

Computational Physics

Microcanonical simulated annealing: Massively parallel Monte Carlo simulations with sporadic random-number generation

M. Bernaschi^a, C. Chilin^{b,c}, L.A. Fernandez^b, I. González-Adalid Pemartín^a, E. Marinari^{c,d}, V. Martin-Mayor^b, G. Parisi^{e,c,d}, F. Ricci-Tersenghi^{c,d}, J.J. Ruiz-Lorenzo^{f,g}, D. Yllanes^{h,i,j,*}

^a Istituto per le Applicazioni del Calcolo, CNR, Rome, Italy

^b Departamento de Física Teórica, Universidad Complutense, Madrid, 28040, Spain

^c Dipartimento di Fisica, Sapienza Università di Roma, Rome, 00185, Italy

^d CNR-Nanotec, Rome Unit, and INFN, Sezione di Roma, Rome, 00185, Italy

^e International Research Center of Complexity Sciences, Hangzhou International Innovation Institute, Beihang University, Hangzhou, 311115, China

^f Departamento de Física, Universidad de Extremadura, Badajoz, 06006, Spain

^g Instituto de Computación Científica Avanzada (ICCAEx), Universidad de Extremadura, Badajoz, 06006, Spain

^h Fundación ARAID, Diputación General de Aragón, Zaragoza, 50018, Spain

ⁱ Instituto de Biocomputación y Física de Sistemas Complejos (BIFP), Universidad de Zaragoza, Zaragoza, 50018, Spain

^j Zaragoza Scientific Center for Advanced Modeling (ZCAM), Zaragoza, 50018, Spain

ARTICLE INFO

Editor: Prof. Weigel Martin

Keywords:

Monte Carlo simulation

Ising machines

CUDA

Parallel computing

Spin glasses

ABSTRACT

Numerical simulations of models and theories that describe complex systems such as spin glasses are becoming increasingly important. Beyond fundamental research, these computational methods also find practical applications in fields like combinatorial optimization. However, Monte Carlo simulations, an important subcategory of these methods, are plagued by a major drawback: they are extremely greedy for (pseudo) random numbers. The total fraction of computer time dedicated to random-number generation increases as the hardware grows more sophisticated, and can get prohibitive for special-purpose computing platforms. We propose here a general-purpose microcanonical simulated annealing (MicSA) formalism that dramatically reduces such a burden. The algorithm is fully adapted to a massively parallel computation, as we show in the particularly demanding benchmark of the three-dimensional Ising spin glass. We carry out very stringent numerical tests of the new algorithm by comparing our results, obtained on GPUs, with high-precision standard (*i.e.*, random-number-greedy) simulations performed on the Janus II custom-built supercomputer. In those cases where thermal equilibrium is reachable (*i.e.*, in the paramagnetic phase), both simulations reach compatible values. More significantly, barring short-time corrections, a simple time rescaling suffices to map the MicSA off-equilibrium dynamics onto the results obtained with standard simulations.

1. Introduction

Monte Carlo algorithms have become a ubiquitous tool in science and technology, with an amazing variety of applications ranging from lattice gauge theories in particle physics to high-frequency trading in stock markets. Most of the time, massively parallel computations are needed to achieve the expected goals. A common bottleneck to these parallel computations is in the generation of (pseudo-) random numbers, which has been recognized as critical for large-scale Monte Carlo simulations [1]. The difficulties encountered are due to two mutually conflicting goals. On the one hand, a large number of long streams of (sufficiently) independent random numbers are needed for the parallel update of the spins. On the other hand, generating random numbers of

sufficiently high quality is computationally intensive (to the point of being a significant limiting factor for numerical simulations).

Modern efforts in high-performance computing have directed a great deal of attention to the numerical simulation of deceptively simple-looking disordered magnetic alloys, known as spin glasses [2–6]. Interest has grown to the point that hardware is being specifically designed to deal with a “spin-glass Hamiltonian” through a variety of algorithms and/or physical principles (see, *e.g.*, Refs. [7–15]). The choice of this seemingly exotic problem as a framework for more general computations is also connected to the crucial notion of NP-completeness [16], which we briefly recall next.

The point is here that one of the interesting applications of Monte Carlo methods is to minimize a cost function, for example, by

* Corresponding author.

E-mail address: david.yllanes@unizar.es (D. Yllanes).

<https://doi.org/10.1016/j.cpc.2026.110182>

Received 12 November 2025; Received in revised form 5 April 2026; Accepted 24 April 2026

Available online 27 April 2026

0010-4655/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

incrementally reducing the “temperature” control parameter, starting from a very high value and going down close to zero (where the “Boltzmann” measure concentrates at the minimum of the cost function). This algorithm is the well-known simulated annealing [17], which can be used to solve, under some conditions, optimization problems in a very effective way [18,19].

The particular problem at hand is formalized in terms of a cost function that one aims to minimize. For the spin glass, the cost function is the energy of the spin configuration, while, for the traveling salesman, the cost function is the total travel distance to visit a set of cities. One wants to find the assignment that minimizes the cost function, often with the help of a computer. If the computational resources (*i.e.*, computer time, memory, etc.) grow with N faster than any polynomial for all known algorithms, the problem is generally regarded as *hard*. A great attention is paid to a small subcategory of these hard problems, termed NP-complete, because if a polynomial scaling algorithm were discovered for any of the NP-complete problems, then a vast family of hard optimization problems would become *easy*. Finding the minimal energy state—the ground state (GS)—of an Ising spin-glass Hamiltonian on a non-planar graph is one of such NP-complete problems [20,21].

Understandably, considerable effort has been devoted to the efficient numerical simulations of spin glasses, both on unconventional dedicated hardware (see above) or through the writing of code carefully optimized for its use on GPUs [22–26] or CPUs [27–29]. Perhaps unsurprisingly, the generation of random numbers turns out to be a major headache in both computational contexts.

Here we propose a widely usable microcanonical simulated annealing (MicSA) formalism, which is both extremely frugal in the number of needed random numbers and fully adapted to a massively parallel computation. The method is demonstrated in the rather demanding case of the three-dimensional Ising spin glass. We carry out our simulations on GPUs to demonstrate that our algorithm nicely parallelizes. We start from a microcanonical ensemble for an extended configuration space that comprises the degrees of freedom we are really interested in (the spins) and some auxiliary degrees of freedom traditionally called daemons. We shall term these auxiliary variables *walkers* if they are allowed to wander through the system or *daemons* if each is kept tied to a single lattice site; should the distinction be immaterial, which happens at many points in the discussion, we shall name them daemons/walkers. The reader will find useful analogies between our Monte Carlo formalism and molecular dynamics simulations coupled with a thermostat [30].

A microcanonical framework with just *one* daemon was proposed by Creutz [31], precisely to minimize the need for random numbers. Unfortunately, Creutz’s ensemble leads to a rather significant dynamic slowing down in the context of spin glasses [32]. To overcome this dynamic slowdown in non-disordered systems, Rummukainen and other authors proposed using a number of daemons proportional to the number of spins [33–36]. We shall follow a similar approach, although those authors emphasized equilibrium physics—our focus, instead, is on non-equilibrium dynamics—and they were not particularly concerned with massively parallel simulations. Lustig’s microcanonical ensemble [37] also contains a number of daemons proportional to the number of spins. Lustig’s framework is very effective for studying first-order phase transitions [38], also for disordered systems [39], which suggests a potential effectiveness in the study of spin glasses (see also [33] for a study of first-order transitions in non-disordered systems). This formalism, however, has three major inconveniences: (i) it is unsuitable for parallel computations, (ii) it is *not* frugal in its usage of random numbers and (iii) it eventually approaches the equilibrium state corresponding to the microcanonical ensemble (which coincides with the canonical ensemble only in the limit of an infinite system size). We cure problems (i) and (ii) by simply skipping the integration of the daemons that Lustig carries out and keeping them instead as explicit dynamic variables. In order to cure problem (iii), we generalize the simulated annealing algorithm [17] to the microcanonical context. We argue (and show by example) that it is sufficient to lower the total energy of the system by acting solely on

the daemons at a logarithmic time rate (this is the only part of the algorithm that involves random numbers) to obtain results indistinguishable (after a simple time rescaling) from those produced using a canonical simulation employing the Janus II special-purpose computer [15]. Although rescalability will be considered in much deeper detail below, as a starter, the reader may wish to check Fig. 1.

In a canonical simulation, one may define an effective temperature (for instance) by comparing the instantaneous value of the energy with that of an equilibrated system at the effective temperature. Only in equilibrium does the effective temperature coincide with the thermal-bath temperature. In close analogy, see the more detailed discussion in Section 2.7, the refresh of the daemons/walkers effectively cools the system only as long as thermal equilibrium has not been reached. In equilibrium, the net energy flow between spins and daemons/walkers cancels.

The rescalability evinced in Fig. 1 can be related to analytical work on microcanonical dynamics—see, for instance, the Chapter entitled “Critical Dynamics” in Ref. [40] for model-C critical dynamics (*i.e.*, non-conserved order parameter and conserved energy). If the coherence length grows with time without bound—which is the case for a spin glass at $T < T_c$, T_c being the critical temperature—and if the specific heat does not diverge in that limit (which is also the case for spin glasses), one expects that the canonical and the microcanonical dynamics can be superimposed at long times after a simple time rescaling. Although the algorithm proposed here is *not* microcanonical, it alternates between microcanonical and non-microcanonical phases. Hence, the field-theoretical prediction is reassuring.

The structure of the rest of the paper is as follows: We recall the Edwards-Anderson model and explain our algorithm in Section 2. Then, we compare the microcanonical and canonical dynamics for the Edwards-Anderson model in three dimensions in Section 3, showing that both belong to the same dynamic universality class. We conclude by summarizing the results and discussing potential future applications of this algorithm in Section 4. We also provide additional information in three appendices. Appendix A provides technical details about our time-rescaling procedures, Appendix B provides some basic information about our publicly available CUDA code. Finally, Appendix C briefly describes the application of the method in equilibrium simulations for the specific example of a small sample of a three-dimensional spin glass with Gaussian couplings at low temperatures.

2. The model and the algorithm

We have chosen to benchmark our algorithm in the rather demanding context of the three-dimensional Edwards-Anderson model [41,42], which is the standard for spin glasses. While describing the model in Section 2.1, we shall emphasize its separability, which is the crucial feature that makes parallel computation feasible. A similar separation between mutually non-interacting degrees of freedom is a necessary condition in any envisaged application of the solution proposed herein.

We have chosen to organize this section in several separate paragraphs. After the somewhat introductory Sections 2.1 and 2.2, each paragraph will discuss a specific decision we needed to make to design the algorithm. Indeed, efficient implementations require tuning Monte Carlo algorithms to the user’s specific needs, which are not only dictated by the peculiarities of the problem under study, but also by those of the hardware platform that will be used in the simulation. We hope that making our tuning process explicit will help the reader adapt the algorithm to their needs. We have summarized the defining details of the algorithms that we have numerically tested against a canonical simulation in Section 2.8.

2.1. The model (and its separability)

We consider N Ising spins $s_x = \pm 1$ that occupy the nodes, denoted by x , of a cubic lattice with linear size L ($N = L^3$). The lattice is endowed

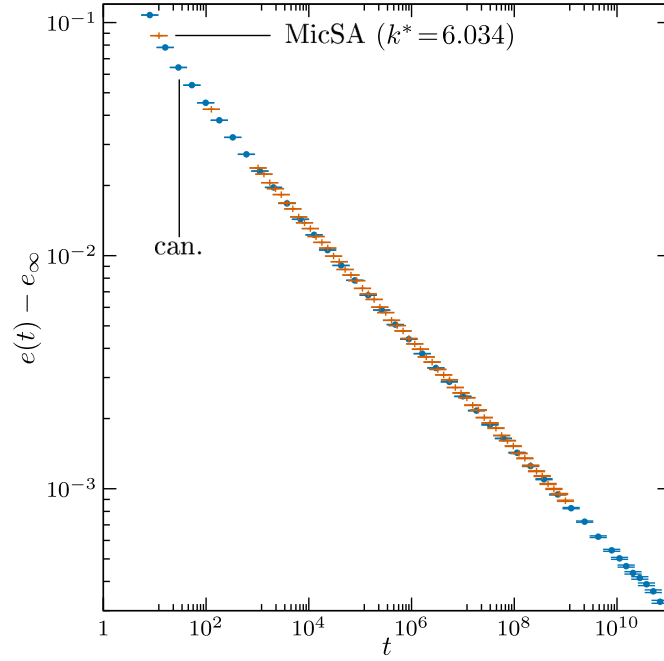


Fig. 1. After a time rescaling, our microcanonical simulated annealing reproduces canonical dynamics. We study the energy relaxation towards its equilibrium value by plotting the excess energy per spin over its equilibrium value, $e(t) - e_\infty$, versus simulation time t , as computed for a spin-glass sample on a $160 \times 160 \times 160$ lattice with binary couplings at temperature $T = 0.7 \approx 0.64T_c$ (T_c is the critical temperature for the phase transition separating the paramagnetic phase at high temperatures from the spin-glass phase at low temperature). The sample is investigated using two dynamics: a standard, random-number greedy Metropolis algorithm (the label *can* in the plot stands for canonical; these data were obtained from the Janus II supercomputer) and the new microcanonical simulated annealing (MicSA) algorithm introduced in this work. Specifically, we are using the version of the algorithm with 6 walkers per spin, see Section 2.8. We plot both data sets, subtracting the same estimate for e_∞ , which was obtained from the canonical simulation as explained in Appendix A. After the MicSA time is rescaled by the factor $k^* = 6.034$ (see Table 1), $t^{\text{plot}} = k^* \times t^{\text{MicSA}}$, both energy relaxations are indistinguishable (the naive rescaling would have been $k^{\text{naif}} = 6$, because of the 6 walkers per spin).

with periodic boundary conditions. The spins interact with their lattice nearest neighbors through the Hamiltonian

$$H^{\text{EA}} = - \sum_{\langle x,y \rangle} J_{x,y} s_x s_y. \quad (1)$$

The quenched couplings $J_{x,y}$ are independent and identically distributed random variables that are drawn at the beginning of the simulation and kept fixed. The couplings are typically drawn from a symmetric distribution. Our choice will be $J_{x,y} = \pm 1$ with 50% probability (Gaussian-distributed couplings are a fairly popular choice, as well). A specific realization of the coupling matrix $\{J_{x,y}\}$ is named a *sample*. With binary couplings $J = \pm 1$, the system undergoes a continuous phase transition at $T_c = 1.1019(29)$ [43].

In principle (but see below for an alternative point of view), one is interested in the canonical ensemble, where spin configurations are found as dictated by Boltzmann's probability at inverse temperature $\beta = 1/T$:

$$p(\{s_x\}) = \frac{e^{-\beta H^{\text{EA}}(\{s_x\})}}{Z_{\text{EA}}}, \quad Z_{\text{EA}} = \sum_{\{s_x\}} e^{-\beta H^{\text{EA}}(\{s_x\})}. \quad (2)$$

It is fairly common in spin glasses that thermal equilibrium will not be achievable in a reasonable computer time (nor in reasonable human time, if we are talking about a laboratory experiment). Many important experimental effects can still be investigated through an out-of-equilibrium simulation [6,44–52], which is carried out with a local algorithm, such as Metropolis or heat bath (see, for instance, Refs. [53,54]) that has Eq. (2) as its stationary distribution.

A crucial feature of the Hamiltonian H^{EA} that makes parallel computations feasible is separability, which is itself a consequence of the limited interaction range (lattice nearest neighbors in our case). Here, separability means that the N spins can be split into disjoint sets such that spins in the same set are mutually non-interacting (entailing that

they can be updated simultaneously). The number of non-interacting sets could be larger than two, but it should be small to make parallel computing efficient. In our case, a very natural possibility is to group the spins into *even* and *odd* sublattices, depending on the parity of the sum of their Cartesian coordinates. This choice results in a *checkerboard* decomposition of the lattice which has proven instrumental in many computational and theoretical studies for decades (Ref. [55] provides an early example).

As we discuss in Section 2.2, separability is endangered by the (in so many other ways fully satisfactory) Lustig microcanonical formalism.

2.2. A generalization of Lustig's microcanonical ensemble

Initially, we shall follow Lustig and expand the configuration space, introducing new degrees of freedom that we shall name hereafter either daemons or walkers (the distinction is discussed in Section 2.6). Specifically, every spin will be endowed with a number γ of these auxiliary degrees of freedom $n_{x,\alpha} \geq 0$, $\alpha = 0, 1, 2, \dots, \gamma - 1$. The $n_{x,\alpha}$ could be either real or integer. Our total Hamiltonian will be [H^{EA} was defined in Eq. (1)]

$$H = H^{\text{EA}} + H^{\text{aux}} \quad (3)$$

with

$$H^{\text{aux}} = \sum_x \sum_{\alpha=0}^{\gamma-1} n_{x,\alpha}. \quad (4)$$

The corresponding canonical (p_c) and microcanonical (p_{mic}) probability densities are

$$p_c(\{s_x, n_{x,\alpha}\}; \beta) = \frac{\exp[-\beta H(\{s_x, n_{x,\alpha}\})]}{Z_c}, \quad (5)$$

$$p_{\text{mic}}(\{s_x, n_{x,\alpha}\}; \mathcal{E}) = \frac{\delta[\mathcal{E} - H(\{s_x, n_{x,\alpha}\})]}{Z_{\text{mic}}}, \quad (6)$$

where δ stands for Dirac's delta function and the partition functions Z_c and Z_{mic} ensure correct normalization as in Eq. (2).

If we were to follow Lustig, our next step would be to integrate out the auxiliary degrees of freedom in Eq. (6):

$$p_{\text{mic}}(\{s_x\}; \mathcal{E}) = \frac{e^{(\gamma N - 1) \log[\mathcal{E} - H^{\text{EA}}(\{s_x\})]}}{\hat{Z}_{\text{mic}}} \theta(\mathcal{E} - H^{\text{EA}}(\{s_x\})), \quad (7)$$

where the Heaviside step function ensures that $\mathcal{E} > H^{\text{EA}}(\{s_x\})$, and the modified partition function \hat{Z}_{mic} ensures normalization. Eq. (7) is well adapted for a serial Monte Carlo simulation, where spins are updated one by one (an independent random number would be necessary per every spin update). Unfortunately, the effective action in Eq. (7), $\log(\mathcal{E} - H^{\text{EA}})$ is highly non-linear in H^{EA} , which spoils the separability property discussed in Section 2.1. Indeed, even an effective action quadratic in H^{EA} would be non-separable because every pair of spins would be mutually interacting. With some ingenuity, a limited parallel computation can be set for an all-to-all interaction [7], but we regard this as a very undesirable feature. Hence, we shall depart from Lustig's approach at this point.

2.3. Creutz's move saves the day

Fortunately, the obstacle described in the previous paragraph can be cleared by adapting to our needs the basic move from Creutz's algorithm [31]. In this way, three different problems get solved at once: (i) we obtain an algorithm that verifies detailed balance for both the canonical and the microcanonical probabilities in Eqs. (5) and (6), (ii) the move does not require any random number, and (iii) separability is not endangered (hence parallel simulations are feasible).

Let us describe the elementary move of our modified form of Creutz's algorithm. Pick spin s_x and one of its daemons/walkers, $n_{x,\alpha}$, and compute the change in H^{EA} , Eq. (1), if we flip the spin at site x (that is, $s_x \rightarrow -s_x$) while keeping all other spins fixed, ΔH^{EA} . If we update

$$n_{x,\alpha} \longrightarrow n_{x,\alpha} - \Delta H^{\text{EA}}, \quad (8)$$

the total Hamiltonian (3) remains unchanged because the total energy gained (or lost) by the spin system is exactly compensated by the $n_{x,\alpha}$. However, to obtain a valid new configuration, we should have $n_{x,\alpha} \geq 0$ after the update in Eq. (8). The spin is flipped, and the $n_{x,\alpha}$ is updated only if the new configuration is legal. Up to now, the only difference from Creutz lies in that he worked with a single auxiliary degree of freedom that is shared by the whole system. In our Lustig-inspired setting, every spin has its own set of γ daemons/walkers $n_{x,\alpha}$.

Let us now briefly explain why our elementary move solves the three problems mentioned above.

The proof of detailed balance—for both the canonical (5) and the microcanonical weight (6)—lies in the reversibility of our elementary move. Indeed, if we pick exactly the same spin s_x and $n_{x,\alpha}$ after we carried out the elementary move—all other spins and daemons/walkers unchanged—and try to update them again, we would exactly undo the spin flip and the update (8) because ΔH^{EA} changes sign after the spin flip. And, of course, the newly reached configuration (which is the configuration we started with) would be legal. Using the language of the Metropolis algorithm [53], we have a symmetric proposal of change in the configuration space that leaves the total energy invariant [implying that the probabilities for the starting and the final configurations are identical, either in the canonical, Eq. (5), or in the microcanonical ensemble, Eq. (6)]. If such a move is accepted with probability one, detailed balance is fulfilled.

Of course, no random numbers are involved in the elementary move. However, it is amusing to notice that if a fresh value of the $n_{x,\alpha}$ were drawn for every spin-flip attempt,¹ our elementary move would become

¹ This is easy to do, although it would require a random number per spin update, because in the canonical ensemble, Eq. (5), the $n_{x,\alpha}$ are independent exponentially-distributed random variables.

the textbook Metropolis update—see, e.g., Refs. [53,54]—for the spin system, recall Eq. (2). Baroque as it may seem, this disguising of the textbook Metropolis algorithm has resulted in a very efficient multispin coding implementation [56].

As for separability, notice that ΔH^{EA} depends only on spins with parity opposite to the parity of our site x . Hence, if we consider two (say) even sites y and z , the two determinations on whether s_y and s_z can be flipped depend solely on odd spins and their chosen n_{y,α_y} and n_{z,α_z} . If the spins at odd sites are held fixed, it is immaterial which of the two sites, y or z , is updated first. A parallel update of all sites with a common lattice parity is, then, feasible.

2.4. Continuous or discrete daemons/walkers?

A computationally convenient feature of the $J = \pm 1$ choice in H^{EA} , Eq. (1), is that the energy changes that a single spin-flip may cause are limited. In three spatial dimensions, $\Delta H^{\text{EA}} = 0, \pm 4, \pm 8$ or ± 12 . Hence, in our particular case, it is practical to limit the possible values of the daemons/walkers to integers multiple of four $n_{x,\alpha} = 4m$, $m = 0, 1, 2, 3, \dots$. For Gaussian couplings, it would probably be more natural to keep continuous daemons/walkers, or have them take values on a finer mesh of energies.

In the canonical ensemble only, the γN auxiliary variables $n_{x,\alpha}$ are statistically independent from each other and from the spins. In the continuous case, the $n_{x,\alpha}$ are exponentially distributed:

$$\text{Prob}(0 < n_{x,\alpha} < R) = 1 - \exp(-\beta R). \quad (9)$$

In Lustig's original ensemble the $n_{x,\alpha}$ enter quadratically in the H^{aux} [37], so that their distribution is Gaussian. For the case with discrete daemons/walkers, see Eq. (11) below.

2.5. Daemons/walkers with a limited capacity

To further parallelize our computation we employ multispin coding [57]. We use the different bits in a memory word to represent different physical spins. Hence, we have chosen to restrict the possible values to $n_{x,\alpha} = 4m$ with $m = 0, 1, 2, 3$. Thus m is representable with two bits only, and we can use two 32-bit words to represent 32 different daemons/walkers. Note that these choices are problem specific. In other situations, different choices are possible, see, for instance, Appendix C.

However, this limitation in the maximum value of $n_{x,\alpha}$ comes with costs and benefits. On the one hand, the configuration obtained after Creutz's move will be legal, meaning that the spin flip will be accepted, if and only if

$$0 \leq n_{x,\alpha} - \Delta H^{\text{EA}} \leq 12. \quad (10)$$

In other words, we should reject any spin-flip that lowers the energy by an amount larger than what the daemon/walker responsible for that update can currently accommodate. On the other hand, having daemons/walkers representable with just two bits reduces the total memory needed. Furthermore, in a multispin coding simulation, the whole algorithm should be implemented employing solely Boolean operations [this includes checking condition Eq. (10) or performing the update in Eq. (8)], which is certainly simpler if the daemons/walkers are representable with only two bits. Refs. [34–36] also considered daemons with limited capacity and managed them accordingly.

Daemons/walkers will be initialized—or refreshed—according to their canonical weight, Eq. (5). That is to say, as independent, identically distributed random variables drawn from the distribution

$$p(n = 4m) = e^{-4m\beta} \frac{1 - e^{-4\beta}}{1 - e^{-16\beta}}, \quad m = 0, 1, 2, 3. \quad (11)$$

We conclude this paragraph with a word of caution. The daemons/walkers should be allowed to take values at least as large as the largest possible $|\Delta H^{\text{EA}}|$ (the problem is in the negative values of H^{EA}). In this way, it will be ensured that no spin (or the equivalent elementary variable) will lock forever at local energy minima too deep to be evaded from.

2.6. Daemons or walkers?

One might worry about the presence of local conservation laws. Local energy conservation is needed if we want to use a parallel updating algorithm, but it could cause dynamic slowing down. This is why we have considered two variants of the algorithm.

If the auxiliary variable $n_{x,\alpha}$ is forever bound to a lattice site x , it will be named a *daemon*. Yet, if it is allowed to wander through the lattice, it will be called a *walker*.

We have designed a wandering procedure that runs efficiently on a GPU because it can be implemented using a memory-access pattern that fulfills the memory coalescing conditions. Consider one of the lattice directions (say Y) and a moving direction (backward or forward). A forward-Y wandering is a two-timestep cycle that can be implemented only for even lattice size L :

1. At the starting time of the cycle, all the α th walkers sitting on even lattice sites exchange their values with the α th walkers sitting on their lattice nearest neighbors along the positive Y direction (periodic boundary conditions are employed; the neighbor will be an odd site).
2. At the second timestep in the cycle, the walkers at odd lattice sites will exchange their values with their lattice nearest-neighbor along the positive Y direction.

Note that after $L/2$ cycles, or L timesteps, every forward-Y wandering walker will return to its starting point, although it will have been modified along its way through the updates in Eq. (8).

The backward-Y wandering operation differs from the forward-Y counterpart just in that the lattice nearest-neighbor will be sought along the *negative* Y direction.

Interestingly enough, Refs. [34,35] implemented their own daemon wandering (although they refreshed their daemons far more often than we do, see Section 2.7).

2.7. The microcanonical simulated annealing

The operations described so far do not allow us to approach the equilibrium canonical distribution, Eq. (5). Indeed, with our microcanonical update, the total energy \mathcal{E} remains equal to its starting value $\mathcal{E}_{\text{start}}$. Eq. (3) tells us that $\mathcal{E}_{\text{start}}$ is the sum of two contributions, the starting value of the spin Hamiltonian H^{EA} (1)—typically much larger than the equilibrium value at inverse temperature β , $\langle H^{\text{EA}} \rangle_\beta$ —and the energy of the daemons/walkers. Although the daemons/walkers will be directly initialized in thermal equilibrium through Eq. (11), it is clear that $\mathcal{E}_{\text{start}}$ will be too high. As a result, the microcanonical updates will lower the energy of the spins—which is desirable—at the cost of heating the daemons/walkers—which is undesirable. It is easy to assign an instantaneous temperature to the daemons/walkers, because the β -dependent equilibrium expectation value can be easily computed from Eq. (11). We have γ daemons/walkers per each of the N spins, as

$$\frac{\langle H^{\text{aux}} \rangle}{\gamma N} = \frac{4e^{-4\beta}}{1 - e^{-4\beta}} - \frac{16e^{-16\beta}}{1 - e^{-16\beta}}. \quad (12)$$

Hence, one may simply invert Eq. (12) to obtain the effective temperature from the instantaneous value of the daemon/walker energy. Since the microcanonical dynamics will increase the value of H^{aux} by relaxing the spin system, one may aptly describe the process by saying that the effective temperature of the daemons/walkers is raised because of the energy absorbed from the spins.

Now, if we aim to reach the canonical equilibrium, Eq. (5), it is clear that we need to find a way to *cool* the daemons/walkers. In some ways, our solution will be similar to the strategy known as simulated annealing [17]. At specified times, see the discussion below, we throw away the daemon/walker configuration—leaving spins untouched—and generate a new, statistically independent one from Eq. (11). By construction, the new daemon/walker configuration will have the correct effective temperature dictated by Eq. (12). In other words, the nominal

daemon/walker temperature T will coincide with the effective temperature. Hence, we shall have decreased the total energy \mathcal{E} of our *annealing* system. When equilibrium is finally reached, the energy gains and losses suffered by the spins finally compensate, and, although the effective temperature of the daemons/walkers fluctuates, it remains stable (on average).

From the above discussion, it should be clear that refreshing the daemon/walker configuration at regular times, say every 1000 timesteps (see Section 2.8 for the precise definition of timestep), should eventually lead any system of finite size to canonical equilibrium. This would dilute the need for random numbers by a factor of 1000, which would probably make the computational weight of generating random numbers irrelevant for most (if not all) applications. However, we shall try to be even thrifter with our random numbers. Indeed, the energy relaxation of spin glasses at and below their critical temperature decreases very slowly with time t . The excess energy over the equilibrium value decays as a power law, $H^{\text{EA}}(t) - \langle H^{\text{EA}} \rangle_\beta \propto N/t^{\epsilon(T)}$, where $\epsilon(T)$ is a fairly small exponent [47,58] (the exponent is so small that the power law resembles a logarithmic behavior). In fact, for a $1/\log(t)$ law, the time required to have a prefixed energy factor decay f grows exponentially with $1/f$. We decided accordingly to refresh our daemons/walkers with a schedule that is linear with the logarithm of the time. At each of those timesteps, the spin system will experience an abrupt, but tiny, energy change as a reaction to the cooler daemons/walkers. This is inappreciable at the scale of Fig. 1, but can be observed with the more refined analysis in Fig. 2-d.

2.8. Putting all pieces together: The algorithms we tested

We have tested the algorithm described above with $\gamma = 1, 2, 4$ and 6 daemons per spin. We have also tried the one with $\gamma = 1$ and 6 walkers per spin.

In all cases, a timestep consisted of γ full-lattice sweeps (*i.e.*, the flip of every spin was considered γ times within a single timestep). A sweep consists of updating the $N/2$ sites with even lattice parity, followed by the update of the $N/2$ sites with odd parity. On the α th sweep in a timestep ($\alpha = 0, 1, \dots, \gamma - 1$), the daemon/walker $n_{x,\alpha}$ was used in the update of the site x .

When applied, the walker-wandering part of the algorithm occurred once every timestep (the parity of the time index t defines which of the two wandering operations is used, recall Section 2.6). In the case with just $\gamma = 1$ walker per spin, we employed forward-X wandering. In the case with $\gamma = 6$ walkers per spin, we used

$$\begin{aligned} \alpha = 0 &\longrightarrow \text{forward-X,} \\ \alpha = 1 &\longrightarrow \text{forward-Y,} \\ \alpha = 2 &\longrightarrow \text{forward-Z,} \\ \alpha = 3 &\longrightarrow \text{backward-X,} \\ \alpha = 4 &\longrightarrow \text{backward-Y,} \\ \alpha = 5 &\longrightarrow \text{backward-Z.} \end{aligned} \quad (13)$$

In all cases, the refresh of the daemons/walkers took place at timesteps $t_s = 2^s$, as specified in Section 2.7. Observe that our definition of timestep ensures that the number of sweeps that a particular subfamily of daemons/walkers, α , completes between consecutive refreshing steps is independent of γ .

2.8.1. The algorithm in the form of a procedure

1. Initialize $s = 0$ and the time counter variable $t = 0$.
2. Get a fresh daemons/walkers configuration $\{n_{x,\alpha}\}$ from Eq. (11).
3. Carry out 2^s steps. Each step consists of:
 - (a) Cycle for $\alpha = 0, \dots, \gamma - 1$:
 - i. For every site in the even sublattice, x , compute ΔH^{EA} for the spin-flip of s_x and use $n_{x,\alpha}$ to check condition in Eq. (10). If the condition is met, flip the spin and update the daemon/walker as in Eq. (8).

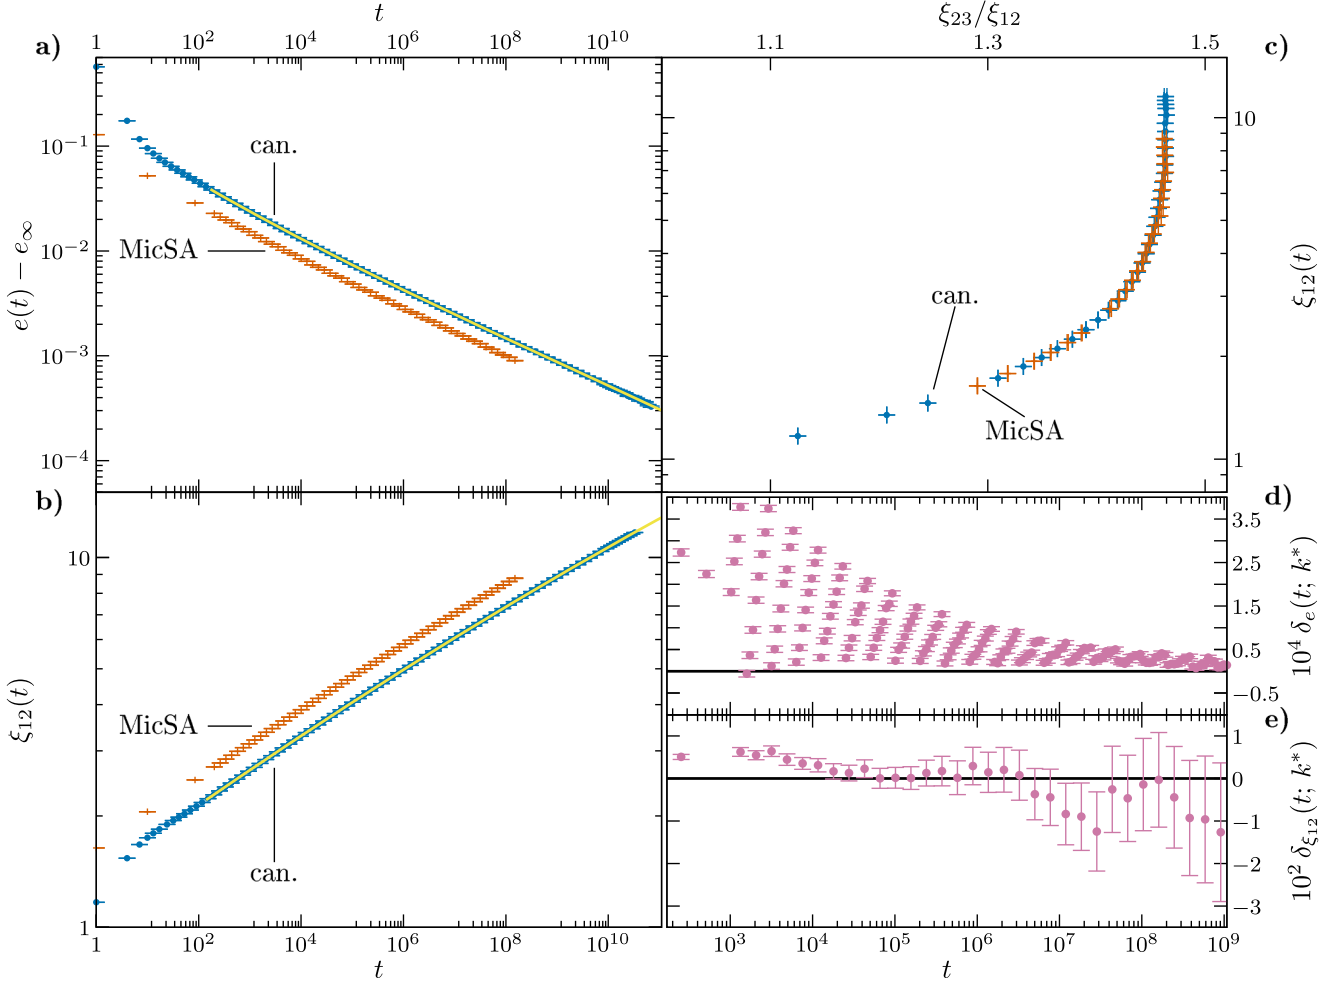


Fig. 2. Comparison of Metropolis and MicSA dynamics with 6 walkers in the spin-glass phase at $T=0.7 \approx 0.64T_c$. **a)** Time evolution of the difference of the energy at time t and the equilibrium value obtained through a fit to the canonical data. The solid yellow line is for the best fit described in Table 2. **b)** The coherence length, ξ_{12} , as a function of the simulation time for both models. The solid line is again for the best fit described in Table 2. **c)** Coherence length, $\xi_{12}(t)$, as a function of the ratio between the coherence lengths $\xi_{23}(t)$ and $\xi_{12}(t)$. **d)** Difference between the results from the two algorithms for the energy, Eq. (21), as a function of the simulation time. **e)** As in **d)** but for the coherence length. To compute these differences, we have considered the best fit to our canonical data—solid lines in **a)** and **b)**—and the rescaling coefficient $k^* = 6.034$ (see Appendix A for more information on the estimation of k^*). Error bars are smaller than data points. For the sake of clarity points are plotted at regular time intervals in logarithmic scale.

ii. Same as (i) for the odd sublattice.

(b) If the $\{n_{x,\alpha}\}$ are walkers, carry out the part of the wandering cycle dictated by the parity of t , see Section 2.6 (for daemons, do nothing).

(c) Increase t by one unit.

4. Increase s by one unit and go back to (2).

We insist on updating all sites with a given parity first, and then the sites with the opposite lattice parity, just because we intend to carry out a parallel computation. Readers willing to work with a serial update may access the lattice in whatever order they prefer.

3. Results

We have tested the new set of algorithms under three very different dynamical conditions. In all three cases, we start the simulation from a fully disordered configuration (which mimics infinite temperature). At the starting time, $t = 0$, the system is instantaneously quenched to the working temperature T , where it is allowed to relax. The features of the relaxation are very much dependent on T . In the paramagnetic phase, $T > T_c$, thermal equilibrium can eventually be reached. Instead, equilibrium is unreachable in a large lattice if one works at the criti-

cal point $T = T_c$, or in the spin-glass phase $T < T_c$. This impossibility to equilibrate is, exactly, the situation encountered in experiments [6]. For $T \leq T_c$ the dynamics consists in the endless growth of glassy magnetic domains of linear size $\xi(t, T)$ (when comparing data at the same temperature we shall use the lighter notation $\xi(t)$; a similar convention will be used for other quantities). The time growth of the coherence length $\xi(t, T)$ turns out to be excruciatingly slow, and slower the lower T is. For $T > T_c$, instead, $\xi(t, T)$ eventually saturates to its equilibrium value $\xi_{eq}(T)$. In order to be representative of the experimental situation, one should have $L \gg \xi(t, T)$ at all times [59]. We shall not feel constrained by the $L \gg \xi(t, T)$ condition, because we want to test our algorithms also when L and $\xi(t, T)$ become comparable.

To investigate the domain-growth process briefly described above, it is mandatory to use replicas, namely statistically independent system copies $\{s_x^{(a)}\}$ that evolve under the very same Hamiltonian H_x^{EA} , Eq. (1), i.e., with the same coupling matrix $\{J_{x,y}\}$. The superscript $a = 1, 2, \dots, N_R$ labels the N_R replicas in a simulation. For any quantity, O , we shall use $\langle O \rangle$ for its average over the initial conditions. In particular, for different replica indices $a \neq b$, one has

$$\langle s_x^{(a)}(t) s_y^{(a)}(t) s_x^{(b)}(t) s_y^{(b)}(t) \rangle = \langle s_x(t) s_y(t) \rangle^2, \quad (14)$$

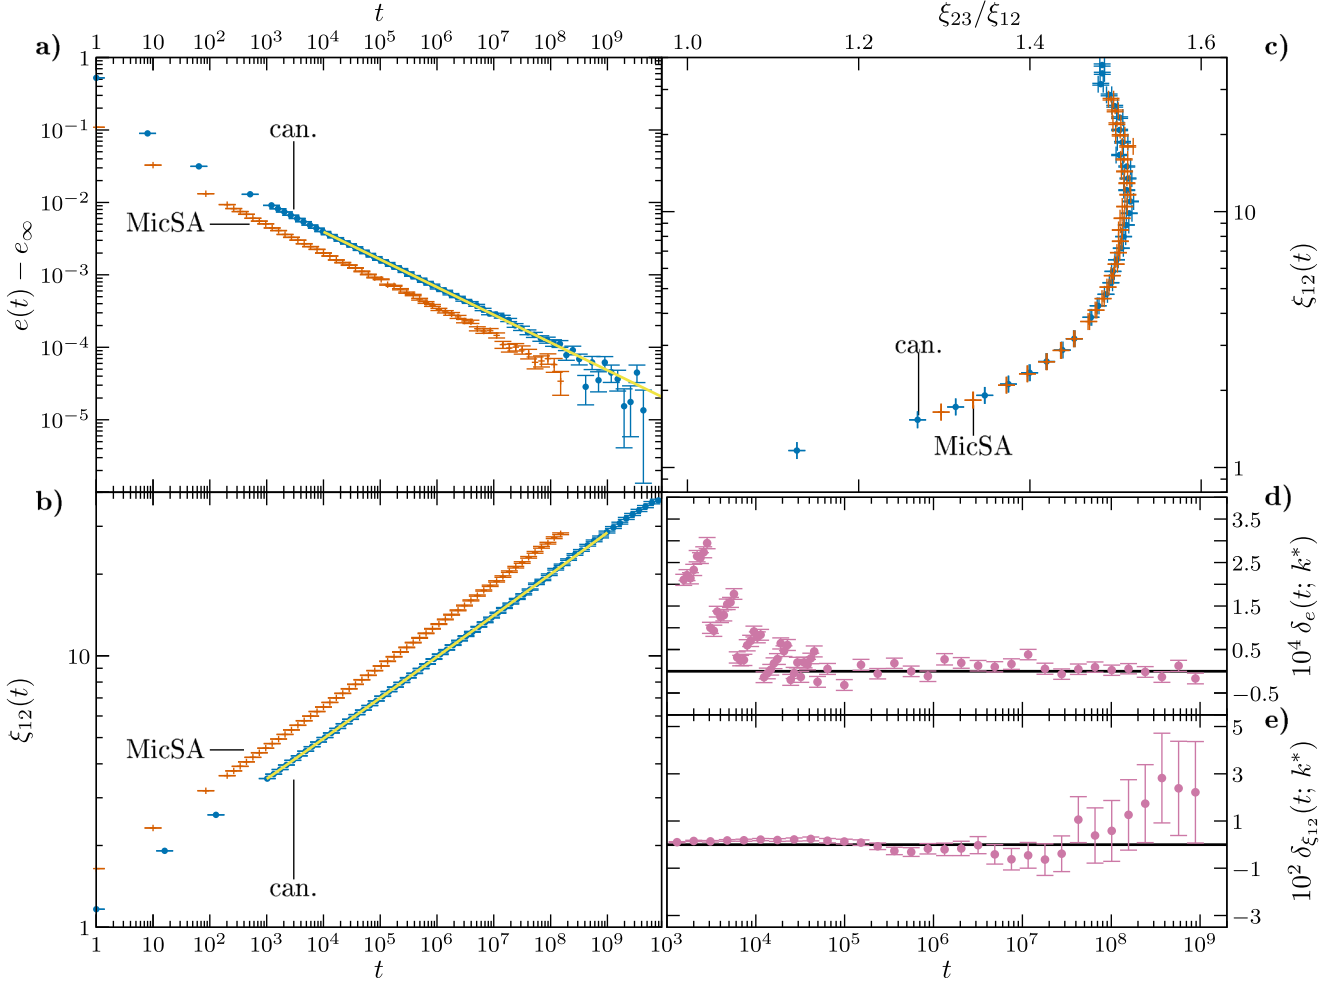


Fig. 3. Six walkers at $T = 1.1 \approx T_c$. As in Fig. 2, but with the time-rescaling factor $k^* = 5.876$ (see Table 1).

Table 1

Estimation of the best time-rescaling factor k^* as obtained from a fit to Eq. (20) for the different number of daemons/walkers per spin γ and temperatures (we mark with a symbol \dagger the algorithm based on walkers and not on daemons). We denote the statistical uncertainty by Δk^* . We also report the fits figure of merit $\chi^2/\text{d.o.f}$ as computed from the diagonal part of the covariance matrix (d.o.f. meaning degrees of freedom). The choice of using either daemons or walkers (and how many) results only in a rescaling of the time units, but not in a change of the dynamical critical exponents.

T	γ	k^*	Δk^*	$[\xi_{\min}, \xi_{\max}]$	$\chi^2/\text{d.o.f.}$
0.7	6^\dagger	6.034	0.007	[4, 12.5]	34.1/108
T_c	1	0.3984	0.0012	[8.5, 14]	25.8/33
T_c	1^\dagger	0.967	0.003	[8, 12]	15.4/25
T_c	2	1.051	0.002	[7, 12]	21.3/36
T_c	4	2.667	0.004	[6, 14]	13.2/59
T_c	6	4.79	0.04	[18, 30]	8.4/31
T_c	6^\dagger	5.876	0.011	[7, 30]	66.0/106
1.4	6^\dagger	5.904	0.014	$t > 10^3$	84.5/95

because the different replicas are independent, identically distributed random variables.

The main quantities we have used to investigate the dynamics are the energy per spin $e(t)$ and the correlation function $C_4(\mathbf{r}, t)$. For both quantities, we give the formal definition and, after the \approx sign, the estimators

that one can compute with a finite number of replicas N_R :

$$e(t) = \frac{1}{N} \langle H^{\text{EA}}(\{s_x(t)\}) \rangle, \quad (15)$$

$$\approx \frac{1}{N N_R} \sum_{a=1}^{N_R} H^{\text{EA}}(\{s_x^{(a)}(t)\}),$$

$$C_4(\mathbf{r}, t) = \frac{1}{N} \sum_{\mathbf{y}=\mathbf{x}+\mathbf{r}} \langle s_x(t) s_y(t) \rangle^2, \quad (16)$$

$$\approx \frac{1}{N} \sum_{a \neq b} \sum_{\mathbf{y}=\mathbf{x}+\mathbf{r}} s_x^{(a)}(t) s_y^{(a)}(t) s_x^{(b)}(t) s_y^{(b)}(t),$$

where the normalizing factor $\tilde{N} = N N_R (N_R - 1)/2$ takes into account the different ways of choosing two different replica indices. We also have employed the periodic boundary conditions when computing $\mathbf{y} - \mathbf{x}$. Rotational invariance [58] suggests choosing the displacement vector \mathbf{r} along one of the lattice axes. In fact, we average over the three different choices $\mathbf{r} = (r, 0, 0)$, $(0, r, 0)$ or $(0, 0, r)$ and name the result simply $C_4(\mathbf{r}, t)$.

Both $e(t)$ and $C_4(\mathbf{r}, t)$ are self-averaging quantities, which means that their sample-to-sample fluctuations vanish for large L as a power of the ratio $\xi(t)/L$ —recall that a sample is a realization of the coupling matrix $\{J_{x,y}\}$ in Eq. (1). In fact, experiments are typically performed on a single sample. We have decided to mimic this approach and work with just one sample, obviously the same one for all algorithms (a sample already studied in previous investigations [46,47,49–52,60]), because our target is obtaining very precise comparisons. Eliminating the noise induced by the sample-to-sample fluctuations helps us achieve this purpose and provides a much stronger test of the equivalence between formalisms.

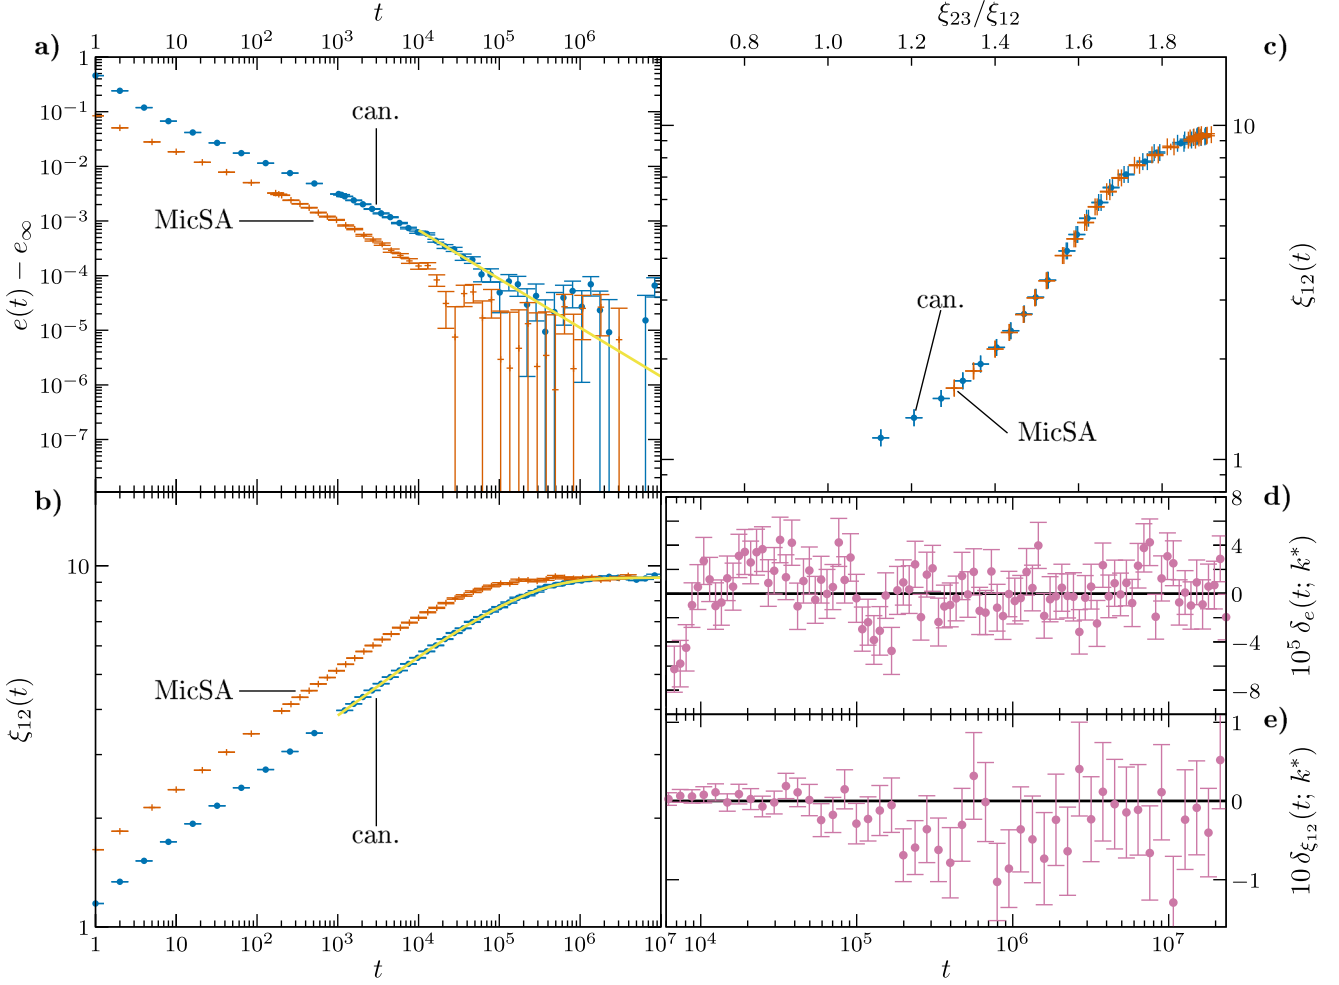


Fig. 4. Six walkers in the paramagnetic phase at $T = 1.4 \approx 1.27T_c$. As in Fig. 2, but with the time-rescaling factor $k^* = 5.904$ (see Table 1).

However, while comparing the curves $e(t)$ from different algorithms is relatively straightforward, there is too much information in the $C_4(r, t)$. To simplify the analysis, we squeeze the r -dependence of the correlation function into two numbers, the coherence length and a dimensionless ratio of characteristic lengths. Specifically, we follow Refs. [58,59] and compute the integrals

$$I_n = \int_0^\infty dr r^n C_4(r, t). \quad (17)$$

The integral converges because at large distances we have $C_4(r, t) \sim f(r/\xi(t))/r^\theta$, where $f(x)$ is a cut-off function decaying super-exponentially in x . For technical details about estimating the integrals I_n , we refer the reader to Refs. [46,61]. The characteristic lengths $\xi_{n,n+1}$ are computed from the I_n ,

$$\xi_{n,n+1} = \frac{I_{n+1}}{I_n} \quad (18)$$

We shall follow the (by now) standard practice of identifying the spin-glass coherence length with $\xi_{12}(t)$. As for the overall r dependence of $C_4(r, t)$ we consider the dimensionless ratio ξ_{23}/ξ_{12} [27,62]. $C_4(r, t)$ is expected to reach a scale-invariant limit at long times [i.e., large $\xi_{12}(t)$] when $T \leq T_c$, which implies that the ratio ξ_{23}/ξ_{12} will reach a long-time limit as well [46].

Our estimates for $e(t)$, $\xi_{12}(t)$ and the dimensionless ratio $\xi_{23}/\xi_{12}(t)$ are displayed as a function of time in Fig. 2 ($T = 0.7 \approx 0.64T_c$), Fig. 3 ($T = 1.1 \approx T_c$) and Fig. 4 ($T = 1.4 \approx 1.27T_c$). For all three temperatures, we observe that, when $\xi_{12}(t)$ is plotted parametrically as a function of $\xi_{23}/\xi_{12}(t)$, the canonical and the MicSA data fall onto a single curve. This

finding indicates that it suffices to consider $\xi_{12}(t)$ when trying to find the correspondence between the $C_4(r, t)$ computed with the two dynamics. The figures suggest as well that a shift in $\log t$ (i.e., a simple rescaling for t) suffices to superimpose the MicSA data with the canonical curve for both $e(t)$ and $\xi_{12}(t)$. We shall attempt to make this observation quantitatively precise.

The technical difficulty to address is that $e(t)$ and $\xi_{12}(t)$ are known only on a discrete time grid. We overcome this difficulty by interpolating the curves from the canonical simulations with a simple, physically motivated continuous function of t . The interpolating function is found from a fit to

$$\xi_{12}^{\text{can}}(t, T) = F_{\text{can}}^\xi(t, T), \quad e^{\text{can}}(t, T) = F_{\text{can}}^e(t, T). \quad (19)$$

The form of the interpolating functions varies very much with temperature, as explained in Appendix A, where details on the above fits are provided (in particular, for $T \leq T_c$ it is simpler to find t as a function of ξ_{12} than the other way around). The reader should take neither $F_{\text{can}}^\xi(T, t)$ nor $F_{\text{can}}^e(T, t)$ too far, because they are mere interpolating devices that are not good for extrapolations. Among other problems, these functions represent the results of a single sample, including times when—for $T \geq T_c$ —the condition $L \gg \xi(T, t)$ is no longer met.

Having promoted $\xi_{12}^{\text{can}}(t, T)$ to a continuous function, we can simply fit the MicSA curves to

$$\xi_{12}^{\text{MicSA}}(t, T) = F_{\text{can}}^\xi(k t, T), \quad (20)$$

with the rescaling factor k as the only fitting parameter. The best-fit rescaling parameter is named k^* in Table 1 and in Figs. 2–4. In the

Table 2

Results of the different fits to our canonical data shown in Figs. 1–4. These fits correspond to the interpolation used to estimate $\delta_e(t, k)$ and $\delta_\xi(t, k)$ [see Eqs. (21) and (22)] explained in Section 3. We specify the time interval considered in the fits as $[t_{\min}, t_{\max}]$ (* indicates that all times $t \geq t_{\min}$ are included in the fit).

\mathcal{O}	T	$F_{\text{can}}^\#(t)$	$[t_{\min}, t_{\max}]$	Fit	$\chi^2/\text{d.o.f.}$
e	0.7	$e_\infty + a_1 t^{-z_1} + a_2 t^{-z_2}$	$[1.7 \times 10^2, *]$	$e_\infty = -1.7708874(19)$, $a_1 = 0.0884(7)$, $a_2 = 0.1060(3)$, $z_1 = 0.2235(5)$, $z_2 = 0.445(3)$	242.4/248
e	T_c	$e_\infty + a_1 t^{-z_1}$	$[10^4, *]$	$e_\infty = -1.7022747(19)$, $a_1 = 0.1325(12)$, $z_1 = 0.3820(9)$	152.4/149
e	1.4	$e_\infty + a_1 t^{-z_1}$	$[10^4, *]$	$e_\infty = -1.607026(4)$, $a_1 = 2.5(9)$, $z_1 = 0.90(4)$	91.7/83
ξ_{12}	0.7	$a_0 + a_1 t^z + a_2 \log(t)$	$[1.4 \times 10^2, *]$	$a_0 = -2.71(3)$, $a_1 = 4.10(3)$, $a_2 = -0.1179(14)$, $z = 0.05930(17)$	197.6/235
ξ_{12}	T_c	$a_0 + a_1 t^z$	$[10^3, 10^9]$	$a_0 = 0.159(8)$, $a_1 = 1.158(4)$, $z = 0.1543(2)$	126.4/157
ξ_{12}	1.4	$C \left(1 - \sum_{i=0}^7 a_i e^{-t^i/\tau}\right)$	$[10^3, *]$	$C = 9.29(2)$, $\tau = 8(3) \times 10^6$, $b = 4.0(4)$, $a_0 = 0.007(4)$, $a_1 = 0$, $a_2 = 0.11(3)$, $a_3 = 0.166(17)$, $a_4 = 0.136(9)$, $a_5 = 0.123(9)$, $a_6 = 0.095(7)$, $a_7 = 0.069(8)$	49.4/104

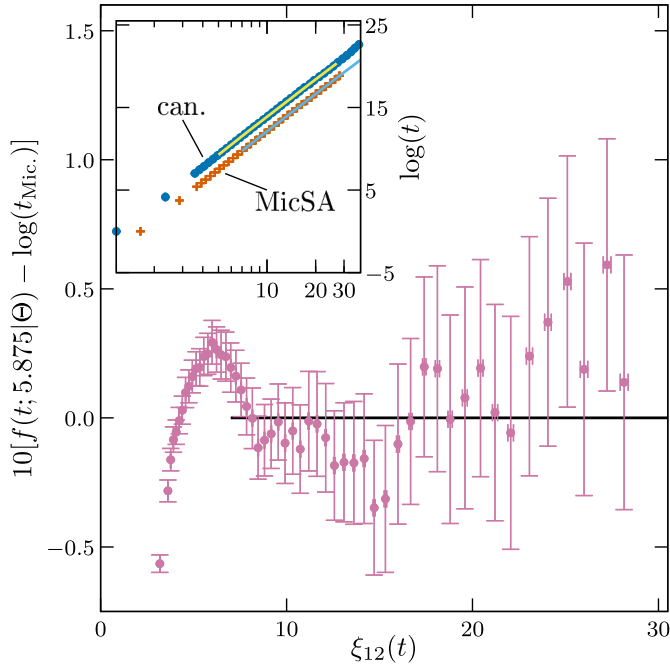


Fig. 5. Estimation of the best time shift k^* at T_c . In the main panel, the difference between the logarithm of the original measured time in the MicSA data and the interpolating function $f(k^*t)$ as a function of the coherence length ξ_{12} . The horizontal solid line at zero indicates the interval of the fit used to estimate k^* (see Table 1 for details). Inset: $\log(t)$ as a function of ξ_{12} for both algorithms. Solid lines indicate the best fit for our data. In the case of the MicSA data, the only fitting parameter is k , while the rest of the parameters $\Theta = \{a_0, a_1, z\}$ come from the fit to the canonical data described in Appendix A.

figures, we have used for the energy the same value of k^* computed with ξ_{12} . Each of the different variants of the MicSA algorithm will have its own k^* that slightly depends on the temperature. For further details on the computation of k^* , see Fig. 5 and Appendix A.

Ideally speaking, we would hope k^* to be close to γ , because in a single step of the microcanonical simulated annealing algorithm, every spin is updated γ times (to be compared with just one spin update per canonical timestep). A direct inspection of Table 1 reveals that only the variants of the algorithm with walkers get close to this ideal. When one works with daemons rather than walkers, k^* turns out to be significantly smaller than γ . Yet, as we shall discuss in the Conclusions, in order to choose the most convenient algorithmic variant, one should also take into account the computational cost of the different parts of the algorithm (see Appendix B).

To better ascertain the quality of the time rescaling, we show in the panels **d** and **e** of Figs. 2–4 the differences

$$\delta_e(t; k) = e^{\text{MicSA}}(t) - F_{\text{can}}^e(k t, T), \quad (21)$$

$$\delta_\xi(t; k) = \xi_{12}^{\text{MicSA}}(t) - F_{\text{can}}^\xi(k t, T), \quad (22)$$

as computed for $k = k^*$. These figures show that, when looked at with enough accuracy, $e^{\text{MicSA}}(t)$ is not a smooth function of time (this is most clear for $T \leq T_c$). Indeed, these curves have abrupt (but tiny) changes every time the walkers get refreshed. This is why we have not attempted to get the rescaling factor from the energy. These abrupt changes, as well as the difference $\delta_e(t)$ itself, strongly decrease as the time t increases.

4. Conclusions

We have proposed a new Monte Carlo algorithm, microcanonical simulated annealing (MicSA), that is both fully adapted to massively parallel computations and extremely frugal in its use of random numbers. The method is fully general and has a certain number of adaptable features. Chief among these are the number γ of daemons/walkers per elementary degree of freedom (spins in our case) and the choice between using daemons or walkers. Walkers have the advantage of avoiding local conservation laws and their resulting slower dynamics, but one can decide whether to use daemons or walkers, and how many, on the view of the application at hand and the hardware platform that will be used.

The method performs fully satisfactorily in the demanding setting of the three-dimensional Ising spin glass, where we have shown that it can reproduce the results of a standard (random-number-greedy) Metropolis algorithm both in and out of equilibrium. In particular, we have shown that the out-of-equilibrium dynamics derived from MicSA can be rescaled onto the one from traditional Metropolis by a simple change in the time unit. The fact that, in the case of walkers, the rescaling factor is nearly equal to the ratio of the number of spin updates in a single timestep for both methods (which coincides with γ) tells us that a single step of MicSA is worth γ Metropolis sweeps, without requiring a single random number.

Our CUDA implementation demonstrates the parallel-friendly nature of the algorithm. The code, see Appendix B, is publicly available and may run both standard Metropolis and MicSA. It achieves a performance that could previously only be provided by the custom-built Janus II computer.² Table 3 shows the time required for the execution of the two main kernels of the standard Metropolis algorithm and the MicSA algorithm on a lattice with 160^3 spins and 128 replicas using a single walker. Using more walkers, the spin-update time increases linearly as expected (data not reported).

² In the context of quantum spin glasses, performance competitive with Janus II has already been obtained on GPU [25].

Table 3

Metropolis and MicSA timings for a full sweep of 128 replicas of a lattice with 160^3 spins. The test was executed on an Nvidia H100 GPU. The total time of the Metropolis sweep is 1.9 times longer than that of MicSA.

Method	Kernel	Time (ns)
Metropolis	random-number generation	301 363
	spin update	113 510
MicSA	spin update	121 554
	walking procedure	97 378

By adding the timings for the two kernels of the new algorithm and taking into account the total number of spins, we see that the spin-update time is ≈ 0.4 picoseconds. This value is much lower than the 3.17 picosecond per spin-update obtained on a single Janus II FPGA. However, the above two numbers should be compared with care for two reasons. First, the computing power on the GPU is spread over 128 replicas, while on Janus II updates are concentrated on only two replicas. Hence, for a high-accuracy simulation spanning a limited time range (say up to $t \sim 10^9$), the GPU is faster. If, however, one needs to reach times $t \sim 10^{11}$ or larger (which is not uncommon [6]), Janus II still fares better (the GPU might make more spin updates per second by running many replicas, but would need a wall-clock time of many months to finish each simulation). Second, we are, somewhat unfairly, comparing 2025 high-end GPUs with the 2013 Virtex 7 FPGAs that Janus II is made of.

Since MicSA is particularly suitable for high-end FPGAs or even dedicated hardware (based on an ASIC, for instance), we expect that users interested in extreme-scale Monte Carlo simulations will find a significant advantage in our proposal.

The new algorithm raises some intriguing questions, some of which are more theoretical. Let us mention a few of them.

The performance of MicSA for system sizes L small enough to be brought to thermal equilibrium also in the spin-glass phase, $T < T_c$, needs to be ascertained. In particular, MicSA needs to be combined with Parallel Tempering [63]. We are already testing this algorithm combination for small systems, see Appendix C and Ref. [64].

The dynamical properties of MicSA in small systems are likely to be quite peculiar. If the daemon/walkers are refreshed at a very low rate, the system may have enough time to equilibrate in the time lapse between the refreshing steps. Since the equilibrium values of the different observables differ in the canonical and the microcanonical ensemble, it is not obvious that it will always be possible to map MicSA dynamics onto canonical dynamics for small systems. Of course, in the long run, canonical equilibrium values will be recovered by MicSA because the daemons/walkers' refresh will vary the total energy.

The performance of MicSA should be investigated in problems where the elementary excitation (which is a spin flip in our case) will cause a continuously distributed energy change. The simplest example would be, perhaps, a spin glass with Gaussian-distributed couplings. As Appendix C shows, choosing the daemons/walkers as continuous, exponentially distributed random variables fully solves the problem.

Another interesting issue regards the optimal refreshing rate for the daemons/walkers. As we have already explained, if the refresh occurs at a constant time rate, we expect dynamical rescalability with Metropolis. If we refresh at every step with daemons/walkers of infinite capacity, MicSA is the Metropolis algorithm (if the capacity is limited, the resulting algorithm is just an unusual version of Metropolis). Of course, if the refreshing rate is too low, microcanonical equilibration will occur in between the refreshing steps, which would complicate the dynamical analysis. However, refreshing once every (say) 1000 steps would bring to microcanonical equilibrium only a fairly small system (and the cost of generating random numbers would be divided by a factor of 1000 as compared to a traditional simulation).

One may wonder how far the extreme strategy followed here can be pushed (recall that we refreshed the daemons/walkers in an exponentially spaced time schedule, namely at times $t_s = 2^s$ for $s = 0, 1, 2, 3, \dots$). Given that in, critical systems with a divergent specific heat (e.g., the three-dimensional ferromagnetic Ising model), the microcanonical dynamics cannot be rescaled to coincide with the canonical one [40], it is conceivable that the minimal daemon/walker time schedule allowing dynamic scalability will be less spaced than for the spin glass (perhaps $t_s \propto s^a$, with a some appropriate exponent, would do the job for the three-dimensional Ising model).

Data availability

Our CUDA code and the data required to reproduce the figures are publicly available at <https://github.com/IsidoroGlez/micSA-EA-MC/releases/tag/v1.1>

CRediT authorship contribution statement

M. Bernaschi: Writing – original draft, Software, Funding acquisition, Conceptualization; **C. Chilin:** Software, Investigation, Conceptualization; **L.A. Fernandez:** Investigation, Conceptualization; **I. González-Adalid Pemartín:** Writing – original draft, Software, Investigation, Conceptualization; **E. Marinari:** Writing – original draft, Funding acquisition, Conceptualization; **V. Martin-Mayor:** Writing – original draft, Software, Funding acquisition, Conceptualization; **G. Parisi:** Writing – review & editing, Conceptualization; **F. Ricci-Tersenghi:** Writing – review & editing, Funding acquisition, Conceptualization; **J.J. Ruiz-Lorenzo:** Writing – original draft, Investigation, Funding acquisition, Conceptualization; **D. Yllanes:** Writing – original draft, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Given his role as Specialist Editor for Condensed Matter Physics in this journal, co-author Massimo Bernaschi had no involvement in the peer review of this article and had no access to information regarding its peer review. Full responsibility for the editorial process for this article was delegated to another journal editor. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank Martin Hasenbusch for calling our attention to Refs. [33–36] and for kindly explaining to us some of the details in these works. This work was partially supported by **Ministerio de Ciencia, Innovación y Universidades (Spain)** and by the **European Regional Development Fund (MCIU/AEI/10.13039/501100011033/FEDER, UE)** through grants no. **PID2022-136374NB-C21**, **PID2024-156352NB-I00**, **PID2024-158623NB-C22** and **RED2022-134244-T**; by **Junta de Extremadura (Spain)** through grant no. **GR24022**; and by funding from the 2021 first FIS (Fondo Italiano per la Scienza) funding scheme (FIS783 - SMaC - Statistical Mechanics and Complexity) from Italian MUR (Ministry of University and Research). We acknowledge the use of the CESAR computational resources at the BIFI Institute (University of Zaragoza). IG, MB and FRT acknowledge the support of the National Center for HPC, Big Data and Quantum Computing, Project CN_00000013 – CUP E83C22003230001, CUP B93C22000620006 and CUP B83C22002940006, Mission 4 Component 2 Investment 1.4, funded by the European Union – NextGenerationEU.

Appendix A. Computation of the best time shift

As we have explained in Section 3, to estimate the best value of k^* , we use the coherence length ξ_{12} . A similar analysis can be performed with the energy, but, as it is clear from panel **d**) in Fig. 2, the energy in the MicSA algorithm does not decrease to its equilibrium value monotonically. Instead, for short times, it is sensitive to the daemons' refreshing process. This behavior makes it more challenging to estimate k^* from these data. Nonetheless, the relation between the coherence length and time in the spin-glass phase, which can be extended to the critical point, offers a straightforward way to estimate k^* .

In particular, at the critical point, we can use the following relation [46]

$$\log(t) = a_0 + z \log(\xi_{12}) + \frac{a_1}{\xi_{12}^\omega} \quad (\text{A.1})$$

where ω is the exponent characterizing the first corrections to leading scaling and it takes the value $\omega = 1.12(10)$ [43].

For $T < T_c$, the leading behavior is given by $\omega = \theta$, where θ is the replicon exponent (we used $\omega = 0.35$ in the fits [46]).

As the reader will notice, it is easy to include in Eq. (A.1) a time shift k that connects our MicSA data with the canonical curve:

$$\log(t) = a'_0 - \log(k) + z' \log(\xi_{12}) + \frac{a'_1}{\xi_{12}^{\omega'}}. \quad (\text{A.2})$$

In this situation, we propose the following method to extract the best time shift, k^* :

1. Fit the canonical data to Eq. (A.1) in order to obtain parameters a_0 , a_1 , and z .
2. With parameters fixed to the values obtained in the previous step ($a'_0 = a_0$, $a'_1 = a_1$, and $z' = z$), we perform a new fit to the MicSA data to Eq. (A.2).

This process guarantees that the only fitting parameter in step 2 is k , obtaining k^* . Table 1 summarizes the results of $T \leq T_c$. Fig. 5 reproduces this procedure. In the main panel, we report the difference between our MicSA data $\log(t_{\text{MicSA}})$ and the interpolation $f(t; k^* | \Theta)$, see Eq. (A.2), with the best time shift $k^* = 5.875$ and the parameters $\Theta = (a_0, a_1, z)$ obtained following this process for $\gamma = 6$ and $T = T_c$. The inset shows the non-subtracted data.

This approach is, however, not useful at $T > T_c$, because Eq. (A.1) does not hold. In this regime, we interpolate $\xi_{12}(t)$ as a function of time using the following expression

$$F_{\text{can}}^\xi(t) = C \left(1 - \sum_{i=0}^7 a_i e^{-t b^i / \tau} \right). \quad (\text{A.3})$$

In this situation, we can use a similar protocol to extract the value of k^* :

1. We fit our canonical data to Eq. (A.3) to obtain the values of the parameters (see the fit result in Table 2).
2. Using the parameters obtained in the previous step, we fit our MicSA data to $F_{\text{can}}^\xi(k t)$, where k is the only fitting parameter.

As in the case of $T \leq T_c$, this process allows us to obtain directly the best time shift k^* , even in the case of $T > T_c$. Notice that, in this situation, t is the independent variable. Then, in Table 1 we reported the fitting range in t , instead of ξ_{12} . Notice that, although we use a slightly different process to obtain the value of k^* for $T > T_c$, panels **d** and **e** in Fig. 4 reported a value of δ_e and $\delta_{\xi_{12}}$ compatible with zero at long times.

Appendix B. The CUDA implementation

In the following, we provide basic technical information about the CUDA implementation of the MicSA algorithm described in Section 2 applied to the three-dimensional Ising spin glass. The code implements a two-level parallelism: we use multispin coding to pack 64 spins in 64-bit words (`long long int`) and then CUDA threads for the concurrent

update of *even* and *odd* spins. Multiple *replicas* and *samples* can be simulated in a single execution; the actual number of both depends on the resources (mainly memory) available on the GPU (we successfully tested up to 512 replicas). The code also implements the standard Metropolis algorithm using random numbers generated on the GPU with the procedure described in Ref. [25]. The refresh of the daemons/walkers is carried out on the CPU asynchronously to the microcanonical update (we recall that the refresh takes place according to an exponential schedule) using CPU threads. The walking procedure and the microcanonical update are implemented in two separate CUDA kernels to increase flexibility. The main reason is that the microcanonical update is carried out separately for each daemon/walker, whereas the walking procedure is completed for all the daemons in a single kernel execution.

The walking procedure takes $\approx 44\%$ of the execution time on the GPU. In contrast, random-number generation takes the majority ($\approx 72\%$) of the time in the Metropolis algorithm. We argue that the higher the peak performance of the platform, the greater the advantage provided by minimizing the use of random numbers. In particular, the design of an *Application Specific Integrated Circuit* (ASIC) would be greatly simplified under the assumption that generating random numbers is not an issue.

Further details about our code can be obtained by direct inspection of the source available from [65].

Appendix C. A Parallel Tempering example (walkers in equilibrium simulations)

In this section, we provide an example of the possible use of the walkers for a simulation at equilibrium, which radically differs from the off-equilibrium dynamics studied in the main text. Bringing a spin-glass sample to equilibrium is possible only if the linear size of the system is tiny (particularly at low temperatures). Furthermore, special simulation methods, such as Parallel Tempering [63], are mandatory. The particular problem that we shall investigate here is a simple variant of Eq. (1) in which the couplings are drawn from a Gaussian distribution $J \sim \mathcal{N}(0, 1)$. The rationale for choosing Gaussian-distributed couplings is that we wish to test our methods in a demanding situation, namely at the very low temperature $T = 0.2 \approx 0.21 T_c$ (T_c is the critical temperature). Indeed, continuously distributed couplings are usually preferable for very-low-temperature studies.

Specifically, we shall study in detail a single sample of a 3D system with lattice linear size $L = 10$. For this sample, the walkers method reproduces the equilibrium values of the energy per spin [*i.e.*, the spin Hamiltonian in Eq. (1) divided by the number of spins N] and of the overlap between the instantaneous value of two replicas

$$q = \frac{1}{N} \sum_{\mathbf{x}} s_{\mathbf{x}}^{(a)} s_{\mathbf{x}}^{(b)}. \quad (\text{C.1})$$

In order to do this, we adapted the walker method to the Parallel Tempering procedure enhanced with Houdayer moves [66], which has been proven to be (mildly) effective also for the equilibrium simulation of three-dimensional spin glasses [67], even though our schedule is a bit different [64]. We are then running $N_R = 2$ replicas of $N_\beta = 40$ clones with temperatures uniformly distributed between $T_{\min} = 1/\beta_{\max} = 0.2$ and $T_{\max} = 1/\beta_{\min} = 2$.

In this case, our method requires the use of continuous walker variables, distributed as in Eq. (5). Additionally, we do not need to limit the capacity of the walkers as in Section 2.5, but we can take a step back and use the elementary move described in Section 2.3.

For each replica/clone, we simulate a system with one walker per spin ($\gamma = 1$) that gets refreshed at each Elementary Monte Carlo Step (EMCS) as described below.

The EMCS is structured as follows:

1. For each clone and for each replica, perform a fixed-temperature move:
 - (a) Generate the walkers according to Eq. (5).

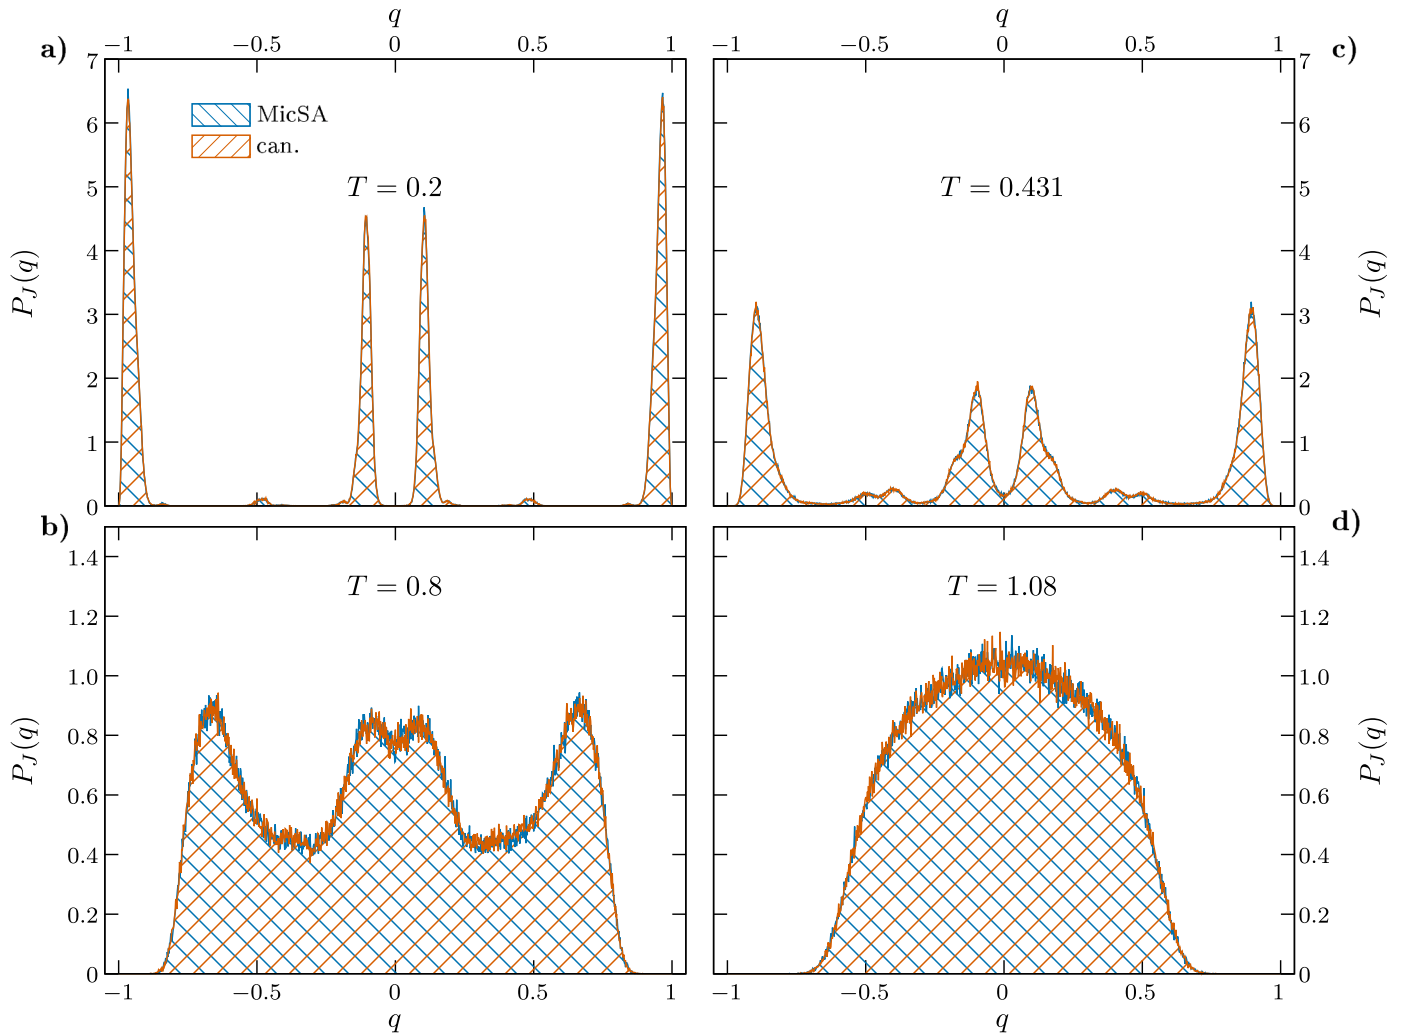


Fig. C.1. Overlap distribution of some selected temperatures. The curves for each method are obtained by simulating 200 times the same sample for 250000 EMCS (half of which have been used for prethermalization, meaning that we keep data only from the second half of the simulation) and later averaging over the 200 curves obtained. The simulation parameters are $L = 10$, $N_R = 2$, $N_\beta = 40$, temperatures uniformly distributed between $T_{\min} = 1/\beta_{\max} = 0.2$ and $T_{\max} = 1/\beta_{\min} = 2$.

Table C.1

Measures of the equilibrium expectation values for $|q|$, e_∞ as extracted from the simulations described in Appendix C, at the lowest temperature $T_{\min} = 0.2$. For each of the two algorithms, we carried out 200 independent runs. Errors are straightforwardly computed by considering the fluctuations between runs.

	$ q $	error	e_∞	error
Walkers	0.743 8	0.012 9	-1.696 330 17	0.000 001 02
Metropolis	0.741 6	0.010 4	-1.696 331 13	0.000 000 89
Difference	0.002 19	0.016 7	0.000 000 96	0.000 001 36

- (b) Alternate a full-lattice sweep with a wandering of the walkers towards the forward-X direction. Repeat 10 times.
2. Perform Houdayer's move.
3. Perform a Parallel Tempering step between the N_β clones belonging to the same replica group.

Fig. C.1 qualitatively shows how the probability distribution of q computed in this manner reproduces the results of using a canonical Metropolis update instead of point 1 of the above algorithm.

As a quantitative test, we have also evaluated the average energy per spin and the average overlap (in absolute value) yielded by both sim-

ulations. The obtained values are summarized in **Table C.1**, which displays good agreement between the two algorithms. Furthermore, from the ratio between the errors we can estimate the ratio of the integrated autocorrelation time for both algorithms [53]. We obtain for $|q|$ the ratio $\tau_{\text{walkers}}/\tau_{\text{Metropolis}} \simeq 1.7$, meaning that, in this setting, one EMCS with the walkers is roughly equivalent to 0.59 steps with the usual Metropolis. We can understand which of the two ways is more effective only by measuring also the CPU time for one EMCS, which depends not only on hardware, but also on parameters of the simulation such as the linear size L or the number of temperatures N_β of the Parallel Tempering.³

References

- [1] M.H. Kalos, P.A. Whitlock, *Monte Carlo Methods*, 2nd Edition Wiley-VCH, Reading, Massachusetts, 2008.
- [2] J.A. Mydosh, *Spin Glasses: an Experimental Introduction*, Taylor and Francis, London, 1993 <https://doi.org/10.1201/9781482295191>
- [3] A.P. Young, *Spin Glasses and Random Fields*, World Scientific, 1998. <https://doi.org/10.1142/3517>

³ For instance, our partially multi-spin coded program runs about twice as fast with walkers as with Metropolis when $L = 18$ on an AMD EPYC 7763 64-Core Processor; the computational advantage of the walker algorithm seems to grow with L .

- [4] P. Charbonneau, E. Marinari, M. Mézard, G. Parisi, F. Ricci-Tersenghi, G. Sicuro, F. Zamponi (Eds.), *Spin Glass Theory and Far Beyond*, World Scientific, Singapore, 2023. <https://doi.org/10.1142/13341>
- [5] G. Parisi, Nobel lecture: multiple equilibria, *Rev. Mod. Phys.* 95 (2023) 30501. <https://doi.org/10.1103/RevModPhys.95.030501>
- [6] E.D. Dahlberg, I. González-Adalid Pemartín, E. Marinari, V. Martín-Mayor, J. Moreno-Gordo, R.L. Orbach, I. Paga, G. Parisi, F. Ricci-Tersenghi, J. Ruiz-Lorenzo, D. Yllanes, Spin-glass dynamics: experiment, theory and simulation, *Rev. Mod. Phys.* 97 (2025) 045005. <https://doi.org/10.1103/ctp2-zwyr>
- [7] M. Aramon, G. Rosenberg, E. Valiante, T. Miyazawa, H. Tamura, H.G. Katzgraber, Physics-inspired optimization for quadratic unconstrained problems using a digital annealer, *Front. Phys.* 7 (2019) 00048. <https://doi.org/10.3389/fphy.2019.00048>
- [8] H. Goto, K. Tatsumura, A.R. Dixon, Combinatorial optimization by simulating adiabatic bifurcations in nonlinear Hamiltonian systems, *Sci. Adv.* 5 (2019) eaav2372. <https://doi.org/10.1126/sciadv.aav2372>
- [9] S. Matsubara, M. Takatsu, T. Miyazawa, T. Shibasaki, Y. Watanabe, K. Takemoto, H. Tamura, Digital annealer for high-speed solving of combinatorial optimization problems and its applications, in: 2020 25th Asia and South Pacific Design Automation Conference (ASP-DAC), IEEE, 2020, pp. 667–672. <https://doi.org/10.1109/ASP-DAC47756.2020.9045100>
- [10] T. Takemoto, M. Hayashi, C. Yoshimura, M. Yamaoka, A 2 × 30k-spin multi-chip scalable CMOS annealing processor based on a processing-in-memory approach for solving large-scale combinatorial optimization problems *IEEE J. Solid-State Circuits*, 55 (1) 2020. <https://doi.org/10.1109/JSSC.2019.2949230>
- [11] C. McGeoch, P. Farré, The D-wave advantage system: an overview, 2022, (D-Wave Technical Report Series). URL: https://www.dwavesys.com/media/3xvdipecn/14-1058a-a-advantage_processor_overview.pdf.
- [12] A.D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A.J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud, et al., Quantum critical dynamics in a 5,000-qubit programmable spin glass, *Nature* (2023) 61–66. <https://doi.org/10.1038/s41586-023-05867-2>
- [13] P.L. McMahon, A. Marandi, Y. Haribara, R. Hamerly, C. Langrock, S. Tamate, T. Inagaki, H. Takesue, S. Utsunomiya, K. Aihara, R.L. Byer, M.M. Fejer, H. Mabuchi, Y. Yamamoto, A fully programmable 100-spin coherent Ising machine with all-to-all connections, *Science* 354 (6312) (2016) 614–617. <https://doi.org/10.1126/science.aah5178>
- [14] M. Baity-Jesi, R.A. Baños, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, A. Gordillo-Guerrero, M. Guidetti, D. Iniguez, A. Maiorano, F. Mantovani, E. Marinari, V. Martín-Mayor, J. Monforte-García, A. Muñoz Sodupe, D. Navarro, G. Parisi, M. Pivanti, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, B. Seoane, A. Tarancon, P. Tellez, R. Tripiccone, D. Yllanes, Reconfigurable computing for Monte Carlo simulations: Results and prospects of the Janus project, *Eur. Phys. J. Special Topics* 210 (2012) 33. <https://doi.org/10.1140/epjst/e2012-01636-9>
- [15] M. Baity-Jesi, R.A. Baños, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, A. Gordillo-Guerrero, D. Iniguez, A. Maiorano, F. Mantovani, E. Marinari, V. Martín-Mayor, J. Monforte-García, A. Muñoz Sodupe, D. Navarro, G. Parisi, S. Perez-Gaviro, M. Pivanti, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, B. Seoane, A. Tarancon, R. Tripiccone, D. Yllanes, Janus II: a new generation application-driven computer for spin-system simulations, *Comp. Phys. Comm.* 185 (2014) 550–559. <https://doi.org/10.1016/j.cpc.2013.10.019>
- [16] C. Papadimitriou, K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*, Dover Publications, 2013.
- [17] S. Kirkpatrick, C.D. Gelatt Jr, M.P. Vecchi, Optimization by simulated annealing 220 (4598) 1983 671–680. <https://doi.org/10.1126/science.220.4598.671>
- [18] M.C. Angelini, F. Ricci-Tersenghi, Monte Carlo algorithms are very effective in finding the largest independent set in sparse random graphs, *Phys. Rev. E* 100 (1) (2019) 13302. <https://doi.org/10.1103/PhysRevE.100.013302>
- [19] M.C. Angelini, F. Ricci-Tersenghi, Limits and performances of algorithms based on simulated annealing in solving sparse hard inference problems, *Phys. Rev. X* 13 (2) (2023) 21011. <https://doi.org/10.1103/PhysRevX.13.021011>
- [20] F. Barahona, On the computational complexity of Ising spin glass models, *J. Phys. A Math. Gen.* 15 (1982) 3241. <https://doi.org/10.1088/0305-4470/15/10/028>
- [21] S. Istrail, Statistical mechanics, three-dimensionality and NP-completeness: I. universality of intracatability for the partition function of the Ising model across non-planar surfaces (extended abstract), in: Proceedings of the thirty-second annual ACM symposium on theory of computing, 2000, pp. 87–96. <https://doi.org/10.1145/335305.335316>
- [22] T. Yavorskii, M. Weigel, Optimized gpu simulation of continuous-spin glass models, *Eur. Phys. J. Special Topics* 210 (2012) 159. <https://doi.org/10.1140/epjst/e2012-01644-9>
- [23] S. Feng, Y. Fang, K.-M. Tam, Z. Yun, J. Ramanujam, J. Moreno, M. Jarrell, Three dimensional edwards-anderson spin glass model in an external field, 2014. [arXiv:1403.4560](https://arxiv.org/abs/1403.4560)
- [24] M. Lulli, M. Bernaschi, G. Parisi, Highly optimized simulations on single- and multi-GPU systems of 3D ising spin glass model, *Comput. Phys. Commun.* 196 (2015) 290–303. <https://doi.org/10.1016/j.cpc.2015.06.019>
- [25] M. Bernaschi, I. González-Adalid Pemartín, V. Martín-Mayor, G. Parisi, The QISG suite: high-performance codes for studying quantum Ising spin glasses, *Comp. Phys. Comm.* 298 (2024) 109101. <https://doi.org/10.1016/j.cpc.2024.109101>
- [26] M. Bernaschi, I. González-Adalid Pemartín, V. Martín-Mayor, G. Parisi, The quantum transition of the two-dimensional Ising spin glass, *Nature* 631 (2024) 749–754. <https://doi.org/10.1038/s41586-024-07647-y>
- [27] L.A. Fernández, V. Martín-Mayor, Testing statics-dynamics equivalence at the spin-glass transition in three dimensions, *Phys. Rev. B* 91 (2015) 174202. <https://doi.org/10.1103/PhysRevB.91.174202>
- [28] A. Billoire, L.A. Fernandez, A. Maiorano, E. Marinari, V. Martín-Mayor, J. Moreno-Gordo, G. Parisi, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, Numerical construction of the Aizenman-Wehr metastate, *Phys. Rev. Lett.* 119 (2017) 37203. <https://doi.org/10.1103/PhysRevLett.119.037203>
- [29] A. Billoire, L.A. Fernandez, A. Maiorano, E. Marinari, V. Martín-Mayor, J. Moreno-Gordo, G. Parisi, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, Dynamic variational study of chaos: spin glasses in three dimensions, *J. Stat. Mech. Theory Exp.* 2018 (3) (2018) 33302. <https://doi.org/10.1088/1742-5468/aaa387>
- [30] D. Frenkel, B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications*, Academic Press, London, 3rd edition, 2023
- [31] M. Creutz, Microcanonical Monte Carlo simulation, *Phys. Rev. Lett.* 50 (1983) 1411. <https://doi.org/10.1103/PhysRevLett.50.1411>
- [32] J. Ruiz-Lorenzo, C. Ullod, Study of a microcanonical algorithm on the $\pm J$ spin glass model in $d = 3$, *Comput. Phys. Commun.* 125 (1) (2000) 210–220. [https://doi.org/10.1016/S0010-4655\(99\)00490-7](https://doi.org/10.1016/S0010-4655(99)00490-7)
- [33] K. Rummukainen, Multicanonical cluster algorithm and the two-dimensional 7-state Potts model, *Nucl. Phys. B* 390 (3) (1993) 621–636. [https://doi.org/10.1016/0550-3213\(93\)90491-7](https://doi.org/10.1016/0550-3213(93)90491-7)
- [34] M. Caselle, R. Fiore, F.F. Gliozzi, M. Hasenbusch, K. Pinn, S. Vinti, Rough interfaces beyond the gaussian approximation, *Nucl. Phys. B* 432 (3) (1994) 590–618. [https://doi.org/10.1016/0550-3213\(94\)90035-3](https://doi.org/10.1016/0550-3213(94)90035-3)
- [35] M. Hasenbusch, K. Pinn, Comparison of Monte Carlo results for the 3D Ising interface tension and interface energy with (extrapolated) series expansions, *Physica A Stat. Mech. Appl.* 203 (2) (1994) 189–213. [https://doi.org/10.1016/0378-4371\(94\)90152-X](https://doi.org/10.1016/0378-4371(94)90152-X)
- [36] V. Agostini, G. Carlino, M. Caselle, M. Hasenbusch, The spectrum of the 2 + 1-dimensional gauged Ising model, *Nucl. Phys. B* 484 (1) (1997) 331–352. [https://doi.org/10.1016/S0550-3213\(96\)00539-1](https://doi.org/10.1016/S0550-3213(96)00539-1)
- [37] R. Lustig, Microcanonical monte carlo simulation of thermodynamic properties, *J. Chem. Phys.* 109 (1998) 8816. <https://doi.org/10.1063/1.477552>
- [38] V. Martín-Mayor, Microcanonical approach to the simulation of first-order phase transitions, *Phys. Rev. Lett.* 98 (2007) 137207. <https://doi.org/10.1103/PhysRevLett.98.137207>
- [39] L.A. Fernández, A. Gordillo-Guerrero, V. Martín-Mayor, J.J. Ruiz-Lorenzo, First-order transition in a three-dimensional disordered system, *Phys. Rev. Lett.* 100 (2008) 57201. <https://doi.org/10.1103/PhysRevLett.100.057201>
- [40] J. Zinn-Justin, *Quantum Field Theory and Critical Phenomena*, Clarendon Press, Oxford, 4th edition, 2005
- [41] S.F. Edwards, P.W. Anderson, Theory of spin glasses, *J. Phys. F Met. Phys.* 5 (1975) 965. <https://doi.org/10.1088/0305-4608/5/5/017>
- [42] S.F. Edwards, P.W. Anderson, Theory of spin glasses II, *J. Phys. F* 6 (10) (1976) 1927. <https://doi.org/10.1088/0305-4608/6/10/022>
- [43] M. Baity-Jesi, R.A. Baños, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, A. Gordillo-Guerrero, D. Iniguez, A. Maiorano, F. Mantovani, E. Marinari, V. Martín-Mayor, J. Monforte-García, A. Muñoz Sodupe, D. Navarro, G. Parisi, S. Perez-Gaviro, M. Pivanti, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, B. Seoane, A. Tarancon, R. Tripiccone, D. Yllanes, Critical parameters of the three-dimensional Ising spin glass, *Phys. Rev. B* 88 (2013) 224416. <https://doi.org/10.1103/PhysRevB.88.224416>
- [44] M. Baity-Jesi, E. Calore, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, A. Gordillo-Guerrero, D. Iniguez, A. Maiorano, E. Marinari, V. Martín-Mayor, J. Monforte-García, A. Muñoz Sodupe, D. Navarro, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, B. Seoane, A. Tarancon, R. Tripiccone, D. Yllanes, A statics-dynamics equivalence through the fluctuation-dissipation ratio provides a window into the spin-glass phase from nonequilibrium measurements, *PNAS* 114 (8) (2017) 1838–1843. <https://doi.org/10.1073/pnas.1621242114>
- [45] M. Baity-Jesi, E. Calore, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, A. Gordillo-Guerrero, D. Iniguez, A. Maiorano, E. Marinari, V. Martín-Mayor, J. Monforte-García, A. Muñoz Sodupe, D. Navarro, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, B. Seoane, A. Tarancon, R. Tripiccone, D. Yllanes, Matching microscopic and macroscopic responses in glasses, *Phys. Rev. Lett.* 118 (2017) 157202. <https://doi.org/10.1103/PhysRevLett.118.157202>
- [46] M. Baity-Jesi, E. Calore, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, A. Gordillo-Guerrero, D. Iniguez, A. Maiorano, E. Marinari, V. Martín-Mayor, J. Moreno-Gordo, A. Muñoz Sodupe, D. Navarro, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, B. Seoane, A. Tarancon, R. Tripiccone, D. Yllanes, Aging rate of spin glasses from simulations matches experiments, *Phys. Rev. Lett.* 120 (2018) 267203. <https://doi.org/10.1103/PhysRevLett.120.267203>
- [47] M. Baity-Jesi, E. Calore, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, A. Gordillo-Guerrero, D. Iniguez, A. Lasanta, A. Maiorano, E. Marinari, V. Martín-Mayor, J. Moreno-Gordo, A. Muñoz Sodupe, D. Navarro, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, B. Seoane, A. Tarancon, R. Tripiccone, D. Yllanes, The Mpemba effect in spin glasses is a persistent memory effect, *PNAS* 116 (2019) 15350–15355. <https://doi.org/10.1073/pnas.1819803116>
- [48] M. Baity-Jesi, E. Calore, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, I. I. González-Adalid Pemartín, A. Gordillo-Guerrero, D. Iniguez, A. Maiorano, E. Marinari, V. Martín-Mayor, J. Moreno-Gordo, A. Muñoz Sodupe, D. Navarro, I. Paga, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, B. Seoane, A. Tarancon, R. Tripiccone, D. Yllanes, Temperature chaos is present in off-equilibrium spin-glass dynamics, *Commun. Phys.* 4 (1) (2021) 74. <https://doi.org/10.1038/s42005-021-00565-9>
- [49] Q. Zhai, I. Paga, M. Baity-Jesi, E. Calore, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, I. Gonzalez-Adalid Pemartín, A. Gordillo-Guerrero, D. Iniguez, A. Maiorano, E. Marinari, V. Martín-Mayor, J. Moreno-Gordo, A.M. Muñoz Sodupe, D. Navarro, R.L. Orbach, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, D.L. Schlagel, B. Seoane, A. Tarancon, R. Tripiccone, D. Yllanes, Scaling law de-

- scribes the spin-glass response in theory, experiments, and simulations, *Phys. Rev. Lett.* 125 (2020) 237202. <https://doi.org/10.1103/PhysRevLett.125.237202>
- [50] M. Baity-Jesi, E. Calore, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, I. Gonzalez-Adalid Pemartin, A. Gordillo-Guerrero, D. Iñiguez, A. Maiorano, E. Marinari, V. Martin-Mayor, J. Moreno-Gordo, A. Muñoz Sudupe, D. Navarro, I. Paga, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, B. Seoane, A. Tarancon, D. Yllanes, Memory and rejuvenation effects in spin glasses are governed by more than one length scale, *Nat. Phys.* 19 (2023) 978–985. <https://doi.org/10.1038/s41567-023-02014-6>
- [51] I. Paga, Q. Zhai, M. Baity-Jesi, E. Calore, A. Cruz, C. Cummings, L.A. Fernandez, J.M. Gil-Narvion, I.G.-A. Pemartin, A. Gordillo-Guerrero, D. Iñiguez, G.G. Kenning, A. Maiorano, E. Marinari, V. Martin-Mayor, J. Moreno-Gordo, A. Muñoz Sudupe, D. Navarro, R.L. Orbach, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, D.L. Schlagel, B. Seoane, A. Tarancon, D. Yllanes, Superposition principle and nonlinear response in spin glasses, *Phys. Rev. B* 107 (2023) 214436. <https://doi.org/10.1103/PhysRevB.107.214436>
- [52] I. Paga, J. He, M. Baity-Jesi, E. Calore, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, I. Gonzalez-Adalid Pemartin, A. Gordillo-Guerrero, D. Iñiguez, A. Maiorano, E. Marinari, V. Martin-Mayor, J. Moreno-Gordo, A. Muñoz Sudupe, D. Navarro, R.L. Orbach, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, D.L. Schlagel, B. Seoane, A. Tarancon, D. Yllanes, Quantifying memory in spin glasses, *Phys. Rev. Lett.* 133 (2024) 256704. <https://doi.org/10.1103/PhysRevLett.133.256704>
- [53] A.D. Sokal, Monte Carlo methods in statistical mechanics: Foundations and new algorithms, in: C. DeWitt-Morette, P. Cartier, A. Folacci (Eds.), *Functional Integration: Basics and Applications* (1996 Cargèse School), Plenum, N.Y., 1997, pp. 131–192. https://doi.org/10.1007/978-1-4899-0319-8_6
- [54] D.P. Landau, K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics*, Cambridge University Press, Cambridge, 2nd edition, 2005
- [55] M. Barma, B. Sriram Shastry, D -dimensional Hubbard model as a $(D+1)$ -dimensional classical problem, *Phys. Lett. A* 61 (1) (1977) 15–18. [https://doi.org/10.1016/0375-9601\(77\)90248-1](https://doi.org/10.1016/0375-9601(77)90248-1)
- [56] N. Ito, Y. Kanada, Monte Carlo simulation of the Ising model and random number generation on the vector processor, in: *Proceedings SUPERCOMPUTING '90*, 1990, pp. 753–763. <https://doi.org/10.1109/SUPERC.1990.130097>
- [57] L. Jacobs, C. Rebbi, Multi-spin coding: a very efficient technique for Monte Carlo simulations of spin systems, *J. Comput. Phys.* 41 (1981) 203. [https://doi.org/10.1016/0021-9991\(81\)90089-9](https://doi.org/10.1016/0021-9991(81)90089-9)
- [58] F. Belletti, A. Cruz, L.A. Fernandez, A. Gordillo-Guerrero, M. Guidetti, A. Maiorano, F. Mantovani, E. Marinari, V. Martín-Mayor, J. Monforte, A. Muñoz Sudupe, D. Navarro, G. Parisi, S. Perez-Gaviro, J.J. Ruiz-Lorenzo, S.F. Schifano, D. Sciretti, A. Tarancon, R. Tripiccion, D. Yllanes, An in-depth look at the microscopic dynamics of Ising spin glasses at fixed temperature, *J. Stat. Phys.* 135 (2009) 1121. <https://doi.org/10.1007/s10955-009-9727-z>
- [59] F. Belletti, M. Cotallo, A. Cruz, L.A. Fernandez, A. Gordillo-Guerrero, M. Guidetti, A. Maiorano, F. Mantovani, E. Marinari, V. Martín-Mayor, A.M. Sudupe, D. Navarro, G. Parisi, S. Perez-Gaviro, J.J. Ruiz-Lorenzo, S.F. Schifano, D. Sciretti, A. Tarancon, R. Tripiccion, J.L. Velasco, D. Yllanes, Nonequilibrium spin-glass dynamics from picoseconds to one tenth of a second, *Phys. Rev. Lett.* 101 (2008) 157201. <https://doi.org/10.1103/PhysRevLett.101.157201>
- [60] M. Baity-Jesi, E. Calore, A. Cruz, L.A. Fernandez, J.M. Gil-Narvion, I. González-Adalid Pemartin, A. Gordillo-Guerrero, D. Iñiguez, A. Maiorano, E. Marinari, V. Martin-Mayor, J. Moreno-Gordo, A. Muñoz Sudupe, D. Navarro, I. Paga, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J.J. Ruiz-Lorenzo, S.F. Schifano, D. Seoane, A. Tarancon, D. Yllanes, Multifractality in spin glasses, *PNAS* 121 (2) (2024) e2312880120. <https://doi.org/10.1073/pnas.2312880120>
- [61] L.A. Fernández, E. Marinari, V. Martín-Mayor, G. Parisi, J.J. Ruiz-Lorenzo, An experiment-oriented analysis of 2D spin-glass dynamics: a twelve time-decades scaling study, *J. Phys. A Math. Theor.* 52 (2019) 224002. <https://doi.org/10.1088/1751-8121/ab1364>
- [62] L.A. Fernández, E. Marinari, V. Martín-Mayor, G. Parisi, J. Ruiz-Lorenzo, Out-of-equilibrium 2D Ising spin glass: almost, but not quite, a free-field theory, *J. Stat. Mech. Theory Exp.* 2018 (10) (2018) 103301. <https://doi.org/10.1088/1742-5468/aae2e1>
- [63] K. Hukushima, K. Nemoto, Exchange Monte Carlo method and application to spin glass simulations, *J. Phys. Soc. Jpn.* 65 (1996) 1604. <https://doi.org/10.1143/JPSJ.65.1604>
- [64] C. Chilin, E. Marinari, V. Martín-Mayor, G. Parisi, J.J. Ruiz-Lorenzo, D. Yllanes, *In preparation*, (2026)
- [65] Data and code available at <https://github.com/IsidoroGlez/micSA-EA-MC/releases/tag/v1.1>. accessed: October 2025
- [66] J. Houdayer, A cluster Monte Carlo algorithm for 2-dimensional spin glasses, *Eur. Phys. J. B* 22 (4) (2001) 479–484. <https://doi.org/10.1007/PL00011151>
- [67] Z. Zhu, A.J. Ochoa, H.G. Katzgraber, Efficient cluster algorithm for spin glasses in any space dimension, *Phys. Rev. Lett.* 115 (2015) 77201. <https://doi.org/10.1103/PhysRevLett.115.077201>