Quasirandomness via rotor-routers

Jim Propp
(U. Mass. Lowell)

March 18, 2008

(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://jamespropp.org/CCR.pdf`

I use the term “quasirandom” in the sense introduced by Niederreiter et al. in the late ’70s (e.g., quasirandom sets of sample points for derandomized Monte Carlo integration), not in the sense introduced by Chung, Graham, and Wilson in the late ’80s (e.g., quasirandom graphs).

Theme of talk: One can design a deterministic process that mimics desired aspects of a stochastic process by replacing random sequences of choices by low-discrepancy sequences of choices.

I. Derandomization of finite Markov chains

Recall that the archetype for discrete randomness is an "unpredictable" fair coin.

The archetype for discrete *quasirandomness* is the deterministic sequence H, T, H, T, H, T, ... (or, equally good, T, H, T, H, T, H, ...).

This is almost as far from random as it can be, but it's still "fair"!

After N tosses of a "quasirandom coin", the number of heads is $N/2 + O(1)$; i.e., the empirical estimate of the bias is $1/2 + O(1/N)$.

Consider a strongly connected finite directed graph in which each vertex has outdegree 2. A bug moving through the directed graph chooses which way to go at each vertex by using a quasirandom coin sitting at that vertex.

For fixed vertex v, let s_N be the number of times the bug visits v during the first N steps of its quasirandom walk.

FACT: $|s_N - Np| = O(1)$, where p is the steady-state probability associated with v under random walk. That is, $|s_N/N - p| = O(1/N)$.

Note that $O(1/N)$ is the best one could hope for (and that $O(1/\sqrt{N})$ is what one gets from ordinary random simulation).

More generally, the simplest kind of quasirandom variable with m different equally probable values is a sequence that rotates through the m allowed values in some fixed order:

$$\begin{array}{l} 1, 2, 3, \ldots, m-1, m, \\ 1, 2, 3, \ldots, m-1, m, \\ 1, 2, 3, \ldots, m-1, m, \\ \ldots \end{array}$$

We call this a rotor.

If these m values are the m arcs emanating from v, we call this quasirandom variable a **rotor-router**, and we picture it as an arrow that points at the neighbors of v in some fixed cyclic order.

We advance the rotor at a vertex *before* we move the bug.

Thus, the rotor at an unoccupied site that has been visited before always points in the direction in which the bug left the vertex on its most recent visit.

FACT: Quasirandom walk on any strongly connected finite directed graph gives discrepancy

$$|s_N - Np| = O(1)$$

where s_N is the number of times the bug visits vertex v during the first N steps of its quasirandom walk, and p is the steady-state probability associated with v under random walk.

This generalizes to finite Markov chains with rational transition probabilities in a straightforward way.

A similar fidelity property holds for quantities such as expected hitting time (if a finite Markov chain starts in state x, what is the expected time until it first enters state y?) and hitting probability (if a finite Markov chain starts in state x, what is the probability that it enters state y before it enters state z?).

II. Derandomization of random walk in two dimensions

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 $(0,0)$ ("escape") 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 the number of escapes divided by the number of trials (call this denominator n) converges to $\pi/8$ with probability 1, with error falling like $1/\sqrt{n}$.

Equivalently, the number of escapes 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 $|D_n| \approx 50$.

We can get good approximations to $\pi/8$ faster if we replace random choices by low-discrepancy choices. Specifically, we insure that for each vertex v, the choices we make each successive time we visit v (regarding where to go next) form a low-discrepancy sequence ("control of local discrepancy").

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.)

Under quasirandom simulation, with rotor-routers, the n trials aren't independent, or even random — yet D_n is provably $O(\log n)$ (rather than $O(\sqrt{N})$) and indeed seems to be bounded!

See demo at

`http://jamespropp.org`
`~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 escapes 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!

Extra wrinkle: In the standard theory of random walk, the probability that a walker who starts at $(0, 0)$ will never reach $(1, 1)$ or return to $(0, 0)$ is 0 and can be ignored. In the quasirandom theory, such paths can occur in a sequence of trials (and hence cannot be ignored), but their frequency provably must go to 0.

III. 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.)

Puzzle: Suppose n units of mass flow in $\{0, 1, 2, ...\}$ starting at 1 and moving with flow-proportions $p(k, k-1) = \frac{1}{3}$ and $p(k, k+1) = \frac{2}{3}$ for all $k > 0$, where mass that arrives at 0 stays there. If mass is infinitely divisible, the amount of mass that eventually arrives at 0 is $\frac{n}{2}$ (corresponding to the fact that a "two-to-one-rightward-biased" random walker who starts at 1 has a 50% chance of ever reaching 0), but suppose mass must always be a whole number (non-integer quantities $\frac{m}{3}$ or $\frac{2m}{3}$ are immediately rounded to the next closest integer). For what n does the amount of mass that eventually arrives at 0 equal $\frac{n}{2}$? $(2, 8, 48, 50, 200, \ldots)$

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 rotor-router blob after n steps lies inside a ball of volume $n + o(n)$ and contains a ball of volume $n - o(n)$.

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 other derandomized random objects (such as derandomized random tilings).

For more information, see

`http://jamespropp.org/quasirandom.html`