

Low-energy excitations in spin glasses from exact ground states

Matteo Palassini

University of California, 3333 California Street, Suite 415, San Francisco, California 94118, USA

Frauke Liers and Michael Juenger

Institut für Informatik, Universität zu Köln, D50969 Cologne, Germany

A. P. Young

Physics Department, University of California, Santa Cruz, California 95064, USA

(Received 22 December 2002; revised manuscript received 1 May 2003; published 13 August 2003)

We investigate the nature of the low-energy, large-scale excitations in the three-dimensional Edwards-Anderson Ising spin glass with Gaussian couplings and free boundary conditions, by studying the response of the ground state to a coupling-dependent perturbation introduced previously. The ground states are determined *exactly* for system sizes up to 12^3 spins using a branch-and-cut algorithm. The data are consistent with a picture where the surface of the excitations is not space filling, such as the droplet or the trivial-non-trivial (TNT) picture, with only minimal corrections to scaling. When allowing for very large corrections to scaling, the data are also consistent with a picture with space-filling surfaces, such as replica symmetry breaking. The energy of the excitations scales with their size with a small exponent θ' , which is compatible with zero if we allow moderate corrections to scaling. We compare the results with data for periodic boundary conditions obtained with a genetic algorithm, and discuss the effects of different boundary conditions on corrections to scaling. Finally, we analyze the performance of our branch-and-cut algorithm, finding that it is correlated with the existence of large-scale, low-energy excitations.

DOI: 10.1103/PhysRevB.68.064413

PACS number(s): 75.50.Lk, 75.40.Mg, 05.50.+q

I. INTRODUCTION

There is still considerable debate about the nature of the spin-glass state in finite-dimensional spin glasses. Two principal theories have been investigated: the “droplet theory” proposed by Fisher and Huse¹ (see also Refs. 2 and 3), and the replica symmetry-breaking picture of Parisi.^{4–6} In the droplet theory, the lowest-energy excitation of length scale l (a “droplet”) has energy of the order of l^θ , where θ is a positive exponent. Furthermore, the droplets have a surface with fractal dimension d_s less than the space dimension d .

Replica symmetry breaking (RSB) is well established in mean-field theory, but it remains to be proven in finite dimensions. The precise nature of RSB in finite dimensions is not uniquely defined but it is generally agreed that a key feature of RSB is the existence of excitations whose energy, unlike that of droplets, remains of the order of unity even as their size tends to infinity. Furthermore, upon the creation of such a large-scale, finite-energy excitation, a finite fraction of the bonds changes state (from satisfied to unsatisfied, or vice versa) or, equivalently, the surface of these excitations is space filling, i.e., $d_s = d$.

Recently, Krzakala and Martin⁷ (KM), and two of us⁸ [Palassini and Young (PY)], have argued, on the basis of numerical calculations at zero temperature, in favor of an intermediate scenario where there are large-scale excitations whose energy does not increase with size, as in RSB, but which have a surface with $d_s < d$. Following KM, we shall denote this the trivial-non-trivial (TNT) scenario. In the TNT scenario, it is necessary to introduce *two* exponents that describe the growth of the energy of an excitation of scale L : (i) $\theta (> 0)$ such that L^θ is the typical change in energy when the

boundary conditions are changed, for example, from periodic to antiperiodic, in a system of size L , and (ii) θ' , which characterizes the energy of clusters excited within the system for a *fixed* set of boundary conditions (θ' was called θ_g in Ref. 7). The TNT picture has been challenged (although in opposite senses) by Marinari and Parisi⁹ and by Middleton.¹⁰ Subsequently, low-temperature Monte Carlo simulations¹¹ have found results consistent with the TNT scenario. The RSB, droplet, TNT and some other scenarios have been also studied by Newman and Stein.^{12,13} For some recent related work, see Refs. 14 and 15.

The work of KM and PY determined the ground state with and without a certain perturbation (which was different in the two cases), designed so that the ground state of the perturbed system is a large-scale excitation of the original system. They used *heuristic* algorithms, i.e., algorithms which are not *guaranteed* to give the exact ground state, although both KM and PY argue that they do find the exact ground state in most cases.

In this paper, we reconsider the problem of determining θ' and d_s , following the perturbation approach of PY, described in Sec. II, but we apply an *exact* algorithm, known as “branch and cut,”¹⁶ so we are guaranteed that the true ground state is reached every time. Exact optimization algorithms have been used before for spin glasses, see, e.g., Refs. 17–19, but, to our knowledge, their use in three dimensions has been restricted to smaller sizes than studied here, and they were not used to investigate the real-space structure of the low-energy excitations.

Our implementation of the branch-and-cut technique can handle significantly larger sizes for free boundary conditions (BC) than for periodic BC,²⁰ so we use free BC here. We

consider a different (and enlarged) set of observables than PY, in the attempt to gain a fuller understanding of what picture fits better the whole set of observables. We also perform a similar analysis of the data of PY, who used periodic BC, in order to investigate the effects of different types of boundary conditions. The various pictures discussed refer to the large volume limit, while the sizes that can be currently reached are rather small. We will, therefore, pay particular attention to properly take into account *corrections to scaling*. In particular, we will try to determine what values of the parameters θ' and $d-d_s$ fit the data in the more “natural” way, namely, with the smallest corrections to scaling for the range of sizes considered.

A summary of our results is as follows. We find that for *periodic* BC, a simple scaling ansatz fits the results in a natural way, i.e., with negligible corrections to scaling and no adjustable parameters besides $d-d_s$ and θ' . This gives $d-d_s=0.42\pm 0.03$; $\theta'=-0.01\pm 0.03$ (the meaning of the error bars will be explained later), which agrees with the results of PY, and is compatible with the TNT picture. We cannot rule out crossover to either the droplet or the RSB picture at length scales larger than our system sizes, but these scenarios, especially the latter, would require larger corrections to scaling than the TNT picture.

For *free* BC, *all* forms of fitting require some corrections to scaling. The most natural scenario, in the sense explained above, gives $d-d_s=0.45\pm 0.02$; $\theta'=0.18\pm 0.03$, with small corrections (of the order of 3%), which is compatible with the droplet picture. Allowing somewhat larger corrections (of the order of 10%), the data are also compatible with $\theta'=0$, namely, with the TNT picture. Finally, if we allow for very large corrections, the data are also consistent with the RSB picture.

In the second part of the paper, we analyze the performance of the branch-and-cut algorithm. We find that the number of elementary operations required to find the ground state increases exponentially with the size, as expected since computing a ground state of a three-dimensional spin-glass system is an \mathcal{NP} -hard problem.²¹ We also find, interestingly enough, that the CPU time is larger for samples in which there is an excited state close in energy to the ground-state energy, yet different from the ground state in the orientation of a large number of spins. We are not aware of any previous quantitative measures of this trend, which we expect to be common to other algorithms as well.

The rest of this paper is organized as follows. In Sec. II, we describe the method of perturbing the ground states to get information about low-energy excitations, introduced by PY. Our results for the nature of the large-scale, low-energy excitations are given in Sec. III. A discussion and summary of the results are given in Sec. III D. A short description of the branch-and-cut algorithm used is given in Sec. IV, and the performance of the algorithm is analyzed in Sec. V. Our conclusions are summarized in Sec. VI.

II. GROUND-STATE PERTURBATION METHOD

The Hamiltonian of the spin-glass model is given by

$$\mathcal{H} = - \sum_{\langle i,j \rangle} J_{ij} S_i S_j, \quad (1)$$

where the sites i lie on a simple cubic lattice with $N=L^3$ spins in dimension $d=3$, $S_i = \pm 1$, and the J_{ij} are nearest-neighbor interactions chosen from a Gaussian distribution with zero mean and standard deviation unity. Free boundary conditions are applied in all directions.

For a given set of bonds, we determine the exact ground state using a branch-and-cut algorithm discussed in Sec. IV. Let $S_i^{(0)}$ be the ground-state spin configuration. As discussed by PY, we then perturb the couplings J_{ij} by an amount proportional to $S_i^{(0)} S_j^{(0)}$ in order to increase the energy of the ground state relative to the other states and therefore possibly induce a change in the ground state. This perturbation, which depends upon a positive parameter ϵ , is defined by

$$\Delta \mathcal{H}_\epsilon = \frac{\epsilon}{N_b} \sum_{b \langle i,j \rangle} S_i^{(0)} S_j^{(0)} S_i S_j, \quad (2)$$

where $N_b = dL^{d-1}(L-1)$ is the number of bonds in the Hamiltonian. We denote the unperturbed ground-state energy by $E^{(0)}$ and the perturbed energy of the *same* state by $E_\epsilon^{(0)}$. The energy of the unperturbed ground state will thus increase exactly by an amount $\Delta E^{(0)} \equiv E_\epsilon^{(0)} - E^{(0)} = \epsilon$. The energy of any other state, say α , will increase by the lesser amount $\Delta E^{(\alpha)} \equiv E_\epsilon^{(\alpha)} - E^{(\alpha)} = \epsilon q_l^{(0,\alpha)}$, where $q_l^{(0,\alpha)}$ is the “link overlap” between the states “0” and α , defined by

$$q_l^{(0,\alpha)} = \frac{1}{N_b} \sum_{\langle i,j \rangle} S_i^{(0)} S_j^{(0)} S_i^{(\alpha)} S_j^{(\alpha)}, \quad (3)$$

in which the sum is over all the N_b nearest-neighbor bonds. Note that the *total* energy of the states changes by an amount of the order of unity.

As we apply the perturbation, the energy *difference* between a low-energy excited state and the ground state decreases by the amount

$$\Delta E^{(0)} - \Delta E^{(\alpha)} = \epsilon(1 - q_l^{(0,\alpha)}). \quad (4)$$

If there is at least one excited state such that $E^{(\alpha)} - E^{(0)} < \Delta E^{(0)} - \Delta E^{(\alpha)}$, then one of these excited states will become the ground state of the perturbed Hamiltonian. We denote the new ground-state spin configuration by $\tilde{S}_i^{(0)}$, and indicate by q_l and q , with no indices, the link and the spin overlap between the new and old ground states $S_i^{(0)}$ and $\tilde{S}_i^{(0)}$, where q is defined as usual by $q = 1/N \sum S_i^{(0)} \tilde{S}_i^{(0)}$.

Due to the spin-flip symmetry of the Hamiltonian (1), the ground state is doubly degenerate, and therefore the distribution of q is symmetric²² around $q=0$. Hence, in the rest of the paper we will restrict ourselves to $q \geq 0$ without loss of information.

Consider the probability $P(\epsilon, L)$ (with respect to the random couplings) that q is less than unity, namely that $S_i^{(0)}$ and $\tilde{S}_i^{(0)}$ differ in a finite fraction of the spins. As discussed by PY, we assume that $P(\epsilon, L)$ is dominated by those samples in which $S_i^{(0)}$ and $\tilde{S}_i^{(0)}$ differ by flipping a single connected cluster of spins, with linear size L . Deviations from this assumption give rise to corrections to scaling, as pointed out by Middleton,¹⁰ and will be analyzed in Sec. III. There are two

energy scales in the problem: the typical energy above the ground state of such an excitation, which scales as $L^{\theta'}$ ($\theta' = \theta$ in the droplet picture), and the threshold energy of Eq. (4), which scales as $\epsilon L^{-(d-d_s)}$ since $1-q_l$ is proportional to the surface of the excitation, $1-q_l \sim L^{-(d-d_s)}$. Hence, the dimensionless probability $P(\epsilon, L)$ is a function of the ratio of these two energy scales:

$$P(\epsilon, L) = g(\epsilon L^{-\mu}), \quad (5)$$

where $g(x)$ is a scaling function and

$$\mu \equiv \theta' + d - d_s. \quad (6)$$

From this we obtain scaling relations for various observables. For example, since $1-q \sim O(1)$ and $1-q_l \sim L^{-(d-d_s)}$, we obtain⁸

$$\langle 1-q \rangle = F_q(\epsilon L^{-\mu}), \quad (7)$$

$$\langle 1-q_l \rangle = L^{-(d-d_s)} F_{q_l}(\epsilon L^{-\mu}), \quad (8)$$

where $\langle \dots \rangle$ is the average with respect to the random couplings. By measuring $\langle 1-q \rangle$ and $\langle 1-q_l \rangle$ we can then determine $d-d_s$ and θ' , the two exponents discriminating the various pictures of the spin-glass phase discussed in Sec. I.

For small ϵ , we expect the probability that the ground state changes to be proportional to ϵ (for fixed L), which implies $g(x) \sim x$ for $x \rightarrow 0$. Hence $F_q(x)$ and $F_{q_l}(x)$ also vary linearly for small x , and the *asymptotic scaling* behavior for $L \gg \epsilon^{1/\mu}$ is

$$\langle 1-q \rangle \sim \epsilon L^{-\mu}, \quad (9)$$

$$\langle 1-q_l \rangle \sim \epsilon L^{-\mu_l}, \quad (10)$$

where

$$\mu_l \equiv \theta' + 2(d-d_s). \quad (11)$$

In the RSB case, $d-d_s = \theta' = 0$, and therefore $\mu = \mu_l = 0$. The scaling relations in Eqs. (7) and (8) reduce in this case to

$$\langle 1-q \rangle = F_q(\epsilon), \quad \langle 1-q_l \rangle = F_{q_l}(\epsilon) \quad (\text{RSB}), \quad (12)$$

and the asymptotic scaling behavior to

$$\langle 1-q \rangle \sim \epsilon, \quad \langle 1-q_l \rangle \sim \epsilon \quad (\text{RSB}). \quad (13)$$

We see that both scaling and asymptotic scaling are in a sense trivial in RSB, since the L dependence disappears. Nevertheless, we will still use the term *scaling*.

It is also convenient to analyze just those samples in which the unperturbed and perturbed ground states are very different, i.e., where $q \leq q_{\max}$, a threshold value. Denoting such restricted averages by $\langle \dots \rangle_c$, we have

$$\langle 1-q_l \rangle_c = L^{-(d-d_s)} F_{q_l}^c(\epsilon L^{-\mu}). \quad (14)$$

This is of the same form as in Eq. (8), but, for sufficiently small q_{\max} , the behavior of the scaling functions $F_{q_l}(x)$ and $F_{q_l}^c(x)$ at small argument will be different for the following reason. If we average over all samples, we need to include the probability $P(\epsilon, L)$ that the perturbation generates an excitation with $q < 1$. This is proportional to $\epsilon L^{-\mu}$ for $\epsilon L^{-\mu} \ll 1$, which is the reason why $F_q(x) \sim x$ for small x . However, this factor is automatically taken into account in the *selection* of the samples in the restricted average in Eq. (14), and so should not be included again when performing the average. As a result, $F_{q_l}^c(x)$ tends to a constant for $x \rightarrow 0$, therefore the asymptotic scaling is

$$\langle 1-q_l \rangle_c \sim L^{-(d-d_s)}. \quad (15)$$

In particular, in RSB this becomes

$$\langle 1-q_l \rangle_c \sim \text{const} \quad (\text{RSB}). \quad (16)$$

Note that in both cases the asymptotic scaling is *independent* of ϵ .

When analyzing the numerical data, we must be aware that there are corrections to both (simple) scaling and asymptotic scaling that occur when L is not large enough. Corrections to simple scaling take the form of *additive* corrections to relations such as Eqs. (5), (7), (8), and (14), whose amplitude is characterized by a correction-to-scaling exponent ω . For example, including the leading correction, Eq. (14) becomes

$$\langle 1-q_l \rangle_c = \frac{1}{L^{d-d_s}} \left\{ F_{q_l}^c(\epsilon L^{-\mu}) + \frac{1}{L^\omega} G_{q_l}(\epsilon L^{-\mu}) \right\}. \quad (17)$$

For $\epsilon L^{-\mu} \rightarrow 0$, this gives the correction to asymptotic scaling corresponding to Eq. (15),

$$\langle 1-q_l \rangle_c = \frac{1}{L^{d-d_s}} \left(a + \frac{b}{L^\omega} \right). \quad (18)$$

For the RSB case, this goes over to

$$\langle 1-q_l \rangle_c = a + \frac{b}{L^\omega}, \quad (19)$$

rather than Eq. (16).

Even when these corrections to (simple) scaling are negligible and the scaling *form*, such as Eq. (14), is valid, the argument of the scaling function may not be sufficiently small for asymptotic scaling to hold. In this regime, when fitting the data to asymptotic scaling we have to consider further corrections to (asymptotic) scaling, whose form is obtained by expanding the scaling function in its argument. For example, the leading correction to Eq. (15), coming from expanding the $F_{q_l}^c$ in Eq. (14) to second order, will be

$$\langle 1-q_l \rangle_c = \frac{1}{L^{d-d_s}} \left(a + b \frac{\epsilon}{L^\mu} \right), \quad (20)$$

TABLE I. Number of independent realizations of the disorder (samples) used in the computations.

L	$\epsilon/\tau=\frac{1}{4}$	$\epsilon/\tau=\frac{1}{2}$	$\epsilon/\tau=1$	$\epsilon/\tau=2$	$\epsilon/\tau=4$
4	50 000	50 000	50 000	50 000	50 000
6	20 000	20 000	20 000	20 000	20 000
8	15 000	13 467	13 467	6000	6000
10	10 000	7440	6000	4918	4000
12	5670		4202		

which goes over to $\langle 1 - q_l \rangle_c = a + b\epsilon$ in RSB. In general, both types of corrections need to be borne in mind when fitting the data.

III. RESULTS

We applied the perturbation method described in the preceding section to systems of size $L=4,6,8,10$, and 12. For each size, we considered five values of the perturbation strength ϵ , namely, $\epsilon/\tau=\frac{1}{4},\frac{1}{2},1,2$, and 4, where $\tau=\sqrt{6}$ is the mean-field transition temperature, except for $L=12$ for which only $\epsilon/\tau=\frac{1}{4}$ and 1 were considered. We choose this value of τ so that we can compare our results with the results of PY for periodic BC. In order to discriminate between the different pictures, it is important to have high statistics. Table I reports the number of samples computed for each size. Note that the number of samples necessary to achieve a given statistical error increases as ϵ decreases, since the fraction of samples in which $\tilde{S}^{(0)} \neq S^{(0)}$ decreases.

A. Spin and link overlap

1. Qualitative analysis

We start with a qualitative analysis of the results. In Fig. 1, we show scatter plots in the (q, q_l) plane for $L=4,10$ and $\epsilon/\tau=\frac{1}{4},1$, where each point represents one of 2000 randomly generated samples. Clearly, the link and spin overlaps are strongly correlated. We note that, as ϵ decreases, there are less points with small q , and that as L increases the data shift towards larger values of q_l . Similar plots²³ for periodic BC show that q_l is significantly lower than for free BC. While q has a large variance (the points are distributed along the whole q axis), the link overlap has a much smaller variance, which decreases as L increases, suggesting that in the thermodynamic limit q_l either tends to 1 or becomes a well-defined function of q .

To quantify this, in Fig. 2 we show the standard deviation of q_l ,

$$\sigma = \sqrt{\langle q_l^2 \rangle_c - \langle q_l \rangle_c^2}, \quad (21)$$

restricted over samples with $q \leq q_{\max}$, as a function of the system size for $\epsilon/\tau=1$. We take $q_{\max}=0.2$, since we are interested in the region of small q , which corresponds to large-scale excitations. A power law $\sigma = aL^{-\delta}$ fits well the data with $\delta=0.52 \pm 0.03$ ($\chi^2=1.80$, the best fit is shown in Fig. 2). Here and in the following, unless stated otherwise, the error bars on the fit parameters are purely statistical in

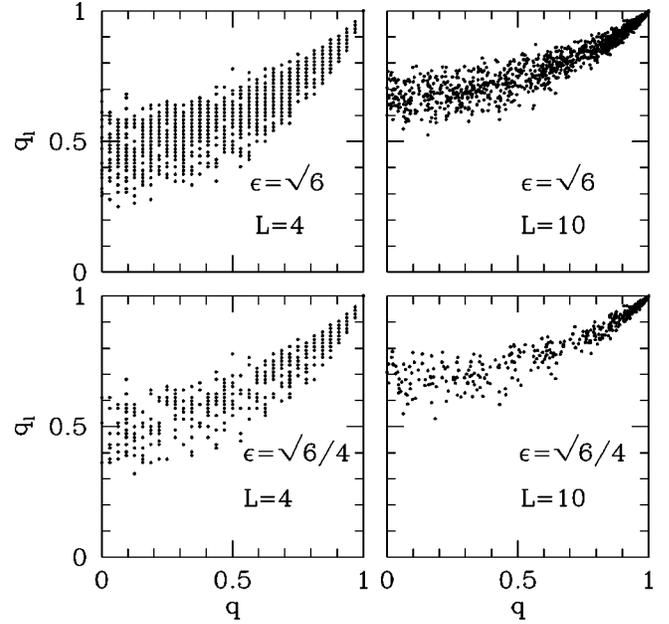


FIG. 1. Scatter plots in the (q, q_l) plane for different values of the size L and perturbation strength ϵ .

relation to the fitting form considered.²⁵ Restricting the average in Eq. (21) to different intervals of q gives results also compatible with a power law. A vanishing σ in the thermodynamic limit is consistent with RSB, which predicts that q_l is a (nontrivial) function of q . It is also trivially consistent with the droplet model or the TNT picture, where $q_l=1$ for all q .

We also measured how the standard deviation of q varies with L , finding that it varies between 0.28 and 0.32. It can be fitted both to a constant (as expected in RSB) or to a power law with a small exponent around 0.1. However, the error bars are very large, hence the fits are not very informative.

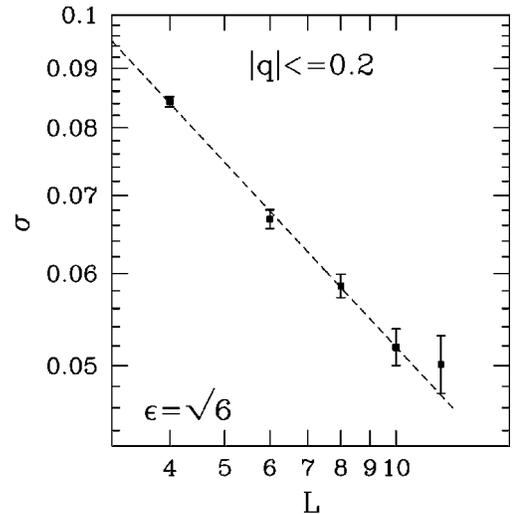


FIG. 2. A plot of the standard deviation of the link overlap $\sigma = \langle (q_l^2)_c - \langle q_l \rangle_c^2 \rangle^{1/2}$, where the average is restricted to samples such that $q \leq 0.2$. The line represents a power-law fit with exponent $\delta = 0.52$.

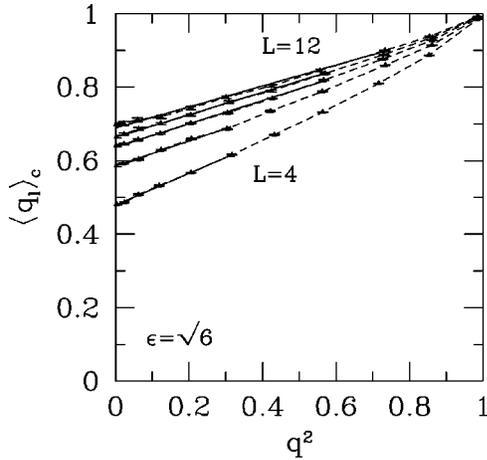


FIG. 3. Average of the link overlap restricted to intervals of q of width 0.1, as a function of q for different sizes L (from bottom to top, $L=4,6,8,10,12$). The continuous lines represent quadratic fits including values of q up to where the lines end. The dashed lines are a guide to the eye.

Under the RSB hypothesis, it is interesting to study the functional relationship between q and q_l . In Fig. 3, we show the *average* value of q_l , restricted to intervals $q \in [q_{\min}, q_{\max}]$, as a function of the mean value of q in each interval²⁴ for $\epsilon/\tau=1$. For fixed L , a quadratic form $q_l = \alpha(L) + \beta(L)q^2$, motivated by the infinite-range Sherrington-Kirkpatrick model where $q_l = q^2$, fits well the data for q less than some cutoff value which increases with L (see Fig. 3). The quadratic fit works well also for other values of ϵ , and $\alpha(L)$ and $\beta(L)$ show a modest variation with ϵ . We also tried global fits including data for all values of q and L , obtaining similar results.

Extrapolating $\alpha(L)$ and $\beta(L)$ to $L \rightarrow \infty$ with fits of the form $\alpha(L) = \alpha + b/L^c$, $\beta(L) = \beta + b'/L^{c'}$, we obtain

$$q_l = (0.77 \pm 0.02) + (0.27 \pm 0.03)q^2, \quad (22)$$

where again the errors are purely statistical for the functional form considered. This nontrivial relation between q and q_l in the thermodynamic limit is consistent with RSB, while in the droplet or TNT pictures the data in Fig. 3 would shift to $q_l \equiv 1$ in this limit.

The power-law form $1 - \alpha(L) = b/L^c$, $\beta(L) = b'/L^{c'}$, which implies $q_l = 1$ in the large volume limit, fits poorly the data if we include all sizes. However, if we exclude the $L=4$ data, the quality of the fit becomes as good as that of the RSB fit just discussed. Hence, allowing for small corrections to scaling, a scenario with $d-d_s > 0$, as in the droplet and the TNT pictures, is consistent with the data.

This already shows that care must be taken to properly consider corrections to scaling when comparing the fits to various pictures. In the following, we will investigate in detail the validity of the various pictures by considering several observables and explicitly discussing the corrections to scaling.

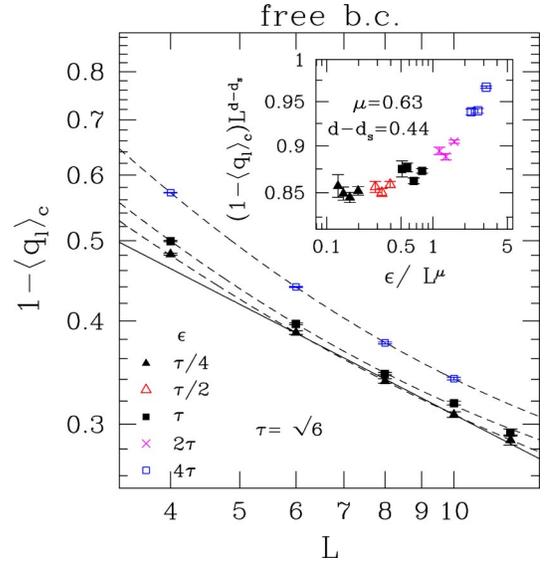


FIG. 4. (Color online) Logarithmic plot of the average $\langle 1 - q_l \rangle_c$, restricted to samples with $|q| \leq 0.4$, as a function of the system size L . Only three values of ϵ are displayed for clarity. The lower continuous line is the best fit with a power-law form III in Eq. (23) for $\epsilon/\tau = \frac{1}{4}$, where $\tau = \sqrt{6}$, and the $L=4$ data have been excluded from the fit. The dashed lines are the best fits with form I in Eq. (23). The inset shows a scaling plot of the data in the main figure, excluding the $L=4$ data, according to Eq. (14). Here and in the following figures, note that the data for various ϵ are correlated, since the samples used for large ϵ were also used for small ϵ .

2. Determination of $d-d_s$

We start with the determination of $d-d_s$ from various observables. We will show that for all observables, a wide range of values of $d-d_s$ fits well the data when allowing corrections to scaling, but that for all observables the smallest corrections are attained for a value of $d-d_s$ around 0.44 in agreement with the results of PY.

The main part of Fig. 4 plots $\langle 1 - q_l \rangle_c$ as a function of L , for various values of ϵ and $q_{\max} = 0.4$ ($q_{\max} = 0.2$ gives essentially the same results). First, note that, independently of what picture holds in the $L \rightarrow \infty$ limit, the data deviate significantly from asymptotic scaling, see Eq. (15), in which the various ϵ values should collapse on a single curve. Second, the data have a noticeable positive (upward) curvature for all values of ϵ . In Sec. III C, we will show that in the case of periodic BC, the data have a much smaller dependence on ϵ and a much smaller curvature (see Fig. 12).

In order to determine how the various pictures fit the data of Fig. 4, we start by considering, following Ref. 9, the following three functional forms:

$$\text{form I, } \langle 1 - q_l \rangle_c = a + b/L^c,$$

$$\text{form II, } \langle 1 - q_l \rangle_c = a + b/L + c/L^2, \quad (23)$$

$$\text{form III, } \langle 1 - q_l \rangle_c = b/L^c.$$

Form I corresponds to the RSB prediction including the leading correction to scaling, see Eq. (19), with $c \equiv \omega$. Form

TABLE II. Fits to $\langle 1 - q_l \rangle_c$ with $q_{\min}=0$ and $q_{\max}=0.4$. The four groups of data refer, from top to bottom, to the three fitting functions I, II, and III in Eq. (23), and form IV in Eq. (25) with $d - d_s = 0.44$. For form III, data for $L=4$ were not included in the fit. The errors are the standard errors of a nonlinear fitting routine (Ref. 25), and Q is the goodness-of-fit parameter.

Form	ϵ/τ	χ^2	Q	a	b	c
I	0.25	0.014	0.99	0.171(2)	1.068(5)	0.890(8)
	0.5	0.015	0.90	0.185(2)	1.14(1)	0.96(1)
	1	3.04	0.22	0.201(14)	1.26(7)	1.04(7)
	2	0.39	0.52	0.206(6)	1.42(3)	1.08(3)
	4	0.27	0.60	0.215(4)	1.73(3)	1.14(2)
II	0.25	0.037	0.98	0.182(1)	1.29(2)	-0.40(5)
	0.5	0.010	0.92	0.189(1)	1.23(1)	-0.14(3)
	1	3.01	0.22	0.198(8)	1.16(10)	0.16(27)
	2	0.47	0.49	0.199(4)	1.21(5)	0.31(14)
	4	0.42	0.52	0.202(3)	1.31(4)	0.64(11)
III	0.25	1.40	0.49		0.85(2)	0.44(1)
	0.5	1.63	0.20		0.87(3)	0.45(2)
	1	9.81	0.007		0.88(3)	0.44(2)
	2	6.75	0.009		0.95(4)	0.47(2)
	4	11.4	0.0007		1.09(5)	0.51(2)
IV	0.25	0.55	0.76	0.808(4)	7(8)	3.6(8)
	0.5	0.65	0.42	0.811(6)	6(7)	3.2(8)
	1	5.91	0.05	0.828(7)	18(37)	4.0(1.5)
	2	2.80	0.09	0.844(9)	5(5)	2.9(7)
	4	3.01	0.08	0.87(1)	3.8(1.6)	2.3(3)

II is a different parametrization of the corrections to scaling. Form III corresponds to the asymptotic behavior of both the TNT and droplet pictures *without* corrections to scaling, see Eq. (15), with $c \equiv d - d_s$.

The results of these fits (performed by χ^2 minimization) are reported in Table II. From the table we see that forms I and II, appropriate to RSB, fit well the data with a low χ^2 and $a > 0$ outside the error bars. The best fits with form I are shown by the dashed lines in Fig. 4. The variation of a between forms I and II provides a measure of the systematic error associated with the unknown corrections to scaling. Within this error, a is independent of ϵ , as predicted by RSB. Therefore, the data for $\langle 1 - q_l \rangle_c$ are compatible with RSB, and our central estimate under the RSB hypothesis is

$$\lim_{L \rightarrow \infty} \langle 1 - q_l \rangle_c = 0.20 \pm 0.02 \quad (\text{RSB}), \quad (24)$$

where the error takes into account also the uncertainty in the *form* of the corrections to scaling, assuming that the corrections considered in either form I or II describe well the data in the whole range of sizes considered. Marinari and Parisi⁹ fitted form I (II) to their data for periodic boundary conditions, $L \leq 14$ and $\epsilon/\tau = 4$, and obtained $a = 0.24$ ($a = 0.30$), from which we estimate a central value $a = 0.27 \pm 0.05$. This is just in agreement with our estimate above for free BC, suggesting that the infinite volume limit of $\langle 1 - q_l \rangle_c$, if non-

zero, may be independent of the boundary conditions, although we do not have an argument why this should be the case.

The power-law form, III, appropriate to the droplet model or the TNT scenario, does not fit well the data if we include all the sizes, but if we exclude $L=4$, it does fit well for $\epsilon/\tau < 2$, and the fit parameters b and c are almost independent of ϵ . The quality of the fit of form III is worse than that of forms I and II, but still acceptable. The main point we want to stress, however, is that the worse fit of form III alone does *not* necessarily favor the RSB picture, since form III does not include corrections to scaling, while forms I and II do. In other words, forms I and II are rather “forgiving” with the RSB picture, allowing corrections of magnitude 100%–200% of the predicted L -independent asymptote, while form III demands that the power-law scenario fits with corrections smaller than the (very small) statistical errors. By looking at Fig. 4, it is clear that the data are closer to a power law than to an L -independent behavior.

Therefore, in order to try a comparison that puts the various pictures on an equal footing, we performed fits with the following more general functional form:

$$\text{form IV, } L^{d-d_s} \langle 1 - q_l \rangle_c = a + b/L^c, \quad (25)$$

where we fix $d - d_s$ and minimize χ^2 with respect to a, b, c , repeating the procedure for different values of $d - d_s$. For $d - d_s = 0$, form IV reduces to form I, while for $d - d_s > 0$, it corresponds to form III plus a correction-to-scaling term, with correction-to-scaling exponent $\omega = c$. We find that, as we might have expected from the previous discussion, form IV fits well the data for a wide range of values of $d - d_s$. For example, for $\epsilon/\tau = \frac{1}{4}$, a value of $d - d_s$ between 0 and 0.45 gives a goodness-of-fit parameter $Q \geq 0.43$, which is entirely acceptable.

This shows that, when allowing for corrections to scaling for *all* pictures, the droplet and TNT pictures are as good as RSB as far as the statistical quality of the fits is concerned. However, within the interval of acceptable values of $d - d_s$, clearly the larger is $d - d_s$ the smaller are the corrections to asymptotic scaling. For example, for $\epsilon/\tau = \frac{1}{4}$ and $d - d_s = 0.42$, the correction term b/L^c in form IV amounts to only 0.1% of the total for $L=12$, while for $d - d_s = 0$ it amounts to 43%. Hence, a large value of $d - d_s$ may be regarded as more natural in this range of sizes.

If we impose that the correction to scaling for $L \geq 8$ is less than a (arbitrary) limit of 3%, we obtain the estimate

$$d - d_s = 0.44 \pm 0.03, \quad (26)$$

where the error is purely statistical within this assumption. In Table II, we show the fits obtained with form IV imposing this value. This value agrees with the estimate $d - d_s = 0.42 \pm 0.02$ of PY for periodic BC (see also Sec. III C of this paper), which is reassuring, since $d - d_s$ should not depend on the boundary conditions.

Note that for $d - d_s = 0$, corresponding to form I, corrections within 3% from the asymptotic limit would only be attained for a size $L \approx 200$. We also note that, as we discussed in Sec. II, even in the regime where corrections to

scaling are negligible, asymptotic scaling sets in only for $L \gg \epsilon^{1/\mu}$. This explains why, if $d-d_s \approx 0.44$, the quality of the power-law fit in Table II becomes progressively worse as ϵ increases. In particular, the deviation from asymptotic scaling is large for $\epsilon/\tau=4$, hence from the data of $\epsilon/\tau=4$ alone one should not necessarily conclude⁹ that an asymptotic power-law behavior is ruled out. This is seen also in the inset of Fig. 4, which shows that, if we exclude $L=4$, the data are compatible with the scaling relation, Eq. (8), where the exponent μ is independently determined below.

PY determined $d-d_s$ from the ratio $R = \langle 1-q_l \rangle / \langle 1-q \rangle$ which has the same scaling behavior as the quantity $\langle 1-q_l \rangle_c$ used here, namely, $R = L^{-(d-d_s)} F_R(\epsilon/L^\mu)$ with $F_R(x) \sim \text{const}$ as $x \rightarrow 0$. Middleton¹⁰ observed that, in two dimensions, small droplets introduce significant corrections to scaling, and suggested that this may be the case also in three dimensions, possibly invalidating the conclusions of PY. The quantity $\langle 1-q_l \rangle_c$ is less sensitive to these corrections, since, by restricting to small q , small droplets should give a smaller contribution, because we have eliminated the part at large q where one can have *only* small droplets. Hence, to investigate these corrections, we fitted our data for R with forms I–IV above (with R replacing $\langle 1-q_l \rangle_c$). The results we find are very similar to those for $\langle 1-q_l \rangle_c$: forms I and II fit well the data with a low χ^2 , giving $a = 0.27 \pm 0.03$ independent of ϵ within the error bars. A power law fits well the data if we exclude $L=4$, with an exponent $d-d_s = 0.43 \pm 0.03$ nearly independent of ϵ and in agreement with Eq. (26). The residual dependence on ϵ is well accounted for by a scaling plot similar to the inset in Fig. 4. Form IV also fits well the data for a wide range of values of $d-d_s$. Again, a power law is more natural in the sense that corrections to scaling are smaller, and the smallest corrections are obtained for $d-d_s$ around 0.43 as for $\langle 1-q_l \rangle_c$.

We interpret the fact that the two quantities give the same value of $d-d_s$ as evidence that corrections due to small droplets are indeed not important in three dimensions in this range of sizes. In Sec. III C, we will show that this is also the case for periodic BC.

To summarize this part, the data for both R and $\langle 1-q_l \rangle_c$ are compatible with a wide range of values of $d-d_s$ between zero (corresponding to RSB) and ≈ 0.44 , but a value at the higher end of this range describes the data in a more natural way, in the sense that the corrections to scaling are smaller.

3. Determination of θ'

Next, we turn to the exponent μ defined in Eq. (6), from which we will extract θ' which is the other exponent, with $d-d_s$, characterizing the spin-glass phase. To this end, we consider the ratio

$$B = \frac{\langle 1-q_l \rangle^2}{\langle (1-q_l)^2 \rangle}, \quad (27)$$

which follows the scaling law

$$B = F_B(\epsilon L^{-\mu}). \quad (28)$$

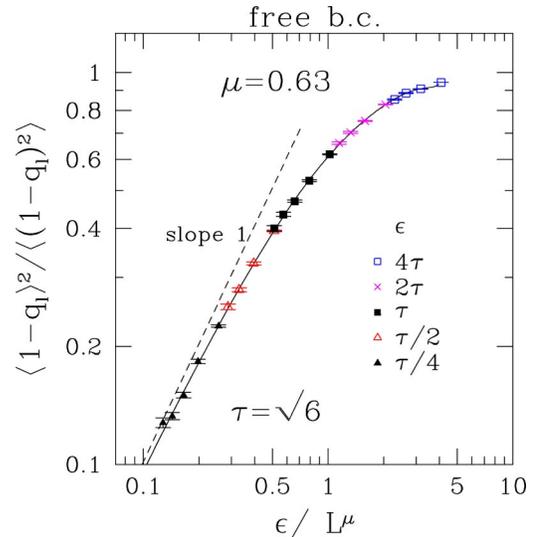


FIG. 5. (Color online) Scaling plot of the ratio $B = \langle 1-q_l \rangle^2 / \langle (1-q_l)^2 \rangle$ according to Eq. (28). The continuous line is a polynomial fit of order $n=6$, which gives χ^2 per degree of freedom (d.o.f) of 1.09 and a goodness-of-fit parameter $Q=0.35$. The dashed line is the linear term of the polynomial fit, corresponding to the asymptotic scaling for $L \rightarrow \infty$.

The factor L^{d-d_s} does not appear here since it cancels between numerator and denominator of Eq. (27), thus allowing us to determine μ independently of $d-d_s$. Following the arguments in Sec. II, we expect $F_B(x) \sim x$ for small x since both $L^{d-d_s} \langle 1-q_l \rangle$ and $L^{2(d-d_s)} \langle (1-q_l)^2 \rangle$ vary as ϵ/L^μ for $\epsilon/L^\mu \rightarrow 0$, hence the asymptotic scaling of B is $B \sim \epsilon/L^\mu$.

To determine μ , we fit the scaling law, Eq. (28), to our data assuming a polynomial form of the order of n for $F_B(x)$, namely, $F_B(x) = \sum_{i=0,n} c_i x^i$, with c_0 set to zero in order to satisfy the asymptotic scaling $F_B(x) \sim x$ as $x \rightarrow 0$. We repeat the fit in an interval of values for μ , and determine the value of μ which gives the best fit, varying n until χ^2 of the best fit becomes approximately constant. In this way, we obtain

$$\mu = 0.63 \pm 0.03, \quad (29)$$

where the error is purely statistical, under the assumption that the corrections to scaling are smaller than the statistical errors of the data. As shown in Fig. 5, scaling is quite satisfactory, with all the data collapsing on one curve, although the data for different ϵ overlap only slightly. The best fit for $n=6$ is displayed by the continuous line. We emphasize that this scaling plot is obtained with only *one* adjustable parameter, μ . Note that in the asymptotic scaling regime the data should follow a straight line (power law), while the data in the figure show a pronounced curvature. Significant corrections to asymptotic scaling must be expected for large ϵ , since B must satisfy the inequality $B \leq 1$. The dashed line in Fig. 5 represents the linear term of $F_B(x)$, corresponding to asymptotic scaling, and the deviation from it gives a measure of the corrections to asymptotic scaling. The $\epsilon/\tau = \frac{1}{4}$ data are quite close to asymptotic scaling, while the data for large ϵ deviate significantly from it. Another manifestation of these

corrections is that, if we fit the data with a power law $B = b/L^{\tilde{\mu}}$, the effective exponent $\tilde{\mu}$ varies strongly with ϵ , converging towards 0.63 for $\epsilon \rightarrow 0$.

In RSB, $B \sim \epsilon$ as $L \rightarrow \infty$ since $\mu = 0$. To test the RSB prediction, we performed fits of B using form I and form II in Eq. (23), (where $\langle 1 - q_l \rangle_c$ is replaced by B). form I gives unphysical (negative) values of a , while form II gives an acceptable fit with a positive a roughly proportional to ϵ . Therefore, the data for B cannot rule out RSB. Note that, if RSB holds asymptotically, the data in Fig. 5 would deviate from the scaling curve for larger L , saturating to a constant value for small values of ϵ/L^μ . The good data collapse we observed, therefore, would be entirely accidental.

We believe that the observed data collapse is a good indication towards the validity of a scaling scenario with large μ . Certainly this scenario is more natural, since it fits the data with (almost) no corrections to (simple) scaling, while the corrections for RSB are very large as apparent from Fig. 5. A similar conclusion was reached in the determination of $d - d_s$.

As a further test, we can obtain a second estimate of μ from the quantity $\langle 1 - q_l \rangle$ (q unrestricted), whose scaling and asymptotic scaling are given in Eqs. (8) and (10). We find results similar to those for B : A power-law fit $\langle 1 - q_l \rangle = b/L^{\tilde{\mu}_l}$ (form III) gives acceptable fits for $\epsilon/\tau \leq \frac{1}{2}$ and for all values of ϵ if we exclude $L = 4$. As for B , the effective exponent $\tilde{\mu}_l$ changes significantly with ϵ , due to corrections to asymptotic scaling, and by extrapolating it to $\epsilon = 0$ we obtain $\mu_l = 1.15 \pm 0.08$. This gives $\mu = \mu_l - (d - d_s) = 0.60 \pm 0.09$ which agrees with the estimate $\mu = 0.63 \pm 0.03$ from B . As for B , the data collapse reasonably well on one curve for $\mu = 0.64 \pm 0.05$ according to Eq. (8), although the quality of the scaling is somewhat worse than that of Fig. 5. To check the RSB prediction, we fitted the data to forms I and II (where now $\langle 1 - q_l \rangle_c$ is replaced by $\langle 1 - q_l \rangle$), finding that they both fit well the data, with a roughly proportional to ϵ as expected in RSB, although, for small ϵ , a is also compatible with zero. Therefore, as for B , the data are also consistent with RSB, but this scenario requires large corrections to scaling, while the hypothesis $\mu = 0.63$ fits the data with almost no corrections to (simple) scaling.

In the analysis so far, we have determined the exponents μ and $d - d_s$ using just the link overlap q_l . By contrast, PY determined μ (for periodic BC) from the scaling of the spin overlap q . An advantage of q_l is that its variance is much lower, as shown in Fig. 1. In any event, we have verified that the scaling relation, Eq. (7), fits well the data for q , giving $\mu = 0.65 \pm 0.02$, in agreement with the estimates from B and $\langle 1 - q_l \rangle$.

Summarizing this part, we find that the data for *all* the quantities considered, namely, B , $\langle 1 - q_l \rangle$, and $\langle 1 - q \rangle$, are consistent with the RSB prediction that $\mu = 0$ asymptotically, but large corrections to scaling are required in the fit, similarly to what we observed in the determination of $d - d_s$. The data are also fitted very well by a scaling scenario with $\mu \approx 0.63$, with almost negligible corrections to scaling (but with sizable corrections to asymptotic scaling, which instead

were small for the observables considered for $d - d_s$). Under the natural assumption of small corrections to scaling, from the estimates of μ and $d - d_s$ in Eqs. (29) and (26), we obtain

$$\theta' = \mu - (d - d_s) = 0.19 \pm 0.06, \quad (30)$$

where, again, the error is purely statistical subject to the condition of having small (less than 3%) corrections to scaling. This result agrees with the droplet theory which predicts that $\theta' = \theta > 0$, and is compatible with the value of $\theta \approx 0.2$ characterizing the energy of domain walls induced by a change in boundary conditions.²⁶ By contrast, for periodic BC and under the same assumption of small corrections to scaling, θ' is compatible with zero (see work of PY in Ref. 8 and Sec. III C). In Sec. III D, we will analyze the origin of this discrepancy, and show that by allowing small (of the order of 10%) corrections to scaling the free boundary-condition data can be reconciled with $\theta' \approx 0$.

B. Box overlap

So far we have analyzed the link and spin overlaps which are computed on the whole system (bulk). We now turn to a different observable, the *box overlap* defined as

$$q_B = \frac{1}{L_B^d} \sum_i S_i^{(0)} \bar{S}_i^{(0)}, \quad (31)$$

where the sum runs over the sites contained in a central cubic box of *fixed* size $L_B = 2$. In the following, we will only consider the absolute value $|q_B|$, which we still call q_B for simplicity. When a large-scale cluster of spins is flipped, for large L the probability that its surface goes across the central box is proportional to the ratio of its surface area, $\sim L^{d_s}$, to the volume, L^d . Therefore, $1 - q_B \sim L^{-(d - d_s)}$ from which we obtain the scaling laws

$$\langle 1 - q_B \rangle = L^{-(d - d_s)} F_{q_B}(\epsilon L^{-\mu}), \quad (32)$$

$$\langle 1 - q_B \rangle_c = L^{-(d - d_s)} F_{q_B}^c(\epsilon L^{-\mu}), \quad (33)$$

where, as for the corresponding scaling functions for q_l , $F_{q_B}(x) \sim x$ and $F_{q_B}^c(x) \sim \text{const}$ for small x . Hence, the asymptotic scaling for $L \rightarrow \infty$ is

$$\langle 1 - q_B \rangle \sim \epsilon L^{-\mu_l}, \quad (34)$$

$$\langle 1 - q_B \rangle_c \sim L^{-(d - d_s)}. \quad (35)$$

In RSB, this reduces to $\langle 1 - q_B \rangle \sim \epsilon$ and $\langle 1 - q_B \rangle_c \sim \text{const}$. An advantage of q_B over q_l is that the former, being measured away from the boundaries, should have smaller corrections to scaling and be less sensitive to boundary conditions. Indeed, Monte Carlo simulations^{27,28} show that q_B has rather small corrections to scaling. This is likely to be particularly important for the free boundary conditions used here.

Figure 6 shows the *restricted* average $\langle 1 - q_B \rangle_c$, with $q_{\max} = 0.4$, as a function of L for two values of ϵ . The data are clearly decreasing with L , are essentially independent of ϵ , as expected from Eq. (35), and are close to a straight line on the logarithmic plot. This indicates that the power-law fit,

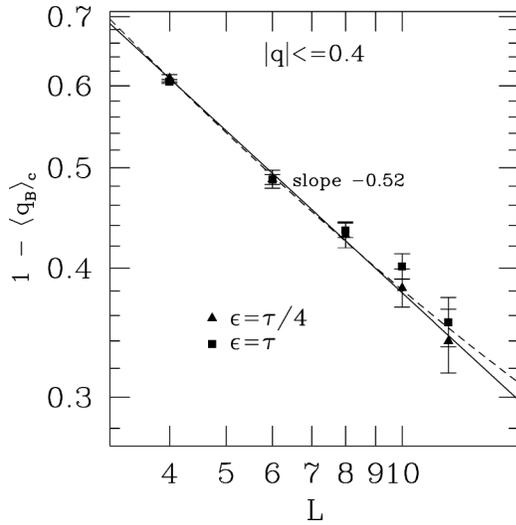


FIG. 6. (Color online) Logarithmic plot of the average box overlap, restricted to samples such that $q \leq 0.4$. We show the data for just two values of ϵ for clarity. The data for other values of ϵ are superimposed. The lower continuous line is a power-law fit for $\epsilon/\tau=4$. The dashed line is the fit with form II in Eq. (23), with q_B replacing q_l . The slope gives the exponent $d-d_s$.

form III, appropriate to the droplet and TNT scenarios, should work well and indeed it does, even for the largest value of ϵ (we note, however, that the statistical errors are larger than for the link overlap, hence the fits are less sensitive to corrections to scaling). The exponent is almost independent of ϵ , varying between 0.48 and 0.52, and from this we obtain the estimate

$$d-d_s = 0.48 \pm 0.03, \quad (36)$$

which is in agreement with the estimates $d-d_s = 0.44 \pm 0.03$ from $\langle 1-q_l \rangle_c$ and $d-d_s = 0.43 \pm 0.03$ from R .

Forms I and II (with q_B replacing q_l) also fit well the data, with a between 0.14 and 0.36 (with no discernible trend with ϵ). Hence, the data are also compatible with RSB, and under the RSB hypothesis, we estimate

$$\lim_{L \rightarrow \infty} \langle 1-q_B \rangle_c = 0.25 \pm 0.10 \quad (\text{RSB}). \quad (37)$$

As usual, we note that the RSB scenario requires rather large corrections to scaling, while the power law fits the data with no corrections.

Figure 7 shows the unrestricted average $\langle 1-q_B \rangle$ multiplied by τ/ϵ , which asymptotically should be independent of ϵ . The data show a small curvature and a significant ϵ dependence, indicating that for this quantity we are not yet in the asymptotic scaling regime (similarly to what we observed for the quantity B). The data are fitted well by a power law, with an exponent that changes with ϵ and tends towards $\mu \approx 0.63$ for $\epsilon \rightarrow 0$. Fits using forms I and II give a compatible with zero. We also determined μ from the scaling relation, Eq. (32), by fixing $d-d_s = 0.44$ and using the same fitting procedure as for B (which assumes no corrections to scaling), finding

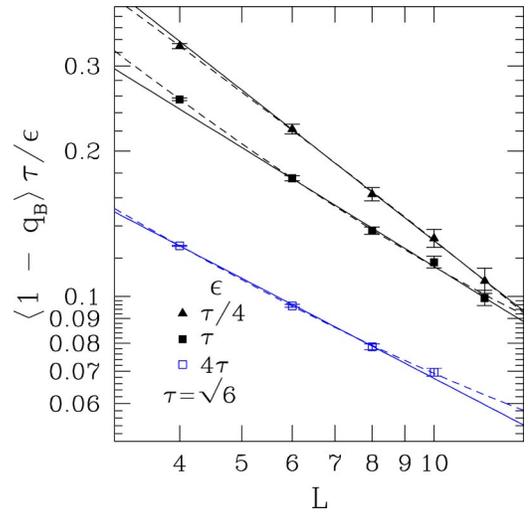


FIG. 7. (Color online) Logarithmic plot of the average box overlap, multiplied by τ/ϵ in order to highlight the deviation from the asymptotic behavior of Eq. (34) in which the data for various ϵ should collapse on a single curve. The continuous lines represent fits with the power-law form III excluding $L=4$. The dashed lines represent fits with form I in Eq. (23).

$$\mu = 0.62 \pm 0.04 \quad (38)$$

which agrees with the various estimates of μ obtained from B , $\langle 1-q_l \rangle$, and $\langle 1-q \rangle$. Figure 8 shows the corresponding scaling plot, in which the data collapse is fairly good.

To conclude this section, the data for box overlaps can be fitted with smaller corrections to scaling than the data for the bulk link and spin overlaps. A fit to the generic scaling picture, with no corrections to scaling, gives results for the exponents $d-d_s$ and μ in agreement with those from the bulk quantities analyzed in the previous sections. However, as

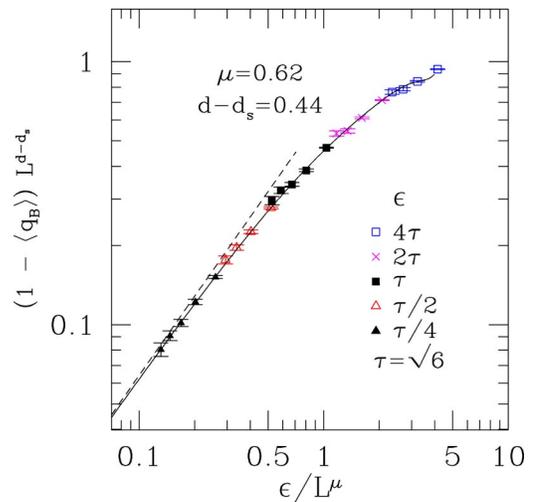


FIG. 8. (Color online) Scaling plot of the box overlap according to Eq. (32). The continuous line is a polynomial fit of the order $n=6$, which gives $\chi^2/\text{d.o.f.} = 0.63$ and a goodness-of-fit parameter $Q = 0.85$. The dashed line is the linear term of the polynomial fit, corresponding to the asymptotic behavior for $L \rightarrow \infty$.

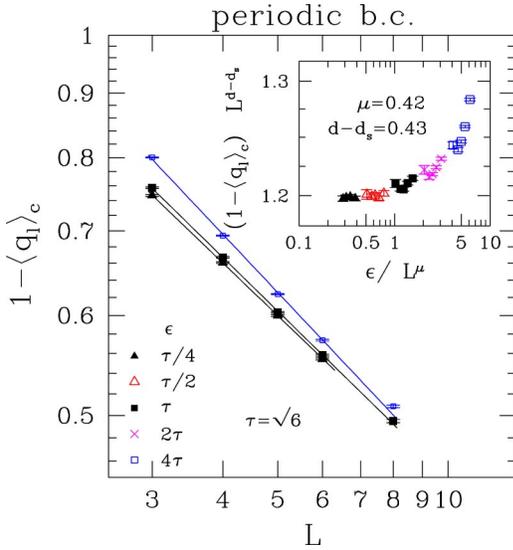


FIG. 9. (Color online) Same as Fig. 4 but for periodic boundary conditions, using the data of PY (Ref. 8).

with the bulk observables, assuming large corrections to scaling, the data can also be fitted to the RSB picture.

C. Comparison with periodic boundary conditions

In order to assess the effect of different boundary conditions, we have repeated part of the analysis above (with the exclusion of box overlaps) for the data of PY (Ref. 8) for periodic boundary conditions and $L \leq 8$. The ground states were obtained using a hybrid genetic algorithm as described by PY. This does not guarantee to find the true ground state, but the systematic errors due to occasionally missing it are smaller than the statistical errors.⁸

Figures 9 and 10 show the equivalent for periodic BC of Figs. 4 and 5 for free BC. The data in Fig. 9 show much less curvature and also a smaller dependence on ϵ than the cor-

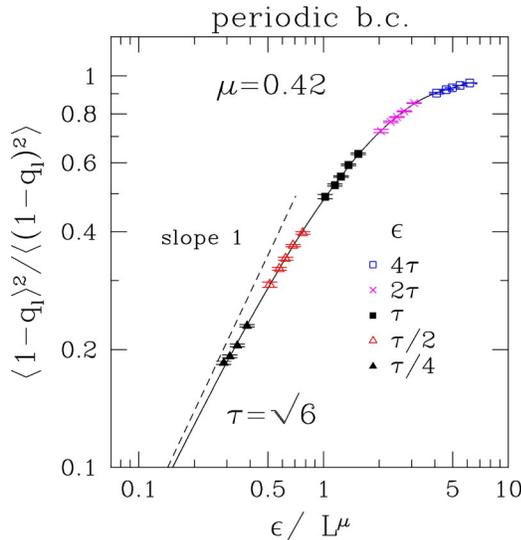


FIG. 10. (Color online) Same as Fig. 5 but for periodic boundary conditions, using the data of PY (Ref. 8).

TABLE III. Fits to $\langle 1 - q_l \rangle_c$ with $q_{\min}=0$ and $q_{\max}=0.4$ for periodic boundary conditions. The three groups of data refer, from top to bottom, to the three fitting functions I, II, and III in Eq. (23), respectively, and form IV in Eq. (25) with $d - d_s = 0.43$.

Form	ϵ/τ	χ^2	Q	a	b	c
I	0.25	0.003	0.99	-0.076(7)	1.256(6)	0.384(4)
	0.5	0.92	0.62	0.05(4)	1.16(2)	0.47(3)
	1	1.58	0.45	0.10(3)	1.16(2)	0.52(3)
	2	2.20	0.33	0.12(3)	1.18(2)	0.54(4)
II	0.25	0.33	0.56	0.279(7)	1.90(5)	-1.5(1)
	0.5	5.22	0.073	0.28(1)	1.90(9)	-1.5(2)
	1	0.69	0.71	0.280(4)	1.89(3)	-1.40(6)
	2	0.04	0.98	0.283(1)	1.90(1)	-1.33(2)
III	0.25	0.02	0.887		1.204(2)	0.433(1)
	0.5	0.35	0.838		1.193(3)	0.427(2)
	1	5.14	0.076		1.21(2)	0.434(6)
	2	7.82	0.020		1.24(2)	0.440(8)
IV	0.25		$2 \cdot 10^{-6}$		1.31(2)	0.46(1)
	1	3.59	0.16	1.205(4)	0.3(1.5)	3(4)
	2	4.69	0.09	1.214(7)	0.2(5)	2(2)
	4	8.38	0.01	1.231(8)	0.7(5)	2.3(7)

responding data for free BC in Fig. 4, indicating that corrections to scaling are smaller for periodic BC. Table III reports the best fits using the three functional forms of Eq. (23). Form I fits well the data, but a varies significantly with ϵ , and for small ϵ it is compatible with zero. Form II also fits well, with a independent of ϵ within the statistical errors. From this fit we estimate

$$\lim_{L \rightarrow \infty} \langle 1 - q_l \rangle_c = 0.28 \pm 0.03 \quad (\text{RSB}) \quad (39)$$

[see the comment after Eq. (24) as to the meaning of the error bar], which agrees with the estimate 0.24 of Marinari and Parisi,⁹ and is just consistent with our estimate 0.20 ± 0.02 for free BC.

The power-law fit with no corrections, form III, fits well the data for the two smallest values of ϵ and, if we exclude $L=3$, for all but the largest value of ϵ . The exponent $c \equiv d - d_s$ is nearly independent of ϵ and gives

$$d - d_s = 0.43 \pm 0.02 \quad (\text{periodic BC}). \quad (40)$$

This result agrees with the estimate $d - d_s = 0.42 \pm 0.02$ of PY obtained from the ratio R defined above, confirming that corrections due to small droplets should not be important in three dimensions, and with our estimate $d - d_s = 0.44 \pm 0.03$ for free BC, indicating that $d - d_s$ does not depend on boundary conditions.

We also performed fits with form IV which includes corrections to scaling. As for free BC, a wide range of values of $d - d_s$ from zero to around 0.44 give a good fit, with the largest values giving the smallest corrections to scaling. The

results of the fit for $d-d_s=0.43$ are shown in Table III. For the two smaller values of ϵ , the fits are difficult because corrections to scaling are very small, hence they are not shown.

We determined the exponent μ from the ratio B using Eq. (28) and the fitting procedure described for free BC, obtaining

$$\mu=0.42\pm 0.03 \quad (\text{periodic BC}), \quad (41)$$

where, as for the estimate of $d-d_s$ above, the errors are purely statistical with the assumption that corrections to scaling are smaller than the statistical errors of the data. Scaling is rather satisfactory as shown in Fig. 10. This value agrees with the estimate $\mu=0.44\pm 0.02$ of PY from the scaling of the spin overlap but incompatible, within the statistical error bars, with the result $\mu=0.63\pm 0.03$ for free BC. We will return in Sec. III D to the origin of the discrepancy between free and periodic BC. The inset of Fig. 9 shows that, with these values of μ and $d-d_s$, the scaling form for $\langle 1-q_l \rangle_c$, Eq. (14), is also well satisfied. Finally, we verified that, if $d-d_s=0.43$, the unrestricted average $\langle 1-q_l \rangle$ satisfies scaling, giving $\mu=0.45\pm 0.02$ in agreement with the estimate from B .

Combining Eqs. (40) and (41), we obtain the estimate of θ' for periodic BC:

$$\theta' = \mu - (d-d_s) = -0.01 \pm 0.03 \quad (\text{periodic BC}). \quad (42)$$

This is compatible with zero and, within the error bars, incompatible with the value $\theta' = \theta \approx 0.2$, where θ characterizes the energy of domain walls induced by boundary-condition changes. A scenario in which $\theta' = 0$ and $d-d_s > 0$ is consistent with the TNT picture. Finally, we note that, although our analysis of the data of PY uses different quantities to extract exponents, our results agree with those given by PY.

D. Discussion and summary of the results

In the previous sections, we have seen that for both free and periodic BC, the analysis of all the different observables considered gives consistent results for the exponents $d-d_s$ and θ' under the assumption of minimal corrections to scaling. However, while the results for $d-d_s$ for free and periodic BC agree with each other, the results for θ' apparently do not, having found $\theta' = -0.01 \pm 0.03$ for periodic BC and $\theta' = 0.19 \pm 0.06$ for free BC. Since θ' , like $d-d_s$, should not depend on the type of boundary conditions, the discrepancy must be due to different corrections to scaling for the two boundary conditions.

Therefore, it is important to analyze further the corrections to scaling. First, we recall that the scaling plots for the quantity B in Figs. 5 (free BC) and 10 (periodic BC), from which we have determined μ (and hence θ'), are obtained by imposing that the *whole* data set (namely, all values of ϵ and L) satisfies scaling with corrections to scaling smaller than the statistical errors, which are less than 1%. Clearly, this is a very stringent requirement. If we relax this require-

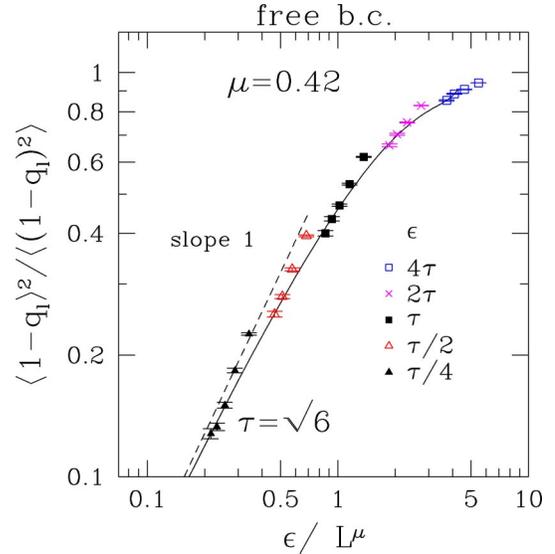


FIG. 11. (Color online) Scaling plot of the ratio $B = \langle 1-q_l \rangle_c^2 / \langle (1-q_l)^2 \rangle$ according to Eq. (28). The continuous line is a polynomial fit of the order of $n=5$, excluding the data with $L=4$ and 6, which gives $\chi^2/\text{d.o.f} = 1.26$. The dashed line is the linear term of the polynomial fit, corresponding to the asymptotic scaling for $L \rightarrow \infty$.

ment, allowing some corrections to (simple) scaling, we can accommodate a larger range of values for μ .

This is shown in Fig. 11, which gives a scaling plot for free BC, analogous to Fig. 5 but assuming the value $\mu = 0.42$ determined from *periodic* BC. The polynomial fitting curve was obtained by excluding from the fit the data points for $L=4$ and 6. One can see that for larger L , the data collapse reasonably well on the curve. The deviation of the $L=4, 6$ data from the curve, less than 10%, is a measure of the corrections to scaling. Therefore, we see that corrections to scaling of less than 10% for the two smallest sizes are sufficient to remove the discrepancy in the value of θ' between free and periodic BC. We verified that also the other quantities considered, namely, $\langle 1-q_l \rangle_c$, R , $\langle 1-q_l \rangle$, $\langle 1-q \rangle$, can be fitted in a similar way.

We also tried the converse operation, namely, a scaling plot of the data for periodic BC but using the value $\mu = 0.63$. We find that one can get a relatively good data collapse excluding the sizes $L=3, 4$, and 5, which deviate from the scaling curve by less than 10%. However, now the data for a given ϵ approach the scaling curve from the right side instead of from the left side as in Fig. 11, but since they have an upward curvature, the correction to scaling should change curvature twice as L increases, which is not very plausible. Hence, we believe that it is more natural to conclude that the correct value of μ is closer to 0.42 than to 0.63, namely, that corrections to scaling are smaller for periodic than for free BC.

Indeed, in general, it is reasonable to expect that corrections are larger for free BC, because these BC have a free surface on which lie a fraction of sites which is quite substantial for moderate sizes. In Fig. 12, we plot together the $\langle 1-q_l \rangle_c$ data of Figs. 4 and 9 for free and periodic BC. The data for free BC lie significantly below those for periodic

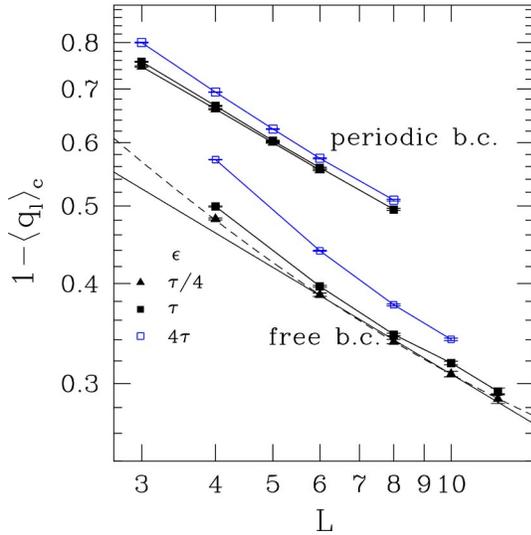


FIG. 12. (color online) This plot shows together the data of Fig. 4 for free boundary conditions and Fig. 9 for periodic boundary conditions.

BC, indicating that the surface of the excitations is smaller for free BC. For periodic BC, the domain wall has to “bend” to return to the same point on the “top surface” as it had on the “bottom surface.” This may be the source of the extra surface area.

Under the hypothesis $d - d_s \approx 0.44$, Fig. 12 then shows that the corrections to *asymptotic* scaling are larger for free BC, since the free boundary-condition data show a more marked deviation from the asymptotic ϵ -independent behavior, and display a larger curvature. This is further indication that free BC have larger corrections.

Evidence that free BC have larger corrections was also found recently in Monte Carlo simulations,²⁹ where some evidence was observed that the free BC data might have a crossover from dropletlike to either TNT- or RSB-like behavior at large sizes.

Incidentally, note that if RSB is the correct asymptotic picture *and* the $L \rightarrow \infty$ limit of $\langle 1 - q_l \rangle_c$ is the same for periodic and free BC, then Fig. 12 would indicate that the corrections are *smaller* for free BC (since the data are closer to their nonzero asymptotic value), which is not very plausible. Note, however, that we do not have an argument why in the thermodynamic limit $\langle 1 - q_l \rangle_c$ should be independent of boundary conditions.

To summarize the first part of the paper, we have analyzed several quantities for periodic and free BC. For both types of boundary conditions, all the data are well described by a general scaling picture involving only two scaling exponents $d - d_s$ and θ' with only small corrections to scaling. Some observables show significant corrections to asymptotic scaling, which are larger for free boundary conditions. Fitting this scaling picture to our data, we obtain comparable values of $d - d_s$ for periodic (0.43 ± 0.02) and free boundary conditions (0.44 ± 0.03).

By imposing that corrections to scaling are less than the statistical errors of about 1%, for periodic boundary conditions we obtain $\theta' \approx 0$, which fits well the TNT scenario

($d - d_s > 0$, $\theta' = 0$), while for free boundary conditions we obtain $\theta' = 0.19 \pm 0.06$, which fits well the droplet picture ($d - d_s > 0$, $\theta' > 0$). By relaxing this requirement and allowing larger corrections to scaling of the order of 10%, the data for free BC can be also fitted by a scenario with $\theta' \approx 0$. Therefore, the data for free BC are also consistent with the TNT picture provided moderate corrections to scaling are allowed, larger than those for periodic BC. We have also provided direct evidence that indeed free BC have larger corrections to scaling.

Data for the box overlap for free BC indicate smaller corrections to asymptotic scaling, which is reasonable since the box is away from the surface, and are consistent with the scenario described above.

For both free and periodic BC, the data are also fitted well by the RSB picture ($d - d_s = 0$, $\theta' = 0$), but only if we allow very large corrections to scaling. In this case, the good scaling behavior we observed for all the observables considered would only be a finite-size artifact, and would disappear at larger sizes. To test this possibility, large system sizes will be needed.

This concludes the first part of the paper, dedicated to the physical results. In the second part, we will describe the branch-and-cut algorithm employed, and analyze its performance in our computations.

IV. THE BRANCH-AND-CUT ALGORITHM

Branch and cut is, to our knowledge, the fastest exact method for determining ground states of spin glasses in three dimensions. To apply this technique, we transform the problem of minimizing the Hamiltonian in Eq. (1) into a standard combinatorial optimization problem known as the *maximum cut* problem. (For a detailed description of optimization and related topics, see Ref. 30.) Consider the interaction graph $G = (V, E)$ associated with the spin-glass Hamiltonian, where G contains vertices $1, \dots, L^3 \in V$ associated with the spin sites and edges $(ij) \in E$ with weight $c_{ij} = -J_{ij}$ associated with the couplings.

Given a partition of V into two sets, $W \subset V$ and its complement $V \setminus W$, the *cut* $\delta(W)$ associated with W is defined as the set of edges with one end point, i say, in W and the other end point, j say, in $V \setminus W$. In formulas, $\delta(W) = \{(ij) \in E | i \in W, j \in V \setminus W\}$. The *weight* of a cut $\delta(W)$ is defined as the sum of the weights of the cut edges $\sum_{(ij) \in \delta(W)} c_{ij}$. A *maximum cut* is a cut with maximum weight among all partitions W . It is easy to show that minimizing the Hamiltonian, Eq. (1), is equivalent to finding a maximum cut in G , see Ref. 19. If we know a maximum cut with node partitions W and $V \setminus W$, the corresponding ground-state spin configuration can be read off by assigning the value up to the spins in W and down to the spins in $V \setminus W$, or vice versa.

The branch-and-cut algorithm solves the maximum cut problem through simultaneous lower and upper bound computations. By definition, the weight of any cut gives a *lower bound* on the optimal cut value. Thus, we can start from any cut and iteratively improve the lower bound using deterministic heuristic rules (local search and other specialized heuristics, see Ref. 31 for details). How do we decide when a cut

is optimal? This can be done by additionally maintaining *upper bounds* on the value of the maximum cut. Upon iteration of the algorithm, progressively tighter bounds are found, until optimality is reached.

Since the availability of upper bounds marks the difference between a heuristic and an exact solution, we now summarize how the upper bound is computed (for more details, see Ref. 31). To each edge (ij) , we associate a real variable x_{ij} and to each cut $\delta(W)$ an *incidence vector* $\chi^{\delta(W)} \in \mathbb{R}^E$ with components $\chi_{ij}^{\delta(W)}$ associated with each edge (ij) , where $\chi_{ij}^{\delta(W)} = 1$ if $(ij) \in \delta(W)$ and $\chi_{ij}^{\delta(W)} = 0$ otherwise. Denoting by $P_C(G)$ the convex hull of the incidence vectors, it can be shown that a basic optimum solution³² of the linear program

$$\max \left\{ \sum_{(ij) \in E} J_{ij} x_{ij} \mid x \in P_C(G) \right\} \quad (43)$$

is a maximum cut. In order to solve Eq. (43) with linear programming techniques, we would have to express $P_C(G)$ in the form

$$P_C(G) = \{x \in \mathbb{R}^E \mid Ax \leq b, 0 \leq x \leq 1\} \quad (44)$$

for some matrix A and some vector b . Whereas the existence of A and b is theoretically guaranteed, even the subsets of $Ax \leq b$ known in the literature contain a huge number of inequalities that render a direct solution of Eq. (43) impractical.

Instead, the branch-and-cut algorithm proceeds by optimizing over a *superset* P containing $P_C(G)$, and by iteratively tightening P , generating in this way progressively better upper bounds. The supersets P are generated by a *cutting plane* approach. Starting with some P , we solve the linear program $\max\{\sum_{(ij) \in E} J_{ij} x_{ij} \mid x \in P\}$ by Dantzig's simplex algorithm.³² Optimality is proven if either of two conditions is satisfied: (i) the optimal value equals the lower bound; (ii) the solution vector \bar{x} is the incidence vector of a cut.

If neither is satisfied, we have to tighten P by solving the *separation problem*. This consists in identifying inequalities that are valid for all points in $P_C(G)$, yet are violated by \bar{x} , or reporting that no such inequality exists. The inequalities found in this way are added to the linear programming formulation, obtaining a new tighter partial system $P' \subset P$ which does not contain \bar{x} . The procedure is then repeated on P' and so on.

At some point, it may happen that (i) and (ii) are not satisfied, yet the separation routines do not find any new cutting plane. In this case, we *branch* on some fractional edge variable x_{ij} (i.e., a variable $x_{ij} \notin \{0,1\}$), creating two subproblems in which x_{ij} is set to 0 and 1, respectively. We then apply the cutting plane algorithm recursively for both subproblems.

V. PERFORMANCE OF THE BRANCH-AND-CUT ALGORITHM

In this section, we study the performance of our current implementation of the branch-and-cut algorithm, in particular, the dependence of the number of computer operations on

TABLE IV. Mean CPU time per sample in seconds for the calculation of the unperturbed ground state, averaged also over the different machines.

L	Mean CPU time per sample
4	0.065
6	0.662
8	10.11
10	338

the system size. The results for size $L=12$ were obtained with a more efficient version of the code, so performance for this size cannot be compared with that for the smaller sizes. Hence, in this section, we shall just consider sizes up to $L=10$.

Finding the ground state of the Hamiltonian, Eq. (1), in three dimensions is an \mathcal{NP} -hard problem,²¹ and all known algorithms to solve this class of problems require a number of operations that grows exponentially on the size of the input, in the worst case.

However, depending on the problem, the number of operations for *typical* instances (for the spin-glass problem, an instance is a realization of the random couplings, or sample) can grow considerably more slowly than the worst case exponential behavior. Furthermore, the number of operations can vary significantly from one instance to another. It is, therefore, useful to investigate experimentally the performance of the algorithm for typical instances, in order to try to extrapolate the computational resources necessary to go to larger sizes, and possibly to identify which parameters of the problem affect most the performance. De Simone *et al.*¹⁹ measured the average CPU time used by the branch-and-cut algorithm to find the ground state of the two-dimensional $\pm J$ spin glass with periodic BC, up to $L=70$, showing that the average CPU time was approximated by a function proportional to L^6 .

Here we analyze the performance of the branch-and-cut algorithm for the three-dimensional spin glass with free BC and Gaussian couplings. In order to do this, we first need a good measure of the performance. For a complex algorithm such as branch-and-cut, a simple and absolute measure of the number of operations is not available. Two possible measures are the CPU time and the number of linear programs solved during the run of the algorithm. In Table IV, we summarize the average running time needed for calculating an unperturbed ground state for the different system sizes.

The CPU time is not an accurate measure, since it depends on the machine architecture and load. Furthermore, our computations were carried out on several different machines, so the CPU time is not useful here. We take instead the number of linear programs solved, n_p , because (i) it is a well-defined and machine-independent quantity; (ii) we have observed that about 95% of the time is spent in solving linear programs; (iii) for a fixed system size, n_p correlates strongly, and almost linearly, with the CPU time. This is shown in Fig. 13, which plots the CPU time versus n_p for 1000 randomly generated samples with $L=10$, computed on the same machine. Note that since the *size* of the linear programs is also

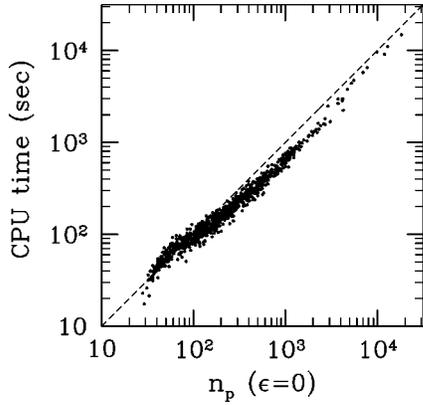


FIG. 13. Scatter plot of the CPU time to find the unperturbed ground state ($\epsilon=0$) versus the corresponding number of linear programs solved (n_p). Each point represents a randomly generated sample with $L=10$. All the computations for this set of samples were run on the same machine. The dashed line indicates a linear behavior.

growing with the system size, the CPU time per linear program increases strongly with L : the average (median) CPU time goes from 0.0077 (0.044) sec for $L=4$ to 0.833 (0.784) sec for $L=10$.

Hence, n_p severely underestimates the rate at which the number of operations increases with L .

From Fig. 13, we also note that the distribution of n_p (and CPU times) is very broad, extending over three orders of magnitude. The histogram distribution of n_p for different system sizes is shown in Fig. 14. In addition to shifting to larger n_p , the distribution broadens as L increases. Also, there is some evidence of a double-peak structure. For $L=10$, we verified that the peak at smaller n_p corresponds to samples that could be solved without branching, while the other peak corresponds to samples where branching was necessary. Since in each branching step the number of subproblems to be solved doubles, the number of linear programs increases rapidly and the second peak is at significantly larger n_p .

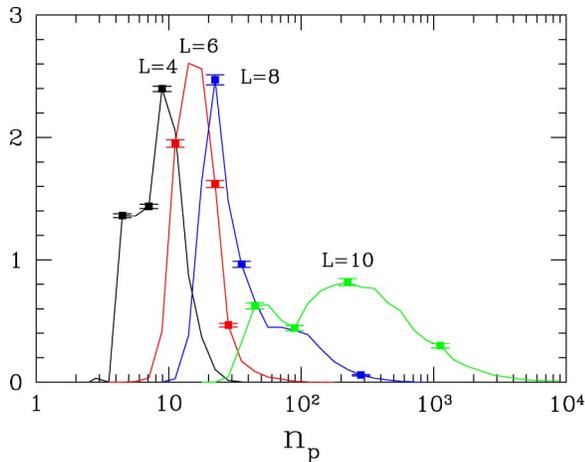


FIG. 14. (Color online) Histogram of the number of linear programs solved by the branch-and-cut algorithm to find the unperturbed ground state for different system sizes.

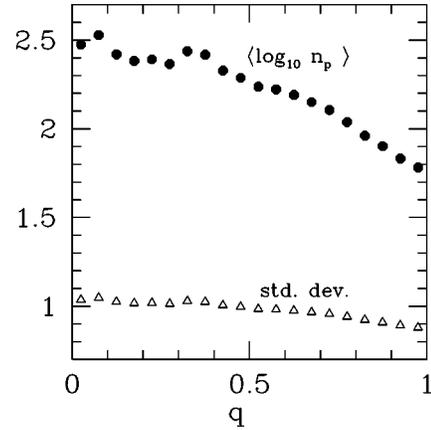


FIG. 15. The circles are a plot of $\langle \log_{10} n_p \rangle$, where n_p is the number of linear programs solved to compute the unperturbed ground state S^0 , versus the overlap between S^0 and the perturbed ground state $\tilde{S}^{(0)}$. The data is for $\epsilon/\tau=4$ and the samples were selected from a set of randomly generated samples with $L=10$, in such a way that the same number of samples is plotted for each consecutive q interval of length 0.1, in order to sample equally all regions of q . The triangles show the standard deviation, among samples, of $\log_{10} n_p$ as a function of q .

In order to identify which parameters of the problem, in addition to the size, affect the performance, we ask whether n_p correlates with the physical observables we measure. No significant correlation was observed with the ground-state energy. Figure 15 plots $\langle \log_{10} n_p \rangle$ for the unperturbed ground state ($\epsilon=0$) and $L=10$ versus the overlap between this state and the perturbed ground state with $\epsilon/\tau=4$. We observe a distinct correlation between n_p and q : for small q , more linear programs are needed than for large q . The figure

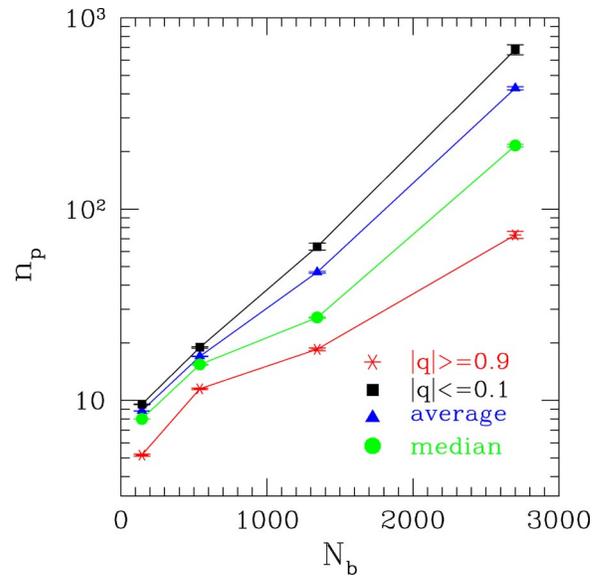


FIG. 16. (Color online) Average n_p , median n_p , and conditional averages of n_p restricted to $|q| \leq 0.1$ and to $|q| \geq 0.9$, as a function of the number of bonds N_b . The data for n_p are for $L=10$ and $\epsilon=0$ (unperturbed ground state), and q is the overlap between the $\epsilon=0$ and $\epsilon/\tau=4$ ground states.

shows that the typical number of linear programs is close to an order of magnitude larger if $q \approx 0$ than if $q \approx 1$. We observed a similar correlation for other values of ϵ as well, and also between the CPU time and q . Again, the distribution of n_p is quite broad as shown by the data for the standard deviation of $\log_{10} n_p$ in Fig. 15.

In order to quantify how the correlation between n_p and q changes with the system size, we show in Fig. 16 the average and median of n_p as a function of N_b , as well as the conditional averages of n_p restricted to samples with large ($|q| \geq 0.9$) and small ($|q| \leq 0.1$) overlaps. We take the number of bonds, N_b , as a measure of the input size, since the maximum cut problem is specified in terms of the edge variables in the graph. From Fig. 16 we see that, first, all measures show an approximately exponential increase with N_b , with corrections for small N_b , and second, the difference between the conditional averages with small and large q seems to increase with the system size, and is about one order of magnitude for $L = 10$.

A qualitative difference between samples with small and large overlaps is that samples with a small $|q|$ have a rougher “energy landscape,” namely, states with an energy close to the ground-state energy yet a spin configuration very different from the ground state. It is then intuitively clear why one would observe a correlation between q and the running time for a stochastic algorithm employing local search heuristics, such as simulated annealing, since when the algorithm encounters one of these configurations with small overlap, it must retrace its steps by a large amount.

For the branch-and-cut algorithm, the reason for the correlation between n_p and q is less obvious, but some insight is provided by an analysis of “reduced cost fixing.” This is a feature of the branch-and-cut algorithm speeding up the computations. In every iteration of the algorithm, reduced cost fixing gives us a sufficient condition to decide which variables (corresponding to the edges in the graph) have already attained their optimal value. Therefore, we can fix the variables with “optimal” status to their current value for all the subsequent iterations of the algorithm, resulting in less overall computational effort. The more variables that can be fixed, the faster the algorithm is in practice.

Since the samples with small overlap have “almost optimal” solutions with spin configurations very different from the ground state, a smaller number of variables can be fixed. Here we do not have the “correct” edge values available until the end. As an example, we checked that for $L = 10$ and $\epsilon = \tau$, for 100 randomly chosen samples with small overlap ($|q| \leq 0.1$), in average 409 ± 39 of the 2700 edge variables could be fixed in the first sub problem, i.e., before branching takes place. In contrast, for 100 randomly chosen samples with big overlap ($|q| \geq 0.9$), 921 ± 34 of the edge variables could be fixed in the first sub problem, about twice as many. Of course, the less variables that can be fixed in the first sub problem, the more overall branching is necessary, resulting in more overall computational effort for samples with small overlap.

A consequence of the broad distribution of the CPU time and of its correlation with the physical observables of interest is that a cutoff in the CPU time produces a systematic error in these quantities. One has, therefore, to ensure that the cutoff is large enough so that the systematic error is smaller than the statistical error.

It is interesting to try to extrapolate the running time needed to deal with larger sizes. The average CPU time in Table IV varies approximately as $\sim \exp(\alpha N_b)$ with α somewhere between 0.0024 and 0.003. Extrapolating to $L = 14$ ($N_b = 7644$), this gives an average CPU time of around $10^{8 \pm 1}$ sec per sample, which is clearly very demanding. Furthermore, memory limitations will set in before we can reach this size. Again, note that n_p increases much more slowly with N_b . The data for $|q| \leq 0.1$ in Fig. 16, for example, vary approximately as $\sim \exp(\alpha N_b)$ with a smaller α around 0.0017, showing that the dominant limiting factor is the solution of the linear programs. Note that the program used for $L = 12$ is significantly faster than that used in this extrapolation. This long extrapolated running time gives us the further motivation to continue our research on the improvement of this algorithm.

VI. CONCLUSIONS

Using an *exact* “branch-and-cut” optimization algorithm, we have studied the large-scale, low-energy excitations in the Ising spin glass in three dimensions with *free* boundary conditions, and compared the results with those obtained earlier by PY for periodic boundary conditions.

In the first part of the paper, we have discussed in detail how the whole set of observables analyzed is fitted by a general scaling picture characterized by two exponents $d - d_s$ and θ' and how the values of these parameters predicted by the various physical pictures proposed for the spin-glass phase fit our data. Our conclusions have been summarized at the end of Sec. III D.

In the second part of the paper, we have analyzed the performance of the branch-and-cut algorithm, finding that the performance is worse when there is a low-energy excited state close in energy to the ground state but far away in configuration space, and have given a quantitative analysis of this effect.

ACKNOWLEDGMENTS

We would like to thank A. J. Bray, G. Parisi, M. Mézard, D. S. Fisher, and M. A. Moore for helpful discussions and correspondence. A.P.Y. acknowledges support from the NSF through Grant No. DMR 0086287. MP would like to thank A. J. Bray for a useful suggestion on the data analysis. We thank the Regional Center of Computing of the University of Cologne for the allocation of computer time. Over the years, Giovanni Rinaldi and Gerd Reinelt contributed much to the algorithm.

- ¹D.S. Fisher and D.A. Huse, *J. Phys. A* **20**, L997 (1987); D.A. Huse and D.S. Fisher, *ibid.* **20**, L1005 (1987); D.S. Fisher and D.A. Huse, *Phys. Rev. B* **38**, 386 (1988).
- ²A.J. Bray and M.A. Moore, in *Heidelberg Colloquium on Glassy Dynamics and Optimization*, edited by L. Van Hemmen and I. Morgenstern (Springer-Verlag, Heidelberg, 1986).
- ³W.L. McMillan, *J. Phys. C* **17**, 3179 (1984).
- ⁴G. Parisi, *Phys. Rev. Lett.* **43**, 1754 (1979); *J. Phys. A* **13**, 1101 (1980); **13**, 1887 (1980); **13**, L115 (1980); *Phys. Rev. Lett.* **50**, 1946 (1983).
- ⁵M. Mézard, G. Parisi, and M.A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
- ⁶K. Binder and A.P. Young, *Rev. Mod. Phys.* **58**, 801 (1986).
- ⁷F. Krzakala and O.C. Martin, *Phys. Rev. Lett.* **85**, 3013 (2000).
- ⁸M. Palassini and A.P. Young, *Phys. Rev. Lett.* **85**, 3017 (2000).
- ⁹E. Marinari and G. Parisi, *Phys. Rev. Lett.* **86**, 3887 (2001).
- ¹⁰A.A. Middleton, *Phys. Rev. B* **63**, 060202(R) (2001).
- ¹¹H.G. Katzgraber, M. Palassini, and A.P. Young, *Phys. Rev. B* **63**, 184422 (2001).
- ¹²C.M. Newman and D.L. Stein, *Phys. Rev. B* **46**, 973 (1992); *Phys. Rev. Lett.* **76**, 515 (1996); *Phys. Rev. E* **57**, 1356 (1998).
- ¹³C.M. Newman and D.L. Stein, *Phys. Rev. Lett.* **87**, 077201 (2001).
- ¹⁴J. Lamarcq, J.-P. Bouchaud, O.C. Martin, and M. Mezard, *Europhys. Lett.* **58**, 321 (2002).
- ¹⁵M. A Moore, cond-mat/0203469 (unpublished).
- ¹⁶M. Jünger, G. Reinelt, and S. Thienel, in *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, edited by W. Cook, L. Lovasz and P. Seymour (American Mathematical Society, Providence, 1995), Vol. 20.
- ¹⁷S. Kobe and A. Hartwig, *Comput. Phys. Commun.* **16**, 1 (1978).
- ¹⁸M. Alava, P. Duxbury, C. Moukarzel, and H. Rieger, in *Phase Transition and Critical Phenomena*, edited by C. Domb and J.L. Lebowitz, *Combinatorial Optimization and Disordered System* Vol. 18 (Academic Press, Cambridge, 2000), pp. 141–317.
- ¹⁹C. De Simone, M. Diehl, M. Jünger, P. Mutzel, G. Reinelt, and G. Rinaldi, *J. Stat. Phys.* **80**, 487 (1995).
- ²⁰F. Liers and M. Jünger, *Int. J. Mod. Phys. C* **11**, 589 (2000).
- ²¹F. Barahona, *J. Phys. A* **15**, 3241 (1982).
- ²²In reality, our implementation of the branch-and-cut algorithm works as follows: If we have found the optimal configuration, i.e., two sets of spins where all spins in one set point up and all spins in the other point down, we assign the value “up” to the set with the smaller cardinality. Therefore, we do not produce a symmetrical distribution. This does not affect the analysis, since we are interested only in the absolute value of the overlap.
- ²³M. Palassini (unpublished).
- ²⁴This average does not coincide with the midpoint of the interval because the perturbation method does not sample uniformly in q .
- ²⁵The statistical errors on the fit parameters are given by the value of the parameter that increases χ^2 by 1 with respect to the best fit (lowest χ^2), when the other parameters are kept fixed. Hence they do not take into account correlations between the various fit parameters, and can underestimate the actual statistical uncertainty. Furthermore, the range of acceptable parameters is usually larger than the condition $\Delta\chi^2=1$, since larger values of χ^2 can still be statistically acceptable.
- ²⁶A.K. Hartmann, *Phys. Rev. E* **59**, 84 (1999); A.J. Bray and M.A. Moore, *J. Phys. C* **17**, L463 (1984); W.L. McMillan, *Phys. Rev. B* **30**, 476 (1984).
- ²⁷E. Marinari, G. Parisi, F. Ricci-Tersenghi, and J.J. Ruiz-Lorenzo, *J. Phys. A* **31**, L481 (1998).
- ²⁸M. Palassini and A.P. Young (unpublished).
- ²⁹H.G. Katzgraber and A.P. Young, *Phys. Rev. B* **65**, 214402 (2002).
- ³⁰*Computational Combinatorial Optimization*, edited by M. Jünger and D. Naddef, *Lecture Notes in Computer Science*, Vol. 2241 (Springer-Verlag, Heidelberg, 2001).
- ³¹F. Barahona, M. Grötschel, M. Jünger, and G. Reinelt, *Oper. Res.* **36**, 493 (1988).
- ³²V. Chvátal, *Linear Programming* (Freeman, New York, 1983).
- ³³The perturbation method does not generate a uniform distribution of q , therefore Fig. 15 was produced by selecting 1000 samples from a random ensemble, such that there is the same number of samples in each consecutive q interval of length 0.1 in the range $q \in [-1, 1]$.