

Quasirandom processes

by Jim Propp (UMass Lowell)

November 7, 2011

Slides for this talk are on-line at

<http://jamespropp.org/cvc11.pdf>

Acknowledgments

Thanks to David Cox 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 Matt Cook, Dan Hoey, Rick Kenyon, Michael Kleber, Oded Schramm, Rich Schwartz, and Ben Wieland.

Quasirandom processes

Consider the sequence $(x_1, x_2, x_3, \dots) = (.618, .236, .854, \dots)$ whose n th term is the fractional part of n times $(1 + \sqrt{5})/2$.

Nobody would ever call this sequence random, or even *pseudorandom*. But it would be considered *quasirandom* for some purposes, because it's uniformly distributed in $[0, 1]$.

In fact, it's more evenly spread out than a random sequence would typically be: for an interval I in $[0, 1]$ of length L , the discrepancy

$$\#\{1 \leq k \leq n : x_k \in I\} - nL$$

is of magnitude $O(\log n)$ rather than of typical magnitude $O(\sqrt{n})$. In low dimensions, quasirandom sampling gives more accurate estimates of integrals than random sampling.

(There's another usage of the word “quasirandom” current among graph-theorists, tracing back to Chung, Graham and Wilson (1989), but that is a different story.)

Quasirandom processes

I'm a probabilist, so the “integrals” that interest me most are probabilities and expected values, and the measure with respect to which I'm integrating is the probability measure associated with a random process.

(Example: the measure space is the sequence of outcomes of infinitely many flips of a fair coin, where each initial string of length n has probability measure 2^{-n} .)

I'm also a combinatorialist, so the kinds of probabilistic systems I like best are discrete ones, like Markov chains. And I want to quasirandomize these processes by using very simple combinatorial constructions.

Rotors

Quasirandom analogue of a fair coin-flip process:

H, T, H, T, H, T, ...

Quasirandom analogue of a fair die-roll process:

1, 2, 3, 4, 5, 6, 1, 2, 3, 4, 5, 6, ...

(This generalizes to arbitrary discrete probability distributions, including ones with infinitely many values and/or irrational probabilities; see “Discrete low-discrepancy sequences” by Angel, Holroyd, Martin, and Propp, [arXiv:0910.1077](https://arxiv.org/abs/0910.1077).)

Rotor-routers

A Markov chain can be thought of as a random walk on a graph, where a walker at vertex u has probability $p(u, v)$ of moving to a vertex v in the next time-step, regardless of her previous history.

Instead of making these choices randomly, we can use a rotor situated at each vertex u to tell the walker which v to go to next.


E.g., if $p(u, v_1) = p(u, v_2) = 1/2$, then we use a simple 2-way rotor at u , and the walker follows the rule “Do whatever you **didn't** do the last time you were at u .”

Gambler's ruin

A gambler starts with \$1, and makes a sequence of fair bets, each of which results in his purse going either up by \$1 or down by \$1.

The gambler stops when he reaches either \$0 or \$3.

We can view this as a random walk on a path of length three, with a source at \$1 and sinks at \$0 and \$3.



What is the probability that, starting with \$1, the gambler reaches \$3?

Gambler's ruin, repeated

The probability that, starting with \$1, the gambler reaches \$3, is $1/3$.

So, if the gambler does this routine n times, starting from the source \$1 and returning to \$1 after each arrival at \$0 or \$3, he will reach \$3 rather than \$0 about $n/3$ times.

To quasirandomize this, replace the random gambles by rigged gambles, where the gambler **wins** his current gamble if and only he **lost** his gamble the last time he had the exact same amount of money in his purse.

Gambler's ruin with rotor-routing

See <http://www.cs.uml.edu/~jpropp/rotor-router-model/>;
select The Applet (tab at top of page) and set Graph/Mode to
Walk on Finite Graph A.

(This applet was written by Hal Canary and Yutai Wong while they
were undergraduates at the University of Wisconsin.)

The color at a site conveys the same information as the rotor there.

The rotor-walker ends up with \$3 (“success”) one-third of the
time, just like the random-walker.

But for a random walker, the number of successes in the first n
trials typically differs from $n/3$ by $O(\sqrt{n})$, while for a rotor-walker,
the number of successes in the first n trials differs from $n/3$ by at
most a **constant**.

This generalizes to arbitrary finite-state Markov chains, and some
infinite-state Markov chains as well.

Confluence

Instead of putting the same walker through the system n times, we can put n walkers through the system one at a time. The walkers share the same rotors, so the path travelled by the k th walker changes the rotor-states and thereby affects the path travelled by the $k + 1$ st rotor, for all $k < n$.

Instead of requiring that for all $k < n$ the k th walker must arrive at a sink before the $k + 1$ st walker can leave the source, we can permit walkers to take steps in arbitrary order.

It can be shown that if all the walkers are **indistinguishable**, then the outcome of simultaneous rotor-walk for several walkers is *independent* of the order in which the walkers take steps. This is the **confluence property** (aka **Abelian property**) of rotor-routing. See "**Chip-Firing and Rotor-Routing on Directed Graphs**" by Holroyd, Levine, Mészáros, Peres, Propp, and Wilson.

Random walk on the two-dimensional grid

Starting from $(0, 0)$, the walker takes random steps in the set $\{E, W, N, S\} = \{(1, 0), (-1, 0), (0, 1), (0, -1)\}$, stopping upon arriving at either $(0, 0)$ or $(1, 1)$.

Starting from anywhere, the walker ends up in $\{(0, 0), (1, 1)\}$ with probability 1.

The probability that a random walker who starts at $(0, 0)$ ends up at $(1, 1)$ rather than $(0, 0)$ (“success”) is $\pi/8$.

Rotor-walk on the two-dimensional grid

The walker successively goes $\dots, N, E, S, W, N, E, S, W, \dots$ upon leaving a particular vertex.

Rule: The rotor advances and the walker then moves in the direction indicated by the *current* (updated) state of the rotor; thus, if a vertex u been visited, and the rotor at that site points in a direction, this is the direction that was travelled by the walker after the walker's most recent visit to u .

Set Graph/Mode to 2-D Walk.

Under suitable initial settings of the rotors, it can be shown that the number of successes in the first n trials differs from $n\pi/8$ by at most $C \log n$ for some constant C . (For a proof, see “Rotor Walks and Markov Chains” by Holroyd and Propp.)

Two open problems

Can we **tighten** the preceding result, so that $O(\log n)$ is replaced by something smaller? (Empirically, it seems that $O(\log \log n)$ or even $O(1)$ might be closer to the truth.)

Can we **broaden** the preceding result, so that it applies to a wider class of initial rotor-settings?

Bringing in more symmetry

In the preceding example, after each visit to $(1, 1)$ the walker gets sent back to $(0, 0)$.

Let's change this, so that there's a rotor at $(1, 1)$ that sends the walker to the four neighbors of $(1, 1)$.

To make things even more symmetrical, let's have all rotors initially point in the same direction.

What happens? That is, what is the path travelled by a rotor-walker that starts from $(0, 0)$?

To infinity and beyond

The walker heads off to infinity, never hitting either $(0, 0)$ or $(1, 1)$!

But we can still make sense of the setting of the rotors “at time ω ”, since each vertex gets visited only finitely often (indeed, at most once).

So we can restart the process with a second walker, using the new rotor-settings.

After the second walker has gone off to infinity, we can restart with a third walker. And so on.

What do the rotors look like after n walkers have passed through the system?

An alternative viewpoint

If you don't like processes indexed by transfinite times $\tau = a\omega + b$, we can rephrase the preceding question by invoking a transfinite version of the confluence property:

Start with n walkers, and let them each move through the system, taking moves in arbitrary order.

As long as every walker takes an infinite number of steps, all of them will wander off to infinity.

Moreover, each rotor will be changed only finitely many times, so it makes sense to speak of the states of the rotors at time infinity (the “eventual settings” of the rotors).

The eventual settings of the rotors are independent of the choices one makes along the way (deciding which walker gets to take the next step).

Quasirandom walk in the plane

<http://jamespropp.org/2drotorwalk.pdf> (a still taken from a dynamic animation created by [Lionel Levine](#)) shows what happens: the n th frame shows the state of the rotors after n walkers have gone through the system.

(Note that the color scheme and rotation-scheme are different from what was used in the “ $\pi/8$ machine” example.)

Eventually each walker heads off to infinity going North, and these columns are contiguous, so that the $n + 1$ st walker heads off to infinity either via the first column to the Right of the n previously visited columns or via the first column to the Left of the n previously visited columns.

Conjecture: The sequence of Rights and Lefts associated with the preceding observation repeats with period 8.

The quasirandom quincunx

Modify the random walk process so that the walker can only go East or South. Rotated by 45 degrees, this is the quincunx or Galton board process.

In the quasirandom version we put a 2-way rotor at each site; assume all the 2-way rotors all point the same way at the start.

To see what happens, view the animated gif file

<http://jamespropp.org/quincunx.gif> or the “movie-version”

<http://jamespropp.org/Galton.swf>

Diffusion-Limited Aggregation

We place a particle at $(0, 0)$, to serve as a seed for aggregation.

Another particle in the plane “wanders in from infinity” until it hits the seed, and then it sticks there.

A third particle also wanders in from infinity until it hits one of the two stuck particles, and it too joins the aggregate.

Etc. (Here I am ignoring a few technical issues, so this is not the exact right definition, but it's close enough for a talk of this kind.)

See <http://jamespropp.org/BigDLA2.gif> (taken from <http://classes.yale.edu/fractals/Panorama/Physics/DLA/BigDLA2.gif>) for a view of the sort of dendritic structure DLA creates.

Quasirandom Diffusion-Limited Aggregation

<http://jamespropp.org/rotorDLA.gif> (created by Oded Schramm) shows one way of using rotor-walk instead of random walk to create something that resembles a DLA cluster.

As far as I know, nobody has followed up on this.

Directed DLA: random and quasirandom

Tobias Friedrich created an animation showing what happens for directed DLA with absorption occurring along a segment rather than at a point.

In his simulations (one **random** and one **quasirandom**), particles that don't hit the segment disappear from the system.

The n th frame of each simulation shows the state of the system after the n th particle has joined the aggregate.

No theorems here, but some intriguing patterns!

Internal DLA

Diaconis and Fulton (1991 and 1993) devised an inside-out variant of DLA in which the particle starts from a source-point rather than infinity.

At stage 0, the blob is empty.

At stage 1, the blob consists of just the source $(0,0)$.

To turn the stage- n blob into the stage- $(n + 1)$ blob, a particle starts from the source and does random walk until it reaches a vertex that isn't in the blob; the new vertex gets added to the blob.

Roundness of Internal DLA

Internal DLA dynamics are much more stable than DLA dynamics.

In particular, Lawler, Bramson, and Griffeath (1992) showed that the internal DLA cluster of size n , rescaled by $\sqrt{n/\pi}$, converges almost surely to a disk of radius 1.

Lawler (1995) showed that the inward and outward fluctuations from roundness are $O(n^{1/6})$, and it was finally shown in 2010 (by two different research groups working independently and using different methods) that the fluctuations are $\Theta(\log n)$.

Quasirandom Internal DLA

We can replace the random walkers by rotor-walkers.

To see what quasirandom Internal DLA looks like using the [rotor-router-model applet](#), set Graph/Mode to 2-D Aggregation; to see what one has after a million particles have joined the aggregate, see <http://jamespropp.org/million.gif>.

Roundness of quasirandom Internal DLA

Levine and Peres proved in “**Strong Spherical Asymptotics for Rotor-Router Aggregation and the Divisible Sandpile**” that when $n = \pi r^2$ particles have been joined the aggregate, the inradius of the set of occupied sites is at least $r - O(\log r)$, while the outradius is at most $r + O(r^\alpha)$ for any $\alpha > 1/2$.

But empirically we observe that the blobs are much rounder.

E.g., for a rotor-router blob of cardinality $n = 2^{32}$ (and radius $r = \sqrt{n/\pi} \approx 36975$), the inradius and outradius of the blob measured from the point $(1/2, 1/2)$ (which is believed to be the limiting location of the center of mass of the blob on both empirical and theoretical grounds) differ by only $0.366 \approx r/10^5$.

Systematic deviations from circularity

Matt Cook noticed that the (small!) deviations from circularity are systematic: in certain directions the blob tends to deviate inward, and in certain directions the blob tends to deviate outward.

We may phrase this equivalently in terms of the stage at which a site joins the blob. A site at distance r from the center should join the blob at stage $\approx \pi r^2$. Sites that join the blob before this time are “early” and sites that join the blob after this time are “late”.

<http://jamespropp.org/early-late.gif> shows when sites joined the blob. Early points are colored blue and late points are colored red.

Stacks

A unified way of viewing random and quasirandom Internal DLA is by having *stacks* at each vertex. Each stack is an infinite sequence of N's, S's, E's, and W's. When a particle visits a site, it pops off the top element of the stack and heads in that direction.

For random Internal DLA, the elements of the stack are random.

For rotor-router Internal DLA, the elements of the stack are periodic with period 4.

Internal DLA with random low-discrepancy stacks

An interesting hybrid is gotten by using random low-discrepancy stacks: The stack at each site is a concatenation of blocks NSEW, NSWE, SNEW, SNWE, EWNS, EWSN, WENS, WESN chosen independently and uniformly at random.

The resulting blobs are much more like rotor-router Internal DLA blobs than like (random) Internal DLA blobs, in terms of the discrepancy from roundness.

Lesson: For ensuring that “global discrepancy” is low, what matters most is that “local discrepancy” is low. Randomness versus non-randomness is of secondary importance.

Internal DLA with non-random high-discrepancy stacks

One can also use deterministic stacks like

N, E, E, S, S, S, W, W, W, W, N, N, N, N, N, . . .

in which the discrepancy between $n/4$ and the number of N 's (or E 's or S 's or W 's) seen in the first n elements of the stack is $O(\sqrt{n})$ (which is the sort of discrepancy one would see for purely random stacks).

Curiously, these processes behave quite differently from any of the processes discussed above; we have no idea why. For examples, see <http://jamespropp.org/TF-A.gif> , <http://jamespropp.org/TF-B.gif>

Fast simulation

The naive method of constructing the rotor-router Internal DLA blob of cardinality n takes $\Theta(n^2)$ steps. Friedrich and Levine have found clever shortcuts that let them construct the rotor-router Internal DLA blob of size n much more quickly (experimentally, in time about $n \log n$).

They have implemented their method, creating blobs so big that the only way to study them is via a Google Maps interface that allows the user to navigate between different scales. See <http://rotor-router.mpi-inf.mpg.de/>.

Friedrich's website shows rotor-router blobs associated with not just the standard style of rotor (rotating clockwise) but other styles of rotor as well.

Interesting spots

There are places in the picture where many adjacent sites have the same color, forming a monochromatic *patch*.

Matt Cook and Dan Hoey independently noticed that if one normalizes the blob to be the inside of the unit disk in the complex plane, the patches occur precisely at complex numbers of the form $(A + Bi)^{-1/2}$ with A, B integers (not both 0), though if A or B is large, one needs n to be quite large before the patch becomes visible.

One also finds patches on which the coloring is not constant but periodic of small period. E.g., on the boundary of the disk, one sees this behavior near points $(a + bi)/\sqrt{a^2 + b^2}$ with a, b integers (not both 0).

Interesting curves

In some of the pictures, the eye detects curves, especially the “lrdu” picture (though it’s hard to say, mathematically speaking, exactly what the eye is detecting along those curves).

Rick Kenyon pointed out that some of these curves appear to be the images of circles of radius $1/2$ centered at points $a + bi$ under the map $z \mapsto 1/\sqrt{z}$; see <http://jamespropp.org/RRcircles.pdf>.

Prospects

To the extent that quasirandom processes **share** properties with their random counterparts, they teach us that many of the theorems of probability theory remain true when hypotheses of randomness are replaced by weaker hypotheses of discrepancy.

To the extent that quasirandom processes have properties **different** from their random counterparts, the task of proving that these properties actually prevail offers exciting challenges to theorists, blending combinatorics, probability, and geometry.