

Ground states of two-dimensional $\pm J$ Edwards-Anderson spin glasses

J. W. Landry* and S. N. Coppersmith†

James Franck Institute, University of Chicago, 5640 South Ellis Avenue, Chicago, Illinois 60637

(Received 6 September 2001; revised manuscript received 11 December 2001; published 12 March 2002)

We present an exact algorithm for finding all the ground states of the two-dimensional Edwards-Anderson $\pm J$ spin glass and characterize its performance. We investigate how the ground states change with increasing system size and antiferromagnetic bond ratio x . We find that the ground-state distribution is log normal and that its dependence on x is nontrivial. We also study the breakdown of ferromagnetic order and show how percolating paths are created by the overlap of different ground-state solutions.

DOI: 10.1103/PhysRevB.65.134404

PACS number(s): 75.50.Lk, 75.10.Nr

I. INTRODUCTION

The Edwards-Anderson (EA) $\pm J$ spin glass¹ is a canonical example of a system with competing interactions that give rise to large numbers of low-energy states. Despite extensive investigation, the low-energy landscape of this model is still controversial in both three²⁻¹⁰ and two dimensions.¹¹⁻²⁸

Although the two-dimensional (2D) $\pm J$ Edwards-Anderson model is much simpler than the three-dimensional case (it does not have a phase transition at nonzero temperature²⁹ and individual ground states can be found in a time that scales as a polynomial of the system size³⁰), it still has many metastable states and a complex energy landscape. At low temperatures spin glass relaxation times become very long, complicating investigations using standard Monte Carlo sampling techniques³¹ and also to varying extent more sophisticated sampling methods such as cluster algorithms^{21,32} and multicanonical methods.³³⁻³⁵

For ground-state properties, exploiting optimization algorithms that find exact ground states has proved a powerful approach.^{13,18,36,37} However, these algorithms find a single ground state of a single realization, and one must sample appropriately from the ground states of each realization^{3,38,39} and also perform a reliable realization average to obtain correct results.

In this paper we investigate the low-energy properties of the two-dimensional $\pm J$ EA model by finding exactly *all* the ground states of each realization. We check that our enumeration is exhaustive by comparing the number of ground states that are found to exact results for the partition function obtained using the method of Saul and Kardar.^{11,40}

In the course of developing our algorithm, we made a series of studies of the distribution of ground states. We find that the the ground-state distribution is consistent with a log-normal distribution and that the mean of the distribution increases with increasing x (the fraction of antiferromagnetic bonds in the system) until it saturates for $0.3 < x < 0.5$.

Though our algorithm is based on an existing polynomial-time algorithm that finds individual ground states,⁴¹ it does not run in polynomial time. This is impossible because the time just to enumerate the ground states grows exponentially with system size. Nonetheless, the number of ground states is vastly smaller than the number of spin configurations, and

empirically we find that our algorithm runs in a time roughly linear in the number of ground states. Memory issues limit our current implementation to about 2×10^6 ground states. Because there are huge variations in the number of ground states among realizations, the system sizes that we investigate are rather small. Though the median number of ground states of a 10×10 system in which half the bonds are antiferromagnetic is 10^4 , at this system size 3% of the realizations have greater than 2×10^6 ground states. Therefore, most of the data presented here is for systems of size 10×10 and less.

The advantage of our method is that it produces qualitatively new information because all the ground states are known explicitly and exactly, so that one can compute in detail the relationships between them. Moreover, these exact results can be used to validate sampling methods appropriate for larger systems.

We use our method to investigate the paramagnetic-ferromagnetic phase transition that occurs as x , the fraction of antiferromagnetic bonds in the system, is increased.^{13,30,42} Quantitative analysis is complicated greatly by the fact that many quantities exhibit large, non-Gaussian variability between realizations. Nonetheless, we are able to show how readily identifiable groups of spins in the various ground states together create a percolating path that destroys the ferromagnetic order of the system.

The paper is organized as follows. Section II presents the algorithm, Sec. III presents data on the distribution of ground states, and Sec. IV describes our investigation of the destruction of ferromagnetic order as the fraction of antiferromagnetic bonds is increased. The results are discussed in Sec. V. Appendix A gives a detailed presentation of the algorithm, while Appendix B discusses the performance of the algorithm in detail.

II. MODEL AND METHODS

A. Edwards-Anderson model

We study the two-dimensional Edwards-Anderson model,¹ in which nearest-neighbor Ising spins ($\sigma_i = \pm 1$) on an $L \times L$ square lattice interact either via a ferromagnetic or an antiferromagnetic coupling. The Hamiltonian is

$$H = - \sum_{\langle ij \rangle} J_{ij} \sigma_i \sigma_j, \quad (1)$$

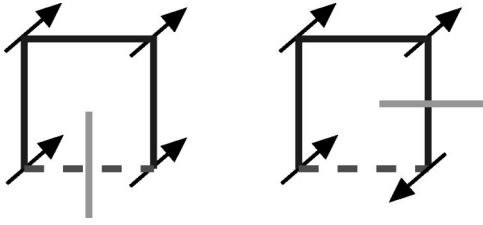


FIG. 1. Two ground states of a frustrated plaquette with four bonds. Ferromagnetic bonds are black lines, while antiferromagnetic bonds are dashed lines. Every configuration of spins produces at least one unsatisfied bond (denoted with a line perpendicular to the bond), and there are four minimum energy configurations.

where the sum $\langle ij \rangle$ is over all pairs of nearest-neighbor spins. Each bond J_{ij} is chosen to be either +1 (ferromagnetic) or -1 (antiferromagnetic). We designate the fraction of antiferromagnetic bonds as x ; $x=0$ is the Ising ferromagnet (no disorder), $x=0.5$ (with equal numbers of ferromagnetic and antiferromagnetic bonds) is the maximally frustrated spin glass (maximum disorder), and $x=1$ is the Ising antiferromagnet (no disorder). Our systems range from $x=0.05$ to $x=0.5$ and have periodic boundary conditions.

B. Calculating all the ground states of the EA model

Our algorithm for finding all the ground states of the EA model first converts the problem of finding ground states into a graphical matching problem, as in Refs. 30 and 43. Next, all possible optimal matching solutions of this problem are found, and finally, these matchings are converted back into spin configurations.

1. Conversion of energy minimization to a matching problem

References 30 and 43 show that the problem of finding a ground state for this spin glass model can be converted to a matching problem in graph theory. Here, we sketch out this conversion and discuss some subtleties that arise from our use of periodic boundary conditions.

A ground state of a spin glass can be described not only in terms of spins and bonds, but also as frustrated plaquettes and paths of broken bonds.³⁰ In a frustrated system, it is not possible for all bonds to be satisfied simultaneously,⁴³ which leads to a natural degeneracy of states. A simple example is shown in Fig. 1.

We denote plaquettes with an odd number of unsatisfied bonds as frustrated, while satisfied plaquettes have an even number of unsatisfied bonds. Frustrated bonds form paths that connect frustrated plaquettes to each other. Because every frustrated plaquette has an odd number of frustrated bonds, it must be the end point of a path. Satisfied plaquettes either have no path through them or are midpoints in a path. This can be seen in Fig. 2, where perpendicular lines have been added to frustrated bonds to show the paths.

We identify the frustrated plaquettes as nodes of a graph and the paths as edges with a weight equal to the number of broken bonds along the path. Ground states have the minimum number of frustrated bonds, so the problem of finding a spin-glass ground state is also the problem of finding those

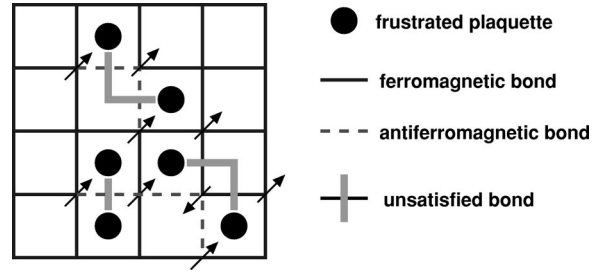


FIG. 2. Sample ground state of a spin glass, showing frustrated plaquettes, unsatisfied bonds, and the corresponding spin and bond configuration.

edges that have the shortest total length. This problem arises in the context of graph theory and is called the *minimum-weight perfect matching problem*.⁴¹ In solutions of this problem, each node is joined to one and only one other node, with the smallest possible total weight (which corresponds to the lowest possible energy). Figure 3 illustrates a sample ground state and its equivalent matching solution.

2. Boundary conditions

References 30 and 44 discuss the relation between spin-glass ground states and solutions to a graphical matching problem and prove a number of results for planar graphs (e.g., free boundary conditions). In these references, the ground-state problem is first converted to a matching problem that can be solved in polynomial time.^{45,46} This matching solution is then shown to correspond always to a spin-glass ground state.

For a periodic lattice, the transformation from spins and bonds to nodes and edges proceeds exactly as in the planar case, and the resulting matching problem can be solved in polynomial time. The issue that distinguishes this problem from the planar case is the conversion of the matching solution back to a ground-state solution. The matching solution found will not always correspond to a ground-state spin configuration for a given toroidal boundary condition. This complication arises because four lattices with four different boundary conditions will produce the same matching problem. These four boundary conditions are periodic on all sides, antiperiodic on the top and bottom, antiperiodic on the left and right, and antiperiodic on all sides.

Boundary conditions can be changed from periodic to antiperiodic on an $L \times L$ system either by setting $J_{iL}^{new} = -J_{iL}$

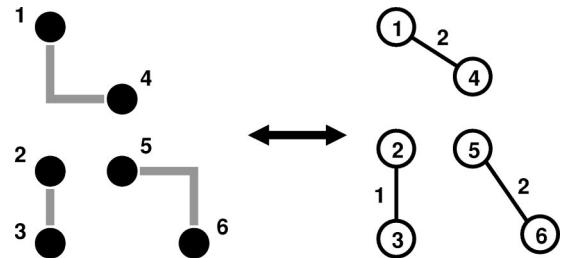


FIG. 3. Correspondence of a spin-glass ground state to solution of the matching problem. The numbers along each edge of the matching solution indicate the weight of that edge.

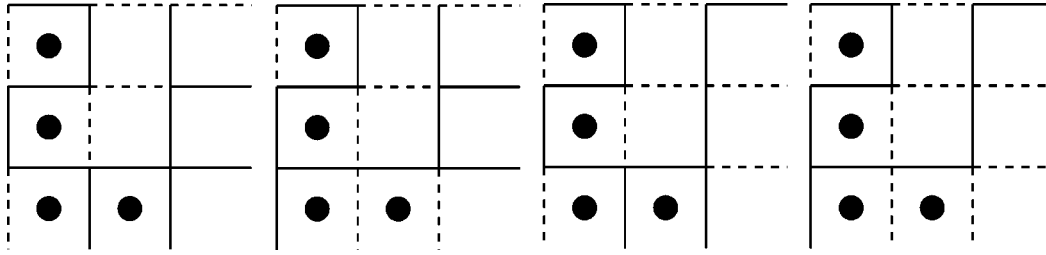


FIG. 4. The four different boundary conditions for a sample lattice. Ferromagnetic bonds are solid lines and antiferromagnetic bonds are dashed lines. Bonds on the right and bottom sides wrap around the lattice to reconnect on the other side. Each boundary condition has the same set of frustrated plaquettes, shown as solid circles.

along the desired edge or by flipping the spins so that $S_{i1}^{new} = -S_{i1}$ along the desired edge. In this study, we flip the bonds. A sample lattice with the four different boundary conditions is shown in Fig. 4.

These four lattices have exactly the same frustrated plaquettes, so they produce the same matching problem. In this sense, the matching problem does not understand boundary conditions. When a matching solution is converted back into spins and bonds, it may not correspond to a ground state for a given boundary condition.

We resolve this ambiguity by converting explicitly each matching into a spin configuration and checking the viability of each spin configuration for each boundary condition. A ground state is only accepted for a given boundary condition if it has a consistent spin configuration. We find numerically that a matching solution always corresponds to a ground-state solution of at least one boundary condition.

This subsection has described the necessary procedures for generating all the ground states in the case of periodic boundary conditions. The algorithm works perfectly for planar graphs without these procedures.

3. Generating all optimal matchings

Our algorithm for finding all the optimal matchings has three parts. The first part finds all edges that make up the optimal solutions. This part exploits the structure of the edges, since the number of edges appearing in the ground states is a small subset of the total number of edges. The second part takes this subset of edges and combines them to find all optimal matchings. The third part converts the optimal matchings into ground-state configurations. The next section describes our algorithm briefly. More detail, including an example, is in Appendix A.

The algorithm uses Edmonds' blossom algorithm,^{45,46} which finds a single optimal solution to a matching problem in polynomial time. We use the Concorde implementation of the algorithm.^{41,47}

C. Finding all edges in all solutions

The algorithm begins by making a list of nodes and possible edges. All frustrated plaquettes are found and designated as nodes. Pairs of nodes that are within a distance of 5 are considered to have edges between them. This restriction controls the combinatorial explosion of possible edges, and optimal solutions involving weights larger than 5 are incred-

ibly rare. Moreover, any errors introduced by this truncation are identified and eliminated at a later stage when the total number of ground states found for a realization is compared to an independent determination of the ground-state degeneracy.

To construct a list of edges that exist in at least one minimal-weight matching, which we designate as viable edges, we start with an empty list. The blossom algorithm is run on the unmodified matching problem and a single matching solution is found. Each edge in this solution is added to the list of viable edges.

To find more viable edges, the nodes are considered successively. For each node, a modified list of edges is created from the original list by deleting those known viable edges connecting to the current node. The blossom algorithm is run on this modified list to find an optimal solution for this new problem. If the solution has the correct path length (i.e., corresponds to a ground state), then the new viable edges that have been found are added to the list of viable edges. The process continues for this node. If the path length of the new solution is too large (i.e., does not correspond to a ground state), then we know that we have found all the viable edges associated with this node. The algorithm then proceeds to the next node in the list. By moving sequentially through the nodes, all viable edges are found.

D. Determining optimal matchings

The next part of the algorithm uses the list of viable edges to find all of the optimal matchings. It picks edges systematically from the list of viable edges until each node is connected by a given edge to one and only one other node. All possible combinations of viable edges in which each node is incident on exactly one edge are examined.

Whenever there is this kind of combination of elements, there is a danger of a combinatorial explosion. In this case, the number of matchings (combinations of edges) is relatively controlled. Section II B discusses this issue in detail.

E. Converting matchings to ground states

All optimal solutions to the matching problem must be converted back into ground-state spin configurations. This conversion is nontrivial because one matching solution can correspond to many different ground states and the same ground state can be represented by different matchings. Simple examples of this phenomenon can be seen in Fig. 5.

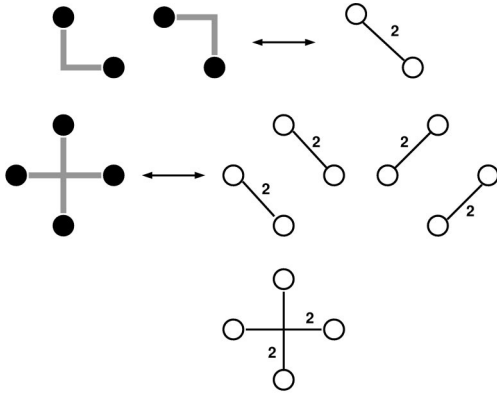


FIG. 5. The relationship between matching solution is not one to one. As shown in the top diagram, two or more different ground states can correspond to the same matching solution. In addition, as seen in the bottom diagram, a single ground state can correspond to multiple matching solutions. The solid circles are frustrated plaquettes, the thick dark lines are unsatisfied bonds, the open circles are nodes, and the thin dark lines are edges with lengths as shown.

To resolve these complications, we keep every ground state we find in memory. Any proposed ground state is checked to see that it does not correspond to a ground state we have already found.

The other important issue is the role of boundary conditions in this conversion from matchings to ground states. We determine the ground state or ground states from the matching by fixing the value of a single spin (in our case, we fix the upper left-hand spin as $+1$). Every other spin follows from this initial spin, because we know the specific bonds of the current boundary condition and their status as satisfied or unsatisfied.

F. Partition function

To check that our algorithm finds every ground state, we also generate the partition function of the realization at $T=0$. This partition function gives the number and energy of the ground states of a given realization. We generate a partition function in polynomial time using Saul and Kardar’s technique,^{11,40} which is a generalization of methods used for finding the partition function at $T=0$ for the two-dimensional Ising model.^{48–51} For reasons of computational efficiency, we consider only $L \times L$ lattices where L is even. Because our methods yield ground states not only for lattices with regular periodic boundary conditions but also for those with antiperiodic boundary conditions, we generate four different partition functions for each possible lattice, corresponding to the four different boundary conditions mentioned above.

We are confident that our algorithm works properly because the number of ground states found by our algorithm agrees with the partition function result for every realization and boundary condition that we have examined.

III. GROUND-STATE DISTRIBUTION

Before presenting the results from our algorithm, we first discuss the distribution of numbers of ground states for dif-

ferent realizations with varying system size L and antiferromagnetic bond ratio x . All of these results were obtained from the partition functions of these realizations, generated using Saul and Kardar’s method.⁴⁰ These ground-state distributions show large sample-to-sample variations for realizations with the same L and x .

Figure 6 shows four histograms of the number of ground states for the four different boundary conditions of 3006 realizations of $L=10$ lattices with $x=0.5$. The solid lines on the figure are fits to a log-normal distribution,^{52–54} where the number of realizations with between m and $m+dm$ ground states is $G(m; \mu, \sigma)d(\log_{10}(m))$, with

$$G(m; \mu, \sigma) = \frac{A}{\sigma \sqrt{2\pi}} \exp\left[-\frac{[\log_{10}(m) - \mu]^2}{2\sigma^2}\right]. \quad (2)$$

Here, μ is the most probable value of $\log_{10}(m)$, σ describes the width of the distribution, and A is a normalization constant. Table I gives the parameter values from a χ^2 fit with the errors for the bin heights taken to be $\sqrt{N_b}$, where N_b is the number of realizations in a given bin. The distributions fit a log-normal distribution extremely well, and all the parameters of the fit for the four different boundary conditions are consistent with each other within error bars.

This log-normal distribution means that the variations in the number of ground states of different realizations are enormous. Sampling a few realizations will not give a meaningful picture of the behavior of the system. Averages over realizations need significant numbers to produce reasonable results, and still may not give sufficient information. Also, because the distribution of ground states is so wide, our methods to find ground states and apply perturbations to them have wildly varying performance on realizations with the same L and x . A few outliers with many ground states will completely dominate the computation time of all algorithms. A change in thinking is necessary—the concept of an average realization or number of ground states is not necessarily useful in considering the physics of this system.

Figure 6 also demonstrates that at $x=0.5$ the boundary condition does not affect the ground-state distribution. All future results in this section will be presented for the fully periodic boundary condition.

Next we study how the distribution of ground states varies with x . Figures 7 and 8 show how the parameters μ and σ characterizing the mean and width of the distribution change

TABLE I. Fits of the boundary condition distributions to a Gaussian with mean μ , standard deviation σ , and amplitude A . μ and σ are given in terms of z , the base-10 logarithm of the number of ground states. Boundary condition has little effect on the distribution on ground states.

Boundary condition	μ	σ	A
Periodic all	3.91 ± 0.066	1.191 ± 0.047	589 ± 32
Antiperiodic NS	3.94 ± 0.065	1.199 ± 0.044	588 ± 31
Antiperiodic EW	3.95 ± 0.056	1.189 ± 0.040	591 ± 28
Antiperiodic All	3.90 ± 0.057	1.201 ± 0.040	591 ± 28

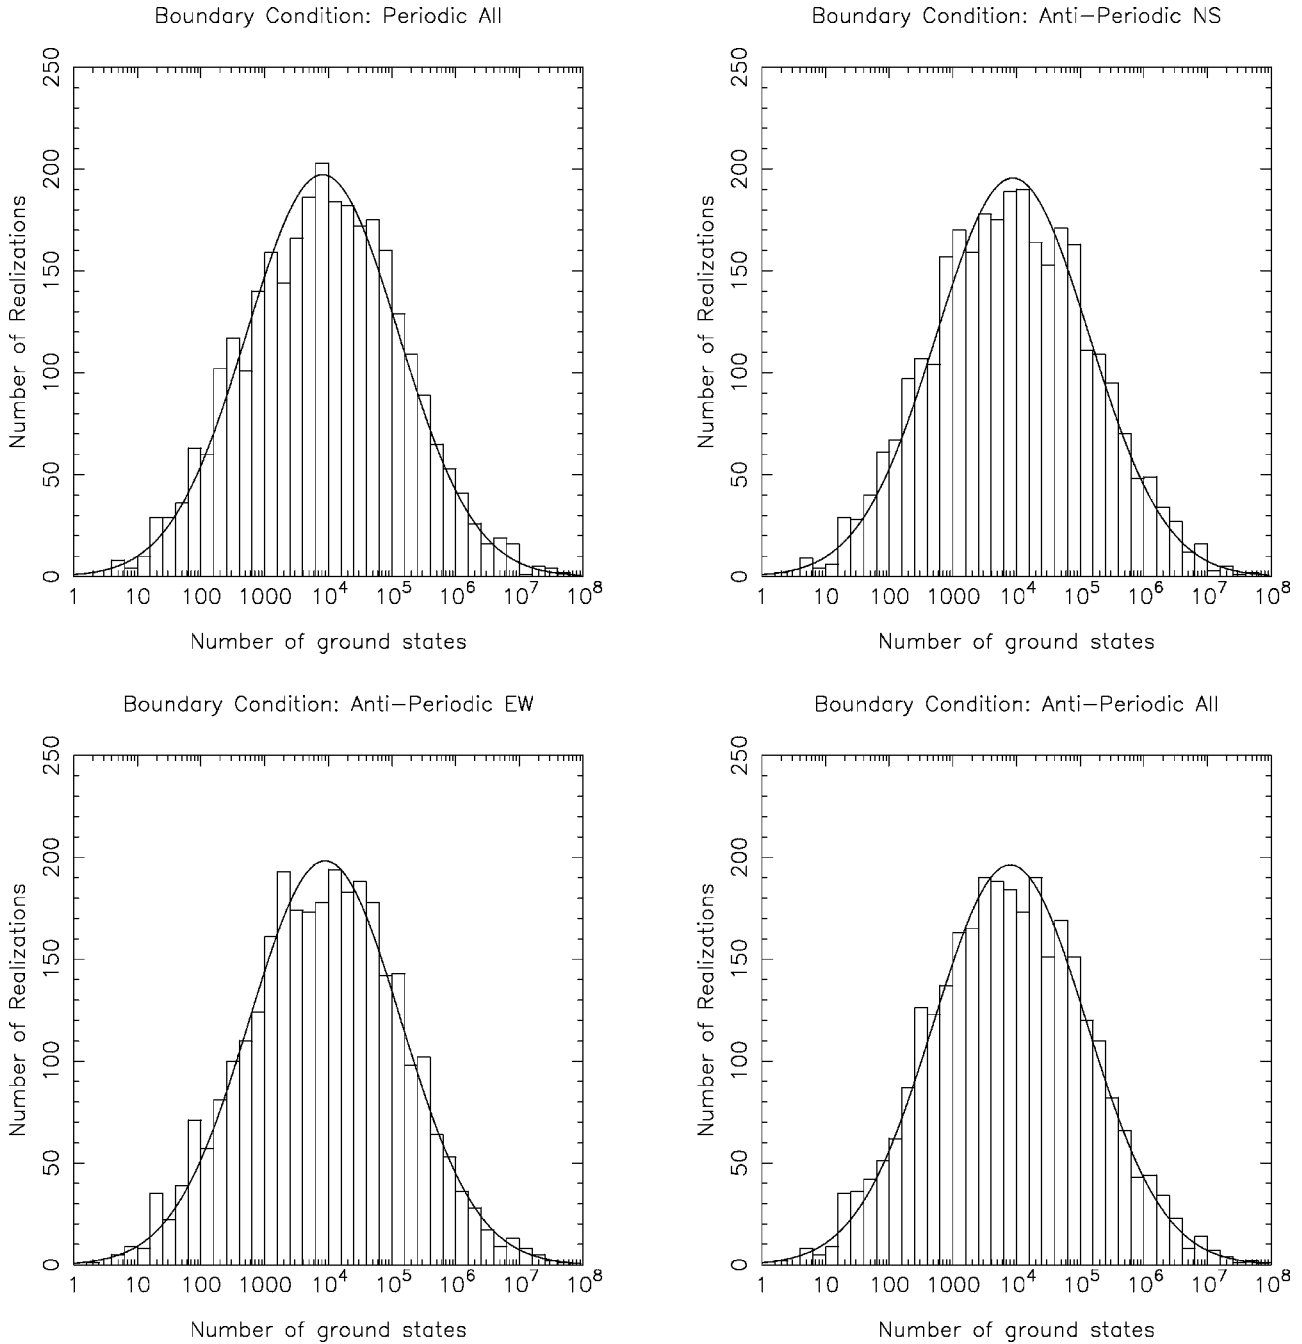


FIG. 6. Histograms of the number of ground states for 3006 realizations with $L=10$ and $x=0.5$ for the four different boundary conditions. The solid lines are χ^2 fits to the log-normal distribution, Eq. (2), with parameters are given in Table I. At this value of x , changing boundary conditions have very little effect on the distribution of ground states.

as x is varied between 0 and 0.5. Both the mean and width tend to increase with x until they saturate between $x=0.25$ and $x=0.3$.

The saturation of the ground-state distribution at $x=0.3$ appears to be completely distinct from the breakdown of ferromagnetic order at $x \approx 0.1$ and seems to be relatively insensitive to changes in system size. Since the distribution of ground states is essentially unchanged from $x=0.3$ to $x=0.5$, this suggests that systems in this range of parameters have no essential physical differences.

Finally, we present the variation of the distribution of

ground states with L at $x=0.5$. As Fig. 9 shows, increasing the system size moves the ground-state distribution over to larger numbers of ground states but does not change the log-normal distribution of the states. Again we fit these distributions to the form of Eq. (2). As seen in Fig. 10, the mean of the ground-state distributions at $x=0.5$ scales exponentially with lattice area L^2 , $\mu \propto a + bL^2$, with $b \approx 0.03$. It is because $b \ll \ln 2$ that our algorithm finds all ground states much more efficiently than an exhaustive search of all configurations.

We also investigate whether the distribution of ground states is self-averaging,⁵⁵ that is, whether the ratio of the

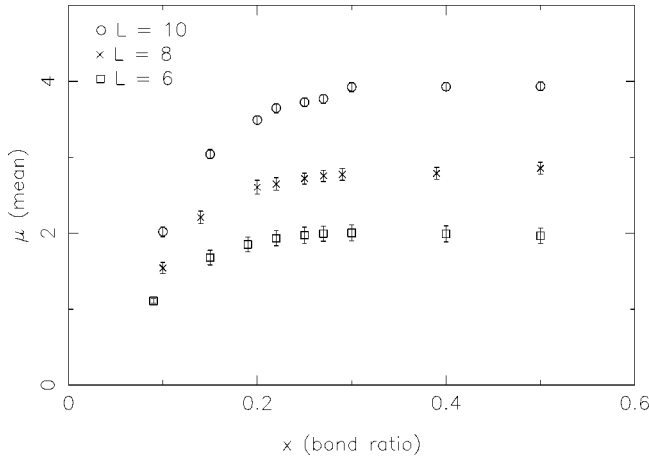


FIG. 7. Plot of the parameter μ defined in Eq. (2) (which is the base-10 logarithm of the maximum of the ground-state distribution) vs x , the fraction of antiferromagnetic bonds. Data for systems with $L=6,8,10$ are shown; the qualitative features do not exhibit a strong L dependence, except for the overall exponential scaling of the number of ground states with system size.

width of the distribution to its mean vanishes in the thermodynamic limit $L \rightarrow \infty$. Figure 11, a plot of σ/μ as a function of L , shows that σ/μ actually increases with L up to $L=8$ (a cautionary note for higher dimensions, where studies of $L > 8$ are extremely difficult), but that σ/μ decreases as L increases for $L > 8$, consistent with self-averaging behavior for the limit $L \rightarrow \infty$. Data for larger lattices would be extremely helpful in determining the behavior of σ/μ as the system size increases.

The wide distribution in numbers of ground states can be rationalized in terms of the matching problem by noting that random arrangements of bonds produce relatively random arrangements of frustrated plaquettes. If we treat the number of possible edge choices for different plaquettes as random

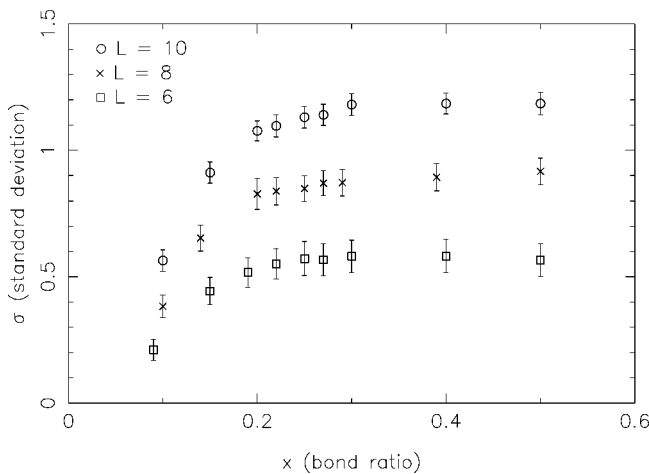


FIG. 8. Plot of the parameter σ defined in Eq. (2), which is a measure of the width of the log-normal distribution of ground states, vs x , the fraction of antiferromagnetic bonds. Data for systems with $L=6,8,10$ are shown; the qualitative features do not exhibit a strong L dependence, except for the overall exponential scaling of the number of ground states with system size.

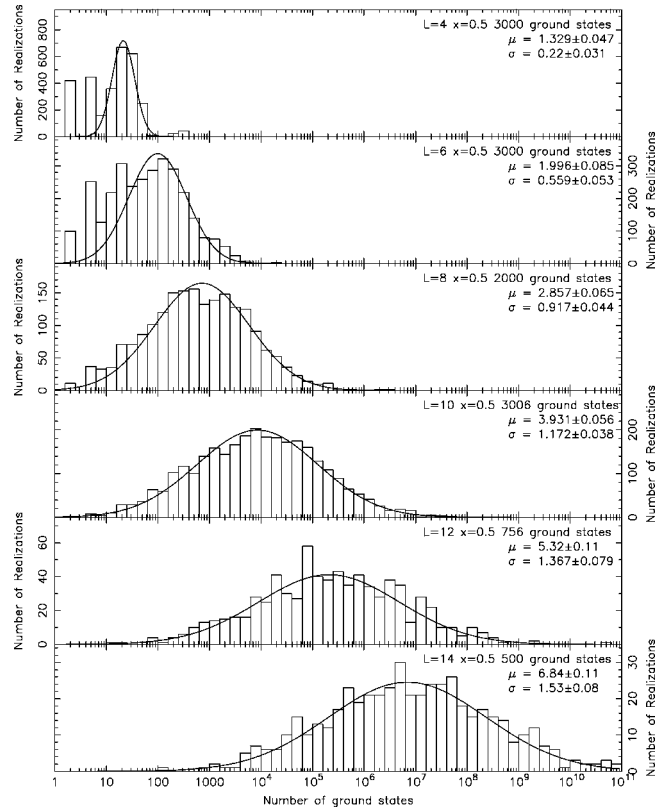


FIG. 9. Histograms of the number of ground states for realizations of different system sizes $L=4, 6, 8, 10, 12,$ and 14 at $x=0.5$. The solid lines are χ^2 fits to log-normal distributions. The parameters μ (which describes the mean) and σ (which describes the width) are defined in Eq. (2); they are given in terms of the base-10 logarithm of the number of ground states. Increasing system size increases both μ and σ .

variables and assume that they are essentially independent from plaquette to plaquette, then the number of total possibilities is multiplicative, and it is well known that multiplicative random processes lead to log-normal distributions.^{52–54} Moreover, since entropy is the logarithm of the number of accessible configurations, it is natural to expect $\log_{10}(m)$ to be normally distributed. This simple argument would also suggest that $\sigma \propto L^2$, but our system sizes are not large enough to address this question definitively.

IV. DESTRUCTION OF FERROMAGNETIC ORDER

In this section we investigate the destruction of ferromagnetic order that occurs as x , the fraction of antiferromagnetic bonds, is increased.^{13,42,56–61}

As Bendisch and collaborators discuss,^{59–61} investigating how the ground-state energy depends on boundary condition is a powerful method for locating the transition at which ferromagnetic order is destroyed. In a system of infinite size, when x is less than the transition point $x_c \approx 0.1$, one expects all the lowest-energy states to occur when the boundary conditions are consistent with ferromagnetic order, while for $x > x_c$, there should be no preference for this type of boundary condition. References 59–61 present calculations supporting

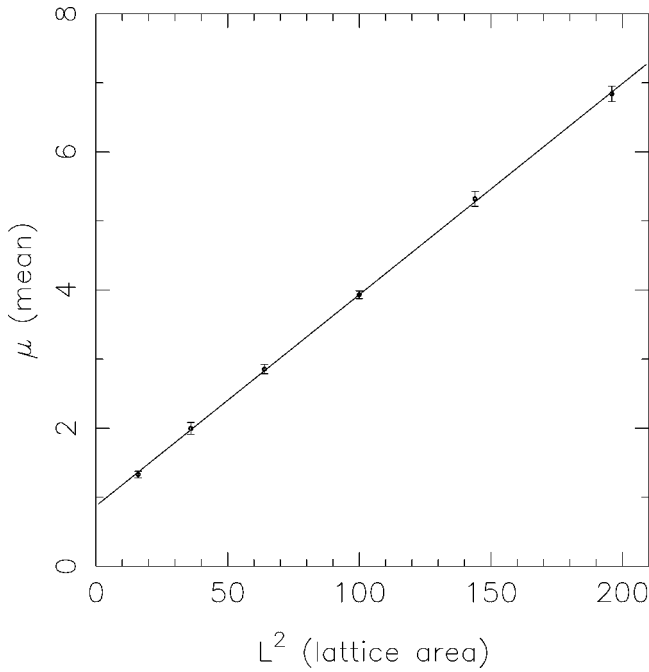


FIG. 10. The parameter μ of the ground-state distribution, defined in Eq. (2), vs system area L^2 for systems with $x=0.5$. The solid line on the graph is a fit to the form $\mu=(a+bL^2)$, with $a=0.874\pm 0.025$ and $b=0.03058\pm 0.00023$. The most probable number of ground states, 10^μ , increases exponentially with L^2 .

this picture for different nontoroidal boundary conditions. Our results obtained by computing the partition function for toroidal boundary conditions also support this picture. However, because of our relatively small system sizes, finite-size effects in our calculations are large. We did not do a finite-

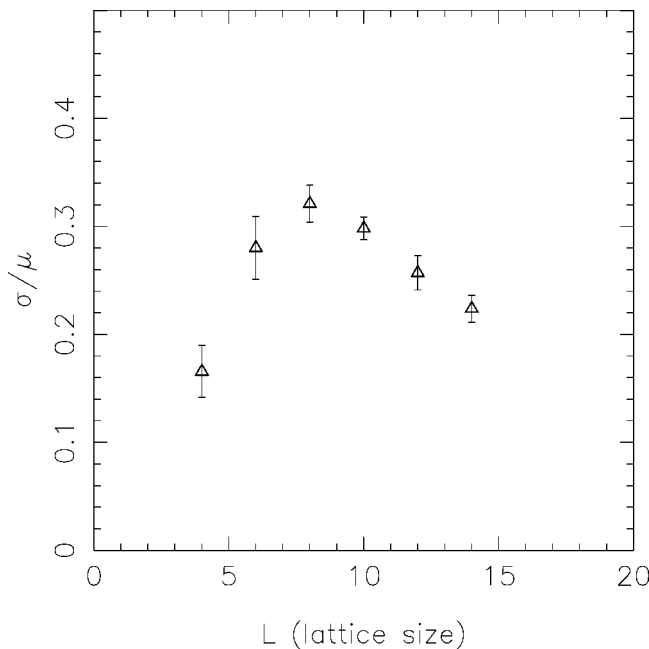


FIG. 11. Plot of the ratio σ/μ vs L . This is a test of self-averaging for the free energy. Note that σ/μ increases initially and then begins to fall off as L increases in a roughly linear way.

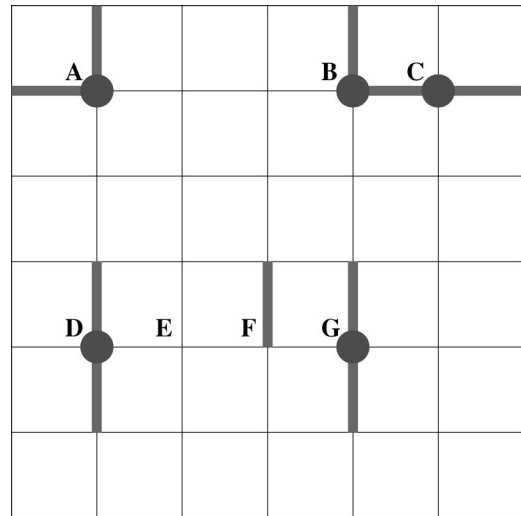


FIG. 12. Three bunches of flippable spins. The flippable spins are denoted by solid circles and the frustrated bonds by thick lines.

size scaling analysis of our data because we do not expect it to yield qualitatively new information about the destruction of ferromagnetic order.

Our algorithm enables us to investigate in detail how the ground states change when ferromagnetic order is destroyed. Barahona *et al.*⁴² present evidence that long-range order in the system is related to whether a set of spins with the same relative orientation in *all* the ground states spans the system. We can refine this picture further by noting that the set of ground states can naturally be subdivided into clusters,⁶² where a cluster is a group of ground states related by a sequence of single spin flips, each of which leaves the energy the same. By definition, all states in the same cluster can be reached from each other by single spin flips without raising the energy, whereas ground states in different clusters can only be reached from each other without raising the energy by making cooperative flips of multiple spins. A realization's ground states may all fall into a single cluster or populate many distinct clusters. For our systems, the number of clusters is moderate—we have observed up to 12 clusters in a single 10×10 realization. It is natural to ask whether the destruction of ferromagnetic order corresponds to growth of the number of spins contributing to the individual clusters or whether the relationship between the clusters plays a vital role.

To address this question, it is useful to focus on the spatial relations between the ground states in each cluster. To do this, we first define a flippable spin as one with an equal number of satisfied and unsatisfied bonds; all the states in a cluster are related by sequential flips of flippable spins. We now introduce the concept of a bunch and define it to be a group of flippable spins in a given cluster whose flippability does not depend on the state of other flippable spins in the system. Figure 12 shows three spin bunches.

The first bunch, spin A, is just a single isolated spin. The second bunch consists of spins B and C. Note that if spin B is flipped, then spin C is no longer flippable, and conversely.

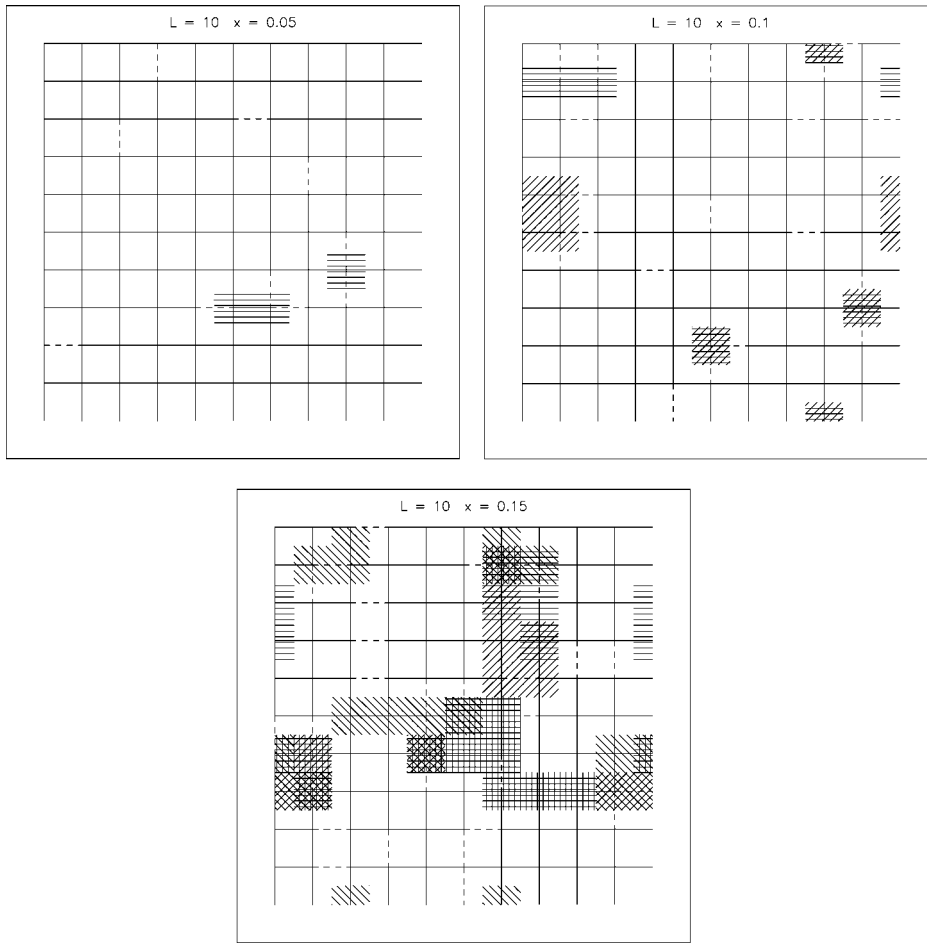


FIG. 13. Three different 10×10 realizations at $x=0.05$, $x=0.1 \approx x_c$, and $x=0.5$. All bunches in each cluster are shown with hatched marks, with different angles signifying different clusters. Multiple clusters allow bunches to combine to span the space and destroy ferromagnetic order.

The third bunch consists of spins D , E , F , and G . The bunch contains all four spins, because if G is flipped, then F becomes flippable. If D and F are both then flipped, then E becomes flippable.

Identification of bunches gives a complete picture of the clusters. Different clusters cannot have all the same bunches, though a given bunch can appear in more than one cluster. Bunches are useful because within a given cluster they are independent, so their contribution to the ground-state degeneracy is multiplicative.

We extract from the complete set of ground states all the bunches of a system using an algorithm described in Ref. 63. Figure 13 shows bunches from three realizations with $x=0.05 < x_c$, $x=0.1 \approx x_c$, and one with $x=0.15 > x_c$. One can see that the bunch structure for a single cluster does not change drastically as one crosses the transition, but when $x > x_c$ multiple clusters exist and the overlap of all the bunches from the different clusters spans the system. We believe that the key element governing the destruction of ferromagnetism is whether the overlap between the different bunches in different clusters is such that the union of all the bunches forms a path that percolates across the system. Thus, the “rigidity” transition discussed by Barahona *et al.*⁴² is governed by overlap of the bunches composing different clusters.

Figure 14 shows the single cluster fraction as a function of x for $L=6, 8$, and 10 . The number of realizations used to generate this figure is shown in Table II. The presence of

multiple clusters is strongly correlated with the destruction of ferromagnetism.

We can investigate further the bunch overlap as x is increased through the spin-glass transition. We do this by defining an overlap fraction for our realizations. To do this, we

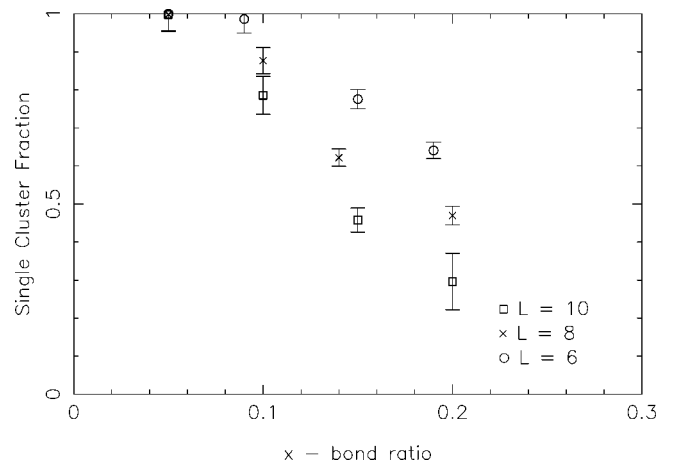


FIG. 14. Plot of the single cluster fraction vs x for $L=6, 8$, and 10 . Error bars are estimated as $\sqrt{N_r}$, where N_r is the number of realizations. At $x=0.05$, in the spin-glass phase, the single-cluster fraction is essentially 1. As x increases and passes through the ferromagnetic transition, the single-cluster fraction drops precipitously.

TABLE II. Runs used to generate Fig. 14. N_l is the number of distinct lattices, N_r is the number of ground-state realizations, and N_{scr} is the number of ground-state realizations with a single ground-state cluster.

L	x	N_l	N_r	N_{scr}
6	0.0556	500	515	515
6	0.0972	500	729	719
6	0.1528	500	1212	940
6	0.1944	500	1386	888
8	0.0547	500	507	506
8	0.1016	500	754	661
8	0.1484	500	1187	738
8	0.2031	281	775	364
10	0.05	500	504	503
10	0.1	200	317	249
10	0.15	199	452	207
10	0.2	18	54	16

first sum up the number of spins in bunches for each cluster in the realization, n_{sb} . Then, all bunches in all clusters of a given realization are overlaid, and the total number of spins covered by bunches is counted, o_{sb} . We define the overlap fraction o_f as

$$o_f = \frac{n_{sb} - o_{sb}}{n_{sb}}. \tag{3}$$

If no spins overlap between the bunches in different clusters, $o_{sb} = n_{sb}$ and $o_f = 0$. If all the spins between different bunches in N different clusters overlapped (an impossibility), then o_f would tend toward 1 as $N \rightarrow \infty$. The overlap fraction o_f is not defined for a realization with only a single cluster.

Figure 15 shows histograms of o_f for various x values for lattices with $L=8$. As the number of realizations with multiple clusters rises, the overlap fraction for these clusters also increases. Not only are there more clusters, but the variability of the bunches in these clusters also increases.

Thus, the appearance of multiple clusters for a single realization and the increase of overlaps between the bunches of these clusters both seem to be closely correlated to the destruction of ferromagnetic order. We believe these additional markers of the transition will aid in the understanding of this transition.

V. DISCUSSION

This paper investigates the ground states of the two-dimensional $\pm J$ Edwards-Anderson spin glass. An algorithm that finds systematically all the ground states of this model is

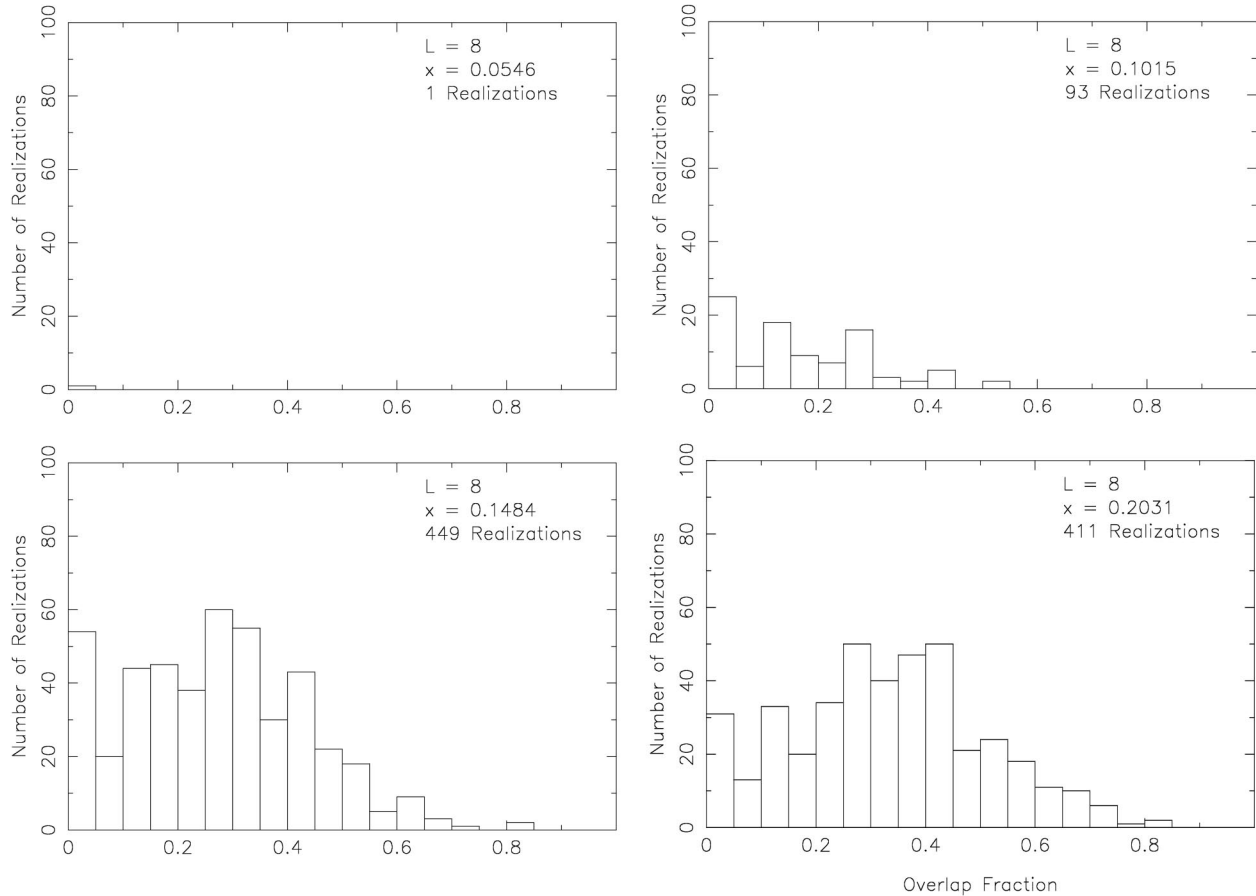


FIG. 15. Plots of the overlap fraction o_f for $L=8$ realizations. As the number of multiple-cluster realizations increases, the structure of the bunches also grows increasingly complicated as the overlap fraction increases.

presented. This algorithm is used to investigate the nature of the destruction of long-range ferromagnetic order as the fraction of antiferromagnetic bonds is increased.

The running time of our algorithm for a given realization is found empirically to scale roughly linearly in the number of ground states for the range that we study (up to 10^7 ground states). Memory is the most serious limitation, because to avoid duplication each new ground state must be compared to those already found. Our method enables the examination of significantly larger systems than would be accessible by exhaustive enumeration of all states—for example, a system with $L=10$ has $2^{100}-10^{30}$ configurations, but only 10^2-10^7 ground states. The advantage of our technique is even greater when we introduce quantum tunneling,⁶³ where the full Hamiltonian of a system with $L=10$ is a $2^{100}\times 2^{100}$ matrix.

We use our method to investigate the transition at which ferromagnetism is destroyed as the fraction x of antiferromagnetic bonds is increased. By comparing the “rigidity” of bonds between different ground-state configurations, we show how the appearance of ground-state clusters and the overlap of bunches is intimately connected to the percolative nature of the ferromagnetic-paramagnetic phase transition.

In general, we find that many quantities exhibit large, non-Gaussian variability between realizations. For example, the number of ground states is well described by a log-normal distribution, with the ratio of the width to the mean of the distribution increasing with system size L up to the relatively large value $L\sim 8$. This relatively large length scale is unsettling when one keeps in mind the relatively small sizes that can be investigated numerically in higher dimensions.^{5,64-67}

The availability of complete information about all the ground states provides a new avenue for probing the nature of the system at low temperatures, and our work can be extended in many directions. We have investigated how the introduction of quantum tunneling and of coupling to a deformable lattice affect the low-energy landscape of the model,⁶³ and the effects of other physical perturbations should be examined. Our work also provides a means for stringent validation of other sampling methods that can be used to study larger systems. We also note that in the course of this work we have shown that one can adapt the standard matching algorithm to study systems with fully toroidal periodic and/or antiperiodic boundary conditions. It would be interesting to compare the effects of changing boundary conditions of this type to those in Ref. 24.

ACKNOWLEDGMENTS

We gratefully acknowledge L. Saul for providing us with the code of his program for finding the partition function of a two-dimensional spin glass and J. Cook and A. Rohe for putting the Concorde version of the blossom algorithm in the public domain. We also wish to thank Scot Shaw for putting his ARPACK++ routines in the public domain. We thank J. Brooke, B. DiDonna, S.R. Nagel, T.F. Rosenbaum, and S. Venkataramani for enlightening conversations and J. Bogan for computer assistance. S.N.C. thanks the Aspen Center for

Physics for hospitality during the preparation of this manuscript. This work was supported by MRSEC Program of the National Science Foundation under Award No. DMR-9808595.

APPENDIX A: DETAILED DESCRIPTION OF THE GROUND-STATE ALGORITHM

1. Goals of the algorithm

These notes are a short description of our algorithm that generates all of the ground states of an Edwards-Anderson 2D spin glass. This algorithm was designed with several goals in mind.

The algorithm should be relatively simple, so that it is easy to understand and implement.

The algorithm should proceed in a completely deterministic and systematic way, so there is no way to get off track or lose information.

The algorithm should be exhaustive. Given infinite time and memory, it should be able to find all the ground states of an arbitrary lattice.

The algorithm should use as much information as possible. It should use the fact that selecting one edge automatically eliminates many other edges, simply because the two plaquettes joined by the edge cannot be linked by any other edges. It should also use the fact that the number of edges in at least one ground state is a small subset of the total number of edges.

2. Short description

The algorithm has two parts. The first is preliminary and the second uses the results of the first to generate all the ground states.

The first part finds all the viable edges: edges that exist in matching solutions of the correct energy. This list of viable edges is then the only thing that needs to be studied, since the other edges do not exist in matchings with the correct ground-state energy.

The second part takes the list of viable edges and systematically goes through the possible combinations (those that connect each plaquette to one and only one other plaquette) and tests them to see if they correspond to ground states. Once all the combinations are exhausted, all the ground states will have been found.

3. Detailed outline of the algorithm

- (1) Given L and x , generate realization with randomly placed bonds.
- (2) Find the partition function of this realization:
 - (a) gives number of ground states,
 - (b) gives energy of ground states.
- (3) Convert this lattice to a matching problem in graph theory:
 - (a) frustrated plaquettes \rightarrow nodes,
 - (b) path of unsatisfied bonds between frustrated plaquettes \rightarrow edges.
- (4) Construct the list of viable edges (edges that appear in at least one ground-state matching):

(a) run blossom algorithm to find a matching solution of the correct energy.

(b) add the edges in this matching solution to the list of viable edges.

(c) build up the list of viable edges, one plaquette at a time, until all the edges corresponding to ground states are found.

(i) remove all edges in the list for the current plaquette from the problem.

(ii) run blossom algorithm to get a new matching solution.

(iii) if the new matching has the correct energy, add the new edges to the list of viable edges, return to (i) and do it again.

(iv) if it does not, there are no new edges for this plaquette. Go to the next plaquette.

(5) Take the list of viable edges and pair up the plaquettes until all ground states are found.

(a) sequentially go through the various combinations of edges until every ground state is found.

(i) pick series of edges until each plaquette is matched up with one and only one other plaquette.

(ii) if matching solution does not have the correct energy, go to (i).

(iii) convert this matching solution to a set of spins and bonds for each boundary condition. If this corresponds to a ground state, enter it for that boundary condition only.

(iv) return to (i) if some edge combinations remain.

(b) number of possible combinations is manageable, because each pick of an edge removes two plaquettes from the list that need to be matched.

(6) When a matching solution is found, it does not necessarily correspond to a ground state, so all ground states are kept in memory.

(a) one matching can represent several ground states. (A simple example is a plaquette diagonally separated by one from another plaquette joined by an edge. This represents two ground states.)

(b) multiple matchings can represent one ground state. (Take the example of four plaquettes in a + formation linked together. Many possible matchings represent the same ground state.)

(c) all ground states are kept in memory, and when a new matching solution is found, this is compared to ground states already found in the system. If it is new, it is added to the list.

(7) Once the algorithm is done, the number of ground states found is compared to the number found in step (2) as a check.

a. Finding all the viable edges

This part of the algorithm uses removal of edges from the matching problem (cutting) to find all the viable edges. It works by repeatedly cutting all the edges involving a specific plaquette until all matchings generated are higher energy. At that point, all viable matchings contain one of the edges already cut.

As an example, suppose we have a 4×4 lattice with 12 plaquettes and 66 edges. The arrangement of bonds and frustrated plaquettes is shown in Fig. 16.

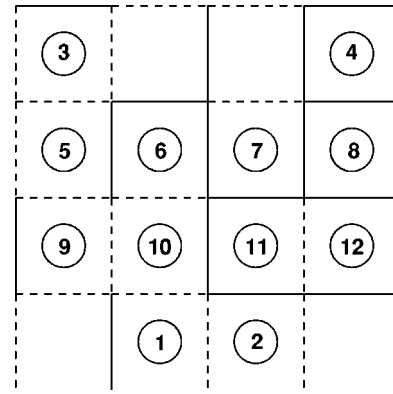


FIG. 16. 4×4 lattice for the algorithm example. Ferromagnetic bonds are dark lines and antiferromagnetic bonds are dashed lines. Frustrated plaquettes are represented by open circles and labeled by number.

We run the blossom program without cutting any edges and get a ground-state matching with edges (9,20,23,34,46,63). Edge 9 links plaquettes 1 and 10, so we put edge 9 in our list under both plaquette 1 and plaquette 10. We would do the same for the five other edges in the matching.

We want to find all the edges that appear in ground-state matchings that connect to plaquette 1. We thus cut edge 9 and feed the problem back into the blossom algorithm. We get the ground-state matching (1,23,34,49,55,63). Edge 1 links plaquettes 1 and 2, so we put edge 1 in our list under both plaquette 1 and plaquette 2. We do the same for the other five edges.

We now cut both edge 9 and edge 1 and feed the result into the blossom algorithm. We get out the result (8,20,23,34,49,56), a matching with an energy larger than the ground-state energy. This means that every ground-state matching contains either edge 1 or edge 9. We can then do the same for plaquette 2, 3, etc., until we have all the edges that appear in ground-state matchings, what we call viable edges.

b. Generating all the ground states

Now that we have the list of viable edges, we restrict our attention to it. All ground states can be found by combining these edges. We combine these edges such every plaquette has one and only one other plaquette connected to it by an edge.

As an example, consider the 4×4 example from above.

The set of possible edges, grouped by plaquette, looks like this

- plaquette 1: (1,9)
- plaquette 2: (1,20)
- plaquette 3: (22,23)
- plaquette 4: (22,34)
- plaquette 5: (23,39,41,42)
- plaquette 6: (39,46,49)
- plaquette 7: (46,52,55)
- plaquette 8: (34,41,52,60)
- plaquette 9: (42,61,63)

plaquette 10: (9,49,61,64)
 plaquette 11: (20,55,64,66)
 plaquette 12: (60,63,66)

We want to generate a matching with six edges. Starting off is easy. We merely pick the first one in plaquette 1's list, edge 1. Edge 1 involves plaquettes 1 and 2, so we skip to plaquette 3's list and select edge 22. That involves edges 3 and 4, so the next highest plaquette is 5. We select edge 23, but that involves plaquette 3 and is forbidden, so we skip to the next one, which is 39. This continues until all the plaquettes are accounted for and we get the matching (1,22,39,52,61,66). We would then convert this matching into possible ground states for each boundary condition. Those potential ground states that have consistent spin configurations and energies are entered into the ground-state lists.

Now we go to the next state. We remove the last two edges from our current state, giving us (1,22,39,52). The second to last plaquette list we looked at was 9, so we choose the edge after edge 61. We get 63, which involves plaquettes 9 and 12. We then skip to plaquette 10 and choose edge 9. That involves plaquette 1, which is already linked. Edge 49 involves plaquette 6 which is taken, and edge 61 involves plaquette 9 which is already taken. Only edge 64 involves the empty plaquettes 10 and 11, so we get the new state (1,22,39,52,63,64).

We continue in this manner to get all of the ground states. If we reach the end of a list for a specific plaquette, we know to erase that entry in the current state and move to the previous plaquette list. If we reach the end of the list for plaquette 1, we know we have completed the algorithm and found all the ground states.

4. Notes on the algorithm

The algorithm currently is memory limited. The blossom algorithm and selection of edges for each matching in part 5 are both very fast. Eventually, as the number of ground states grows with the size of the system and the value of x , the available memory of the computer is exhausted.

The reason every ground state is stored in memory now is that in certain circumstances, distinct matching solutions can correspond to the same ground state. If these circumstances were determined and recorded, then only a small subset of relevant ground states would have to be kept, while the others are written to a file. Whenever a matching solution with a "dangerous" edge or edges came up, it could be compared to the small subset of comparable matching solutions or ground states. It is not clear if this approach is really feasible or not.

Finally, certain combinations of edges that automatically lead to higher-energy matching solutions constantly recur. If a list of these could be quickly generated, it should greatly decrease the number of combinations of viable edges that must be explored.

APPENDIX B: GROUND-STATE ALGORITHM PERFORMANCE

In this section we present the performance of our ground-state algorithm and discuss the steps limiting its perfor-

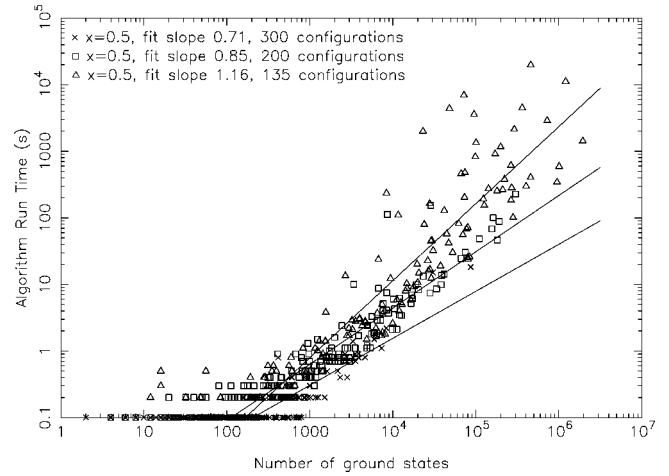


FIG. 17. Algorithm run time vs the number of ground states. The fits are power law fits to each of the separate data sets with more than 100 ground states.

mance. One typically investigates the performance of an algorithm on a finite lattice by showing its scaling behavior relative to system size.⁴⁰ In this instance, such an approach does not provide much understanding because of the log-normal distribution of ground states for a given system size and x . Any average over such a broad distribution would not convey much information.

We instead study our algorithm's performance relative to the number of ground states, which we believe is a much better measure of performance. In addition, because the number of ground states of a realization can be generated quickly from the partition function,⁴⁰ this measure gives a useful predictor of time to completion before the algorithm is run.

For all of the following results, the algorithm was run on a Pentium III 800 MHz machine, with 512 MB of RAM.⁶⁹

Figure 17 is a plot of the algorithm's run time versus number of ground states. The plot also shows power law fits of the results for different system sizes. Though the scatter is substantial, it can be seen from the plot that the run times remain reasonably well behaved as the number of ground states increases. Assuming enough memory is available, we expect that systems with 10^8 ground states could be completed in a week of running time.

Figure 18 demonstrates a more useful way to characterize the scaling behavior of our algorithm. Here we plot the run time versus the number of matchings (sets of edge combinations) explored. Results from different system sizes agree and scale as $\approx O(n_m^{4/5})$, where n_m is the number of matchings. This measure of system performance is not a very good predictor, because it is difficult to determine the number of matchings before running the first part of the algorithm. In the fits for both Figs. 17 and 18, runs with very small numbers of matches and ground states are not included, because there is an initial overhead in computation irrespective of the scale of the problem.

Figures 17 and 18 show data for different lattice sizes L with the same $x=0.5$. Figure 19 demonstrates that the same scaling of algorithm run time with ground states is observed

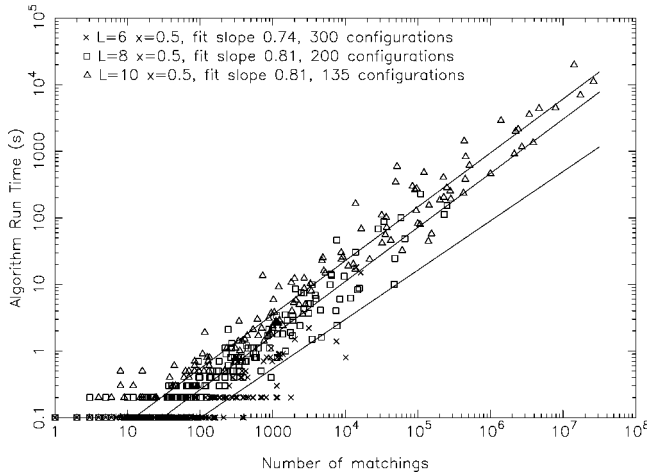


FIG. 18. Algorithm run time vs the number of matchings. The power law fits yield exponents that agree and give a scaling of $O(n^{4/5})$. Lattices with less than 100 matchings were excluded from the fit because of start-up time costs.

if we vary x and keep L fixed.

The algorithm scales as a power law with number of matchings, because the core of the algorithm (its second part) runs sequentially through all possible matchings of edges in ground states. We thus expect it to scale as $O(n_m)$, where n_m is the number of matchings.

The algorithm's scaling with ground states is more complicated, precisely because the algorithm deals with matchings instead of ground states, except when inserting into the list of ground states, which should go like $O(n_{gsbc}^{3/2})$, where n_{gsbc} is the number of ground states for a given boundary condition. In practice, the number of ground states is smaller than the number of matchings by about an order of magnitude. Also, since we observe that ground states are split relatively evenly among the boundary conditions in realizations with large numbers of ground states, we expect n_{gsbc} to be about one and a half orders of magnitude smaller than n_m .

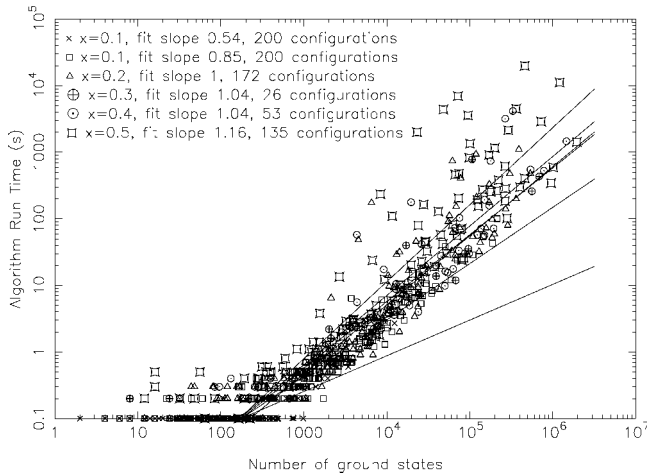


FIG. 19. Algorithm run time as a function of the number of ground states for $L=10$ realizations with various x values. The power law fits are to each of the separate data sets.

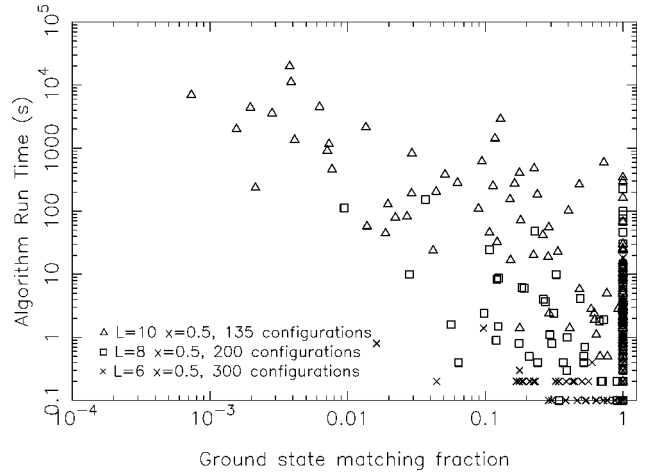


FIG. 20. Algorithm run time vs the fraction of all matchings constructed from viable edges that are ground-state matchings. Small ratios imply long run times.

Since significantly more work is done only on the ground-state matchings, which require an additional investment of searches, the interplay between the time for matchings and for the ground-state conversions is nontrivial.

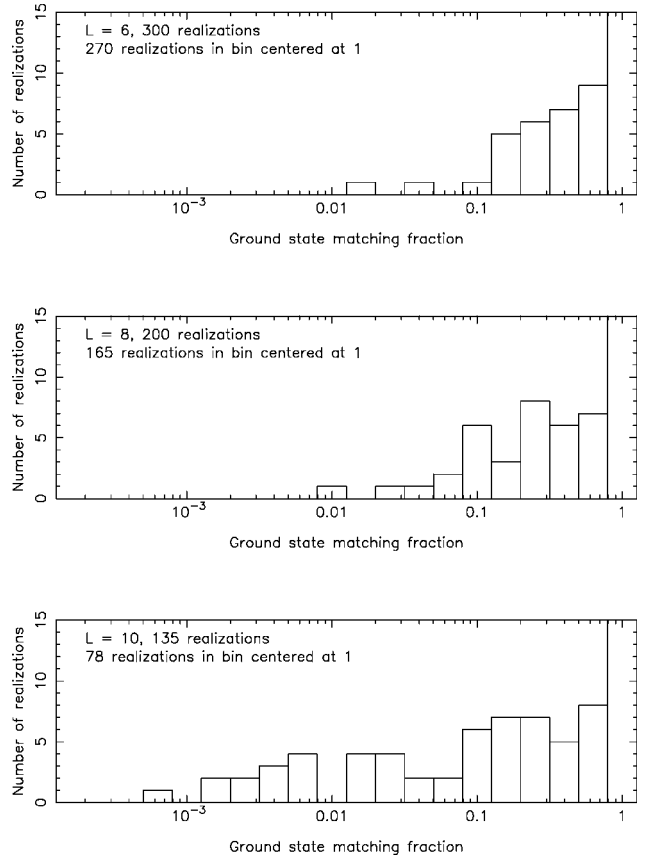


FIG. 21. Histograms of the ground-state matching ratios for realizations with $L=6, 8,$ and 10 . In all cases, at least half of the realizations had ground-state matching fractions equal to 1. Note the decreasing ground-state matching fractions as the system size increases.

The algorithm performs well only if the ground-state matchings are a significant fraction of all possible combinations of viable edges. If two distinct systems have the same number of ground states, but one has 10 times as many matchings as the other, their run time will be very different. We denote those matchings that lead to a ground state as ground-state matchings, and the ratio of ground-state matchings to the number of total matchings of viable edges is inversely correlated with the run time of the algorithm. This is shown in Fig. 20, a plot of run time versus ground-state matching ratio. Realizations with small ground-state matching ratios have large run times, because the algorithm spends most of its time cycling through matchings that do not yield ground states.

To illustrate the behavior of the ground-state matchings, we present in Fig. 21 a histogram of the ground-state matching ratios for different lattice sizes L . Realizations with small ground-state matching ratios occur rarely, but increase in fre-

quency with increasing system size. However, even for $L = 10$, over 78% of the configurations have matching ratios greater than 0.1. We thus expect this ratio to remain appreciable for somewhat larger system sizes.

The above plots do not show ground-state numbers higher than 2×10^6 , because of memory limitations of our hardware. The code currently stores all of the ground states while running to prevent duplication of ground states. With more memory, much larger numbers of ground states could be investigated relatively quickly with this algorithm, as seen in the scaling above. In addition, if the algorithm could be modified so that it did not require all the ground states to be stored, this memory constraint would be removed. This would require some cleverness about how matchings are converted into ground states. One could conceive of keeping only a list of “problematic” ground states that correspond to multiple matchings, if they could be determined easily.

-
- *Present address: Sandia National Laboratory, Albuquerque, NM 87185. Electronic address: jwlandry@uchicago.edu
- [†]Present address: Department of Physics, University of Wisconsin, 1150 University Avenue, Madison, WI 53706.
- ¹S. Edwards and P. Anderson, *J. Phys. F: Met. Phys.* **5**, 965 (1975).
- ²A. Hartmann, *Phys. Rev. E* **59**, 84 (1999).
- ³A. K. Hartmann, *Europhys. Lett.* **45**, 746 (1999).
- ⁴M. Palassini and A. Young, *Phys. Rev. Lett.* **83**, 5126 (1999).
- ⁵F. Krzakala and O. Martin, *Phys. Rev. Lett.* **85**, 3013 (2000).
- ⁶F. Krzakala and O. Martin, *Europhys. Lett.* **53**, 749 (2001).
- ⁷E. Marinari, G. Parisi, F. Ricci-Tersenghi, J. Ruiz-Lorenzo, and F. Zuliani, *J. Stat. Phys.* **98**, 973 (2000).
- ⁸E. Marinari and G. Parisi, *Phys. Rev. B* **62**, 11 677 (2000).
- ⁹J. Lamarcq, J.-P. Bouchaud, O. Martin, and M. Mèzard, cond-mat/0107544 (unpublished).
- ¹⁰M. Palassini and A. Young, *J. Phys. Soc. Jpn.* **69**, 165 (2000).
- ¹¹L. Saul and M. Kardar, *Phys. Rev. E* **48**, R3221 (1993).
- ¹²C. de Simone, M. Diehl, M. Junger, P. Mutzel, G. Reinelt, and G. Rinaldi, *J. Stat. Phys.* **84**, 1363 (1996).
- ¹³N. Kawashima and H. Rieger, *Europhys. Lett.* **39**, 85 (1997).
- ¹⁴T. Shirakura and F. Matsubara, *Phys. Rev. Lett.* **79**, 2887 (1997).
- ¹⁵F. Matsubara, T. Shirakura, and M. Shiomi, *Phys. Rev. B* **58**, R11 821 (1998).
- ¹⁶A. K. Hartmann, *Eur. Phys. J. B* **8**, 619 (1999).
- ¹⁷M. Shiomi, F. Matsubara, and T. Shirakura, *J. Phys. Soc. Jpn.* **69**, 2798 (2000).
- ¹⁸A. Middleton, *Phys. Rev. Lett.* **83**, 1672 (1999).
- ¹⁹A. Middleton, *Phys. Rev. B* **63**, 060202(R) (2001).
- ²⁰M. Palassini and A. Young, *Phys. Rev. B* **60**, R9919 (1999).
- ²¹J. Houdayer, *Eur. Phys. J. B* **22**, 479 (2001).
- ²²C. Newman and D. Stein, *Phys. Rev. Lett.* **84**, 3966 (2000).
- ²³A. Carter, A. Bray, and M. Moore, cond-mat/0108050 (unpublished).
- ²⁴A. Hartmann and A. Young, *Phys. Rev. B* **64**, 180404 (2001).
- ²⁵I. A. Gruzberg, N. Read, and A. W. W. Ludwig, *Phys. Rev. B* **63**, 104422 (2001).
- ²⁶N. Read and A. W. W. Ludwig, *Phys. Rev. B* **63**, 024404 (2001).
- ²⁷J. T. Chalker, N. Read, V. Kagalovsky, B. Horovitz, Y. Avishai, and A. W. W. Ludwig, *Phys. Rev. B* **65**, 012506 (2001).
- ²⁸F. Merz and J. T. Chalker, cond-mat/0106023 (unpublished).
- ²⁹M. Palassini and A. Young, *Phys. Rev. B* **63**, 140408(R) (2001).
- ³⁰L. Bieche, J. P. Uhry, R. Maynard, and R. Rammal, *J. Phys. A* **13**, 2553 (1980).
- ³¹K. Binder and A. Young, *Rev. Mod. Phys.* **58**, 801 (1986).
- ³²N. Persky, I. Kanter, and S. Solomon, *Phys. Rev. E* **53**, 1212 (1996).
- ³³W. Janke, *Physica A* **254**, 164 (1998).
- ³⁴N. Hatano and J. E. Gubarnatis, *Suppl. Prog. Theor. Phys.* **138**, 442 (2000).
- ³⁵K. Bhattacharya and J. Sethna, *Phys. Rev. E* **57**, 2553 (1998).
- ³⁶H. Rieger, in *Lecture Notes in Physics Vol. 501* (Springer-Verlag, Heidelberg, 1998).
- ³⁷A. K. Hartmann, *J. Phys. A* **33**, 657 (2000).
- ³⁸A. Sandvik, *Europhys. Lett.* **45**, 745 (1999).
- ³⁹A. K. Hartmann, *Eur. Phys. J. B* **13**, 539 (2000).
- ⁴⁰L. Saul and M. Kardar, *Nucl. Phys. B* **432**, 641 (1994).
- ⁴¹W. J. Cook and A. Rohe, *INFORMS J. Comput.* **11**, 138 (1999).
- ⁴²F. Barahona, R. Maynard, R. Rammal, and J. P. Uhry, *J. Phys. A* **15**, 673 (1982).
- ⁴³G. Toulouse, *Commun. Phys. (London)* **2**, 115 (1977).
- ⁴⁴I. Bieche, Ph.D. thesis, University of Grenoble, 1979.
- ⁴⁵J. Edmonds, *Can. J. Math.* **17**, 449 (1965).
- ⁴⁶E. Lawler, *Combinatorial Optimization: Networks and Matroids* (Holt, Rinehart, and Winston, New York, 1976).
- ⁴⁷D. Applegate, R. Bixby, V. Chvataland, and W. Cook, computer code CONCORDE, (See Ref. 41), <http://www.math.princeton.edu/tsp/concorde.html>
- ⁴⁸M. Kac and J. Ward, *Phys. Rev.* **88**, 1332 (1952).
- ⁴⁹R. Potts and J. Ward, *Prog. Theor. Phys.* **13**, 38 (1955).
- ⁵⁰P. Burgoyne, *J. Math. Phys.* **4**, 1320 (1963).
- ⁵¹M. Glasser, *Am. J. Phys.* **38**, 1033 (1970).
- ⁵²J. Aitchison and J. Brown, *The Log-normal Distribution* (Cambridge University Press, Cambridge, UK, 1957).
- ⁵³E. Crow and K. Shimizu, *Log-normal Distributions: Theory and Application* (Dekker, New York, 1988).
- ⁵⁴E. Limpert, W. Stahel, and M. Abbt, *BioScience* **51**, 341 (2001).
- ⁵⁵B. Derrida, *Physica D* **107**, 186 (1997).

- ⁵⁶S. Kirkpatrick, Phys. Rev. B **16**, 4630 (1977).
- ⁵⁷G. Grinstein, C. Jayaprakash, and M. Wortis, Phys. Rev. B **19**, 260 (1979).
- ⁵⁸J. Vannimenus, J. Maillard, and L. de Séze, J. Phys. C **12**, 4523 (1979).
- ⁵⁹J. Bendisch, Physica A **245**, 560 (1997).
- ⁶⁰J. Bendisch and H. Trotha, Physica A **253**, 428 (1998).
- ⁶¹M. Achilles, J. Bendisch, and H. Trotha, Physica A **275**, 178 (2000).
- ⁶²A. K. Hartmann, Phys. Rev. E **63**, 016106 (2001).
- ⁶³J. W. Landry and S. N. Coppersmith (unpublished).
- ⁶⁴A. K. Hartmann, Phys. Rev. E **60**, 5135 (1999).
- ⁶⁵J. Houdayer, F. Krzakala, and O. Martin, Eur. Phys. J. B **18**, 467 (2000).
- ⁶⁶H. Katzgraber, M. Palassini, and A. Young, Phys. Rev. B **63**, 184422 (2001).
- ⁶⁷E. Domany, G. Hed, M. Palassini, and A. Young, cond-mat/0104264 (unpublished).
- ⁶⁸J. Bogan (unpublished), <http://honeycomb.uchicago.edu>
- ⁶⁹The machine was a part of a Beowulf Linux cluster [details available online (Ref. 68)]. Since the code was not parallelized, each run was only on a single drone.