Self-Avoiding Walks

Gordon Slade

What Are They?

Self-avoiding walks are discrete paths without self-intersections. For our purposes they live in the d-dimensional integer lattice $\mathbf{Z}^d$, which consists of the points in $\mathbf{R}^d$ whose components are all integers. Elements of $\mathbf{Z}^d$ will be referred to as *sites*. Two sites are called nearest neighbours if they are separated by unit Euclidean distance. An n-step *self-avoiding walk* is then defined to be an ordered set of $n+1$ sites in $\mathbf{Z}^d$ for which each consecutive pair of sites consists of nearest neighbours and in which no site occurs more than once (Fig. 1). It is this last feature that gives the walk its self-avoiding character and turns out to make life difficult. For despite its simple definition, the self-avoiding walk leads to mathematical problems which are simple enough to state to the mathematically uninitiated, but which are very hard to solve and mainly still open. This article gives an introduction to some of these mathematical problems; a detailed account can be found in [23].

The self-avoiding walk has been used for some time as a model of linear polymers. Linear polymers are molecules which form in long chains of basic units called monomers. The chains can get pretty long: some consist of about 10^5 monomers. Polymer scientists want to know how many different configurations an n-monomer chain can adopt, and also how far apart the endpoints of the molecule typically are, assuming each configuration is equally likely. These questions translate into questions about self-avoiding walks, where the self-avoidance constraint models the *excluded volume effect*: No two monomers can occupy the same region in space. For polymers, n is very large, so it is natural to ask about the asymptotic behaviour of properties of the set of n-step self-avoiding walks, in the limit as n goes to infinity. Real polymers live in the continuum rather than on a lattice, but this local issue turns out to be largely irrelevant for long-range questions, and the lattice approximation is a good one for large n. For polymers, the most relevant dimension is $d = 3$, but 2-dimensional self-avoiding walks are also important as a model of polymers constrained between two narrowly separated

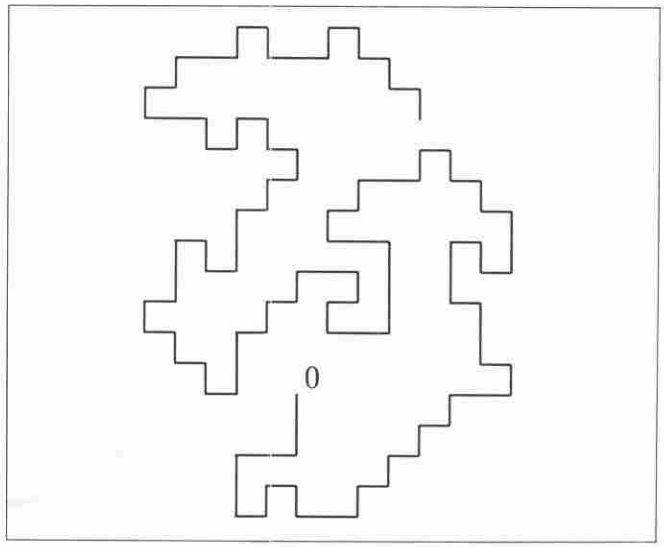

Figure 1. A 2-dimensional self-avoiding walk with 105 steps, beginning at 0. A subsequent step to the left would eventually lead to a trap.

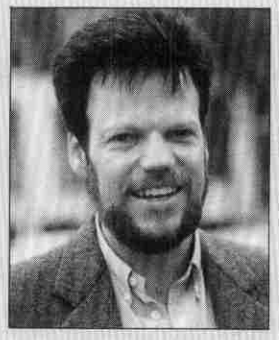

Gordon Slade

Gordon Slade began his university education at the University of Toronto, studying Engineering Science as an undergraduate, and beginning graduate study in mathematics. He received a doctorate in mathematics from the University of British Columbia in 1984. Following postdoctoral work at the University of Virginia, he moved to McMaster University in 1986, where he is currently Associate Professor. He has held visiting positions in Canada, France, Germany, Great Britain, Japan, Switzerland, and the United States. His current research interests centre on the critical behaviour of random systems.

parallel planes. It is perhaps less clear what dimensions 4 and higher have to do with polymers, but it is instructive nevertheless to study how the behaviour of the model changes as a function of the spatial dimension d.

Visualizing a walk as a path in the lattice, the self-avoidance constraint says that the path has no loops. Although a walk is defined to be a sequence of sites, it is useful to think of it as the continuous path formed by joining each pair of consecutive sites by the intervening unit line segment; see Figure 1. Actually, the term "walk" can lead to some confusion, and perhaps "chain" would be better. The self-avoiding walk is not a model of polymers which are growing in time: n is static. It is tempting to try to think of the self-avoiding walk as a kind of non-Markovian stochastic process (a Markov process forgets its past, whereas the self-avoiding walk has to remember its entire history), but the self-avoiding walk is not only non-Markovian — it is also not a process. A walk may be trapped and impossible to extend by another step.

The Connective Constant

Let c_n denote the number of n-step self-avoiding walks which begin at the origin. This measures the number of possible configurations of a polymer of $n+1$ monomers. For simple random walks, which have no self-avoidance constraint, the analogue of c_n is just $(2d)^n$ as there are $2d$ options for the walk at each step. The situation is not as easy for self-avoiding walks, but we can at least say the following. All possible $(n+m)$-step self-avoiding walks can be formed by concatenating n-step self-avoiding walks to m-step walks, but not all such concatenations will be self-avoiding. This means that $c_{n+m} \leq c_n c_m$, or in other words the sequence $\{c_n\}$ is submultiplicative.

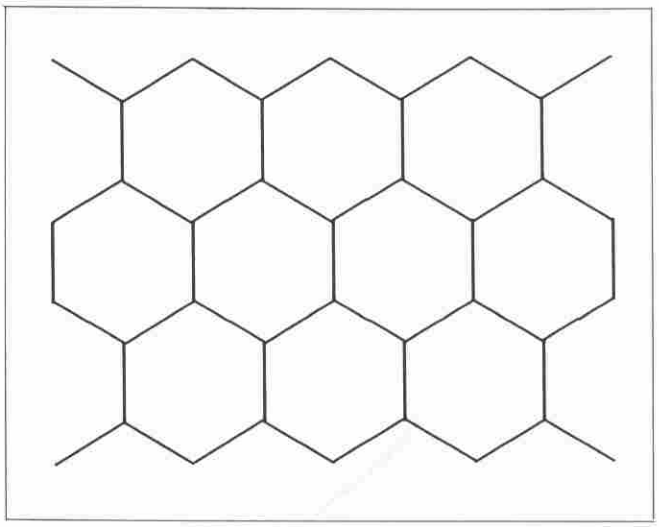

Figure 2. A piece of the honeycomb lattice, the only lattice for which the connective constant is believed to be precisely known.

Table 1. Rigorous lower and upper bounds on the connective constant μ, together with estimates of actual values, for dimensions 2, 3, 4, 5, 6. Estimated errors in the last digit(s) are shown in parentheses.

d	Lower Bound	Estimate	Upper Bound
2	2.620 02[a]	2.638 158 5 (10)[d]	2.695 76[b]
3	4.572 140[c]	4.683 907 (22)[e]	4.756[b]
4	6.742 945[c]	6.772 0 (5)[f]	6.832[b]
5	8.828 529[c]	8.838 6 (8)[g]	8.881[b]
6	10.874 038[c]	10.878 8 (9)[g]	10.903[b]

References: (a) [6], (b) [1], (c) [17], (d) [5, 11], (e) [8], (f) [9], (g) [10].

It is an immediate consequence of submultiplicativity that the sequence $\{(c_{2^k})^{1/2^k}\}$ is nonincreasing. With a bit more work it can be shown that submultiplicativity ensures the existence of the limit $\mu \equiv \lim_{n\to\infty} c_n^{1/n}$ and that $c_n \geq \mu^n$ for all n. This limit μ is known as the *connective constant* and was first shown to exist by Hammersley and Morton [13]. Roughly speaking, c_n is of order μ^n for large n, so that μ measures the average number of possible next steps available for a long self-avoiding walk. Of course, any particular walk can have anywhere from zero to $2d-1$ possible next steps.

It is not hard to see that $d^n \leq c_n \leq (2d)(2d-1)^{n-1}$, and hence $d \leq \mu \leq 2d-1$. The lower bound is due to the fact that walks which only take steps in positive coordinate directions are surely self-avoiding. The upper bound follows from the observation that the set of all walks having no immediate reversals (steps which reverse their predecessor) includes all self-avoiding walks and can be counted by noting that such a walk has $2d$ choices for its first step and $2d-1$ for each subsequent step.

Improving these estimates on μ is an activity that some people enjoy, and Table 1 gives the current status for dimensions 2 through 6. The estimates of the precise value of μ are obtained from series extrapolation methods: Values of c_n are enumerated on a computer (a nontrivial job) for n as large as can be managed, and then μ is estimated from these data. For $d=2$ the current world record is $n=39$, but this is changing rapidly. Table 2 gives the known values in two dimensions.

The precise value of μ is not known in any dimension. Early guesses that in two dimensions μ equals $1+\sqrt{2} = 2.4142\ldots$ or $e = 2.7182\ldots$ have been ruled out by the rigorous bounds. There is one interesting exception where μ is believed to be known exactly: For the 2-dimensional honeycomb lattice (Fig. 2) there is strong evidence [26] from physical arguments that $\mu = \sqrt{2+\sqrt{2}}$. This intriguing value has been confirmed numerically, but not yet by a rigorous proof.

Table 2. Values of c_n on the 2-dimensional square lattice. The most recent additions to this table are from [5].

n	c_n	n	c_n
1	4	21	2 408 806 028
2	12	22	6 444 560 484
3	36	23	17 266 613 812
4	100	24	46 146 397 316
5	284	25	123 481 354 908
6	780	26	329 712 786 220
7	2 172	27	881 317 491 628
8	5 916	28	2 351 378 582 244
9	16 268	29	6 279 396 229 332
10	44 100	30	16 741 957 935 348
11	120 292	31	44 673 816 630 956
12	324 932	32	119 034 997 913 020
13	881 500	33	317 406 598 267 076
14	2 374 444	34	845 279 074 648 708
15	6 416 596	35	2 252 534 077 759 844
16	17 245 332	36	5 995 740 499 124 412
17	46 466 676	37	15 968 852 281 708 724
18	124 658 732	38	42 486 750 758 210 044
19	335 116 620	39	113 101 676 587 853 932
20	897 697 164		

Because of trapping, it is not immediately apparent even that $c_n < c_{n+1}$ for all n, although this strict monotonicity has been proved recently [27]. In fact, one would expect from the definition of the connective constant that c_{n+1}/c_n approaches the limiting value of μ as $n \to \infty$, but this has only been proved for $d \geq 5$ and remains open for $d = 2, 3, 4$. (The case of $d = 1$ is trivial for self-avoiding walks.) Thirty years ago Kesten [19] showed that $\lim_{n\to\infty} c_{n+2}/c_n = \mu^2$ in all dimensions, using a sophisticated argument based on a "pattern theorem," but his proof doesn't work if the subscript $n+2$ is replaced by $n+1$.

The Mean-Square Displacement

The standard measure of the average end-to-end distance of an n-step self-avoiding walk or polymer is the mean-square displacement, which is denoted by $\langle R_n^2 \rangle$ and is defined as the average of the squared Euclidean distance between the endpoints of a walk. The average is taken over all possible n-step self-avoiding walks, with each walk equally weighted. For the simple random walk, which is a Markov process, an elementary probability argument[1] shows that the analogue of the mean-square displacement is exactly equal to n. In contrast, for the self-avoiding walk the most "obvious" bounds on the mean-square displacement remain unproven in low dimensions. For a lower bound, it seems clear that the self-avoidance constraint should force the self-avoiding walk to move away from its starting point at least as fast as the simple random walk, and hence that $\langle R_n^2 \rangle \geq O(n)$. But it remains an open problem to prove this in dimensions 2, 3, and 4. Nor has an upper bound of the form $\langle R_n^2 \rangle \leq O(n^{2-\epsilon})$ been proved in these dimensions for positive ϵ, even though it seems obvious that the ballistic n^2 behaviour cannot be typical above one dimension.

[1] Let X_i $(i = 1, 2, 3, \ldots)$ be independent and identically distributed random variables such that X_i takes values equal to the $2d$ positive and negative unit vectors in $\mathbf{Z}^d$ with equal probabilities $1/2d$. The analogue of $\langle R_n^2 \rangle$ for the simple random walk is the variance of the random variable $X_1 + \cdots + X_n$. Because the variance of a sum of independent random variables is the sum of the variances and because the variance of each X_i is unity, the analogue of $\langle R_n^2 \rangle$ is n.

Critical Exponents

Chemists and physicists can't always afford to wait for rigorous proofs when they need results, and this is the case for the self-avoiding walk. They now have precise conjectures about the behaviour of the number of n-step walks and of the mean-square displacement, and more. Actually some physicists and chemists may be surprised to see their results reported as conjectures, as they regard them rather as facts. Whatever we call them, they are almost certainly correct. The consensus is that there are *critical exponents* γ and ν, and amplitudes A and D, such that

$$c_n \sim A\mu^n n^{\gamma-1}, \tag{1}$$

$$\langle R_n^2 \rangle \sim Dn^{2\nu}. \tag{2}$$

Here the symbol $\sim$ means that the left side is asymptotic to the right side as $n \to \infty$, in the sense that their ratio has limiting value of unity.

These critical exponents are the subject of a considerable body of research, not least because they are believed to be *universal*. Universality means that they should depend on the spatial dimension and on almost nothing else. For example, the exponents are believed to have the same value on all 2-dimensional lattices such as the square or honeycomb. This will certainly not be the case for the connective constant, which describes the average number of next steps available for a long self-avoiding walk and, hence, depends greatly on the lattice. Because they are believed to be universal, the critical exponents are physically more meaningful than the connective constant. Their dependence on dimension makes it natural to study the self-avoiding walk in general dimensions, not just in dimensions 2 and 3 where the application to polymers is most clear.

Numerous tools have been used to arrive at relations (1) and (2). Important confirmation has been provided by numerical studies, including extrapolation of exact enumeration data and Monte Carlo simulation. The renormalization group method has provided a powerful formal tool in the understanding of these relations. Application of the renormalization group has been facilitated by a remarkable connection between the self-

avoiding walk and the theory of ferromagnets, which provides a link with the general theory of critical phenomena and phase transitions. In this approach it has proved useful to employ a representation in which the spatial dimension d can be replaced by a *complex* variable and to study the behaviour of the model as a function of this upgraded dimension. Conformal field theory (the physicist's, not the algebraist's, field theory) has played an important role in two dimensions.

Relations (1) and (2) were recently proved for dimensions 5 and higher by Hara and Slade [15, 16], with the values $\gamma = 1$ and $\nu = 1/2$. These are the same values as for the simple random walk, which indicates that the self-avoidance constraint is not playing a terribly dramatic role in high dimensions. It is, however, known that D is *strictly* greater than 1 for $d \geq 5$: The self-avoidance constraint does push the walk away from the origin faster than simple random walk, but only at the level of the amplitude and not at the level of the exponent. For $d = 5$ the bounds $1.098 \leq D \leq 1.803$ give limits on this effect; as $d \to \infty$ it is known that D approaches the corresponding simple random walk value of 1.

The proof of these results relies on an expansion known as the lace expansion, first introduced by Brydges and Spencer [4]. This expansion has its roots in the cluster expansions of statistical mechanics and constructive quantum field theory. Brydges and Spencer used the lace expansion to study the *weakly* self-avoiding walk in more than four dimensions. The weakly self-avoiding walk involves a measure on simple random walks in which all self-avoiding walks receive the same weight, whereas walks which intersect themselves receive a slightly smaller weight—there is a small penalty for each intersection. The weakly self-avoiding walk is believed on the basis of the renormalization group to have the same critical exponents as the usual self-avoiding walk whenever the penalty for intersections is strictly positive (an example of universality), and the small parameter helps a lot with convergence issues. The Hara–Slade proof involves instead using the inverse dimension as a small parameter. More precisely, the small parameter is proportional to $(d-4)^{-1}$, which is a serious but surmountable hindrance when $d = 5$.

It is perhaps not surprising that high dimensions are easier to treat because in high dimensions it is more difficult for a walk to run into itself, and, therefore, the self-avoidance constraint is weaker. Four dimensions is the borderline case here, as is well known from the theory of simple random walks [21]. For example, the probability that two independent simple random walks of length n do not intersect remains bounded away from zero as $n \to \infty$ for dimensions $d > 4$, but not for $d \leq 4$. Nevertheless, it is still a nontrivial problem to treat the high-dimensional case, and the proof given in [15, 16] is long, technical, and computer-assisted.

For $d = 4$ it is believed on the basis of physical and numerical arguments that the critical exponents are the same as for $d \geq 5$, but that the right sides in relations (1) and (2) should be modified by the insertion of factors equal to the fourth root of the logarithm of n. This is a subtle deviation from simple random walk behaviour, and four dimensions is particularly amenable to renormalization group analysis. Indeed, in the physics literature [22] one finds that three dimensions can be treated by expanding about four dimensions in a complex parameter $\epsilon = 4 - d$ and then setting $\epsilon = 1$. This ϵ-expansion, introduced 20 years ago by Wilson and Fisher, has become a standard tool in the renormalization group. Although it remains unclear how to apply such arguments to study $d = 3$ in a mathematically rigorous way, important steps have recently been taken in making renormalization group arguments rigorous for the weakly self-avoiding walk in four dimensions [2, 3, 18].

For $d = 3$ it is believed that γ is about 1.162 and ν is about 0.588, whereas for $d = 2$ it is believed that γ is exactly equal to the astonishing value $43/32$ and ν is exactly equal to $3/4$. These 2-dimensional values were first predicted by Nienhuis [26] on the basis of connections with the theory of ferromagnetism. Numerical work leaves little room for doubt that these values are correct, but a proof is lacking.

It will likely be some time before things are completely settled in a mathematically rigorous fashion in dimensions 2, 3, and 4. The 30-year-old bounds

$$c_n \leq \begin{cases} \mu^n \exp[Cn^{1/2}], & d = 2 \\ \mu^n \exp[Cn^{2/5} \log n], & d = 3 \\ \mu^n \exp[Cn^{1/3} \log n], & d = 4, \end{cases} \tag{3}$$

which are due to Hammersley and Welsh [14] and Kesten [20], have not been improved, nor has the submultiplicativity bound $c_n \geq \mu^n$ been improved in dimensions 2, 3, or 4. The bounds (3), whose proofs have at their heart an elegant argument relying on submultiplicativity, are still a long way from (1). And as has already been mentioned, the situation is even more embarrassing for the mean-square displacement.

Monte Carlo Methods

Suppose you want to measure the mean-square displacement on a computer, to check the asymptotic relation (2). One way is to compute $\langle R_n^2 \rangle$ exactly for as many small values of n as you can and then extrapolate, and a lot of work has been done in that direction [12].

Another way is to do a Monte Carlo experiment. In a Monte Carlo experiment, you first fix a value of n and generate a lot of (hopefully) independent examples of n-step self-avoiding walks, and then take the average of the square displacements of this sample. If your sample is "typical," then the measured average should be a good approximation to the true mean-square displacement (in which the average is taken over all n-step self-avoiding

walks and not just your particular sample). The process can be repeated for several different values of n, and the results fitted to (2). Interesting algorithmic issues arise.

How do you generate a random sample of 100-step self-avoiding walks? One way that jumps to mind is to do the following. Start by constructing the first step, by picking a neighbour of the origin at random. Then proceed inductively by choosing your next site at random from those neighbours of your current site which have not previously been visited, until you have a 100-site walk. If in this process you become trapped, so that any next step would force an intersection, then discard the walk and start over.

This is certainly a way of generating 100-step walks, but unfortunately they will have the wrong probability distribution. This can be easily seen by looking at a specific example: Consider 4-step walks in two dimensions. The walk NEEE, where N denotes a step to the north (upwards) and E denotes a step to the east (to the right) has probability $\frac{1}{4} \times \frac{1}{3} \times \frac{1}{3} \times \frac{1}{3} = \frac{1}{108}$, whereas the walk NESS has probability $\frac{1}{4} \times \frac{1}{3} \times \frac{1}{3} \times \frac{1}{2} = \frac{1}{72}$. Because these two walks have different probabilities of being generated by the algorithm, taking averages over walks produced by this algorithm will not be the same as taking averages with respect to walks with the desired uniform distribution, in which all walks of the same length are equally likely. Hence, measured quantities like the mean-square displacement cannot be trusted if the walks are generated by this algorithm. Indeed, there are good reasons to believe that the distribution of walks generated by this algorithm (sometimes called "true" or "myopic" self-avoiding walks) has very different properties from the uniform distribution of the self-avoiding walk.

A second natural method is to begin with a particular 100-step walk, such as a straight line, and then to perform a sequence of random modifications of the walk to generate new walks. Perhaps the simplest such modification would be to perform *local moves*, which just change a few (say up to k) contiguous sites of the walk. For example, a small subwalk could be chosen randomly from within a long walk and then replaced by a different self-avoiding subwalk having the same length and endpoints (unless the original subwalk occurs at the beginning or end of the original walk, in which case the unattached endpoint can move). If a self-avoiding walk is produced, then the new walk is kept; otherwise the new walk is rejected and we try again. The set of rules which are used to say how replacements are made defines a specific algorithm. This is usually arranged in such a way that the procedure is reversible, which means that the probability of a particular transition being made is equal to the probability of the reverse transition.

Unfortunately, all such algorithms suffer from the following fundamental drawback. A theorem due to Madras and Sokal [24] states that for any fixed k, any such algorithm can explore only an exponentially small subset of all n-step self-avoiding walks. More precisely, given a reversible local algorithm in which each move changes up to k contiguous sites, define $a_n(k)$ to be the maximum, over all possible n-step initial walks, of the number of walks which can be reached from the initial walk by applying any number of allowed moves to the initial walk. The Madras–Sokal theorem states that for any fixed k,

$$\limsup_{n\to\infty} a_n(k)^{1/n} < \mu.$$

Because the number of walks is at least μ^n, this means that the algorithm explores only an exponentially small subset of all walks.

The proof of this result relies on Kesten's pattern theorem. A *pattern* is defined as a self-avoiding walk of some fixed length. A *proper pattern* is a pattern that can occur arbitrarily often in some self-avoiding walk, or, in other words, one which does not inherently entail traps. The pattern theorem says essentially that any given proper pattern must appear often on all but an exponentially small subset of self-avoiding walks. The theorem about the local algorithms then follows by exhibiting a specific pattern which is frozen, or, in other words, cannot be changed by a local move. The number of occurrences of the frozen pattern remains unchanged under the algorithm. If there are few occurrences of the pattern, then only a small subset of all self-avoiding walks can be explored, by the pattern theorem. On the other hand, if there is a large number of occurrences of the pattern, then again only a small subset of all self-avoiding walks can be explored as only the subwalks which are not part of the occurrences of the frozen pattern are able to change.

An example in two dimensions of a proper pattern which is frozen under local moves changing up to k contiguous sites is shown in Figure 3. A similar pattern has been written down in three dimensions. This has not been done in higher dimensions, and, hence, the theorem has not been completely proved for $d \geq 4$, but this could

Al andar se hace camino,
y al volver la vista atrás
se ve la senda que nunca
se ha de volver a pisar.

Antonio Machado

(As you go you make the way, and looking back, you see the path your feet will never tread again.)

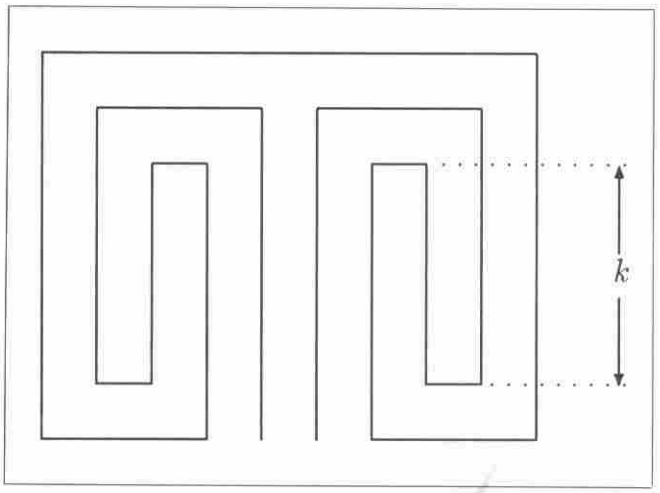

Figure 3. A frozen pattern for local moves.

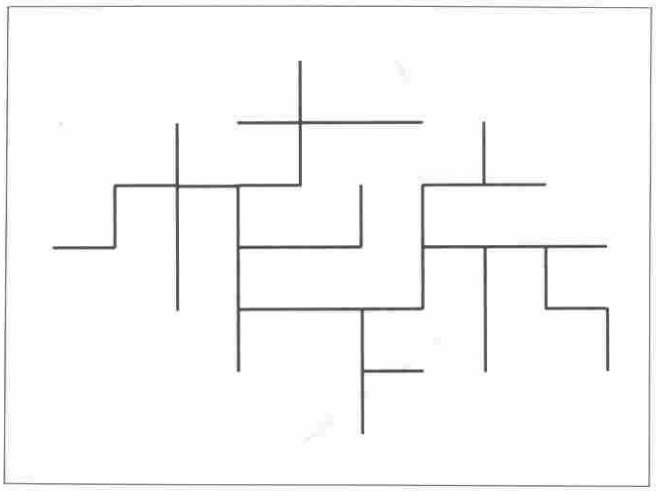

Figure 4. A tree.

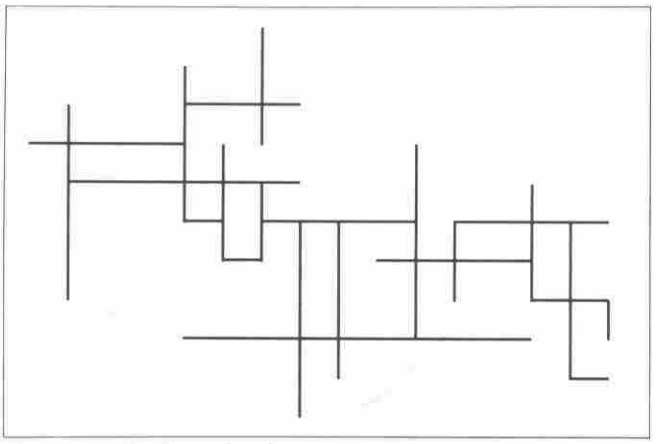

Figure 5. A lattice animal.

likely be done with some effort. The pattern in Figure 3 is the $(10k + 39)$-step walk

$$N^{k+2}W^3S^{k+1}EN^kES^{k+1}W^3N^{k+3}E^9S^{k+3}W^3N^{k+1} \cdot ES^kEN^{k+1}W^3S^{k+2}.$$

An algorithm which does explore the entire sample space of all n-step self-avoiding walks is the *pivot algorithm*. In this algorithm, the first step is to choose a site at random on a self-avoiding walk, thereby dividing the walk into two pieces. Treating this site as the origin of the lattice, one of the pieces is then acted upon by a random lattice symmetry, namely, reflection or rotation. This has the drawback that very often the resulting walk will not be self-avoiding and the trial will therefore be rejected, but this is compensated by the fact that the resulting walk will typically be quite different from the original walk, facilitating a rapid exploration of all corners of the sample space. The pivot algorithm has been studied by many people and used with great success; the most in-depth analysis is in [25].

Related Problems

Not all polymers are linear. Some like to form branches but still remain self-avoiding, and these branched polymers are modelled by trees in the lattice (Figure 4). A *tree* in the lattice is defined to be a finite connected set of nearest-neighbour bonds having no closed loops, where a *nearest-neighbour bond* is a unit line segment joining two neighbouring sites. One can ask similar questions about these, such as how many n-bond trees are there in the lattice, and what is the average radius of gyration of these trees (assuming they are equally likely)? These questions are harder and less studied than the corresponding self-avoiding walk questions.

Another related problem concerns lattice animals. (See Figure 5.) A *lattice animal* is a finite connected cluster of bonds in the lattice, which, unlike a tree, is permitted to form closed loops. If in two dimensions you replace each bond by the square having the bond as diagonal, then a lattice animal models a connected planar beast composed of unit cells. How many are there, and how big are they typically? It is believed that there are critical exponents for trees and animals analogous to those in relations (1) and (2). The most detailed theorems for trees and animals are also for high dimensions, this time (strictly) above eight dimensions, where things appear to be simpler. Very roughly speaking, trees and animals like to be 4-dimensional objects, and life is therefore easier in more than eight dimensions because two 4-dimensional objects generically do not intersect in more than eight dimensions.

Lattice animals are basic in the theory of percolation. In percolation each nearest-neighbour bond in the infinite lattice is assumed to be "occupied" with probability p and "vacant" with probability $1 - p$. Here p is fixed for all bonds and the occupation status of different bonds is independent. The set of occupied bonds decomposes in a natural way into connected clusters, and each cluster will be a lattice animal or its infinite generalization. Percolation provides a model of a porous medium: The bonds which are occupied correspond to pores which admit the

flow of fluid. These pores are to be thought of as microscopic in size, so that fluid flow on a macroscopic scale requires the existence of an infinite connected cluster of occupied bonds.

It is known that for all dimensions greater than or equal to 2 there is a phase transition in this model: There is a critical value $p_c = p_c(d)$ lying strictly between 0 and 1, such that for $p < p_c$ the probability is 0 that there exists an infinite connected cluster of occupied bonds, whereas for $p > p_c$ this probability is 1. In other words, if the bond density p is less than p_c, then there is certainly no fluid flow on the macroscopic scale, but as soon as p is above p_c there certainly is such flow. For any p for which there is an infinite occupied cluster, it is known that with probability 1 the infinite cluster is unique: There cannot be more than one.

Simulations and physical arguments suggest very strongly that at p_c itself there can be no infinite cluster of occupied bonds, but this has been proved rigorously only for $d = 2$ and for very high d. There are several critical exponents that can be defined for percolation; for example, one exponent is defined in terms of the rate of divergence of the expected size of the connected cluster containing the origin of $\mathbf{Z}^d$, as $p \nearrow p_c$. This time it is above six dimensions that existence of critical exponents is best understood — there is still no rigorous proof of their existence otherwise, although the numerical and other evidence leaves little room for doubt that they are out there waiting. Critical exponents and the nature of percolation at and near p_c is a subject which is attracting great attention in the probability community these days; a good introduction is [7].

Acknowledgments

This article was written during a visit to the Department of Mathematics of the University of Virginia, and I am grateful to David Brydges and the Department for their hospitality. I thank David Brydges, Takashi Hara, Neal Madras, and Alan Sokal for teaching me about self-avoiding walks and other topics. I also thank the Natural Sciences and Engineering Research Council of Canada for its support under grant A9351.

References

1. S. E. Alm, Upper bounds for the connective constant of self-avoiding walks, to appear in *Combinatorics, Probability and Computing* **2** (1993).
2. D. Arnaudon, D. Iagolnitzer, and J. Magnen, Weakly self-avoiding polymers in four dimensions. Rigorous results, *Phys. Lett. B* **273** (1991), 268–272.
3. D. Brydges, S. N. Evans, and J. Z. Imbrie, Self-avoiding walk on a hierarchical lattice in four dimensions, *Ann. Probab.* **20** (1992), 82–124.
4. D. C. Brydges and T. Spencer, Self-avoiding walk in 5 or more dimensions, *Commun. Math. Phys.* **97** (1985), 125–148.
5. A. R. Conway, I. G. Enting, and A. J. Guttmann, Algebraic techniques for enumerating self-avoiding walks on the square lattice, *J. Phys. A: Math. Gen.* **26** (1993), 1519–1534.
6. A. R. Conway and A. J. Guttmann, Lower bound on the connective constant for square lattice self-avoiding walks, *J. Phys. A: Math. Gen.* **26** (1993), 3719–3724.
7. G. Grimmett, *Percolation*, Springer-Verlag, Berlin, 1989.
8. A. J. Guttmann, private communication.
9. A. J. Guttmann, On the zero-field susceptibility in the $d = 4$, $n = 0$ limit: analysing for confluent logarithmic singularities, *J. Phys. A: Math. Gen.* **11** (1978), L103– L106.
10. A. J. Guttmann, Correction to scaling exponents and critical properties of the n-vector model with dimensionality > 4, *J. Phys. A: Math. Gen.* **14** (1981), 233–239.
11. A. J. Guttmann and I. G. Enting, The size and number of rings on the square lattice, *J. Phys. A: Math. Gen.* **21** (1988), L165–L172.
12. A. J. Guttmann and J. Wang, The extension of self-avoiding random walk series in two dimensions, *J. Phys. A: Math. Gen.* **24** (1991), 3107–3109.
13. J. M. Hammersley and K. W. Morton, Poor man's Monte Carlo, *J. Roy. Stat. Soc. B* **16** (1954), 23–38.
14. J. M. Hammersley and D. J. A. Welsh, Further results on the rate of convergence to the connective constant of the hypercubical lattice, *Quart. J. Math. Oxford* (2) **13** (1962), 108–110.
15. T. Hara and G. Slade, The lace expansion for self-avoiding walk in five or more dimensions, *Rev. Math. Phys.* **4** (1992), 235–327.
16. T. Hara and G. Slade, Self-avoiding walk in five or more dimensions. I. The critical behaviour, *Commun. Math. Phys.* **147** (1992), 101–136.
17. T. Hara, G. Slade, and A. D. Sokal, New lower bounds on the self-avoiding-walk connective constant, *J. Stat. Phys.* 1993, 479–517.
18. D. Iagolnitzer and J. Magnen, Polymers in a weak random potential in dimension four: rigorous renormalization group analysis, to appear in *Comm. Math. Phys.*
19. H. Kesten, On the number of self-avoiding walks, *J. Math. Phys.* **4** (1963), 960–969.
20. H. Kesten, On the number of self-avoiding walks. II. *J. Math. Phys.* **5** (1964), 1128–1137.
21. G. F. Lawler, *Intersections of Random Walks*, Birkhäuser, Boston, 1991.
22. J. C. Le Guillou and J. Zinn-Justin, Accurate critical exponents from field theory, *J. Phys. France* **50** (1989), 1365–1370.
23. N. Madras and G. Slade, *The Self-Avoiding Walk*, Birkhäuser, Boston, 1993.
24. N. Madras and A. D. Sokal, Nonergodicity of local, length-conserving Monte Carlo algorithms for the self-avoiding walk, *J. Stat. Phys.* **47** (1987), 573–595.
25. N. Madras and A. D. Sokal, The pivot algorithm: A highly efficient Monte Carlo method for the self-avoiding walk, *J. Stat. Phys.* **50** (1988), 109–186.
26. B. Nienhuis, Critical behavior of two-dimensional spin models and charge asymmetry in the Coulomb gas, *J. Stat. Phys.* **34** (1984), 731–761.
27. G. L. O'Brien, Monotonicity of the number of self-avoiding walks, *J. Stat. Phys.* **59** (1990), 969–979.

Department of Mathematics and Statistics
McMaster University
Hamilton, Ontario, Canada L8S 4K1
e-mail: slade@aurora.math.mcmaster.ca