From *SIAM News*, Volume 38, Number 6, July/August 2005

Intriguing Mathematical Picture Merits a Second Look

By Sara Robinson

One afternoon in the spring of 2003, Yuval Peres, a professor of mathematics and statistics at the University of California, Berkeley, and Alexander Holroyd, a visiting research fellow, were poring over a computer printout of an intriguingly complex mathematical picture.

Starting with 300 "centers" (the centers of the concentric circles in Figure 1) distributed uniformly at random in a square, their program, a simulation of a probabilistic model, had partitioned the square into territories, one for each center. The rules that defined the partition were based on simple local preference relations: Centers preferred to have nearby territory, and points preferred to be assigned to the closest possible center. Because the centers had a particular territorial quota, the model guaranteed that the square would be partitioned into equal-sized territories.

Although the rules were simple, conflicts had arisen, causing the resulting territories to vary widely in shape and extent. Some territories consisted of many separate pieces, often scattered quite far from their associated centers. A few regions of the square seemed to consist of a jumble of fine slivers of territory, each belonging to a different center.

All in all, Peres says, the picture was much more interesting than they expected—more intriguing, in fact, than the problem that had prompted it. "We looked at it and said, 'Wow, this is pretty cool,'" adds Holroyd, who is now an assistant professor of mathematics at the University of British Columbia.

Captivated, the researchers completed their original project and then returned to the picture, trying to understand it. If the square were replaced by $\mathbb{R}^d$ scattered with countably many centers, they wondered, how many pieces could a given territory consist of? Would the territories be bounded? If so, how widely could the pieces be scattered? How many different territories could a bounded region of space contain?

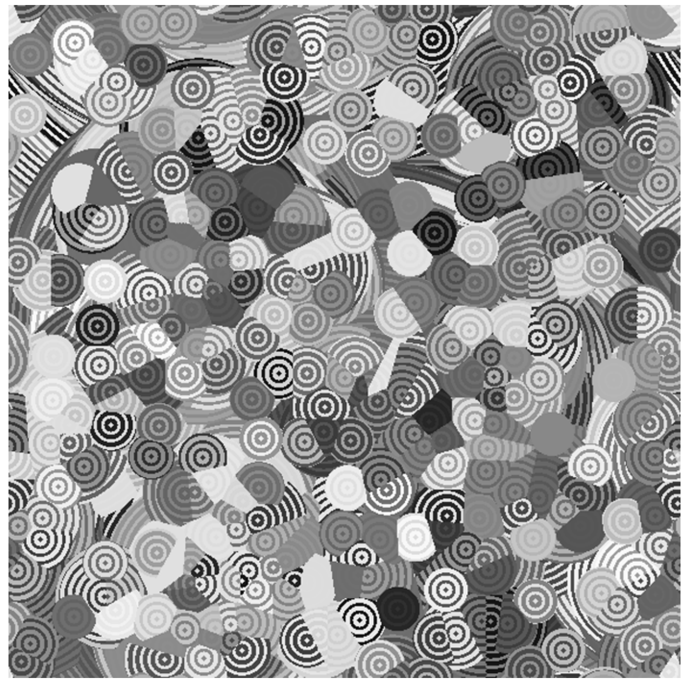

Figure 1. *For a simulation of their probabilistic model, as described in detail in the accompanying article, Alexander Holroyd and Yuval Peres started with 300 points or "centers" distributed independently uniformly at random in a square (actually a 2-torus). Each center has an associated territory described by concentric circles in two alternating colors (the grayscale version shown here only hints at what the researchers found so intriguing in the color version, which can be seen at http://www.stat.berkeley.edu/users/peres/stable/stable.html). All territories have the same area and, together, form a partition of the square. Many of the territories are broken into several pieces, some of which are located a long way from their associated centers (as appreciated much more easily in the color version).*

Working with Christopher Hoffman, a professor of mathematics at the University of Washington, Holroyd and Peres were able to answer some of the questions, though not all of them. Their results were posted in preliminary form on the Math arXiv in June.

Looking for Typical Points

The story behind the picture starts in 1995 in Sweden, at the University of Göteborg, where Peres attended a lecture given by Hermann Thorisson, a probability theorist now at the University of Iceland. Thorisson had just proved a rather surprising result about point processes.

Intuitively, a point process is a countable collection of isolated points scattered randomly (according to some distribution) in d-dimensional space. In practice, point processes are used to model random events—meteorites striking a region, say, or the formation of stars. For such phenomena, it's useful to look at processes that are translation-invariant—that is, the distribution of points is independent of location—and ergodic: Translation-invariant events depending only on the process have probability zero or one.

In this setting, Peres explains, probability theorists are interested in looking at the landscape from the perspective of a "typical" point. For instance, given that a meteorite has struck at one place, how distant on average are the other locations at which meteorites have struck?

But it is not clear how to choose a typical point. With finitely many points, one could be chosen uniformly at random, but this does not work for an infinite set.

Thorisson's result, in essence, provides a way to choose a typical point of a certain type of point process. In particular, his theorem says that there is a particular point p in the process with the property that when the process is translated so that p is at the origin, the rest of the points still follow the initial distribution. Thus, the new distribution is just the starting one conditioned on having a point at the origin. In effect, the distribution has been centered at a typical point, with the rest of the distribution's properties preserved.

As an example, imagine a point process given by a doubly infinite sequence of outcomes of a fair coin flip. In this case, the underlying space is **Z**, and the heads are the points in the process. What Thorisson's theorem indicates is that there is a rule for choosing a head and translating it to the origin so that the new sequence corresponds to fair coin tosses but is conditioned on having a head at the origin.

While Thorisson's theorem has intriguing implications for some specific examples, his proof is non-constructive. He showed that a rule for choosing such a point always exists, but his proof does not distinguish between randomized and deterministic rules, and he did not give a specific rule for any particular process. Following the announcement of his result, other probability researchers thus turned to the problem of finding deterministic rules for choosing a typical point for specific types of point processes.

Several years later, Tom Liggett, a probability theorist at the University of California at Los Angeles, gave a deterministic rule for the coin-flipping example: If a head has the property that there are equal numbers of heads and tails between it and the origin, Liggett showed, then translating it to the origin leaves the same distribution, conditioned on having a head at the origin. That work, published in 2002, was subsequently extended to multidimensional lattices by Liggett and Holroyd, who held a visiting position at UCLA from 1997 to 2000. In a third paper, Peres and Holroyd further generalized some of those results.

In the fall of 2002, Holroyd arrived at Berkeley, and he and Peres continued to think about the problem, trying to find a deterministic rule for translation-invariant, ergodic point processes in $\mathbb{R}^d$. What was needed, they eventually realized, was a translation-invariant way to partition the space into regions of equal area, each containing a point in the process. In this case, the center of the territory containing the origin would be the desired typical point.

Intuitively, the territories could be defined via a ball-growing process: Each center would grow outward at the same rate until reaching its territorial quota. If the intensity, or average density, of the process were one point per unit area, the result would be a partition of the space into equal-sized areas. It was not clear that this process was well defined, however, and the researchers tried for a long time to figure out a way to make the intuitive ball-growing model rigorous. Eventually, Peres recognized that the conditions defining the territories were related to a stability criterion introduced in a celebrated "matching" algorithm from the early 1960s.

A Continuum Stable Marriage Algorithm

David Gale and Lloyd Shapley, mathematicians at UC Berkeley and UCLA, respectively, had considered the problem of pairing n boys and n girls, each of whom has ranked the members of the opposite group in order of preference. They showed that there is always a matching of boys to girls that is "stable," in the sense that no boy and girl prefer one another to their assigned partners in the matching.

Gale and Shapley's proof was in the form of an algorithm, involving iterated proposals and rejections, that is guaranteed to find such a matching. Variants of this algorithm are used in practice to match graduating medical students with slots in residency programs and high school students to public schools (see "How Much Can Matching Theory Improve the Lot of Medical Residents," *SIAM News*, Volume 36, Number 6, July/August 2003).

In round 1 of the algorithm, the boys propose to their first-choice mates. Each girl considers all her proposals and tentatively accepts her favorite and rejects the rest. In round 2, each rejected boy moves on to his second-choice girl, possibly competing with her tentative match from the previous round, and so on. Eventually, this procedure terminates in a stable matching.

The matching produced by the algorithm is "male-optimal," in the sense that all the boys have the best possible mates they could achieve in any stable matching, and "female-pessimal," in that the girls are assigned their worst possible partners in any stable matching. If the girls do the proposing, the properties of the match are reversed.

In general, there can be many (possibly exponentially many) stable matchings for a particular stable marriage instance. The matchings form a lattice with the male- and female-optimal matchings as its extremal points. These results carry over to the case of many-to-one matchings as well.

To make their ball-growing partition rigorous, Holroyd and Peres defined a variant of the Gale–Shapley model in which the roles of boys and girls are played by the centers and the other points in the space. Preferences for both sides are based purely on distance: Points prefer to be assigned to nearby centers, and centers prefer territories composed of nearby points. An allocation of territories to centers is "stable" if there are not both unsated centers and unassigned points, and if no center and point outside that center's territory would prefer one another to their allocated partners.

In their recent write-up, Holroyd, Peres, and Hoffman give two procedures for finding a stable allocation, a "site-proposing" version and a "center-proposing" one. In the site-proposing procedure, sites first apply to the nearest center. When the applicants for a center exceed its appetite, it tentatively accepts applicants out to a radius sufficient to fulfill its quota and rejects those outside that radius. The rejected points then apply to the next closest center. In the center-proposing version, centers first propose to the points in a surrounding ball whose area equals the center's quota. Points with multiple offers tentatively accept only the closest center and reject the rest. After a chunk of territory has rejected a center, the center grows its ball of proposals just large enough to again meet its quota.

Each procedure eventually (after infinitely many steps) results in a stable matching of centers to territories, the researchers showed. In fact, as long as the point process is ergodic, the center-optimal and site-optimal matchings coincide almost everywhere, giving rise to a unique stable allocation. The reason for this, says Holroyd, is that preferences depend only on distance. To illustrate, he points to an analogous, finite version of the situation: Suppose that n men and n women are scattered randomly across a football field, with preferences for one another based on distance. Notice that pairwise distances are distinct with probability one.

In the male-optimal matching, each boy is paired with his closest possible partner in any stable matching; the sum of the distances between couples, therefore, is strictly smaller than in any other stable matching. The same is true of the female-optimal matching, and it thus follows that the male-optimal and female-optimal matchings must coincide. Because the male- and female-optimal matchings are extremal points of the lattice of stable matchings, there must be a unique stable matching.

The Geometry of the Territories

In the simulations, the territories looked as if they might be quite complex, but the researchers were able to show that they are somewhat well behaved: Each is almost surely bounded, and any bounded subset of $\mathbb{R}^d$ intersects only finitely many territories.

The researchers next considered what happens when the intensity of the point process is kept fixed at one point per unit area, but the appetite of the centers varies. Remarkably, the model undergoes a "phase transition" akin to those of percolation models in statistical physics.

If the appetite is less than one, the centers will almost surely be sated, but an infinite volume of space will not belong to any territory. If the appetite is equal to one—the "critical case"—then almost all sites will be claimed and all centers will almost surely be sated, yielding a partition of the space into equal-sized territories. If the appetite is greater than one, almost all the space will be claimed but some centers will be unsated.

The perfect balance in the critical case, Holroyd explains, follows from the ergodicity of the point process. The existence of unclaimed sites or of unsated centers is a translation-invariant event depending only on the process, which, by ergodicity, must have probability zero or one. Unclaimed sites and unsated centers cannot both occur with probability one, he notes, or the allocation would be unstable. Because the distribution of centers is translation-invariant, when the appetite is one it also cannot be the case that sites are unclaimed with probability one while centers are unsated with probability zero, or vice versa. Thus, both events happen with probability zero.

In the critical case when the point process is Poisson, the researchers showed, points are assigned and centers sated at the price of pervasive jealousy: Every center is preferred by an infinite volume of points not assigned to it, and every point is coveted by infinitely many centers that were assigned less desirable territory.

Percolation models are associated with events that have power-law tails at a critical temperature, and exponential tails otherwise. A corresponding event for this model is the extent of the territories. Holroyd, Peres, and Hoffman looked at a random variable X, defined as the distance from the origin to the center of the territory that contains it. (Because the territories are defined by a translation-invariant process, this random variable describes the behavior of all the territories.) X turns out to have exponential tails in the non-critical cases, and at least a power-law tail in the critical case.

Although they believe that the critical case likely follows a power law, the researchers have not been able to find any upper bound on X in the critical case in dimension 2 or higher.

What's Next

Holroyd and Peres, who is visiting UBC this summer, are still trying to prove upper bounds in the critical case for higher dimensions. They are also delving further into the implications of this new continuum version of stable marriage. Having proved that their model results in an almost surely unique stable marriage with probability one for certain random sets of centers, they are now trying to remove the probabilistic aspect of their theorem. Specifically, they hope either to find a deterministic set of centers that corresponds to more than one stable matching, or to prove that no such set exists.

Sara Robinson is a freelance writer based in Los Angeles, California.