $h(\cdot)$ at i gives

$$\sum_j (N(i)p(i,j))(h(j) - h(i)) = 0$$

(where $N(i) = \sum_j N(i,j)$); hence D equals

$$\sum_i \sum_j (N(i,j) - N(i)p(i,j))(h(j) - h(i)).$$

That is, the *global discrepancy* $K - Np$ can be written as the sum of the *local discrepancies*

$$(N(i,j) - N(i)p(i,j))(h(j) - h(i)).$$

If $x_0, x_1, ...$ is given by a random process, the local discrepancies are $\approx N^{1/2}$, so the global discrepancy is $\approx N^{1/2}$ too.

If $x_0, x_1, ...$ has smaller-than-random ($subrandom$) local discrepancies (say $\approx N^{\alpha}$ with $\alpha < 1/2$) then $K - Np$ will be subrandom too.

II. Rotor-routing in general

Suppose we have a (not necessarily finite) Markov chain that is locally finite (for each i, $p(i, j) = 0$ for all but finitely many j) and "rational" (all $p(i, j)$ are rational numbers).

How would we follow these transition probabilities $p(i, j)$ so as to miminimize local discrepancy?

For each state i, in lieu of a d-sided die (where d is a common denominator of the $p(i, j)$'s), we use a periodic process of period d, such that each j occurs with frequency $p(i, j)$.

Whenever we arrive at a state i, we choose whichever j is next in succession.

E.g., if $p(a,b) = \frac{1}{3}$ and $p(a,c) = \frac{2}{3}$, our periodic process could be *cbccbccbc*....

The 1st time we leave a, we go to c;
the 2nd time we leave a, we go to b;
the 3rd time we leave a, we go to c;
etc.

Note that the local discrepancies associated with transitions from a to b and from a to c stay bounded, and indeed vanish whenever the number of visits to a is divisible by 3.

Equivalently, we can imagine that vertex a has a three-state **rotor** associated with it that has one state pointing from a to b and two states pointing from a to c. The states are cyclically ordered. To move the particle forward one step, advance the rotor to its next state and move the particle to the state that the rotor points to.

E.g., if there is a particle at a and the rotor at a currently points from a to b, rotate the rotor so that it points from a to c and then route the particle from a to c.

This is the **rotor-router process**. It is deterministic, but more importantly, it minimizes local discrepancy.

History:

- Eulerian walkers model
- load-balancing
- whirling tours

Recent applications:

- diffusion-limited aggregation (Levine and Peres)
- rumor-spreading (Friedrich)

It's instructive to think about simple Markov chains, such as ordinary random walk on

```
o-----o-----o-----o
0     1     2     3
```

Suppose the initial setting of the rotors is

```
o->   o-> <-o   <-o
0     1     2     3
```

Then a particle that starts at 0 goes

$$010123210123...$$

Note that the proportion of the time that the particle spends in the four states converges to $(1/6, 2/6, 2/6, 1/6)$, which is the stationary measure for the chain.

Theorem (Holroyd-Propp): For any finite rational Markov chain, if $x_0, x_1, \ldots$ is a rotor-router walk on the states of the chain, then

$$|t\pi(v) - n_t(v)| \leq \text{[explicit constant]}$$

for all $t > 0$, where $n_t(v)$ is the number of $0 \leq s < t$ with $x_s = v$.

Intuition for the weaker claim $n_t(v)/t \to \pi(v)$: If we treat $n_t(\cdot)$ as a vector and let $\pi_t = n_t/t$, then $n_t P - n_t$ is bounded because of the definition of rotor-routing, so $\pi_t P - \pi_t \to 0$, so $\pi_t \to \pi$ (by the uniqueness of the stationary measure).

Now use the same graph again, but think of states $b = 0$ and $c = 3$ as absorbing states (or sinks) and state $a = 1$ as the source (gambler's ruin with $n = 3$).

Suppose the initial setting of the rotors is

```
o      o-> <-o      o
0      1      2     3
```

(the rotors at 0 and 3 are unused and hence omitted).

Then a particle that starts at 1 goes

$$1\mathbf{0}12\mathbf{3}1\mathbf{0}121\mathbf{0}12\mathbf{3}1\mathbf{0}121\mathbf{0}12\mathbf{3}\ldots$$

Note that the proportion of the time that the particle arrives at the two sinks converges to 2/3 and 1/3 which are the respective absorption probabilities.

Equivalently, we can imagine that each time a particle gets absorbed at a sink, it stays there, and a new particle is released from the source.

Or: We can imagine that many particles start at the source, and we let them successively travel through the state-diagram of the Markov chain until absorption, under the stricture that a particle can't start to move from the source until its predecessor has been absorbed by a sink.

But is this stricture necessary?

Abelian property: If there are multiple particles on the state-diagram of a Markov chain, so that at each instant you have a choice of which one to advance via rotor-routing, the choices you make don't matter. The end result is the same when all the dust settles (i.e., when all particles have been absorbed).

So the stricture can be dropped.

This gives us intuition for why the absorption frequencies in a rotor-router simulation agree with absorption probabilities for random simulation.

E.g., in our example, when the number of particles N that start at 1 is large, and we move them in tandem (each takes a step, then each takes a step, etc.), $\approx N/2$ of them arrive at 0 in 1 step, $\approx N/8$ of them arrive at 0 in 3 steps, $\approx N/32$ of them arrive at 0 in 5 steps, etc. So

$$\approx N/2 + N/8 + N/32 + \ldots = 2N/3$$

of the particles get absorbed at 0.

(At each stage there is new error introduced, but after $\log 1/\epsilon$ steps all but ϵN of the particles have been absorbed.)

Note that if we look at the states of the rotors at the instant of absorption, only three possibilities keep on occurring:

```
o     o->   o->   o
0     1     2     3
```

and

```
o   <-o   <-o     o
0     1     2     3
```

and

```
o   <-o     o->   o
0     1     2     3
```

The fourth configuration

```
o    o-> <-o    o
0    1     2    3
```

can't occur after both states 1 and 2 have been visited, since that would require that the last time the particle visited 1 it went to 2 and the last time the particle visited 2 it went to 1, which is impossible since the particle is currently at 0 or 3.

Even if we allow ourselves to repeatedly choose between adding a particle at 1 and adding a particle at 2 (letting the added particle walk until it is absorbed), we cannot ever arrive at the fourth rotor-configuration (with all particles absorbed) once states 1 and 2 have both been visited.

We say that the first three rotor-configurations are **recurrent** and the fourth is **transient**.

More precisely, a rotor-configuration is recurrent if it can be obtained from any other rotor-configuration by a succession of operations each of which adds a particle to the state-diagram and lets it do rotor-walk until it is absorbed.

The recurrent rotor-configurations are precisely those that contain no cycles; see "Chip-firing and rotor-routing on finite digraphs" by Holroyd, Levine, Mészáros, Peres, Propp, and Wilson: `arXiv:0801.3306`

This is why we adopt the convention of rotating the rotor before we route the particle, rather than the reverse.

The Holroyd-Propp article gives other examples of the concentration phenomenon for rotor-routing, where rotor-routing for N trials give approximations to probabilistic quantities (e.g. hitting times) that differ from the true values by $O(1/N)$, in contrast to ordinary random simulation, which gives errors on the order of $1/\sqrt{N}$.

Also, if the rotor-and-particle system returns to a configuration it's already visited, then it will behave thereafter in a periodic fashion, and its behavior over the course of one period will give exact values of the probabilistic quantity in question.

(See our earlier examples of four-state Markov chains.)

Some quantities (like mean squared hitting time) do not fit into the rotor-router framework directly, but can be made to do so if we use more than one kind of particle and have more complicated routing rules.

Markov chains with irrational transition probabilities can also be fit into the framework of discrepancy-minimization, but we can no longer use a finite-state router at each vertex.

E.g., if $p(1,2) = \alpha$ and $p(1,3) = 1 - \alpha$ with α irrational, then there is a unique protocol for routing the particle so that after the particle has left 1 for the nth time, the number of times it went to 2 is the integer closest to $n\alpha$ and the number of times it went to 3 is the integer closest to $n(1 - \alpha)$.

For out-degree > 2, things are a little more complicated; see the section of Holroyd-Propp on "stack-walk".

A nice example of rotor-routing is the "goldbug walk" on $\{-1, 0, 1, ...\}$ where states $b = -1$ and $c = 0$ are absorbing, all other states are transient, and $p(i, i-2) = p(i, i+1) = \frac{1}{2}$ for all $i \geq 1$.

This walk has leftward drift, so the probability of absorption in $\{b, c\}$ is 1.

To see what happens when all rotors initially point to the right, run

`http://jamespropp.org/rotor-router-model/`

with Graph/Mode set to 1-D Walk.

If we attend to where the successive particles end up, we see that the whole system, made of infinitely many $(1/2, 1/2)$ rotors, behaves like a single $(\alpha, 1-\alpha)$ rotor, with $\alpha = (-1 + \sqrt{5})/2$.

By the abelian property, we could also put lots of particles at 0 at the start (N, say) and let them do rotor-walk in tandem; $N\alpha \pm O(1)$ of them will be absorbed at -1 and $N(1-\alpha) \pm O(1)$ will be absorbed at 0.

The same is true if the rotors initially point to the left, except that one particle will never get absorbed; it just wanders off to the right forever.

To take full advantage of the abelian property in situations where some of the particles wander off to infinity, it's helpful to define simulation in "transfinite time".

E.g., in the goldbug system we let the first particle wander off to infinity, leaving leftward-pointing rotors in its wake, and "thereafter" continue to add other particles, all of which get absorbed at -1 and 0.

Specifically, we define rotor-simulation indexed by ordinals of the form $m\omega + t$ where m, t are non-negative integers.

This may be possible for some m and not others.

E.g., consider ordinary rotor-walk on $\mathbf{Z}$ where all rotors to the left of 0 point to the right and all other rotors point to the left. We can simulate from time 0 to time ω (the particle goes to $+\infty$) and from time ω to time 2ω (the particle goes to $-\infty$), but from time 2ω to time 3ω, each site gets visited infinitely often so it is impossible to say what state the rotors are in at time 3ω.

On the other hand, consider the Markov chain whose state space is $\{0, 1, 2, \ldots\}$ where state $b = 0$ is absorbing and $p(i, i-1) = \frac{1}{3}$ and $p(i, i+1) = \frac{2}{3}$ for all $i > 0$ (biased random walk with rightward drift-rate $\frac{1}{3}$).

Regardless of the initial setting of the rotors, the particle runs off to infinity during its nth run if and only if it did *not* run off to infinity on the $n-1$st run, for all $n \geq 3$ (but not necessarily $n = 2$), in agreement with the fact that the escape probability is $\frac{1}{2}$.

In this case, the transfinite rotor-walk is defined for *all* times of the form $m\omega + t$.

What controls the well-definedness of transfinite rotor-walk is the following fact:

Recurrence/transience dichotomy:

An infinite rotor-walk on a graph either visits every vertex only finitely often or visits every vertex infinitely often.

(Proof: Every neighbor of a vertex that gets visited infinitely often must be visited infinitely often.)

If a vertex gets visited only finitely often, the rotor at that vertex has a limiting (indeed, eventual) setting; if a vertex gets visited infinitely often, the limiting setting does not exist.

Rotor-walks can be *more recurrent* than random walks: Landau and Levine show that for rotor-walk on an infinite binary tree with source at the root, if we set the rotors so that the first time a particle leaves a vertex v it goes toward the root, the particle will visit the root infinitely often.

Rotor-walks can be more *more transient* than random walks: sample paths can go off to infinity (this a.s. doesn't happen for random walk)

However, if we look at matters in the proper frequentist way, a rotor-walk *cannot* be more transient than its random counterpart, asymptotically.

Theorem (Schramm): For transfinite rotor-walk, let I_n be the number of times the walk goes to infinity before the nth return to a. Then $\limsup_{n \to \infty} I_n/n$ is at most the probability that random walk started from a never returns to a.

In particular, if the walk is recurrent, $I_n/n \to 0$.

III. Rotor-routing on $\mathbf{Z}^2$

For rotor-walk in $\mathbf{Z}^2$, it's natural to have rotors that cycle through the four directions as N,E,S,W,N,E,S,W,...

Cooper and Spencer proved that there is a small finite constant C (less than 10) such that if one starts N particles at the origin and lets them execute T steps of tandem rotor-walk, the discrepancy between

# of particles at site v at time T

and

N times $p^{(T)}(0, v)$

is at most C, for *all* v, N, T and all initial configurations of the rotors!

In fact, Cooper and Spencer showed that an analogous bounded discrepancy property holds for any initial distribution of the N particles on the even sublattice $\{(i,j) \in \mathbf{Z}^2 : i+j \text{ is even}\}$ (not just the distribution where all the particles start at 0).

The bound is *independent* of the initial distribution of the particles.

Boundedness also holds if the rotors cycle in the pattern N,S,E,W,N,S,E,W,... (thought the constant C is different).

Boundedness also holds for $\mathbf{Z}^d$ for all d (though the constants are worse, and might grow quickly with d).

For the rest of this talk I'll focus instead on absorption probabilities.

Let $a = (0, 0)$, $b = (1, 1)$, $c = (0, 0)$.

It is known that the probability that a particle emitted from a arrives at b before it arrives at c is $p = \pi/8$.

To see how closely rotor-walk concentrates around this value, see

`http://jamespropp.org/`
`rotor-router-model/`

with Graph/Mode set to 2-D Walk.

(What sort of stable structures does the rotor-configuration exhibit?)

Theorem (Holroyd-Propp): With rotors that cycle clockwise and initial conditions

```
        .  .  .  .
    N  N  N  N  N  N  N  E
    W  N  N  N  N  N  E  E
 .  W  W  N  N  N  E  E  E  .
 .  W  W  W  N  E  E  E  E  .
 .  W  W  W  W  S  E  E  E  .
 .  W  W  W  S  S  S  E  E  .
    W  W  S  S  S  S  S  E
    W  S  S  S  S  S  S  S
        .  .  .  .
```

we have $D = K - Np = O(\log N)$, where N is the number of particle emitted from a, K is the number of particle absorbed at b, and $p = \pi/8$.

Proof sketch: Recall from the beginning of the talk that the discrepancy D equals

$$\sum_v \sum_w (N(v,w) - N(v)p(v,w))(h(w) - h(v)).$$

The rotor-router protocol guarantees that the differences $N(v,w) - N(v)p(v,w)$ are uniformly bounded, so, using standard facts about the potential kernel for two-dimensional random walk, we can show that the inner sum is on the order of $1/|v|^2$.

If we had to sum over all $v \in \mathbf{Z}^2$, this would diverge like the harmonic series (since the number of points at distance $n \pm \frac{1}{2}$ from the origin is on the order of constant time n).

However, the initial conditions we picked guarantee that when N particles have gone through the system and been absorbed, the sites that have been visited lie in the $2N$-by-$2N$ square centered at $(\frac{1}{2}, \frac{1}{2})$ (the combinatorial details are omitted here), and for sites that have not been visited, the contribution to the discrepancy D is 0.

Hence, the global discrepancy is bounded by the harmonic series truncated after $O(N)$ terms, which is $O(\log N)$.

The method of proof works for any finite target set: Let p be the probability that a random walk in $\mathbf{Z}^2$ that walks from source vertex $(0,0)$ until it hits the finite target set B stops at a particular vertex b in B. If one performs N successive runs of a rotor-router walk in $\mathbf{Z}^2$ from $(0,0)$ to B, the number of runs that stop at b is $Np \pm O(\log N)$.

What's wrong with this theorem:

It's not general enough.

E.g., the concentration phenomenon seems to be just as strong if we use the initial configuration

```
          .   .   .   .
    E   E   E   E   E   E   E   S
    N   E   E   E   E   E   S   S
.   N   N   E   E   E   S   S   S   .
.   N   N   N   E   S   S   S   S   .
.   N   N   N   N   W   S   S   S   .
.   N   N   N   W   W   W   S   S   .
    N   N   W   W   W   W   W   S
    N   W   W   W   W   W   W   W
          .   .   .   .
```

even though the proof given above doesn't apply.

What's wrong with this theorem:

It's not sharp enough.

The observed discrepancy D_N after N trials seems to be a lot less than $\log N$.

In 10,000 trials, $|D_N| < 0.5$ for $5,070$ of the trials. That is, more than half the time, the number of absorptions at b during the first N trials is equal to the integer closest to Np.

We have $|D_N| < 2.05$ for all $N \leq 10^4$.

Is $|D_N|$ bounded? Unknown!

Yuval Peres points out that if the summands in our truncated harmonic series are uncorrelated, we would expect the global discrepancy to behave like the random sum $\pm 1 \pm \frac{1}{2} \pm \frac{1}{3} \pm \ldots$, which is a.s. bounded.

For those who like to code:

What happens if all the rotors are initially lined up?

Some of the trials result in escape to infinity; some result in capture at b; some result in capture at c. We know that these occur with asymptotic frequencies 0, $\pi/8$ and $1 - \pi/8$, but how rapid is the convergence?

IV. Open problems

What about $\mathbf{Z}^3$?

About five years ago I did a transfinite rotor-router simulation of walk on $\mathbf{Z}^3$ with $N = 10^6$, with all rotors initially aligned; it gave the escape probability to four (but not five) significant figures. So in this case discrepancy is almost certainly not bounded, but it might be smaller than $O(\sqrt{N})$.

Joel Spencer and I ask what happens if the rotors are set up in $\{(x, y, z) : x, y, z \in \mathbf{Z}\}$ so that the first time a particle visits (x, y, z) it moves to the neighbor for which $|z|$ is smallest, unless z is already 0, in which case the particle moves to the neighbor for which $|y|$ is smallest, unless y is already 0, in which case the particle moves to the neighbor for which $|x|$ is smallest.

Is this rotor-walk recurrent?

Landau and Levine's result about recurrent rotor-walk on the infinite binary tree requires a rather constrained initial setting of the rotors. If the initial setting of the rotors is random, will the rotor-walk be transient almost surely?

(Side issue: What do we mean by random? What *should* we mean? On finite graphs, the uniform measure on rotor-settings is in many ways less natural that the uniform measure on the set of *recurrent* rotor-settings; how does this idea carry over into the setting of infinite Markov chains?)