Random walk and random aggregation, derandomized

Jim Propp
(U. Wisconsin)

January 30, 2006

(based on articles in progress with Ander Holroyd and Lionel Levine;

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

Slides for this talk are on-line at
`http://www.math.wisc.edu/~propp/NEU.pdf`

I. A motivating example: limit-shape theorems for tilings

An *Aztec diamond* of order n is the union of all the squares $[n, n+1] \times [n, n+1]$ that fall inside the region

$$\{(x, y) : |x| + |y| \leq n + 1\}.$$

A *domino* is a union of two such squares that share an edge. A *domino-tiling* of an Aztec diamond of order n is a set of $n(n+1)$ dominoes whose interiors are disjoint and whose union is the Aztec diamond.

Theorem (Elkies, Kuperberg, Larsen, and Propp): The Aztec diamond of order n has exactly $2^{n(n+1)/2}$ domino-tilings.

The Elkies et al. article contained four proofs, one of which (by Kuperberg and Propp) used a combinatorial construction ("domino-shuffling") that can also be used to generate random domino-tilings of Aztec diamonds.

Theorem (Jockusch, Propp, Shor): As one runs domino-shuffling for infinite time, the boundary of the "temperate zone" (informally, the part of the tiling where horizontal and vertical dominoes are intermixed) of a random domino-tiling of the Aztec diamond of order n converges with probability 1 to a perfect circle.

Therefore, "most" domino tilings of a large Aztec diamond have a temperate zone that is "close to round".

But, how can we generate a specific one that has a close-to-round temperate zone?

And: What deterministic procedure for running domino-shuffling forever is guaranteed to yield a circle in the limit?

We want some way to derandomize the random construction.

Let's switch to a simpler random construction.

II. Derandomization of a simple random walk

Simple random walk on $\mathbf{Z}^2$: For any two vertices $v, w \in \mathbf{Z}^2$, the transition probability $p(v, w)$ (the probability that a particle at v moves to w at the next time step) is $\frac{1}{4}$ if w is one of the four nearest neighbors of v and 0 otherwise.

This random walk is *recurrent*: With probability 1, each vertex in $\mathbf{Z}^2$ gets visited infinitely often.

Fact (Polya?): If a particle starts at $(0,0)$ and does random walk in $\mathbf{Z}^2$ until it either hits $(1,1)$ or returns to $(0,0)$, the probability that it hits $(1,1)$ before returning to $(0,0)$ is exactly $\pi/8$.

Hence, if we modify the walk so that whenever the particle arrives at $(1,1)$ it gets shunted immediately to $(0,0)$, then as $N \to \infty$ the number of visits to $(1,1)$ up to time N divided by the number of visits to $(0,0)$ up to time N converges to $\pi/8$, with probability 1.

How can we make this constructive?

Propp-Schramm: Suppose a particle visits sites $s_1, s_2, s_3, \ldots$ in $\mathbf{Z}^2$, where $s_1 = (0,0)$ and s_{n+1} is a nearest neighbor of s_n for all $n \geq 1$, except in the case where s_n is $(1,1)$, in which instance s_{n+1} is $(0,0)$.

Suppose each site in $\mathbf{Z}^2$ occurs in the sequence $s_1, s_2, s_3, \ldots$ infinitely often.

Suppose moreover that for every site $s \in \mathbf{Z}^2$ other than $(1,1)$ and for every site t adjacent to s, on any n successive visits to s the particle next goes to t $n/4 \pm o(n)$ times ("control of local discrepancy").

Then $(0,0)$ and $(1,1)$ both occur with well-defined density in $s_1, s_2, \ldots$, and the ratio of the latter density to the former is $\pi/8$.

By making the local discrepancy (the "$o(n)$") smaller, we can speed the convergence to $\pi/8$.

The smallest possible local discrepancy is gotten by using a "rotor-router" at each site: e.g., each time the particle leaves a site, it goes in the direction 90 degrees clockwise from whatever direction it went the last time it left that site.

(Physicists invented this rule ten years ago as an example of "self-organized criticality", and computer scientists introduced it as a protocol for load-balancing of processors; but neither group realized that the rotor-walk mechanism is applicable to estimation of properties of random walk.)

Ordinary Monte Carlo: If some quantity of interest, μ, can be expressed as $E(X)$ for some random variable X, then we can estimate μ by

$$\hat{\mu}_n = (X_1 + X_2 + \ldots + X_n)/n,$$

where $X_1, \ldots$ are i.i.d. instances of X.

If $\mathrm{Var}(X) < \infty$, the root-mean-square error of our estimate is $O(1/\sqrt{n})$, and if we have a bound on $\mathrm{Var}(X)$, the central limit theorem will give us asymptotic confidence intervals for μ.

In practice, the X_k are not really random but are generated by a deterministic algorithm (whose output behaves in many respects like the output of a random process). This makes the applicability of the central limit theorem harder to make sense of.

Quasi Monte Carlo: Replace random $X_1, X_2, \ldots$ by "quasirandom" $x_1, x_2, \ldots$ to get $|x_1 + x_2 + \ldots + x_n - n\mu|$ smaller than $\sqrt{n}$ and $|(x_1+x_2+\ldots+x_n)/n-\mu|$ smaller than $1/\sqrt{n}$.

If we can find a constant c such that $|x_1 + x_2 + ... + x_n - n\mu| < cn^{\alpha}$ for all n, with $\alpha < \frac{1}{2}$, then we get better estimates for μ than are given by Monte Carlo, surrounded by certainty intervals rather than confidence intervals.

The best α we can hope for is $\alpha = 0$; in this case the discrepancy $|(x_1 + x_2 + \ldots + x_n) - n\mu|$ is $O(1)$, i.e., *bounded*.

Recall: If a particle starts at (0,0) and does unbiased random walk in the infinite square grid, the probability p that it will arrive at (1,1) before it ever returns to (0,0) is $\pi/8$.

If we do n independent trials, the number of successes divided by the number of trials should be close to $\pi/8$, with an error on the order of $1/\sqrt{n}$.

Equivalently, the number of successes minus $\pi/8$ times the number of trials (write this "global" discrepancy as D_n) should be on the order of $\pm\sqrt{n}$ if we do independent random trials.

For $n = 10^4$, under random simulation, we expect $|D_n| \approx 50$.

Under quasirandom simulation, with rotor-routers, the n trials aren't independent, or even random — yet D_n seems to be bounded!

See demo at

`http://www.math.wisc.edu/~propp/rotor-router-1.0`

In 10,000 trials, $|D_n| < 0.5$ for $5,070$ of the trials. That is, more than half the time, the number of successes after n trials is equal to the integer closest to $p = \pi/8$ times the number of trials.

We have $|D_n| < 2.05$ for all $n \leq 10^4$.

Does $|D_n|$ stay bounded as $n \to \infty$?

Unknown!

III. Highly recurrent walk

For analogous processes in 1 dimension, which have better recurrence properties than 2 dimensional random walk, boundedness of D_n can be proved rigorously using harmonic functions on Markov chains.

Let h be the harmonic function on S defined by $h(s)$ = the probability that the Markov chain started from s hits S_1 before S_2, and define

$$\|\nabla h\| = \sum_x \max\{|h(x) - h(y)| : p(x,y) > 0\}.$$

Theorem (Holroyd and Propp): Suppose

$$\|\nabla h\| < \infty.$$

Then

$$|\hat{\mu}_n - \mu| \leq \|\nabla h\|/n.$$

This gives good results for various kinds of one-dimensional biased and unbiased random walk.

The theorem does not apply to our 2-D example (which has $||\nabla h|| = \infty$).

IV. Quasirandom diffusion

It can be shown that rotor-router walk is parallelizable.

Put some particles in $\mathbf{Z}^d$, where the sites are equipped with rotors. (For technical reasons, the particles must all start out on the same index-2 sublattice.)

Let the particles do rotor-router walk in parallel for n steps.

Cooper and Spencer show that the difference between (1) the number of particles at a site after n steps of rotor-router walk, and (2) the expected number of particles at a site after n steps of random walk, is bounded by a constant C that doesn't depend on n, or on what the original distribution of particles was, or which way the rotors were originally pointing. All it depends on is d, the dimension of the lattice.

See "Simulating a random walk with constant error", by Joshua Cooper and Joel Spencer:

arXiv:`math.CO/0402323`.

When you fully parallelize the rotor-walk algorithm, it essentially become heat flow in fixed precision arithmetic, with a twist: the rotors control the rounding of the least significant bits.

Rotors actually give an improvement over naive methods of simulating heat flow in discrete space and discrete time. (The method might generalize to variants of diffusion that include convection and reaction terms. But this will probably be of only minor interest for PDE, since rounding error isn't as big an issue as error introduced by discretization of space and time.)

V. Quasirandom aggregation

Internal Diffusion-Limited Aggregation (IDLA): To add a new bug to the (initially empty) blob, put the bug at the origin and let it do random walk until it hits an unoccupied site. Adjoin this site to the blob. Repeat.

Theorem (Lawler, Bramson, and Griffeath, 1992): The n-bug IDLA blob in $\mathbf{Z}^2$ is a disk of radius $\sqrt{n/\pi}$, to within radial fluctuation that are $o(n^{1/2})$.

Theorem (Lawler, 1995): We can replace $o(n^{1/2})$ by $O(n^{1/3})$ in the preceding result.

It appears empirically that the radial fluctuations are actually $O(\ln n)$.

IDLA can be derandomized using rotor-routers in the obvious way.

For one-dimensionsional derandomized IDLA, (where the "disk" is an interval), there is an absolute bound on the difference between the inner and outer radius of the blob; this was proved by Lionel Levine under my supervision as a Harvard honors thesis while he was still an undergraduate.

Levine's thesis also contains interesting observations about the dynamics of rotor-router aggregation.

As a graduate student, working with Yuval Peres, Levine succeeded in getting a (highly non-trivial) result for higher dimensions:

Theorem (Levine and Peres): For derandomized IDLA in any dimension, the symmetric difference between the rotor-router blob after n steps and the ball of volume $\sqrt{n}$ has area $o(n)$.

It appears that the radial fluctuations for derandomized IDLA are even smaller than for true IDLA.

E.g., after a million bugs have been added to the system, the inradius is 563.5 and the outradius is 565.1: these figures differ by 1.6 (about three tenths of one percent).

There may be an absolute bound on the difference between the inner and outer radius of the IDLA blob, valid at every time n.

VI. Future goals

Formulate a basic general theory of discrepancy for random walk and random aggregation models, so that standard probabilistic results can be derived as corollaries.

Apply rotor-routers to estimation problems of interest to Quasi Monte Carlo practitioners.

Find the right notion of discrepancy, and the right local mechanisms, for construction of quasirandom tilings.

For more information, see

`http://www.math.wisc.edu/`
`~propp/quasirandom.html`