

Phase transition in three-dimensional Heisenberg spin glasses with strong random anisotropies through a multi-GPU parallelization

M. Baity-Jesi,^{1,2,3,*} L. A. Fernández,^{1,2} V. Martín-Mayor,^{1,2} and J. M. Sanz¹

¹*Departamento de Física Teórica I, Universidad Complutense, 28040 Madrid, Spain*

²*Instituto de Biocomputación and Física de Sistemas Complejos (BIFI), 50009 Zaragoza, Spain*

³*Dipartimento di Fisica, Università La Sapienza, 00185 Roma, Italy*

(Received 6 September 2013; revised manuscript received 18 December 2013; published 21 January 2014)

We characterize the phase diagram of anisotropic Heisenberg spin glasses, finding both the spin and the chiral glass transition. We remark on the presence of strong finite-size effects in the chiral sector. On the spin glass sector, we find that the universality class is that of Ising spin glasses. Our data are compatible with a unique phase transition for the chiral and spin glass sector. We focus on keeping finite-size effects under control, and we stress that they are important to understand experiments. Thanks to large GPU clusters we have been able to thermalize cubic lattices with up to 64^3 spins, over a vast range of temperatures (hence, of relaxation times).

DOI: [10.1103/PhysRevB.89.014202](https://doi.org/10.1103/PhysRevB.89.014202)

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

I. INTRODUCTION

Spin glasses (SGs) are disordered magnetic alloys [1–3]. Their microscopic modelization includes several interactions, such as the RKKY interaction that is invariant over rotations [4–6], and the Dzyaloshinsky-Moriya (DM) interaction that breaks the rotational symmetry [7,8]. Therefore, in theoretical physics SGs are often studied with simplified models that take into account only a few essential characteristics (in particular, quenched disorder and symmetries) [9].

The DM interaction, through a spin-orbit coupling with a third spin, causes the interactions between spins in any SG to have a certain degree of random anisotropy. This implies that real SGs are never fully isotropic (this theoretical limit is named Heisenberg SG). In fact, materials are classified according to the degree of anisotropy in their interactions [10], which turns out to be relevant in their nonequilibrium magnetic response [11]. On one end of the materials' spectrum we find the extremely anisotropic $\text{Fe}_{0.5}\text{Mn}_{0.5}\text{TiO}_3$, which is maybe the best realization of the ideal limit of an Ising SG (Ising SGs correspond to the idealization of uniaxial spins). On the other end, we have very isotropic alloys such as AgMn or CuMn (whose modelization is notoriously difficult [12], due to the presence of short range spin-density wave ordering [13–15]).

Despite the variety of interactions, already in the early 1990s there was general experimental agreement on that SGs undergo a phase transition at sufficiently low temperature [16–18].

On the other hand, theoretical work was less advanced, even though one works with extremely simple models. For the Ising SG there were arguments supporting the existence of a phase transition [19] that were later confirmed numerically [20,21]. In the Heisenberg case, instead, all the attempts carried out during the 1980s and 1990s failed in finding a phase transition at a finite temperature $T_{SG} > 0$ [22–25]. In fact, Matsubara *et al.* argued in 1991 that once a small anisotropic term is added to the Heisenberg Hamiltonian the phase transition becomes visible [25]. This was in agreement with a later domain-wall computation [26]. The accepted picture at the time was that

the lower critical dimension (i.e., the spatial dimension below which there is no phase transition) lies somewhere between three dimensions (3D) and four dimensions (4D) [27].

However, the story was slightly more complicated. Villain and co-workers made a provocative suggestion hypothesizing that, although maybe there was no spin glass transition, a different order parameter called chirality (or vorticity) could be critical [28]. Chirality is a scalar observable that describes vorticity and alignment between neighboring spins (see below the precise definition in Sec. II B). This idea was elaborated by Kawamura in his 1992 *spin-chirality decoupling scenario*: In the ideal case of a purely isotropic system the spin and chiral glass order parameters would be decoupled, but the introduction of any small anisotropy would couple them [29].

Kawamura's scenario was apparently consistent with all the observations until 2003, when Lee and Young employed more efficient simulation algorithms and finite-size scaling techniques to show that the spin glass channel is critical also in the fully isotropic model (i.e., the Heisenberg limit) [30]. Both order parameters seemed to become positive at the same temperature. Further simulations confirmed the existence of a spin glass phase transition, although uncertainty remains on whether the transition is unique [31,32] or chiralities order at a slightly higher temperature T_{CG} [33].

A parallel issue is measuring the chiral order parameter in experiments. Kawamura proposed in 2003 that the extraordinary Hall resistivity is a simple function of the linear and nonlinear chiral-glass (CG) susceptibilities [34]. Experiments based on this proposal observed the chiral transition and measured, for instance, the critical exponent δ [35]. Interestingly enough, the value of δ turned out to be compatible between spin and chiral glass sector. Nonetheless, it was impossible to identify a universality class despite the critical exponents of these systems which had been extensively measured (at least in the SG sector) [10,16,17]: The impression was that they change in a continuous way from the Heisenberg to the Ising limit [36], as we increased the anisotropy.

However, analogy with ferromagnetic materials suggests a different interpretation. Anisotropy would be a relevant parameter in the sense of the renormalization group [37]. There should be a new dominant fixed point, and symmetry considerations lead us to think it should belong to the

*marcobaitjesi@fis.ucm.es

Ising-Edwards-Anderson universality class. Yet, when we add a relevant parameter to the Hamiltonian, there should be some *crossover* effects. In other words, one expects that while the correlation length ξ is small, the critical exponents are closer to the Heisenberg-Edwards-Anderson universality class, and that only for large enough ξ the universality class reveals its nature.

Notwithstanding, it is very hard, both numerically and experimentally, to prepare a SG with a large correlation length, since one should wait very long times (it has been argued that the waiting time t_w required to reach a certain coherence length is proportional to almost its seventh power, see, e.g., Refs. [38] and [39]). To our knowledge, for this reason, the largest measured correlation lengths are of the order of only 100 lattice spacings [11,39]. That is a rather small distance to reveal the true universality class, so it is plausible that experiments will find critical exponents between the two universality classes.

To further complicate things, in experiments one has to take into account at least two relevant crossovers. The first is the competition that we just pointed out between the isotropic and the anisotropic fixed points. It is the one we treat in this paper. The second crossover that we will not address is about short versus long range interactions. In fact, the Hamiltonian we treat is short range, but the DM interaction has been shown to be quasi-long range, in the sense that the interactions are long range, but only until a cut-off distance of the order of some tens of atomic spacings [40].

Aiming to untangle these questions, one of the authors undertook a numerical study of Heisenberg SGs with very weak random anisotropies [41], but the scenario remained even more foggy, since it was observed that

(1) The chiral glass critical temperature T_{CG} was significantly higher than T_{SG} , in disagreement with experiments and expectations.

(2) Apparently, the chiral susceptibility was *not* divergent at T_{CG} . This is surprising and, apparently, in contrast with experiments [35]. Technically, this lacking divergence appeared as a very large anomalous dimension $\eta_{CG} \sim 2$ [42].

(3) Introducing very weak anisotropies changed dramatically T_{SG} . For example, the T_{SG} found by comparing systems of size $L = 6, 12$ was about twice its equivalent on the fully isotropic model. This is surprising, since one expects that the critical temperature would change very little from the isotropic case when D is as small as in Ref. [41].

In this paper we will focus on the uniqueness of the phase transition and on the universality class, proposing that there is a unique transition, belonging to the Ising-Edwards-Anderson (IEA) universality class [9]. We will also give an interpretation to the results of Ref. [41], showing that the apparent inconsistencies are due to scaling corrections that we will try to characterize, since we believe them to be fundamental both in the interpretation of numerical simulations and of experiments.

To do all this, we will study numerically the Heisenberg spin glass model with strong random anisotropies in order to suppress both finite-size effects and traces of the crossover from the isotropic limit.

We simulated on the largest lattices to present (up to $L = 64$), over a wide temperature range [43]. This has been possible thanks to an intense use of graphic accelerators (GPUs) for the computations. We made use of the *Tianhe-1A* GPU cluster

in Tianjin, China [44], and of the *Minotauro* GPU cluster in Barcelona [45].

The rest of the paper is organized as follows. In Sec. II we give an explicit definition of the model, and we introduce the observables we extracted from simulations and analysis. Section III contains details on how we practically conducted the simulations, although much information is relegated to Appendix, where we also discuss the use of GPUs for spin glasses. In Sec. IV we recall some finite-size scaling concepts we used in our analysis, to find the critical temperatures and exponents (some technical details are given in Sec. V). Finally, in Sec. VI we refer to the results obtained in this work, and give our conclusions in Sec. VII.

II. MODEL AND OBSERVABLES

A. The model and its symmetries

We study the model introduced by Matsubara *et al.* [25], which is particularly convenient because of its simplicity. We consider $N = L^3$ three-dimensional unitary vectors $\vec{s}_x = (s_x^1, s_x^2, s_x^3)$ on a cubic lattice of linear size L , with periodic boundary conditions. The Hamiltonian is

$$H = - \sum_{\langle x,y \rangle} \left(J_{xy} \vec{s}_x \cdot \vec{s}_y + \sum_{\alpha\beta} s_x^\alpha D_{xy}^{\alpha\beta} s_y^\beta \right), \quad (1)$$

where $\langle \cdot \rangle$ means the sum goes only over nearest neighbors, and the indices α, β indicate the component of the spins. J_{xy} is the isotropic coupling between sites x and y . D_{xy} is the anisotropy operator: a 3×3 symmetric matrix, where the six matrix elements $D_{xy}^{\alpha\beta}, \alpha \geq \beta$, are independent random variables.

There is quenched disorder, this means that the time scales of the couplings $\{J_{xy}, D_{xy}\}$ are infinitely larger than those of our dynamic variables, so we represent them as constant in time random variables, with $\overline{J_{xy}} = \overline{D_{xy}^{\alpha\beta}} = 0$, $\overline{J_{xy}^2} = 1$, and $\overline{(D_{xy}^{\alpha\beta})^2} = D^2$. The overline $\overline{\dots}$ denotes the averages over the instances of the disorder, while for thermal averages we will use $\langle \dots \rangle$. Each different realization of the couplings $\{J_{xy}, D_{xy}\}$ is called a *sample*. Independent systems with the same couplings are *replicas* of the same sample. We use two replicas per sample.

Notice that if all the matrix elements $D_{xy}^{\alpha\beta}$ are zero we recover the fully isotropic Heisenberg model, with $O(3)$ symmetry. However, if the $D_{xy}^{\alpha\beta}$ are nonvanishing, the only remaining symmetry is time reversal: $\vec{s}_x \rightarrow -\vec{s}_x$ for all the spins in the lattice. Time reversal is an instance of the Z_2 symmetry. This is the symmetry group of the IEA model [9]. Hence we expect that the Z_2 symmetry will be spontaneously broken in a unique phase transition belonging to the IEA universality class (see, e.g., Ref. [26]). Of course, underlying this expectation is the assumption that the anisotropic coupling is a relevant perturbation in the renormalization group sense (as it is the case in ferromagnets [37]). In fact, the infinite-anisotropy limit can be explicitly worked out for a problem with *site* anisotropy [rather than link anisotropy as in Eq. (1)]: one finds an IEA-like behavior [46,47].

It is widely accepted that the universality class does not change with the probability distribution of the couplings [48]. We take advantage of this, and choose a bimodal distribution

for J_{xy} and $D_{xy}^{\alpha\beta}$, $J_{xy} = \pm 1$ and $D_{xy}^{\alpha\beta} = \pm D$. These couplings can be stored in a single bit, which is important because we are using GPUs, special hardware devices where memory read/write should be minimized (Appendix).

We chose the two different values $D = 0.5, 1$. We want to compare our results with those in Ref. [41], where simulations were done on samples with weak random anisotropies. In that work the $D_{xy}^{\alpha\beta}$ did not follow a bimodal distribution, but were uniformly distributed between -0.05 and 0.05 . To make proper comparisons we consider the standard deviation of the distribution. For bimodal distributions it is exactly D , in Ref. [41] it is $(D^2)^{1/2} = 1/\sqrt{1200} \simeq 0.03$.

B. The observables

To define the SG and CG order parameters we use two replicas. The overlap field is $q_x = \vec{s}_x^a \cdot \vec{s}_x^b$, where a and b are replica indices. Its Fourier transform at wave vector \mathbf{k} is $\hat{q}_{\text{SG}}(\mathbf{k}) = \sum_x q_x e^{i\mathbf{k}\cdot\mathbf{x}}/N$.

The chirality represents the oriented volume of the parallelepiped we can construct on three consecutive spins:

$$\zeta_{x,\mu} = \vec{s}_{x+e_\mu} (\vec{s}_x \times \vec{s}_{x-e_\mu}), \quad \mu = 1, 2, 3, \quad (2)$$

where e_μ is the unitary vector in the μ direction. The CG overlap is defined similarly to the SG one, as $\kappa_{x,\mu} = \zeta_{x,\mu}^a \zeta_{x,\mu}^b$. Again a and b indicate the replica. The Fourier transform of the CG overlap field is $\hat{q}_{\text{CG}}^\mu(\mathbf{k}) = \sum_x \kappa_x e^{i\mathbf{k}\cdot\mathbf{x}}/N$.

We define the wave-vector dependent susceptibilities on the two overlap fields as

$$\chi_{\text{SG}} = N \overline{|\hat{q}_{\text{SG}}(\mathbf{k})|^2}, \quad \chi_{\text{CG}} = N \overline{|\hat{q}_{\text{CG}}(\mathbf{k})|^2}, \quad (3)$$

and from each of them we can compute the correlation length of the related field [37]

$$\xi = \frac{1}{2 \sin(k_{\min}/2)} \sqrt{\frac{\chi(0)}{\chi(\mathbf{k}_{\min})} - 1}, \quad (4)$$

being $\mathbf{k}_{\min} = (2\pi/L, 0, 0)$ or permutations. When computing ξ_{CG} , one can choose μ parallel or orthogonal to the wave vector \mathbf{k}_{\min} . As it was already observed in Ref. [31], there is no apparent difference between the two options, so we averaged over all the values of μ to enhance our statistics.

III. SIMULATION DETAILS AND EQUILIBRATION

We used Monte Carlo dynamics throughout all the work. Previous experience advises us to mix several Monte Carlo dynamics [32,57,58]. In fact, our single Monte Carlo step (MCS) consisted of (in successive order): (i) one full lattice sweep with the heat-bath algorithm, (ii) L lattice sweeps of microcanonical over-relaxation algorithm [59], and (iii) one single parallel tempering sweep [60,61]. The combination of the first two, which update one spin at a time, has been shown to be effective in the case of isotropic SGs [75] and other models with frustration [62,63]. Both heat-bath and over-relaxation are directly generalized to the anisotropic case [64].

All the simulations were run on NVIDIA Tesla GPUs. Except $L = 64$, $D = 0.5$, where we parallelized 45 GPUs, each sample was simulated on a single GPU. The interested

TABLE I. Details of the simulations. We show the simulation parameters for each anisotropy D , and lattice size L . N_{samples} is the number of simulated samples. N_T is the number of temperatures that were used in parallel tempering. The temperatures followed a geometric sequence between T_{\min} and T_{\max} , and N_T was chosen so that the parallel tempering's acceptance was around 15%. N_{MCS}^{\min} is the minimum number of MCS for each simulation. The simulation for $L = 64$, $D = 1$ was intended only to locate T_{CG} .

D	L	N_{samples}	N_{MCS}^{\min}	N_T	T_{\min}	T_{\max}
0.5	8	377	2.048×10^4	10	0.588	0.8
0.5	16	377	4.096×10^4	28	0.588	0.8
0.5	32	377	3.28×10^5	45	0.583	0.8
0.5	64	185	4×10^5	45	0.621	0.709
1	8	1024	2.048×10^4	10	0.877	1.28
1	12	716	1.68×10^5	20	0.893	1.28
1	16	1024	4.096×10^4	28	0.877	1.28
1	24	716	1.68×10^5	40	0.900	1.28
1	32	1024	3.28×10^5	45	0.917	1.28
1	64	54	3.44×10^5	45	1.0	1.16009

reader can find in Appendix details on how they were performed.

Table I depicts the relevant simulation parameters. For given L and D , the simulations were all equally long, except for $L = 64$, $D = 0.5$, where we extended the simulation of the samples with the longest relaxation times.

To ensure thermalization we made a *logarithmic data binning*. Each bin had twice the length of the previous, i.e., it contained two times more Monte Carlo steps (MCS), and had twice the measures. More explicitly, let us call i_f the last bin: i_f contains the last half of the Monte Carlo time series, $i_f - 1$ the second quarter, $i_f - 2$ the second octave, and so on. This allowed us to create a sequence of values $\langle O_n(i) \rangle$, for every observable O , where n indicates the sample, and i identifies the bin, that has length 2^i MCS. A set of samples was considered thermalized if $\langle O_n(i) \rangle - \langle O_n(i_f) \rangle$ converged to zero. This test is stricter than merely requesting the convergence of the sequence of $\langle O_n(i) \rangle$, because neighboring blocks are statistically correlated, so the fluctuation of their difference is smaller [65]. Physical results were taken only from the last block.

Since the $L = 64, D = 0.5$ samples were the most GPU consuming, we were more strict with them. To ensure and monitor thermalization, beyond the previous criteria, we measured the integrated autocorrelation time (mixing time) of the random walk in temperatures of each sample [31]. In a thermalized sample, all the replicas stay a significant amount of time at each temperature. We made sure that all the simulations were longer than 10 times this autocorrelation time. The sample-to-sample fluctuations were not extreme, and the autocorrelation times τ spanned between 10 000 to 50 000 MCS, depending on the sample. Finally, we decided to take measures only over the last 64 000 MCS of each simulation.

IV. FINITE-SIZE SCALING

Our simulations were far from the thermodynamic limit, therefore in our analysis we had to take into account finite-size

effects. Finite-size scaling (FSS) consists of comparing results at different lattice size to characterize the critical point. Specifically, we shall be employing phenomenological renormalization, also known as the quotients method [37,66,67].

Since FSS applies irrespectively of the considered order parameter, in the current section we will not distinguish between spin and chiral sector. The generic critical temperature will be called T_c .

If an observable O diverges at the critical temperature as $O \propto |T - T_c|^{x_o}$, then its thermal average close to the critical point can be expressed like

$$\begin{aligned} \langle O(L, T) \rangle = & L^{x_o/\nu} \{ f_o [L^{1/\nu} (T - T_c)] \\ & + L^{-\omega} g_o [L^{1/\nu} (T - T_c)] \\ & + L^{-2\omega} h_o [L^{1/\nu} (T - T_c)] + \dots \}, \end{aligned} \quad (5)$$

where f_o , g_o , and h_o are analytic scaling functions for observable O , while ν is the thermal critical exponent. The exponent $\omega > 0$ is universal, and it expresses the corrections to scaling. The lower dots stand for subleading corrections to scaling. Let us name $\xi_L(T)$ the correlation length in a lattice of finite size L , at temperature T . The case $O = \xi_L(T)/L$ is of special interest, since ν is the critical exponent for the correlation length. Then, Eq. (5) becomes in this case, up to the leading order,

$$\frac{\xi_L}{L} = f_\xi [L^{1/\nu} (T - T_c)] + \dots \quad (6)$$

Therefore we can identify T_c as the temperature where the curves $\xi_L(T)/L$ cross for all L for sufficiently large L . If we let $T^{L,2L}$ be the temperature where $\xi_L(T)/L$ crosses $\xi_{2L}(T)/(2L)$, this regime is reached once the $T^{L,2L}$ has converged. Yet, if ω is small, our lattice sizes may not be large enough, so we will have to take into account the aforementioned corrections to scaling. Including corrections to the order $L^{-2\omega}$, the approach of the crossing temperature $T^{L,2L}$ to the asymptotic value T_c can be written as

$$T^{L,2L} - T_c = AL^{-(\omega+1/\nu)} + BL^{-(2\omega+1/\nu)} + \dots, \quad (7)$$

where A and B are nonuniversal scaling amplitudes.

To compute the critical exponents ν and η we use the quotients' method, taking the quotient of the same observable between different lattice sizes L and $2L$. At the temperature $T^{L,2L}$ we get

$$\frac{\langle O_{2L}(T^{L,2L}) \rangle_J}{\langle O_L(T^{L,2L}) \rangle_J} = 2^{x_o/\nu} + A_{x_o} L^{-\omega} + \dots \quad (8)$$

Again, A_{x_o} is a nonuniversal amplitude, while the dots stand for subleading corrections to scaling. Therefore, if O is the thermal derivative of ξ , we can compute the ν critical exponent through the relation

$$\frac{d\xi_{2L}(T^{L,2L})/dT}{d\xi_L(T^{L,2L})/dT} = 2^{1+1/\nu} + A_\nu L^{-\omega} + \dots \quad (9)$$

To calculate η we use the susceptibility, as $\chi \propto |T - T_c|^{-\gamma}$ and $2 - \eta = \gamma/\nu$, hence

$$\frac{\chi_{2L}(T^{L,2L})}{\chi_L(T^{L,2L})} = 2^{2-\eta} + A_\eta L^{-\omega} + \dots \quad (10)$$

Note that the value of ξ_L/L at the crossing tends as well to a universal quantity:

$$\frac{\xi_L}{L} \Big|_{T^{L,2L}} = \frac{\xi^*}{L} \Big|_{L=\infty} + A_\xi L^{-\omega} + \dots \quad (11)$$

V. INTERPOLATIONS, EXTRAPOLATIONS, AND ERRORS

We have been able to estimate the critical temperature from the crossing of the curves ξ/L at L and $2L$, and the exponents ν and η with the method of the quotients, as described in Sec. IV.

To identify the crossing point between the pairs of curves, we used low-order polynomial fits: For each lattice size, we took the four temperatures in the parallel tempering nearest to the crossing point. We fitted these four data points to a linear or quadratic function of the temperature. The obtained results were compatible within one standard deviation (the values reported in this work come from the linear interpolation). In order to calculate ν we needed the derivative of the correlation length at the crossing point. We extracted it by taking the derivative of the polynomial interpolations.

However, there is a difficulty in the calculation of statistical errors: The fits we had to perform came from strongly correlated data (because of the parallel-tempering temperature swap). Therefore, to get a proper estimate of the error, we made Jack-knife blocks, fitted separately each block, and calculated the Jack-knife error [37].

The whole mentioned procedure was fluid while $T_{SG}^{L,2L}$ fell in our simulated temperature span. Yet, since $T_{SG}^{L,2L}$ was fairly lower than $T_{CG}^{L,2L}$, it occurred in four cases that we did not reach low enough temperatures in our simulations to be able to interpolate the crossing, and we had to recur to extrapolations. This happened with $D = 1$, $T_{SG}^{32,64}$ and $T_{CG}^{32,64}$, and in the lower anisotropy $D = 0.5$, with $T_{SG}^{16,32}$ and $T_{SG}^{32,64}$.

The case of $T_{SG}^{32,64}$ ($D = 1$) and $T_{SG}^{16,32}$ ($D = 0.5$) was not a great issue, because the crossing point was very near to the lowest simulated temperature, so we treated these crossings just like the others.

In the case of $T_{SG}^{32,64}$ ($D = 0.5$), instead, we had to extrapolate at a long distance [see Fig. 1 (top)]. Again, we performed the extrapolation through linear in temperature fits. To make the fit of $L = 64$ more stable, we took into account a progressive number of points (i.e., we fitted to the n lowest temperatures). We increased the number of temperatures, while the crossing temperature was constant. Note that increasing the number of temperatures in the fit results in a smaller statistical error for the crossing temperature. However, $\xi_L(T)/L$ is not a linear function at high T (see Fig. 1). Therefore a tradeoff is needed because when too high temperatures were included in the fit, the crossing temperature started to change, and we knew that curvature effects were biasing it. Our final extrapolation was obtained from a fit performed on the 10 lowest-temperature points. Unfortunately, this approach was not feasible for the SG susceptibility due to its strongly nonlinear behavior. Hence, in the next section we will not give an estimate for $\eta_{SG}(L = 64)$.

In the case of $T_{SG}^{32,64}$ ($D = 1$), the simulation was not devised to reach that crossing point, and we did not extrapolate data.

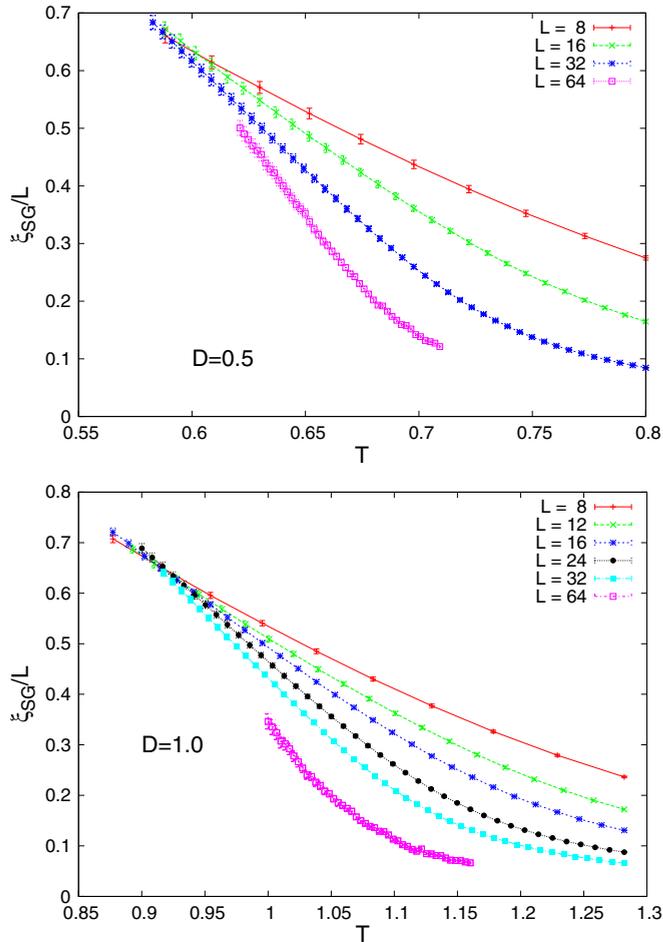


FIG. 1. (Color online) Spin glass correlation length in units of the linear lattice size L for $D = 0.5$ (top) and $D = 1$ (bottom). All the curves cross at about the same temperature for both anisotropies [see Eq. (7)]. The data for $D = 1$, $L = 64$, shown here for the sake of completeness, were only used for the chiral sector.

TABLE II. Determination of the critical quantities for the SG sector. For each anisotropy D , and each pair of lattices $(L, 2L)$, we obtain effective size-dependent estimates for T_{SG} , and the universal quantities ν_{SG} , η_{SG} , and $\xi_L(T_{SG})/L$. The thermodynamic limit, indicated with $L = \infty$, is obtained by means of fits to Eqs. (7), (9), (10), and (11). Exponent ω was not a fitting parameter (we took $\omega_{IEA} = 1.0(1)$ from Ref. [49], see text and Ref. [68]). The line immediately after the extrapolations displays the estimator of the χ^2 figure of merit of each one. $D = IEA$ represents the critical values of the Ising-Edwards-Anderson universality class, taken from Ref. [49]. The numbers in square brackets express the systematic error due to the uncertainty of ω_{IEA} .

D	$(L, 2L)$	T_{SG}	ν_{SG}	η_{SG}	$\xi_{SG}(T_{SG})/L$
0.5	(8,16)	0.602(18)	1.91(27)	-0.388(27)	0.629(48)
0.5	(16,32)	0.577(22)	2.70(63)	-0.449(67)	0.705(76)
0.5	(32,64)	0.596(14)	2.18(45)	-	0.631(56)
0.5	∞	0.591(16)[0]	2.71(82)[3]	-	0.637(87)[1]
	χ^2/DOF	0.55/1	0.47/1	-	0.56/1
1.0	(8,16)	0.910(21)	2.38(25)	-0.410(44)	0.660(34)
1.0	(12,24)	0.927(19)	2.32(28)	-0.370(53)	0.629(36)
1.0	(16,32)	0.910(16)	2.37(28)	-0.400(19)	0.660(35)
1.0	∞	0.917(32)[0]	2.33(67)[0]	-0.391(71)[1]	0.662(83)[0]
	χ^2/DOF	0.66/1	0.030/1	0.37/1	0.55/1
IEA	∞		2.45(15)	-0.375(10)	0.645(15)

VI. RESULTS

A. Spin glass transition

Figure 1 shows the crossings of $\xi_{SG}(T)/L$ for $D = 0.5, 1$. Table II contains the principal results on the SG sector, providing a quantitative description of those figures. As explained in Sec. II A, we expect that the transition belongs to the Ising-Edwards-Anderson (IEA) universality class. This conjecture is supported by the fact that the critical exponents ν_{SG} and η_{SG} , and the height at which the $\xi_{SG}(T)/L$ cross, are compatible with those of the IEA spin glass, indicated in the last line of Table II.

Hence, it is reasonable to extrapolate our results to $L \rightarrow \infty$ by assuming the IEA universality class. We took $\omega_{IEA} = 1.0(1)$ from Ref. [49], and fitted to Eqs. (9), (10), and (11). In those fits we took into account both the anticorrelation in the data [70], and the bias arising from the indetermination of the exponent ω_{IEA} . Notice, from Table II, that the dependence on L of the data is so weak, that this bias is practically negligible. This situation is different from the one encountered in Ref. [41], where the anisotropy fields were extremely small ($D \simeq 0.03$) [71]. There the finite-size effects in the SG sector were huge.

Overall, the strong consistency of our extrapolations to large L with the IEA exponents shows *a posteriori* that our assumption was proper.

B. Chiral glass transition

In the CG channel (Fig. 2 and Table III) the interpretation is slightly more controversial, since finite-size effects are heavy. For the smaller lattice sizes, T_{CG} is consistently larger than T_{SG} , and ν_{CG} is incompatible with the IEA limit. On the other side, when L is larger, T_{CG} approaches noticeably its SG counterpart, and so does ν_{CG} . We notice that η_{CG} marks the distinction between these two regimes. In fact, when L is small, it is very close to 2. This means that the divergence of χ_{CG} is extremely slow ($\chi \sim L^{2-\eta}$) [42], revealing we are still far from the asymptotic limit. When L is larger, η_{CG} is consistently smaller, the divergence of χ_{CG} is less suppressed,

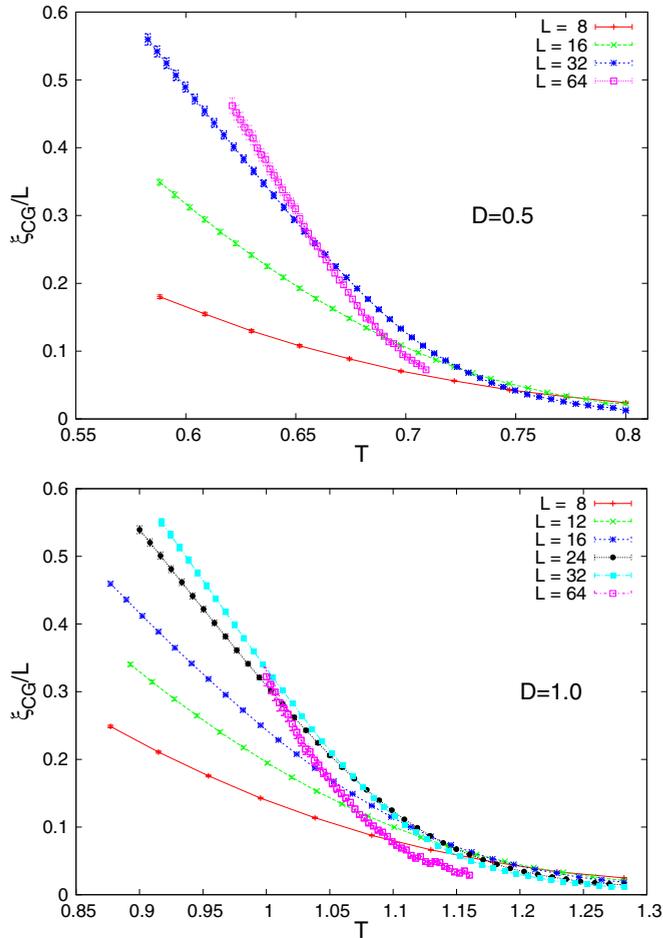


FIG. 2. (Color online) Chiral glass correlation length in units of the lattice size for $D = 0.5$ (top) and $D = 1$ (bottom). When L grows, the crossing temperature shifts significantly towards left.

and we can assume the asymptotic behavior is starting to show up. Consistently with this observation, the value of ξ_{CG}/L at the crossing temperature becomes sizable (indeed, the second-moment correlation length (4) is well defined only if $\eta < 2$, see, e.g., Ref. [37]).

C. Uniqueness of the transition

Although the SG and CG transitions do not coincide yet with our values of L and D , the critical temperatures, as well as ν , become more and more similar as the linear size

TABLE III. Determination of the critical quantities for the CG sector. Same as Table II, but for chirality. In this case the corrections to scaling are significant.

D	$(L, 2L)$	T_{CG}	ν_{CG}	η_{CG}	$\xi_{CG}(T_{CG})/L$
0.5	(8,16)	0.7762(43)	1.45(22)	1.9778(23)	0.0321(22)
0.5	(16,32)	0.7255(29)	1.78(14)	1.8416(98)	0.0735(41)
0.5	(32,64)	0.659(47)	2.40(47)	0.823(68)	0.258(18)
1.0	(8,16)	1.2031(33)	1.205(71)	1.9507(27)	0.0418(12)
1.0	(12,24)	1.1472(40)	1.72(11)	1.8664(51)	0.0691(25)
1.0	(16,32)	1.1046(38)	2.18(10)	1.6995(75)	0.1098(42)
1.0	(32,64)	0.987(22)	2.48(84)	0.53(19)	0.368(58)

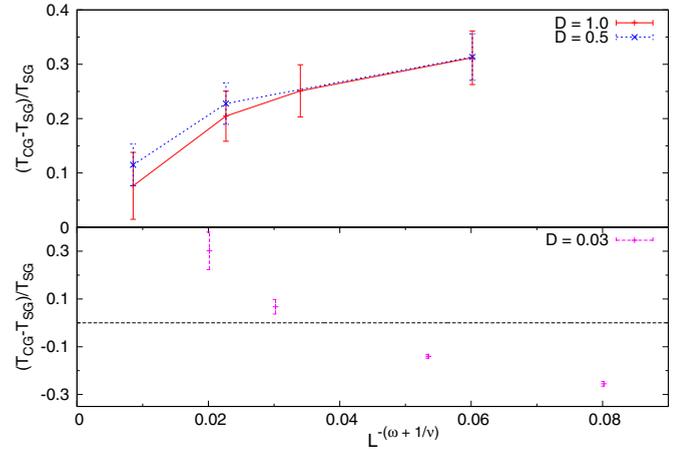


FIG. 3. (Color online) Difference between the chiral crossing T_{CG} and the spin glass transition temperature T_{SG}^{∞} , in units of T_{SG}^{∞} (see Table II for the extrapolations of T_{SG}^{∞}). The exponents ω_{IEA} and ν_{IEA} are taken from Ref. [49]. In the upper plot we represent our data, for $D = 0.5, 1$. The two transitions get closer when we increase L , and the approach appears faster when the lattice size increases. Notice that a linear interpolation between the two largest lattice sizes intercepts the y axis compatibly with a coupling between the two transitions (i.e. $T_{SG} = T_{CG}$). On the bottom plot we show data from Ref. [41], where much lower anisotropies were considered. Here the scenario is completely different, since the critical temperatures drift apart for large enough L . The horizontal dashed line corresponds to $T_{CG} - T_{SG} = 0$.

of the system increases. Moreover, the decrease of η_{CG} as a function of L has not yet stabilized, so it is likely that the chiral quantities will keep changing with bigger lattice sizes.

As explained in Sec. II A, we expect that the transition should belong to the IEA universality class. To confirm this expectation, we make the ansatz of a unique transition, of the IEA universality class, to seek if the two critical temperatures join for $L \rightarrow \infty$. Figure 3 (upper half) shows the difference between the critical temperatures as a function of the natural scale for first-order corrections to scaling $L^{-(\omega_{IEA}+1/\nu_{IEA})}$ [Eq. (7)]. Again, ω_{IEA} and ν_{IEA} are taken from Ref. [49]. Not only Fig. 3 (top) reveals a marked increase of the speed of the convergence for $L = 64$ (to which corresponds the smallest anomalous exponent η_{CG}), but also, a linear interpolation to infinite volume, taking that point and the previous, extrapolates $T_{SG} = T_{CG}$ within the error.

Figure 4 shows how the SG and CG critical temperatures approach each other with L . Again, T_{CG} gets closer to T_{SG} , and the speed of the approach increases with the lattice size. The points in the intercept represent extrapolations to the thermodynamic limit of the T_{SG} . Since the observations are compatible with the ansatz of a unique phase transition, belonging to the IEA universality class, we used the infinite-size limit of T_{SG} to plot the model's phase diagram (Fig. 4, inset) [72].

D. Comparing with weak anisotropies

Both plots of Fig. 3 show the same observable, for different anisotropies. The top plot depicts our data, in the case of strong anisotropies $D = 0.5, 1$. The bottom one represents the case

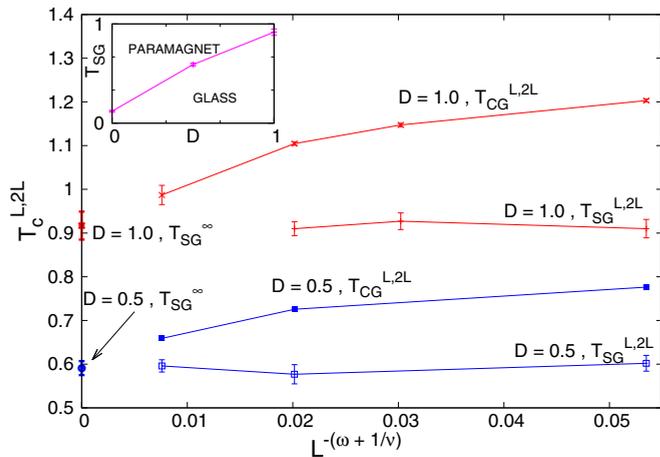


FIG. 4. (Color online) Crossing temperatures as a function of $L^{-(\omega_{IEA} + 1/\nu_{IEA})}$ (large plot). The points on the intercept are the $L \rightarrow \infty$ extrapolations from Table II. The inset shows the phase diagram of the model with these same points, as the most economic interpretation of our data is that in the thermodynamic limit $T_{SG} = T_{CG}$. The $D = 0$ point is borrowed from Ref. [31].

of weak anisotropies ($D \simeq 0.03$) [71], coming from Ref. [41]. The behavior is very different between the two cases. For strong anisotropies, the critical temperatures tend to meet as we increase L . That is qualitatively very different from the weak anisotropy case, where their distance increases. We can ask ourselves where this qualitative difference of behavior comes from.

If we compare same system sizes and different D in Table III, we notice that finite-size effects are larger (and η closer to two) the smaller the anisotropy. These differences in the finite-size effects are appreciable with a factor 2 change in the anisotropy (from $D = 1$ to $D = 0.5$), so it is reasonable that suppressing the anisotropy by a factor 17 or 35 will increase drastically the finite-size effects.

The most economic explanation is then that there is a nonasymptotic effect that disappears with much larger systems or, as we have seen, with larger anisotropies. In other words there is a $L^*(D)$ after which T_{SG} and T_{CG} start joining. For $D \simeq 0.03$, L^* is so large that we observe a growing $T_{CG} - T_{SG}$, while for $D \geq 0.5$ we find $L^* < 8$.

Another peculiarity outgoing from Ref. [41] arises from the SG transition alone. It had been observed that a very weak perturbation on the symmetry of the isotropic system implied huge changes in the critical temperature, while one would expect that the transition line is smooth.

To solve this dilemma, we take advantage of having strong evidence for the universality class of the transition. So, we take the data from Ref. [41], and use once again the exponents ν_{IEA} and ω_{IEA} in Ref. [49] to extrapolate the infinite volume limit with second-order corrections to scaling [Eq. (7)]. The fit is good ($\chi^2/\text{DOF} = 0.70/1$), and, as we show in Fig. 5, its $L \rightarrow \infty$ extrapolation for the critical temperature is compatible with $T_{SG}(D = 0)$ within one standard deviation. Thus, taming the finite-size effects was enough to make the scenario consistent, and the issue reduces to the fact that finite-size effects are extremely strong when the anisotropy is smaller.

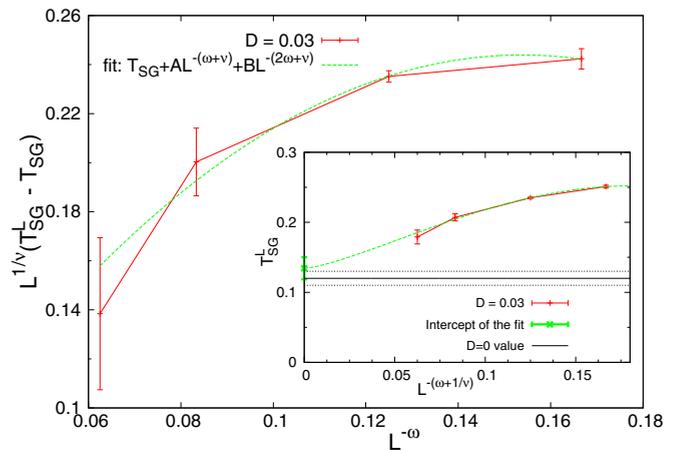


FIG. 5. (Color online) Data from Ref. [41], corresponding to $D \simeq 0.03$ [71], with extrapolations to the thermodynamic limit assuming the Ising-Edwards-Anderson universality class. The data are the same in both plots. The dashed line is a fit of the scaling in L , considering corrections up to the second order [Eq. (7)]. The large figure displays the trend of the scaling variable $L^{1/\nu}(T - T_{SG})$ as a function of $L^{-\omega}$. The inset shows the same data set, plotting $T_{SG}^{L,2L}$ as a function of $L^{-\omega-1/\nu}$, see Eq. (7). The extrapolation to large L (the point in the intercept) is compared with T_{SG} of $D = 0$ from Ref. [31]. The full horizontal line is the central value of $T_{SG}^{D=0}$, and the dashed lines define the error.

VII. CONCLUSIONS

We performed a numerical study of the critical behavior of Heisenberg spin glasses with strong bimodal anisotropies. Our aim was to clarify the role of scaling corrections, as well as the crossover effects between the Heisenberg and Ising universality classes, to be expected when the anisotropic interactions are present. In fact, we show that anisotropic interactions are a relevant perturbation in the renormalization group sense: No matter how small the anisotropy, the asymptotic critical exponents are those of the Ising-Edwards-Anderson model. However, a fairly large correlation length maybe needed to reach the asymptotic regime. This observation is relevant for the interpretation of both numerical simulations [41], and experiments [10].

It is then clear that large system sizes are needed to make progress, something that calls for extraordinary simulation methods. Therefore, we performed single-GPU and multi-GPU simulations to thermalize lattices up to $L = 64$ at low temperatures. As a side benefit, our work provides a proof-of-concept for GPU and multi-GPU massive simulation of spin glasses with continuous degrees of freedom. This topic is elaborated further in the Appendix.

We performed a finite-size scaling analysis based on phenomenological renormalization [66,67]. We imposed scale invariance on the second-moment correlation length in units of the system size ξ_L/L . We followed this approach for both the chiral and spin glass order parameters.

Our results for the spin glass sector were crystal clear: All the indicators of the universality class were compatible with their counterparts in the Ising-Edwards-Anderson model. On the other hand, in the chiral sector scaling corrections were

annoyingly large, despite they decrease upon increasing the magnitude of the anisotropic interactions.

Regarding the coupling of chiral and spin glass transition, our numerical results seem to indicate that the two phase transitions take place at the same temperature (i.e., $T_{CG} = T_{SG}$). However, it is important to stress that we need our very largest lattices to observe this trend. Nevertheless, what we see is in agreement with both Kawamura's prediction and experiments, where the phase transitions are apparently coupled, and the chiral glass susceptibility is divergent [35].

Moreover, we were able to rationalize the numerical results in Ref. [41] with corrections to scaling, by assuming the Ising-Edwards-Anderson universality class.

We remark that there are strong analogies between the interpretation of numerical and experimental data. In both cases, there is a relevant length scale (the correlation length for experiments, the system size for simulations). If that length is large enough, the asymptotic Ising-Edwards-Anderson universality class should be observed. Otherwise, intermediate results between Heisenberg and Ising are to be expected, and indeed appear [10].

The difficulty in reaching the asymptotic regime lies on time: The time growth of the correlation length is remarkably slow ($\xi(t_w) \sim t_w^{1/z}$ with $z \approx 7$ [38,39], where t_w is the waiting time). Indeed, the current experimental record is around $\xi \sim 100$ lattice spacings [11,39], pretty far from the thermodynamic limit [73]. Hence, attention should shift to the study of the intermediate crossover regime. An intriguing possibility appears: One could envisage an experimental study of the crossover effects as a function of the *waiting time*. In fact, t_w varies some four orders of magnitude in current experiments [74], which should result in a factor 4 variation of $\xi(t_w)$.

ACKNOWLEDGMENTS

We thank J. J. Ruiz-Lorenzo and E. Marinari for useful discussions. We were partly supported by MINECO, Spain, through the Research Contract No. FIS2012-35719-C02. M.B.J. was supported by the FPU program (MECD, Spain). The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Grant Agreement No. [247328]. The computations were carried out in the GPU-accelerated clusters Tianhe-1A (Tianjin, China) and Minotauro (Barcelona, Spain). The total amount of time devoted to this project was 2.2×10^5 GPU hours in Tianhe-1A and 2.0×10^5 GPU hours in Minotauro. Access to Tianhe-1A was granted through Research Contract No. 287746 by the EU-FP7. The authors thankfully acknowledge the computer resources, technical expertise, and assistance provided by the staff at the *National Supercomputing Center-Tianjin* and at the *Red Española de Supercomputación-Barcelona Supercomputing Center*.

APPENDIX: SPIN GLASSES ON (MULTIPLE) GPUS

The Appendix is structured as follows. The specific algorithms that we have used are explained in Appendix A 1 with no reference to their implementation. However, implementation *is* crucial: Our simulations are so demanding that we have

used special hardware described in Appendix A 2. This special hardware speeds up the simulations thanks to parallelization, so in Appendix A 3 we give some brief details about it. Finally, we address in Appendix A 4 some issues regarding the generation of pseudorandom numbers.

1. Simulation algorithms

As explained in Sec. III, we used a blend of several Monte Carlo dynamics. Specifically, our single Monte Carlo step (MCS) consisted of (in successive order):

- (1) 1 full lattice sweep with the heat-bath algorithm,
- (2) L lattice sweeps of microcanonical over-relaxation algorithm, and
- (3) 1 parallel tempering sweep [60,61].

Heat-bath by itself would provide correct (but inefficient) dynamics. It actually mimics the natural evolution followed by real spin glasses (that never reach equilibrium near or below the critical temperature). For this reason we enhance it with two more algorithms. However, heat-bath does play a crucial role, since it is irreducible (i.e., the full configuration space is reachable, at least in principle), at variance with over-relaxation, which keeps the total energy constant, and parallel tempering, which changes the temperature but not the spin configuration.

Crucial to perform the heat-bath and over-relaxation dynamics is a factorization property of the Boltzmann weight for the Hamiltonian (1). The conditional probability density for spin \vec{s}_x , given the rest of the spins of the lattice, is

$$P(\vec{s}_x | \{\vec{s}_y\}_{y \neq x}) \propto e^{(\vec{s}_x \cdot \vec{h}_x)/T}, \quad (\text{A1})$$

where \vec{h}_x is the *local field* produced by the lattice nearest neighbors of spin \vec{s}_x (its precise definition is given in Ref. [64]).

In the heat-bath update, a new orientation for spin \vec{s}_x is drawn from the conditional probability (A1), see Ref. [37] for instance.

The over-relaxation update is deterministic. Given a spin \vec{s}_x and its local field, we change the spin as much as possible while keeping the energy constant:

$$\vec{s}_x^{\text{new}} = 2\vec{h}_x \frac{\vec{h}_x \cdot \vec{s}_x^{\text{old}}}{h_x^2} - \vec{s}_x^{\text{old}}. \quad (\text{A2})$$

Contrarily to heat-bath, the order in which the spins are updated is important in over-relaxation. Accessing the lattice randomly increases the autocorrelation time in a substantial way. On the other hand, a sequential update generates a microcanonic wave that sweeps the lattice. The resulting change in the configuration space is significantly larger. A similar microcanonic wave is generated with other types of deterministic lattice sweeps. For instance, one could partition the lattice in a checkerboard way and first update all spins in the black sublattice, updating the white spins only afterwards.

The combination of heat-bath and over-relaxation has been shown to be effective in the case of isotropic spin glasses [75] and other models with frustration [62,63]. However, if one is interested in very low temperatures or large systems, parallel tempering is often useful. For each sample we simulate N_T different copies of the system, each of them at one of the temperatures $T_1 < T_2 < \dots < T_{N_T}$. A parallel tempering

update consists of proposing, as configuration change, a swap between configurations at neighboring temperatures. The exchange has the Metropolis acceptance. Evidently the acceptance is higher if the temperatures T_i are closer to each other, since the energy of the configurations will be similar. Notice that exchanging configurations is equivalent to exchange temperatures, so the data transfer is reduced to a single number.

2. Hardware features

The GPUs we used were of the Tesla generation, produced by NVIDIA, with an SIMD architecture (single instruction, multiple data) [76], optimized for the parallel processing of large amounts of double precision data.

We had access to Tesla M2050 GPUs in the *Tianhe-1A* supercomputer [44] in Tianjin, China, and Tesla M2090 GPUs on the *Minotauro* cluster [45] in Barcelona, Spain. Despite the extremely high performances claimed by NVIDIA (e.g., 665 Gflops in double precision in the case of the M2090 GPUs), it is practically impossible to reach that limit, because the major bottleneck does not reside in the computing speed, but in the memory access. Yet GPUs keep being a valid tool to simulate on spin glasses, as they typically allow the same function to be launched concurrently on thousands of threads. This is exactly what we need, since we can update simultaneously different replicas, and also non-neighboring spins within the same replica, because the interactions are only between nearest neighbors (see Appendix A 3).

More details on the specific hardware and codes will be given in Ref. [77].

3. Parallelization

Our update schemes support two levels of parallelism. Heat-bath and over-relaxation are parallelized within a single lattice. On the other hand, parallel tempering concerns $2N_T$ independent lattices (two replicas, see Sec. II B, at N_T temperatures). Clearly spins in different lattices can be updated simultaneously (between temperature swaps). For small system sizes, the $2N_T$ lattices can be updated efficiently within a single GPU. Yet, for $L = 64$ we have found it convenient to speed up by employing N_T GPUs, each of them simulating two lattices.

a. Single GPU

Our parallelization scheme was not very different from the one described extensively in previous works such as Refs. [78] and [79], so we limit ourselves to remark that we used binary couplings in order to be able to store a full coupling in a single byte. Also, due to the fact that the lattice positions were evaluated with bitwise operations, and to our coalesced memory-reading scheme [77], our program was mostly efficient when the size of the lattice was a power of two, so we favored simulations on those sizes.

b. Multi-GPU

For $L = 64$ and $D = 0.5$ the relaxation times were too long to be able to thermalize on a single GPU. Therefore, we prepared a code that mixed CUDA and MPI, in order to be able to concentrate a major computing capability on

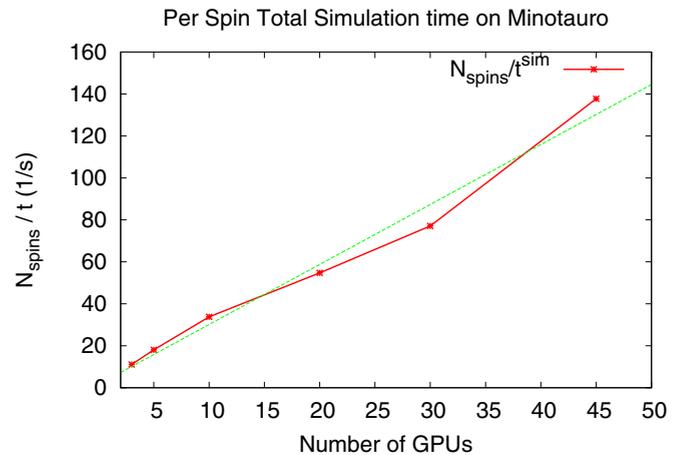


FIG. 6. (Color online) Scaling of the computing time with the number of GPUs N_{GPU} . Benchmark performed on the *Minotauro* GPU cluster [45].

a single sample. We took advantage of the two levels of parallelization that our update algorithms allow. We used $N_{\text{GPU}} = N_T = 45$ GPUs, each updating only two independent lattices with the same couplings, but not necessarily with the same temperature. At the level of the single GPU, the way we swept the lattice with heat-bath and over-relaxation was similar to the single-GPU version. Yet, we had to arrange it in order to get the same thread occupancy as in the single-GPU version. Our choice has been to divide the lattice in rows of eight spins along the x axis. Non-neighboring rows were updated at the same time. A side advantage of this scheme was that we could use for it the same type of coalesced memory reading that we developed for the single-GPU lattice sweeping.

This arrangement resulted in an extremely small overhead when passing from the single to the multiple-GPU algorithm. We were also favored by other factors. Parallel tempering only requires the exchange with the master of a double precision number. Also, the long correlation times allow us to take measurements with low frequency. As a consequence of all this, we obtained a linear scaling of the computing time with the number of GPUs N_T (Fig. 6).

4. Pseudorandom number generator

Pseudorandom number generators (PRNGs) are a critical issue in the implementation of stochastic algorithms [80], but even more in cases like ours, where each of the N_{threads} threads had to carry its own PRNG, and we had a large number of them acting in parallel on the same lattice. This became a major problem especially in the simulations with MPI, where a huge number of PRNGs was concentrated on only two lattices. It was crucial to guarantee the statistical independence of the N_{threads} pseudorandom sequences. We consider three different aspects: (a) the PRNG that each thread uses, (b) the initialization of the generators, and (c) our tests on the generators.

a. The generator

We resorted to a linear combination of Parisi-Rapupano with congruential generators [81].

With the Parisi-Rapuano sequence [82], the n th pseudorandom number P_n is generated through the following relations:

$$\begin{aligned} y_n &= (y_{n-24} + y_{n-55}) \bmod 2^{64}, \\ P_n &= y_n \text{XOR} y_{n-61}, \end{aligned} \quad (\text{A3})$$

where XOR is the exclusive OR logic operator, and y_i are 64-bit unsigned integers. Although some pathologies have been found in the 32-bit Parisi-Rapuano PRNG [83], it looks like its 64-bit version is solid [84].

On the other side, we used a 64-bit congruential generator, where the n th element of the sequence C_n was given by [80,85]

$$C_n = (C_{n-1} \times 3202034522624059733 + 1) \bmod 2^{64}. \quad (\text{A4})$$

Also this generator is not reliable when used alone [81,86].

The final pseudorandom number R_n was obtained by summing P_n and C_n :

$$R_n = (P_n + C_n) \bmod 2^{64}. \quad (\text{A5})$$

b. Initializing the generators

We have found that problems arise if special care is not devoted to the initialization of the random numbers. This is particularly important in the case of multiple GPUs where $N_{\text{threads}} = 32\,768$ threads concurrently update the spins in only two lattices.

We decided to use one seed per node. This seed was used to initialize a 64-bit congruential PRNG [Eq. (A4)]. We

employed it to initialize the state vector of 24-bit Luescher PRNG [87]. The 24-bits words were obtained from three consecutive congruential calls (we kept the most significant byte from each call). As for the Luescher generator, we employed the *full luxury* version, which is fireproof but slow. We took the eight most significant bits from each Luescher call to fill up the state vector of the 64-bit PRNGs in Eq. (A5). We were probably excessively cautious, given the high quality of the full-luxury generator, but initialization takes only a small fraction of the total computing time.

c. Tests

We tested with success our random sequences through the whole battery of tests proposed by Marsaglia in Ref. [88]. To be sure the sequences were reliable also with concurrent threads, we also generated N_{threads} sequences and tested them *horizontally*, i.e., taking first the first number of each sequence, then the second, and so on.

Also, we made simulations with ferromagnetic couplings demanding the energies to be equal, up to the seventh significant digit, to those obtained with an independent CPU program.

Finally, it has been pointed out that local Schwinger-Dyson relations (see, e.g., Ref. [89]) can be useful to assess the quality of PRNGs [83]. The relevant identity here is

$$2T \langle \vec{s}_x \cdot \vec{h}_x \rangle - \langle (\vec{h}_x)^2 \rangle - \langle \vec{s}_x \cdot \vec{h}_x \rangle^2 = 0. \quad (\text{A6})$$

We averaged it over all the sites in the lattice, in order to obtain a more stringent test for the simulations.

-
- [1] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
- [2] J. A. Mydosh, *Spin Glasses: An Experimental Introduction* (Taylor and Francis, London 1993).
- [3] A. P. Young, *Spin Glasses and Random Fields* (World Scientific, Singapore, 1997).
- [4] M. A. Ruderman and C. Kittel, *Phys. Rev.* **96**, 99 (1954).
- [5] T. Kasuya, *Prog. Theor. Phys.* **16**, 45 (1956).
- [6] K. Yosida, *Phys. Rev.* **106**, 893 (1957).
- [7] I. Dzyaloshinsky, *J. Phys. Chem. Solids* **4**, 241 (1958).
- [8] T. Moriya, *Phys. Rev. Lett.* **4**, 228 (1960).
- [9] S. F. Edwards and P. W. Anderson, *J. Phys. F* **5**, 965 (1975).
- [10] D. Petit, L. Fruchter, and I. A. Campbell, *Phys. Rev. Lett.* **88**, 207206 (2002).
- [11] F. Bert, V. Dupuis, E. Vincent, J. Hammann, and J.-P. Bouchaud, *Phys. Rev. Lett.* **92**, 167203 (2004).
- [12] O. E. Peil, A. V. Ruban, and B. Johansson, *Phys. Rev. B* **79**, 024428 (2009).
- [13] J. W. Cable, S. A. Werner, G. P. Felcher, and N. Wakabayashi, *Phys. Rev. Lett.* **49**, 829 (1982).
- [14] J. W. Cable, S. A. Werner, G. P. Felcher, and N. Wakabayashi, *Phys. Rev. B* **29**, 1268 (1984).
- [15] F. J. Lamelas, S. A. Werner, S. M. Shapiro, and J. A. Mydosh, *Phys. Rev. B* **51**, 621 (1995).
- [16] H. Bouchiat, *J. Phys. (Paris)* **47**, 71 (1986).
- [17] L. P. Lévy, *Phys. Rev. B* **38**, 4963 (1988).
- [18] K. Gunnarsson, P. Svedlindh, P. Nordblad, L. Lundgren, H. Aruga, and A. Ito, *Phys. Rev. B* **43**, 8199 (1991).
- [19] S. Franz, G. Parisi, and M. A. Virasoro, *J. Phys. I France* **4**, 1657 (1994).
- [20] M. Palassini and S. Caracciolo, *Phys. Rev. Lett.* **82**, 5128 (1999).
- [21] H. G. Ballesteros, A. Cruz, L. A. Fernández, V. Martín-Mayor, J. Pech, J. J. Ruiz-Lorenzo, A. Tarancón, P. Téllez, C. L. Ullod, and C. Ungil, *Phys. Rev. B* **62**, 14237 (2000).
- [22] W. L. McMillan, *Phys. Rev. B* **31**, 342 (1985).
- [23] J. A. Olive, A. P. Young, and D. Sherrington, *Phys. Rev. B* **34**, 6341 (1986).
- [24] B. M. Morris *et al.*, *J. Phys. C* **19**, 1157 (1986).
- [25] F. Matsubara, T. Iyota, and S. Inawashiro, *Phys. Rev. Lett.* **67**, 1458 (1991).
- [26] M. J. P. Gingras, *Phys. Rev. Lett.* **71**, 1637 (1993).
- [27] B. Coluzzi, *J. of Phys. A: Math. Gen.* **28**, 747 (1995).
- [28] A. Mauger, J. Villain, Y. Zhou, C. Rigaux, N. Bontemps, and J. Ferré, *Phys. Rev. B* **41**, 4587 (1990).
- [29] H. Kawamura, *Phys. Rev. Lett.* **68**, 3785 (1992); **80**, 5421 (1998).
- [30] L. W. Lee and A. P. Young, *Phys. Rev. Lett.* **90**, 227203 (2003).
- [31] L. A. Fernandez, V. Martin-Mayor, S. Perez-Gavero, A. Tarancon, and A. P. Young, *Phys. Rev. B* **80**, 024422 (2009).
- [32] I. Campos, M. Cotallo-Aban, V. Martin-Mayor, S. Perez-Gavero, and A. Tarancon, *Phys. Rev. Lett.* **97**, 217204 (2006).

- [33] D. X. Viet and H. Kawamura, *Phys. Rev. Lett.* **102**, 027202 (2009).
- [34] H. Kawamura, *Phys. Rev. Lett.* **90**, 047202 (2003).
- [35] T. Taniguchi, *J. Phys.: Condens. Matter* **19**, 145213 (2007).
- [36] I. A. Campbell and D. C. M. C. Petit, *J. Phys. Soc. Jpn.* **79**, 011006 (2010).
- [37] D. Amit and V. Martin-Mayor, *Field Theory, the Renormalization Group and Critical Phenomena*, 3rd ed. (World Scientific, Singapore, 2005).
- [38] F. Belletti *et al.*, *Phys. Rev. Lett.* **101**, 157201 (2008); *J. Stat. Phys.* **135**, 1121 (2009).
- [39] Y. G. Joh, R. Orbach, G. G. Wood, J. Hammann, and E. Vincent, *Phys. Rev. Lett.* **82**, 438 (1999).
- [40] A. J. Bray and M. A. Moore, *J. Phys. C: Solid State Phys.* **15**, 3897 (1982).
- [41] V. Martin-Mayor and S. Perez-Gaviro, *Phys. Rev. B* **84**, 024419 (2011).
- [42] Recall that $\gamma_{CG} = \nu(2 - \eta_{CG})$, where γ_{CG} is the critical index for the CG susceptibility, while ν is the correlation-length exponent.
- [43] Of course the limiting factor is in the wide range of *relaxation times*, rather than temperatures. However, relaxation times depend on a variety of implementation-dependent factors (such as the temperature spacing in the parallel tempering, or the number of over-relaxation sweeps). Hence, comparison with other work will be easier in terms of temperatures.
- [44] National Super-Computing Center, Tianjin, China, <http://www.nsc-cj.gov.cn/en/>
- [45] Barcelona Supercomputing Center, Barcelona, Spain, <http://www.bsc.es>
- [46] F. Parisen Toldin, A. Pelissetto, and E. Vicari, *J. Stat. Mech.* (2006) P06002.
- [47] F. Liers, J. Lukic, E. Marinari, A. Pelissetto, and E. Vicari, *Phys. Rev. B* **76**, 174423 (2007).
- [48] Independence from microscopic details such as the disorder distribution has been found for spin glasses [49-52], as well as for other disordered systems like the random field Ising model [53], or disordered ferromagnets (either site [54] or bond [55, 56] diluted).
- [49] M. Hasenbusch, A. Pelissetto, and E. Vicari, *Phys. Rev. B* **78**, 214205 (2008).
- [50] T. Jorg, *Phys. Rev. B* **73**, 224431 (2006).
- [51] H. G. Katzgraber, M. Korner, and A. P. Young, *Phys. Rev. B* **73**, 224432 (2006).
- [52] T. Jorg and H. G. Katzgraber, *Phys. Rev. Lett.* **101**, 197205 (2008).
- [53] N. G. Fytas and V. Martin-Mayor, *Phys. Rev. Lett.* **110**, 227201 (2013).
- [54] H. G. Ballesteros, L. A. Fernandez, V. Martin-Mayor, A. Munoz Sudupe, G. Parisi and J. J. Ruiz-Lorenzo, *Phys. Rev. B* **58**, 2740 (1998).
- [55] P.-E. Berche, C. Chatelain, B. Berche, and W. Janke, *Eur. Phys. J. B* **38**, 463 (2004).
- [56] A. Malakis, A. N. Berker, N. G. Fytas, and T. Papakonstantinou, *Phys. Rev. E* **85**, 061106 (2012).
- [57] M. Baity-Jesi *et al.*, *Eur. Phys. J. Special Topics* **210**, 33 (2012).
- [58] L. W. Lee and A. P. Young, *Phys. Rev. B* **76**, 024405 (2007).
- [59] F. R. Brown and T. J. Woch, *Phys. Rev. Lett.* **58**, 2394 (1987).
- [60] K. Hukushima and K. Nemoto, *J. Phys. Soc. Jpn* **65**, 1604 (1996).
- [61] E. Marinari in *Advances in Computer Simulations*, edited by J. Kerstész and I. Kondor (Springer, Berlin, 1995).
- [62] J. L. Alonso, A. Tarancón, H. G. Ballesteros, L. A. Fernández, V. Martín-Mayor, and A. Muñoz Sudupe, *Phys. Rev. B* **53**, 2537 (1996).
- [63] E. Marinari, V. Martin-Mayor, and A. Pagnani, *Phys. Rev. B* **62**, 4999 (2000).
- [64] It is enough to define the local field as $\vec{h}_x = \sum_{\|x-y\|=1} [J_{xy}\vec{s}_y + D_{xy}\vec{s}_y]$.
- [65] L. A. Fernández, A. Maiorano, E. Marinari, V. Martin-Mayor, D. Navarro, D. Sciretti, A. Tarancón, and J. L. Velasco, *Phys. Rev. B* **77**, 104432 (2008).
- [66] M. P. Nightingale, *Physica A* **83**, 561 (1975).
- [67] H. G. Ballesteros, L. A. Fernández, V. Martín-Mayor, and A. Muñoz-Sudupe, *Phys. Lett. B* **378**, 207 (1996).
- [68] A more recent article from the Janus Collaboration (Ref. [69]) gives a more precise determination of the critical exponents. Using one or the other does not change significantly our results and conclusions.
- [69] Janus Collaboration: M. Baity-Jesi, R. A. Baños, A. Cruz, L. A. Fernandez, J. M. Gil-Narvion, A. Gordillo-Guerrero, D. Iñiguez, A. Maiorano, F. Mantovani, E. Marinari, V. Martin-Mayor, J. Monforte-Garcia, A. Muñoz Sudupe, D. Navarro, G. Parisi, S. Perez-Gaviro, M. Pivanti, F. Ricci-Tersenghi, J. J. Ruiz-Lorenzo, S. F. Schifano, B. Seoane, A. Tarancon, R. Tripiccion, and D. Yllanes, *Phys. Rev. B* **88**, 224416 (2013).
- [70] Some of the points we used for those extrapolations shared some of the data. For example, the crossing of ξ for $L = 8, 16$, had in common the points from size $L = 16$ with the pair $L = 16, 32$. Hence, in the fits we have taken into account the covariance matrix that gave a measure of the anticorrelation between measures that share data.
- [71] In Ref. [41] the D_{ij}^{ab} are uniformly distributed between $\pm D = 0.05$, while in our work we use binary couplings. If we want to compare them, we have to use $D = 1/\sqrt{1200} \simeq 0.03$.
- [72] In the phase diagram we show the $D = 0$ point comes from Ref. [31], where chiral and spin glass transition are assumed to be coupled. There is disagreement on whether $T_{SG} = T_{CG}$ also in the isotropic case. Yet, we do plot it as a single transition because although T_{SG} might be lower than T_{CG} , their best estimates are compatible (and not distinguishable in the plot).
- [73] In a typical system $N = L^3 \sim 6 \times 10^{23} \Rightarrow L \simeq 10^8$.
- [74] G. F. Rodriguez, G. G. Kenning, and R. Orbach, *Phys. Rev. B* **88**, 054302 (2013).
- [75] J. H. Pixley and A. P. Young, *Phys. Rev. B* **78**, 014419 (2008).
- [76] NVIDIA Corporation, *CUDA C Programming Guide*, docs.nvidia.com/cuda/cuda-c-programming-guide/index.html
- [77] M. Baity-Jesi, Ph.D. thesis, Universidad Complutense de Madrid (work in progress).
- [78] M. Bernaschi, G. Parisi, and L. Parisi, *Comput. Phys. Commun.* **182**, 1265 (2011).
- [79] T. Yavors'kii and M. Weigel, *Eur. Phys. J. Special Topics* **210**, 159 (2012).
- [80] D. E. Knuth, *The Art of Computer Programming*, 2nd ed. (Addison-Wesley, Reading, MA, 1981), Vol. 2.
- [81] L. A. Fernandez, V. Martin-Mayor, and D. Yllanes, *Nucl. Phys. B* **807**, 424 (2009).

- [82] V. Parisi, cited in G. Parisi and F. Rapuano, *Phys. Lett. B* **157**, 301 (1985).
- [83] H. G. Ballesteros and V. Martin-Mayor, *Phys. Rev. E* **58**, 6787 (1998).
- [84] L. A. Fernandez, V. Martin-Mayor, D. Sciretti, A. Tarancon, and J. L. Velasco, *Phys. Lett. B* **628**, 281 (2005).
- [85] P. L'Ecuyer, *Math. Comp.* **68**, 249 (1999).
- [86] G. Ossola and A. D. Sokal, *Nucl. Phys. B* **691**, 259 (2004).
- [87] M. Luescher, *Comput. Phys. Commun.* **79**, 100 (1994).
- [88] G. Marsaglia, Diehard Battery of Tests of Randomness, <http://www.stat.fsu.edu/pub/diehard>
- [89] R. J. Rivers, *Path Integral Methods in Quantum Field Theories* (Cambridge University Press, Cambridge, 1990).