Snapshots of modern mathematics
from Oberwolfach

№ 7/2019

Random permutations

Volker Betz

100 people leave their hats at the door at a party and pick up a completely random hat when they leave. How likely is it that at least one of them will get back their own hat? If the hats carry name tags, how difficult is it to arrange for all hats to be returned to their owner? These classical questions of probability theory can be answered relatively easily. But if a geometric component is added, answering the same questions immediately becomes very hard, and little is known about them. We present some of the open questions and give an overview of what current research can say about them.

1 Uniform random permutations

Imagine we have a set A with finitely many, say n, elements. We could number these elements in a certain order by $1, 2, \ldots$ to n. By reordering the elements we obtain what we call a *permutation* of the set A. In other words, a permutation is a *one-to-one map* from the set to itself, that is, a map that assigns to every element from the set exactly one other element and no two elements in the set get assigned to the same element.

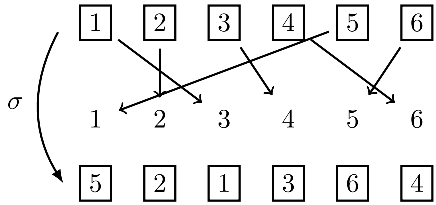

Figure 1: A permutation of the numbers 1 to 6. The arrows show the way σ assigns new places to the numbers and point from j to $\sigma(j)$. The last line shows the reordered set of numbers after having applied σ.

Permutations are interesting to study in various contexts.[1] Here we will add a further ingredient to our setup, namely, randomness. The idea of choosing at random among all permutations of a given set is almost as old as probability theory itself. One of the first examples is what is today known as the 'hat problem', first considered by Pierre Rémont de Montmort (1678–1719) in 1708.

1.1 The hat problem

Assume that n people, each wearing a hat, come to a party, leave their hats by the door, get thoroughly drunk, and pick up a completely random hat when they leave. What is the probability that at least one of them ends up taking their own hat? To find the answer, we assign numbers 1 to n to the people, and also to their respective hats. If we define A to be the set of numbers $\{1, \ldots, n\}$, then the map $\sigma : A \to A$ that maps each person to the hat they have picked up is a permutation. The number of possible ways the people can pick up the hats is the number of arranging n elements, which is exactly $n! = 1 \cdot 2 \cdot \ldots \cdot (n-1) \cdot n$. To see why this is the case, count the number of possible hats the first person could take, multiply it with the number of possible hats for the second person, for the third person, and so on. Since none of the permutations is preferred in any way (the people are so drunk that they pick completely random hats), all $n!$ permutation maps σ occur with equal probability; we also say they are *uniformly distributed*. Phrasing this a little more technically, this means we are studying the *uniform probability measure* on the space $\mathcal{S}(n)$ of permutations of length n.[2]

[1] For example, they are useful when dealing with symmetry groups, as Jay Taylor and Eugenio Giannelli explain in Section 3 of their Snapshot 5/2016 *Symmetry and characters of finite groups*.

[2] The notion of a *probability measure* is central in probability theory. If you are not familiar with this term you may think of it as a function that assigns to every event in a collection of

What is then the probability that at least one person obtains their own hat? This is equivalent to asking: what is the probability that $\sigma(j) = j$ for at least one $j \in A$?

It turns out that the probability $\mathbb{P}(k)$ of precisely k people finding their own hat is equal to

$$\mathbb{P}(k) = \frac{1}{k!} \sum_{j=1}^{n-k} \frac{(-1)^j}{j!}.$$

If the total number of people gets very large, that is, as n tends to infinity, $\mathbb{P}(k)$ converges to $\mathbb{P}(k) = \frac{\mathrm{e}^{-1}}{k!}$. This is the *Poisson distribution* with parameter 1. In particular, for a large party the probability that at least one person picks up their own hat is approximately $1 - 1/\mathrm{e} \approx 0.63$, since the probability that no-one picks up their own hat is e^{-1} and these two probabilities must sum to 1. We will not give a derivation here that the Poisson distribution gives the right probabilities for the hat problem; instead, we focus on an even easier problem.

1.2 The hat problem revisited: telephone chains

Imagine that all of the guests of our party were diligent enough to leave their phone number inside their respective hat. The morning after the party, someone (call her j) wants to get back her own hat. She will call the number in the hat $\sigma(j)$ she took, unless it is her own number (in which case she was lucky enough to grab her own hat in the first place). The person $\sigma(j)$ on the other end will be pleased to learn who picked up their hat, but in most cases j will still not have found her own hat. To find it, she needs to ask the person $\sigma(j)$ to dial the number in the hat they took, and call $\sigma(\sigma(j)) = \sigma^2(j)$. This procedure is then repeated until j gets a call from the person who holds her hat. The question is: how many people do we expect to be involved before j gets a call?

In the language of permutations, we are asking about the length of the cycle containing a given point j. Let us write ℓ_j for the smallest integer so that applying ℓ_j times the map σ to j produces j itself, that is $\sigma^{\ell_j(\sigma)}(j) = j$. For example, in the permutation shown in Figure 1 we have $\ell_2 = 1$, because $\sigma(2) = 2$, whereas $\ell_3 = 5$, because $\sigma^5(3) = 3$ and $\sigma^k(3) \neq 3$ for every $1 \leqslant k < 5$. The (unordered) set

$$O_j(\sigma) = \{j, \sigma(j), \sigma^2(j), \ldots, \sigma^{\ell_j(\sigma)-1}(j)\}$$

of all elements that can be reached from j by applying σ multiple times is called

all possible events a certain probability between 0 and 1, with certain rules for combining probabilities. It is normalized in such a way that if one asks "What is the probability that *something* happens?", we get the answer 1.

the *orbit* of j under σ. The *cycle*

$$C_j(\sigma) = (j, \sigma(j), \sigma^2(j), \ldots, \sigma^{\ell_j(\sigma)-1}(j))$$

is the ordered set containing the elements of $O_j(\sigma)$ in the order in which they appear when applying σ iteratively, but where we do not care about the actual starting point of the cycle. Using the example from Figure 1 again, this means that we identify five possible cycles coming from the orbit $\{3, \sigma(3) = 4, \sigma^2(3) = 6, \sigma^3(3) = 5, \sigma^4(3) = 1\}$, that is, we consider the following cycles to be equivalent:

$$(3,4,6,5,1) = (4,6,5,1,3) = (6,5,1,3,4) = (5,1,3,4,6) = (1,3,4,6,5).$$

The number of elements a cycle contains, that is, the number ℓ_j, is called the *length* of the cycle (and the same applies for an orbit). You can see (think about the example above) that for an orbit of length k there are $k!$ ways to order its elements, but only (k-1)! cycles with the elements of that orbit.

As an answer to our question we will find, for each number k, the probability $\mathbb{P}(\ell_j = k)$ that j is contained in a cycle of length $\ell_j = k$, meaning the probability that exactly k people have to receive a call before j gets her hat back. The problem we have posed can be solved by simple counting.

Let us introduce the notation $\binom{n}{k}$ for the number of ways of picking k elements out of a set of size n when the order of the elements does not matter. It can be shown that $\binom{n}{k} = \frac{n!}{k!(n-k)!}$. We can now easily count the permutations having a cycle of length k containing j. First we pick j, then there are $\binom{n-1}{k-1}$ ways of picking $k-1$ elements other than the fixed element j from the set A, and thus $\binom{n-1}{k-1}(k-1)! = \frac{(n-1)!}{(n-k)!}$ ways of fixing a cycle with k elements that contains j. We do not care what the $n-k$ elements outside the cycle $C_j(\sigma)$ do, and for given $C_j(\sigma)$ there are $(n-k)!$ ways of arranging them. Hence, there are $\binom{n-1}{k-1}(k-1)!(n-k)!$ permutations of n elements which include a cycle of length k containing j. Since we again are assuming that the permutations are uniformly distributed, we only have to divide by the total number of permutations (which is $n!$ as you might remember from before) to obtain

$$\mathbb{P}(\ell_j = k) = \frac{1}{n!}\binom{n-1}{k-1}(k-1)!(n-k)! = \frac{1}{n}. \tag{1}$$

Surprisingly, the result does not depend on k and we thus have found for person j to regain her hat, it is equally likely that k people need to get involved for all $k \leqslant n$. This has the somewhat discouraging consequence that with probability $1/2$, at least half of the party guests will need to be called before person j recovers her hat.

1.3 Yet another problem

Let us now ask a related, but different, question. What is the probability that, for *at least one* of the guests, regaining their hat will involve contacting more than half of the people? In other words, what is the probability that a cycle of length $> n/2$ exists in a uniformly distributed random permutation?

As above, the answer is not difficult and comes down to counting. Let us fix a $k > n/2$. Then the only difference to the reasoning above when counting the permutations is that we do not insist on a fixed element j in the cycle. Hence there are now $\binom{n}{k}$ ways of choosing k elements to form a cycle of length k. Recall from above that there are $(k-1)!$ different cycles containing these elements, and there are $(n-k)!$ ways of completing the permutation outside of the chosen cycle. Thus we find

$$\mathbb{P}(\ell_j = k \text{ for some } j) = \frac{1}{n!}\binom{n}{k}(k-1)!(n-k)! = \frac{1}{k}. \tag{2}$$

A hidden pitfall in this argument is that it would not work if we had considered $k \leqslant n/2$. For example, a permutation having two cycles of length $n/2$ would have been counted twice. But as we are only interested in cycles of length $> n/2$, things are fine.

Observe that generally a permutation σ cannot have a cycle of length k and a cycle of length m at the same time when $k \neq m$ and $k, m > n/2$. In the language of probability theory we say that the events that there exist cycles of length k are *disjoint* for different $k > n/2$. This allows us to simply add up the probabilities $\mathbb{P}(\ell_j = k \text{ for some } j)$ for all $\lfloor n/2 \rfloor + 1 \leqslant k \leqslant n$, where the *floor function* $\lfloor x \rfloor$ of x is defined to be the largest integer smaller than x, to obtain

$$\mathbb{P}(\ell_j > n/2 \text{ for some } j) = \sum_{k=\lfloor n/2 \rfloor + 1}^{n} \frac{1}{k} \overset{n\to\infty}{\longrightarrow} \ln 2 \approx 0.693. \tag{3}$$

As noted in the equation, for a very large party, when n grows to infinity, the probability that there is at least one person for whom more than half of the people will have to get involved before she gets her hat back, is about 0.693.[3] This calculation is at the heart of a very intriguing piece of recreational mathematics, known as the 100 prisoners problem, which you can read about in [12]. Much more is known about uniform random permutations, but we want to go in a different direction now.

---

[3] You can compute this result for example by using integral estimates for the harmonic series that appears here.

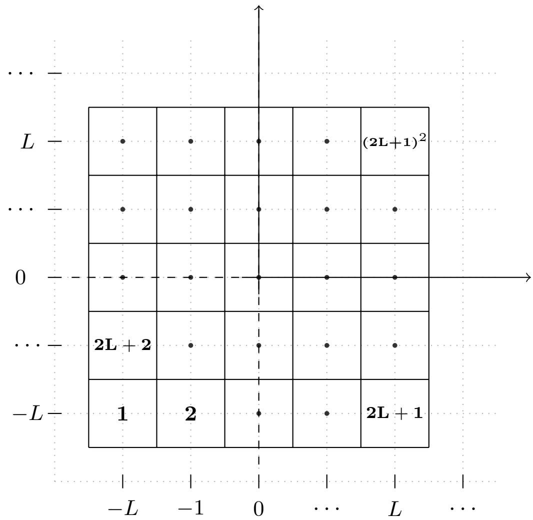

Figure 2: The underlying geometry. The dotted lines show the $\mathbb{Z}_2$-lattice structure; X_L is the 'cutout' defined by the chessboard pattern. Positions x_j are given by the integer coordinates of X_L.

2 Spatial random permutations

Assume that the host of the party takes precautions so that the guests may more easily find their hats. He prepares a large room with a chessboard pattern on the floor, and as many tiles as there are guests. Each guest places their hat on the center of a tile. After the party, the guests go to the tile where they think their hat is. Being drunk, they may be off by a few tiles in any direction, so two or more people may end up claiming the same hat. Such disputes need be resolved, but for our discussion here we will disregard this problem. Eventually everyone should hold precisely one hat, not necessarily their own. The question we pose is the same as above: how many people will need to be called before person j recovers her hat?

In the language of random permutations, the situation can be described as follows. For $L \in \mathbb{N}$ let $X_L = \mathbb{Z}^2 \cap [-L, L]^2$. Each point x_j, $j \leqslant n = (2L+1)^2$, in X_L represents the position of hat number j.

As above, a permutation σ on $\{1, \dots, n\}$ means that person j picked up hat number $\sigma(j)$, but now we do not want to assume that all permutations occur with equal likelihood. Rather, we expect that the guests might still remember roughly where they left their hat and thus are more likely to grab a hat close to their own. A naive way to implement this assumption would be to say person number 1 picks a random hat near (say, no more than 2 tiles from) position x_1, then person 2 picks one of the hats still available near position x_2, and so on. Unfortunately, this will not work. There is a high probability that at some point person would j find that all hats near x_j are already taken, and thus would need to find a hat far from x_j.

A better way is to assign a *cost function* (or "energy") to each person–hat pair $(j, \sigma(j))$ which becomes large when a person picks a hat far from their own tile. For simplicity and concreteness, we pick $\phi_j(\sigma) = |x_j - x_{\sigma(j)}|^2$ for our cost function. Then we say that the sum

$$H(\sigma) = \sum_{j=1}^{n} \phi_j(\sigma) = \sum_{j=1}^{n} |x_j - x_{\sigma(j)}|^2 \tag{4}$$

represents the "total energy" of the system. For a parameter $\alpha \geqslant 0$, and we will see later what the purpose of this parameter is, we define the probability of a given permutation σ to be

$$\mathbb{P}_L(\sigma) = \frac{1}{Z_L(\alpha)}\, \mathrm{e}^{-\alpha H(\sigma)}\,, \quad \text{with } Z_L(\alpha) = \sum_{\sigma \in \mathcal{S}(n)} \mathrm{e}^{-\alpha H(\sigma)}\,. \tag{5}$$

The factor $Z_L(\alpha)$ is the sum over all possible permutations σ. A permutation obtained via the probability measure defined in equation (5), where the probabilities depend on the location of the points x_j, will be called a *spatial random permutation* (SRP for short).

You can see that $\mathbb{P}_L(\sigma)$ gets small very fast for a permutation with large energy $H(\sigma)$ since the probability is directly proportional to the exponential factor $\mathrm{e}^{-\alpha H(\sigma)}$. A permutation with $H(\sigma)$ large is thus very unlikely; this effect is stronger when α is large and weaker when α is small. Indeed, the identity permutation id where every person picks their own hat (and thus $\phi_i(\mathrm{id}) = 0$ for all i) has the smallest total energy, but this does not mean that this situation will occur very often. The reason is that there is only one identity permutation, and the normalization constant $Z_L(\alpha)$ becomes very large when $n = (2L+1)^2$ is large.

For example, let us consider the situation where everyone except two adjacent people obtains their own hat. There are n ways to pick who will have the first misplaced hat, and 4 ways to pick a neighbor (except at the boundary of the chessboard, but this does not make a big difference in our calculation if we assume the party to be quite large). The energy of each such permutation is 2, and so the probability that everyone except two people obtains their own hat is $\frac{2n}{Z_L(\alpha)}\, \mathrm{e}^{-2\alpha}$. If we fix α, then for a very large number of people n, this will always be much larger than $\mathbb{P}_L(\mathrm{id}) = \frac{1}{Z_L(\alpha)}$.

More generally, if we think of not only one pair of misplaced neighboring hats, but k such pairs, where k is a number such that k is still much smaller than n then we have – again roughly and up to boundary effects – a number of $2n \cdot (2n-2) \cdot \ldots \cdot (2n-2k)$ ways to do that. Since $2n \cdot (2n-2) \cdot \ldots \cdot (2n-2k) > (2n-2k)^k$, the probability of σ having k misplaced neighboring pairs is greater than $\frac{(2n-2k)^k}{Z_L} \mathrm{e}^{-2\alpha k}$. For large n and moderate k this is again much larger than the probability of finding one pair.

We see that there is a competition at work: misplacing many hats makes $H(\sigma)$ large, but at the same time there are by far more ways to misplace many hats than there are to misplace only a few. This competition between energy (the size of $H(\sigma)$) and "entropy" (the number of permutations that have a certain property, for example, having precisely k displaced hats) puts the spatial hat problem into the field of *statistical mechanics*.[4]

The parameter α determines the respective importance of the two effects. If α is small, a typical permutation will have many misplacements, while for large α, many people will find their own hat. Figure 3 shows two typical random permutations obtained with the probability measure (5), for different α. Looking at the pictures, one could easily conjecture that for large α, every telephone queue for recovering a given hat will be short, no matter how many guests there are, while for smaller α the length of a typical queue will grow with the number of guests. When translated to the language of permutations, this means that we conjecture long cycles to be present when α is small, but absent when α is large. In fact, it is not a simple thing to decide if this statement is true or false, the answer is a subject of current research. We will come back to the model and some of the recent findings in the last section, but before that, let us look beyond the hat problem.

3 Spatial random permutations: solved and unsolved problems

SRPs would not be so interesting if they were only about a variant of the hat problem. It turns out, however, that they have relationships with many other parts of mathematics and physics. Here we discuss two of these connections.

---

[4] Statistical mechanics is the branch of physics which takes a statistical approach to describing the behavior of very large numbers of particles, for example, the number of atoms in a gas. The behavior of the individual particles is deterministic, but if the number of them becomes sufficiently high, exact calculations of the behavior of the whole system based on the microscopic behavior of the individual particles becomes totally impractical. Considering the system as effectively random has proved to be a very useful technique, and is for example the basis of the theory of *thermodynamics*.

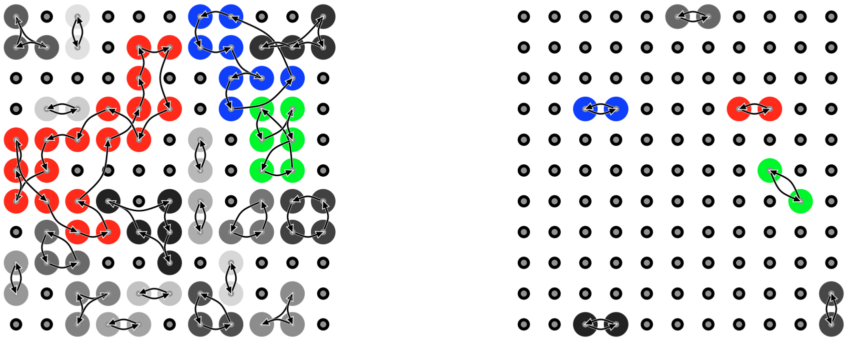

Figure 3: Two examples of a spatial random permutation with side length 11, corresponding to $L = 5$. On the left, $\alpha = 0.7$, on the right, $\alpha = 2$. An arrow from a point x_j to a point y means that $x_{\sigma(j)} = y$, that is that person j, whose hat was at x_j, picked up the hat at the position y.

3.1 SRPs and bond percolation

Firstly, an SRP is a variant of *constrained bond percolation*. Percolation theory has its origins in the mathematical description of the percolation of fluids through porous media, with the basic question being "Does the liquid reach from the center to the boundary if we consider a large piece of material?". This question has given rise to many interesting and difficult mathematical problems, and has found applications in different fields, for instance, epidemiology and geology. The relevant model of bond percolation is the following: in a rectangular grid, such as X_L in Figure 2, toss an unfair coin (with heads probability p, $0 < p < 1$) for each pair of neighboring points. If the coin comes up heads, connect the pair with a straight line, a "bond". After doing this for all pairs, check whether the center of the the grid is connected with its boundary. If this happens with a positive (that is, not 0) probability when the grid grows infinitely large, we say that percolation occurs. It was proved by Harry Kesten in 1980 [7] that bond percolation on the square lattice occurs if and only if $p > 1/2$. We say that there is a *percolation transition* at the critical parameter $p_c = 1/2$.

What does bond percolation have to do with SRPs? We will find out by first considering a variant of the model given in equation (5): instead of allowing (in principle) all permutations σ to occur, let us now forbid those permutations where $|x_i - x_{\sigma(i)}| > 1$. In the context of the hat problem, this means that each guest picks up a hat that is either her own, or lies directly above, below, to the left, or to the right of her own. The set of all permutations that are still allowed is denoted by $\mathcal{S}_1(n)$. The energy $H(\sigma_k)$ of a permutation $\sigma_k \in \mathcal{S}_1(n)$ where k hats are misplaced is then equal to k, and such a permutation has the probability

$$\mathbb{P}_L(\sigma_k) = \frac{1}{\tilde{Z}_L(\alpha)}\,\mathrm{e}^{-\alpha k}\,, \quad \tilde{Z}_L(\alpha) := \sum_{\sigma\in\mathcal{S}_1(n)} \mathrm{e}^{-H(\sigma)}\,. \tag{6}$$

The quantity $\tilde{Z}_l(\alpha)$ is again the sum over all allowed permutations, but note that it is different from $Z_L(\alpha)$ since $\mathcal{S}_1$ is much smaller than $\mathcal{S}$. Note also that thanks to our restriction to nearest neighbor permutations from $\mathcal{S}_1$, the model now makes sense (and is interesting) for all $\alpha \in \mathbb{R}$, not only for $\alpha > 0$.

Why does this variant have anything to do with bond percolation? To answer this question, let us write ω_k for a combination of connected and unconnected neighbors of the set X_L as occurring in bond percolation after having tossed all coins, and with the property that in ω_k precisely k pairs of neighbors have been connected. If m is the number of all pairs of neighbors that exist in X_L, the probability of connecting exactly those of ω_k is given by

$$\mathbb{P}_{L,\mathrm{bp}}(\omega_k) = p^k(1-p)^{m-k} = \Big(\frac{p}{1-p}\Big)^k(1-p)^m =: \frac{1}{Z_{L,\mathrm{bp}}(\beta)}\,\mathrm{e}^{-\beta k} \tag{7}$$

with $Z_{L,\mathrm{bp}} = (1-p)^{-m}$ and $\beta = \ln((1-p)/p)$. By solving for the probability p, we see that in this language, the results of Kesten say that bond percolation occurs if and only if $\beta < 0$. The similarities between equations (6) and (7) should now be obvious: the bond percolation model (7) is of the same form as our special SRP variant (6) when identifying α with β. We could therefore use the latter as a model to analyze bond percolation if we equate the occurence of percolation with the presence of a cycle connecting the point 0 to the boundary of the grid in the SRP model. But there is an important difference. While ω_k can be any arrangement of connections between nearest neighbors, in a a valid arrangement in σ_k for each point x in X_L one of the following three alternatives hold: either x is connected to exactly two of its neighbors, or to none at all, or it is connected to exactly one neighbor y in such a way that the pair $x; y$ is disconnected from all other points.

This difference, and others which we will not address here, have drastic effects. First and foremost, they destroy the "independence" property that is present in the bond percolation model and makes its analysis possible. Here, independence means that the probability of connecting two points is independent of all the other connections already made. This fact leads to the absence of almost any rigorous result for this specific model of SRP. However, numerical investigations (of the closely related model (5)) are possible, and they show that the special rules about which arrangements are possible in our SRP variant makes percolation much harder to achieve than in the original bond percolation model. Numerical simulations [6, 1] suggest that there is a percolation transition, meaning occurence of a cycle connecting 0 to the boundary, only for SRPs in three and more space dimensions. In two dimensions, by contrast, it seems that no matter what α is, the probability of reaching the boundary from the center converges to zero as $n \to \infty$; the expected length of the excursion diverges, though. Of these statements, the only one that we know how to prove is that when α is large enough, there are no long cycles, and no percolation [2].

3.2 SRPs and Bose-Einstein condensation

The other model to which SRPs are related is, perhaps surprisingly, the Bose-Einstein condensate, a state of matter that, unlike the common solid, liquid or gas that we are familiar with in our daily life, can only occur under extreme conditions. Bose-Einstein condensates (BEC for short) were theoretically predicted to exist in 1924/25. Inspired by the work of Satyendra Nath Bose (1894–1974) on the statistics of light particles (photons), Albert Einstein (1879–1955) realized that Bose's mathematics could apply also to atoms and anticipated the existence of a new form of matter, the Bose-Einstein condensate. A BEC is a group of atoms cooled to practically absolute zero. They start to clump together and completely lose their identities as individual atoms, such that the group behaves as though it were a single particle. This state of matter was first experimentally realized in 1995, an achievement which earned its creators a Nobel prize in physics. Although the first theoretical description is almost a century old there are still major open problems to be solved for a complete theoretical description of BEC.

The phenomenon of Bose-Einstein condensation can be linked to an averaged, or "annealed" version of the SRP measure (5), which we will now introduce. Let us pick n completely arbitrary points $\boldsymbol{x} = (x_1, \dots, x_n)$ in a d-dimensional box $\Lambda^d = [-L, L]^d$, and define for every choice of points $\boldsymbol{x}$ the energy in the same way as in (4), with the only difference that the points $\boldsymbol{x}$ are no longer lattice points of $\mathbb{Z}^2$:

$$H(\sigma, \boldsymbol{x}) = \sum_{j=1}^{n} |x_j - x_{\sigma(j)}|^2. \tag{8}$$

The annealed SRP measure is the probability measure on permutations of $\{1, \dots, n\}$ given by

$$\mathbb{P}_L^a(\sigma) = \frac{1}{Z_L(\alpha)} \int_{[-L,L]^d} \mathrm{d}x_1 \dots \int_{[-L,L]^d} \mathrm{d}x_n \, \mathrm{e}^{-\alpha H(\sigma, \boldsymbol{x})}, \tag{9}$$

where now the normalization constant is

$$Z_L(\alpha) = \frac{1}{n!} \int_{[-L,L]^{d\cdot n}} \sum_{\sigma\in\mathcal{S}(n)} \mathrm{e}^{-\alpha H(\sigma,\boldsymbol{x})}\, \mathrm{d}\boldsymbol{x}. \tag{10}$$

Again, we can ask about the typical length ℓ_j of the cycle containing j for fixed j, as the number of points n tends to infinity. Will it stay bounded, or will it grow to infinity?

Whereas we were not able to give a precise answer for the percolation model, this question can be answered in the case of the annealed measure (9). Indeed, the answer matches conjectures made in the percolation model. It was shown in [2] that in two dimensions, for all $\alpha > 0$ the cycle containing j is finite with probability 1 in the limit as n tends to infinity, or, equivalently,

$$\lim_{M\to\infty} \lim_{n\to\infty} \mathbb{P}_L^a(\ell_j > M) = 0.$$

Since this holds for all α, there is no phase transition, that is, there is no critical value of α for which cycles would suddenly have infinite length.

On the other hand, in dimension 3 or higher, there is a phase transition. If we define

$$\alpha_\mathrm{c} := \pi\zeta(d/2)^{-2/d} \tag{11}$$

where $\zeta(z) = \sum_{j=1}^{\infty} j^{-z}$ is the Riemann zeta function, it follows that

$$\lim_{M\to\infty} \lim_{n\to\infty} \mathbb{P}_L^a(\ell_j > M) \begin{cases} = 0 & \text{if } \alpha \geqslant \alpha_\mathrm{c} \\ > 0 & \text{if } \alpha < \alpha_\mathrm{c} \end{cases}.$$

In other words, α_c is exactly the critical value for which the behavior of the cycles changes. For $\alpha < \alpha_c$ we expect to see with a positive probability the appearance of infinitely large cycles containing j, but for $\alpha \geqslant \alpha_c$, the probability of such an event is 0.

The expression (11) is familiar from quantum physics: it is the critical temperature for an ideal Bose gas at density 1, where $\alpha = T/4$ and T is the physical temperature of the quantum gas. This is no coincidence: it turns out that we may describe the free Bose gas as a system of spatial random permutations,[5] and Bose-Einstein condensation then corresponds to the occurrence of infinite cycles in the limit of an infinitely large box. This connection was first observed in a somewhat vague sense by Feynman [4], then made rigorous by Sütő [10], and finally extended and cast into the language of permutations in [2]. The connection between the ideal Bose gas and spatial random permutations is

[5] To make this somewhat more precise: using the Feynman-Kac formula, the density matrix of the gas can be rewritten as such a system of SRPs.

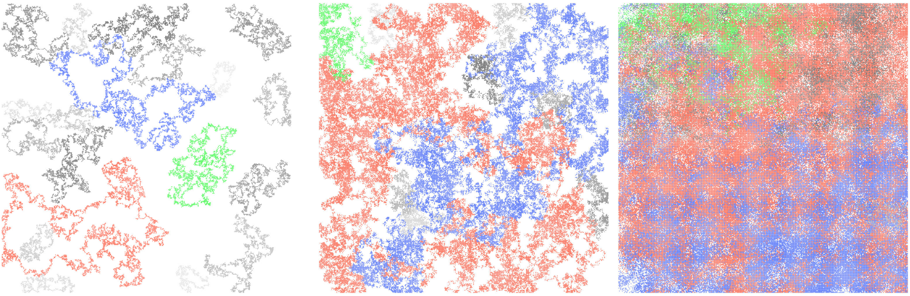

Figure 4: The 20 longest cycles of a SRP of side length 400, having 160000 points. The parameter values are, from left to right, $\alpha = 0.6$, $\alpha = 0.3$ and $\alpha = 0.05$. Points involved in a given cycle are of the same color. For clarity of the image, the arrows indicating where exactly each point is mapped are suppressed, and cycles ranging below 21st in total length are not shown. The holes in the pictures are occupied by such smaller cycles.

not limited to the free Bose gas. As explained in [11] and [3], it works for the interacting Bose gas as well, but leads to a much more complicated model, for which we cannot prove the existence of a phase transition. This is not unexpected; a rigorous understanding of interacting Bose-Einstein condensation at positive temperature is among the most difficult open problems in theoretical physics. What is however interesting is that via SRP, probability theory might be able to contribute to the solution of this problem.

3.3 Spatial random permutations – further links and state of the art

Before we end this article, let us briefly come back to our first SRP model (5) on the rectangular lattice. Even though it looks easier than the annealed model (9) we introduced for the connection to Bose-Einstein condensates, it is much more difficult to handle mathematically, and no rigorous results on the existence of a phase transition are available. The advantage that (5) does have is that it is easy to implement and visualize on a computer, and this gives rise to yet another host of intriguing, but unsolved, problems.

Take a look at Figure 4, which shows the 20 longest cycles of a SRP with $2L + 1 = 400$ for different values of α. You can see that the points belonging to the longest cycles of a SRP are very interesting geometric objects: they look like *fractals*, which are, roughly speaking, objects that reveal fine structure on every scale; no matter how closely we zoom in, there is always finer structure to observe. We can also observe that the long cycles seem to fill the space more densely the smaller the value of α. This intuitive observation coincides with a careful numerical study carried out in [1]. It reveals that the long cycles of SRP in two dimensions indeed appear to have a *fractal dimension*, a characteristic quantity we can view as a measure of how densely the space is filled, in the sense of having "more than" length (1 dimension), but "less than" area (2 dimensions), and that this fractal dimension decreases linearly with α, at least when α is small.

For larger α, SRP cycles appear to undergo what is known as a Kosterlitz-Thouless transition [8, 5]; this means that the probability of finding a point x in a cycle of length K or larger, when viewed as a function of K changes its behavior when α passes a critical value α_c. Finally, there is some evidence that SRP cycles might be related to the traces of Schramm-Löwner curves. Explaining anything about latter object would seriously exceed the capacity of this article, and we refer to [9].

Of the above conjectures, none are proved or disproved at the time of writing, and there is a good chance that all of them will be very difficult. Like many models of statistical mechanics, (5) is easy to describe, but unfortunately difficult to analyze.

References

[1] V. Betz, *Random permutations of a regular lattice*, Journal of Stat. Phys. 155, 1222–1248 (2014)

[2] V. Betz, D. Ueltschi, *Spatial random permutations and infinite cycles*, Commun. Math. Phys. 285, 469–501 (2009)

[3] V. Betz, D. Ueltschi, *Critical temperature of dilute Bose gases*, Phys. Rev. A 81, 023611 (2010)

[4] R. P. Feynman, *Atomic theory of the λ transition in Helium*, Phys. Rev. 91, 1291–1301 (1953)

[5] J. Fröhlich, T. Spencer: *The Kosterlitz-Thouless Transition in Two-Dimensional Abelian Spin Systems and the Coulomb Gas*, Commun. Math. Phys. 81, 527–602 (1981)

[6] D. Gandolfo, J. Ruiz, D. Ueltschi, *On a model of random cycles*, J. Stat. Phys. 129, 663–676 (2007)

[7] H. Kesten: The Critical Probability of Bond Percolation on the Square Lattice Equals 1/2, *Commun. Math. Phys.* 74, 41–59 (1980)

[8] J. M. Kosterlitz, D. J. Thouless, *Ordering, metastability and phase transitions in two-dimensional systems*, Journal of Physics C: Solid State Physics 6: 1181–1203 (1973).

[9] G. Lawler, *Scaling limits and the Schramm-Loewner evolution*, Probab. Surveys 8 (2011), 442–495

[10] A. Sütő, *Percolation transition in the Bose gas*, J. Phys. A 26, 4689–4710 (1993)

[11] D. Ueltschi, *Relation between Feynman Cycles and Off-Diagonal Long-Range Order*, Phys. Rev. Lett. 97, 170601 (2006).

[12] P. Winkler, *Names in Boxes Puzzle*, College Mathematics Journal 37 (4): 260, 285, 289 (2006)

Volker Betz is a professor of mathematics at TU Darmstadt.

Mathematical subjects
Geometry and Topology, Probability Theory and Statistics

Connections to other fields
Physics

License
Creative Commons BY-SA 4.0

DOI
10.14760/SNAP-2019-007-EN

Snapshots of modern mathematics from Oberwolfach provide exciting insights into current mathematical research. They are written by participants in the scientific program of the Mathematisches Forschungsinstitut Oberwolfach (MFO). The snapshot project is designed to promote the understanding and appreciation of modern mathematics and mathematical research in the interested public worldwide. All snapshots are published in cooperation with the IMAGINARY platform and can be found on www.imaginary.org/snapshots and on www.mfo.de/snapshots.

ISSN 2626-1995

Junior Editor
Sabiha Tokus
junior-editors@mfo.de

Senior Editor
Sophia Jahns (for Carla Cederbaum)
senior-editor@mfo.de

Mathematisches Forschungsinstitut
Oberwolfach gGmbH
Schwarzwaldstr. 9 – 11
77709 Oberwolfach
Germany

Director
Gerhard Huisken

Mathematisches Forschungsinstitut Oberwolfach

Member of Leibniz Association

IMAGINARY open mathematics