Aging at the spin-glass/ferromagnet transition: Monte Carlo simulations using graphics processing units

Markus Manssen* and Alexander K. Hartmann
Institute of Physics, Carl von Ossietzky University, 26111 Oldenburg, Germany
(Received 21 November 2014; revised manuscript received 7 May 2015; published 26 May 2015)

We study the nonequilibrium aging behavior of the $\pm J$ Edwards-Anderson model in three dimensions for samples of size up to $N = 128^3$ and for up to 10^8 Monte Carlo sweeps. In particular we are interested in the change of the aging when crossing from the spin-glass phase to the ferromagnetic phase. The necessary long simulation times are reached by employing a CUDA-based GPU implementation, which allows for single-spin flip times as small as 8 ps. We measure typical spin-glass correlation functions in space and time to determine the growing length scale and extract the constituting exponents. We observe a clear signature of the disorder-driven equilibrium transition in the nonequilibrium behavior.

DOI: 10.1103/PhysRevB.91.174433 PACS number(s): 75.50.Lk, 75.40.Mg, 75.10.Hk

I. INTRODUCTION

Spin glasses [1–5] are certain magnetic alloys [6], that possess at low temperatures interesting equilibrium and nonequilibrium behavior which is to a large extent still not well understood. The low-temperature ordered spin-glass phase is characterized by a rough free-energy landscape and by slow glassy dynamics [7,8]. Disorder and frustration in the spin-spin interactions were identified as the underlying principles governing the behavior of spin glasses. Thus models mixing positive, i.e., ferromagnetic and negative (antiferromagnetic) couplings such as the mean-field Sherrington-Kirkpatrick model [9] and the Ising-like short-ranged Edwards-Anderson model [10] were created to understand the spin-glass behavior. A prevailing topic is still whether the replica-symmetry-breaking theory [2] arising from the solution [11,12] of the former model can accurately describe the spin-glass phase of the latter model in three dimensions. The most prominent competitor is the droplet theory [13,14], which centers around the eponymous droplets, low-level excitations from the presumably only two existing pure states. Numerous publications have dealt with simulations in equilibrium [15–20] as well as out of equilibrium [7,21–30].

The standard spin-glass models assume an on average equal fraction of positive and negative couplings. Nevertheless, when decreasing the fraction of negative bonds in the Edwards-Anderson model, it exhibits a phase transition at low temperatures from the aforementioned spin-glass phase to the well known ferromagnetic phase of the Ising model. This transition has been studied in the typical equilibrium approach to phase transitions [31], also via ground-state calculations [32]. Nevertheless, concerning the nonequilibrium "aging" behavior, so far only systems deep in the spin-glass phase have been studied extensively, to the knowledge of the authors. Therefore, the purpose of this study is to determine whether the spin-glass to ferromagnet transition is also visible within the nonequilibrium behavior. Specifically we will be looking at growing correlations in space and time and try to explain them in terms of the dynamical correlation length. The determination and characterization of this growing length scale in the spin-glass phase has been a focus of many previous publications [8,21–23,27,29,33–35], as there are a few stumbling blocks before dependably measuring it. It was quickly found that it appears to follow a power law [21,22] in line with the mean-field theory, though there is discussion [28,29] whether it crosses over into the logarithmic growth expected from droplet theory.

Because reaching sufficiently long simulation times is computationally challenging, even inspiring the adoption of specialized hardware [36], we implemented the model in CUDA [37] to leverage the comparatively high processing power of GPUs. Quite a few pioneering works have already tested the feasibility and performance outlook of this platform for the Ising [38–40] and the Edwards-Anderson model [41,42] but have shown no fruitful application. Our implementation has been carefully optimized for tackling the problem at hand efficiently and with limited resources. This allowed us to study large system sizes of $N = 128^3$ up to long time scales of 10^8 sweeps.

The remainder of this work is organized as follows. In Sec. II we describe the used Edwards-Anderson model and its observables. Section III follows with details on the GPU implementation of the model. The results of the simulation and their analysis are presented in Sec. IV. We close with our conclusions in Sec. V.

II. MODEL

The Edwards-Anderson model [10] describes a D-dimensional cubic system of side length L containing $N = L^D$ Ising spins $S_i = \pm 1$. Its Hamiltonian is given by

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

where the sum runs over nearest neighbors $\langle i,j \rangle$ and the bonds $J_{ij} = \pm 1$ are drawn from a bimodal distribution $P(J) = p\delta(J-1) + (1-p)\delta(J+1)$. We use periodic boundary conditions in all directions. The parameter $p \in [0,1]$ controls the fraction of positive and negative bonds. For $p = 1$ the ferromagnet Ising model is reproduced with a paramagnetic phase at high temperatures and the well-known ferromagnetic phase at small temperatures T for $D > 1$. On the other hand,

*markus.manssen@uni-oldenburg.de

$p = 0.5$ represents the usual spin-glass model with a low temperature spin-glass phase for $D > 2$. We will only be concerned with the case $D = 3$ in the following, which has the transition temperatures $T_{\mathrm{FM}} \approx 4.5115$ ($p = 1$) [43] and $T_{\mathrm{SG}} = 1.102(3)$ ($p = 0.5$) [20], respectively. For intermediate values of p there exists the phase transition from ferromagnet to spin glass at $p \approx 0.77$ ($T \to 0$) [32].

Simulations start with random initial configurations emulating a quench from infinite temperature. We then examine the system at different waiting times t_w (measured in sweeps) after the beginning of the simulation. The spin-glass order parameter is the overlap

$$q = \frac{1}{N} \sum_i q_i, \tag{2}$$

with $q_i = S_i^{(a)} S_i^{(b)}$ the elementwise overlap of two replicas $S^{(a)}$ and $S^{(b)}$ with the same bond configuration J, but different initial configurations and thermal histories. In equilibrium, corresponding to $t_w \to \infty$, the probability distribution $P(q)$ is expected to assume a two peak structure below the transition temperature. In the droplet theory [13,14] this takes the form of two δ peaks at $\pm q_{\mathrm{EA}}$, while the mean-field theory [2] has a wider distribution with a plateau of nonzero probability around $q = 0$.

To measure the growing length scale we make use of the spatial four-point correlation

$$C_4(r,t_w) = \frac{1}{N} \sum_i q_i(t_w) q_{i+r}(t_w) \tag{3}$$

between two points and two replicas. With $i + r$ we denote a spin which has a spatial distance r from spin i.

There exist different approaches to extract a growing coherence (or dynamic correlation) length ξ from the four point-correlation function. The first approach is based on the assumption that C_4 follows the functional form [23]

$$C_4(r,t_w) \propto r^{-\alpha} g\left(\frac{r}{\xi(t_w)}\right). \tag{4}$$

The function g is approximately a stretched exponential $g(x) \sim \exp(-x^{\beta})$. Extracting ξ works by fitting (4) to the data of C_4, for various times t_w. The two unknown exponents ($\alpha \approx 0.5$ and $\beta \approx 1.5$ in the spin-glass phase) complicate the extraction of ξ, which spawned many methods for accomplishing this task [21–23,27].

As an alternative, notably Ref. [33] introduced the use of integral estimators for this problem. One uses the integral

$$I_k(t_w) = \int_0^{L/2} r^k C_4(r,t_w) dr \tag{5}$$

to calculate an estimate for the coherence length

$$\xi_k(t_w) = \frac{I_{k+1}(t_w)}{I_k(t_w)} \sim \xi(t_w). \tag{6}$$

The choice of k determines which regions of the function C_4 contribute most to the estimate. Reference [29] recommends $k = 1$ to get the best trade-off between systematic errors for small values of k and larger influence of statistical error for higher values.

Another observable of interest we use to study the aging behavior around the ferromagnet-spin glass transition is the autocorrelation

$$C(t,t_w) = \frac{1}{N} \sum_i S_i(t_w) S_i(t_w + t) \tag{7}$$

between two points in time separated by a time difference t in reference to the waiting time t_w. It is assumed to split into two parts. The first is a quasiequilibrated part for small $t \ll t_w$, that takes the form [21,22,24,26,28] of a power law

$$C_{\mathrm{eq}}(t) \propto t^{-x}, \tag{8}$$

with another characteristic exponent x. For longer times $t \gg t_w$ the aging part $C_{\mathrm{age}}(t,t_w) = f(\xi(t_w + t)/\xi(t_w))$ can trivially be expected [7,8] to depend only on the ratio of the coherence lengths at the two different times. An additive decomposition $C(t,t_w) = C_{\mathrm{eq}}(t) + C_{\mathrm{age}}(t,t_w)$ is favored by theoretical arguments [8,26,28,44]. But we will present here for the spin-glass phase only the results according to a multiplicative decomposition $C(t,t_w) = C_{\mathrm{eq}}(t) \cdot C_{\mathrm{age}}(t,t_w)$, as this seems to work better with our analysis of the data, even though it is only expected to hold in the critical region [8,26]. For the ferromagnet a plateau $\lim_{t_w \to \infty} C(t,t_w) = C_{\mathrm{eq}}(t) + M^2$ is expected at long waiting times with the equilibrium magnetization M [45,46].

III. IMPLEMENTATION

We implemented a Metropolis Monte Carlo simulation [47] of the model for Nvidia GPUs using the CUDA C programming interface [37], as was first detailed in Ref. [48]. For explanation of the GPU-related terms used in the following we refer to the CUDA Programming Guide [37] or a textbook like Ref. [49]. In order to perform the update of a spin S_i one has to calculate the flipping probability

$$p_{\mathrm{accept}} = \min\left[1, \exp\left(-\frac{2}{T} \sum_{j \in N(i)} J_{ij} S_i S_j\right)\right] \tag{9}$$

incorporating the coupling of the spin i to the local field generated by its direct neighbors $N(i)$ on the lattice. Since GPUs are designed to keep their large number of simple processors busy with many, preferably independent processes at once, a sequential implementation of a single-spin-flip algorithm is ill suited for GPUs. So in order to attain a parallel algorithm suitable for this architecture we adopted a standard checkerboard update scheme. In a first step all "white fields" of the system are updated followed by the other half of the system in a second step. Both add up to a single sweep of the system corresponding to one time step. Each half-step is performed in its own CUDA kernel call to ensure global synchronization. A problem, especially for fits in the ferromagnetic phase we will be performing later, is that the checkerboard metropolis algorithm produces systematic errors in the autocorrelation C at short times t. This is caused by the deterministic access pattern combined with guaranteed flips for certain spins, in particular the ones with zero local magnetic field. Thus these free spins will change direction every time step resulting in a damped oscillation of C around the true value. As a remedy we access the spins in both systems used for calculating C in a

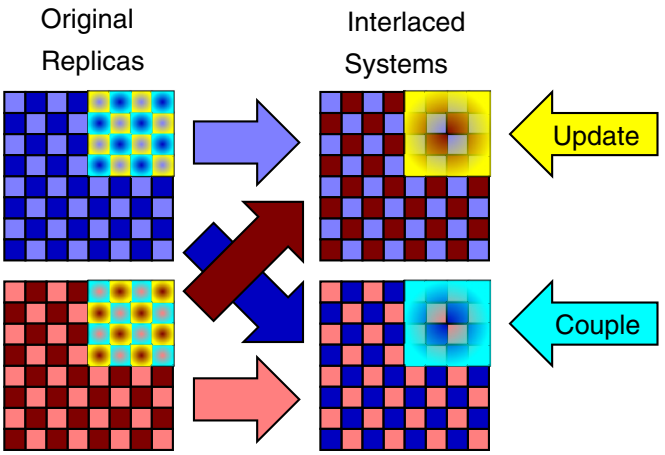

FIG. 1. (Color online) Memory layout in the GPU. Two independent replicas (left), having the same bonds, are simulated in a checker board manner: during each iteration either all "black" or all "white" sites are updated. Thus, within the GPU memory, one memory area (top right) holds all updated sites, e.g., the white sites of replica one and the black sites of replica two, while another memory range (lower right) contains all neighboring interaction partners. After one half-sweep is completed, the role of updated sites and neighboring sites are swapped.

checkerboard scheme and just randomize all free spins, which emulates the result of a step in an algorithm, that flips these spins with probability $1/2$. This can be seen as a correction of first order and removes the effect completely in the spin-glass phase, while greatly reducing it in the ferromagnetic phase.

A straightforward implementation of this update scheme however would be inefficient as the GPU memory is optimized for reading large bulks of neighboring memory cells at once ("memory coalescing"). To circumvent this one could relocate the spins to two different memory regions as was done in Ref. [39]. For a simpler method we instead simulate the two replicas we need for calculating (3) simultaneously and swap the "black spins" between the two lattices to get what we will call the "interlaced checkerboard layout"; see Fig. 1. This way all spins in one lattice can be updated at the same time, while all neighbors they are coupled to reside in the other lattice. Basically the same approach was also used in Ref. [36] just by virtue of the simplifications it introduces. Specifically the spins' indices remain unchanged and we can use the same bonds for both update steps. Since the bonds J_{ij} are symmetric, we only need to store the left/up bonds of a spin and they can be read efficiently via texture memory. The joint neighbors of the updated spins are loaded into shared memory so they can be shared between the members of a thread block to calculate the flip probability (9). We choose dimensions of $32 \times 4 \times 2$ for the thread blocks in the GPU thread hierarchy.

We also employed 64bit-multispin-coding, meaning that the spins taking values ± 1 are coded as single bits and 64 of them are stored together in a 64bit variable. The same applies for the corresponding bonds. We choose spins from the same position in 64 different samples, which is sometimes known as asynchronous multispin coding. Bit operations are used to perform the update for all bit-coded spins at the same time. We only need to differentiate between a few possible cases of spin alignments using boolean logic at the bit level. Then we look up the precalculated flipping probabilities from constant memory and apply them for the matching bits. It is customary to save computational effort by using only one random number for multiple samples. As no suitable and efficient random number generators were available at the time of implementation, we established a 1024bit variant of a Xorshift generator [50]. The generator was optimized for generating a single random number per thread and kernel, as was needed here. With this complete implementation we reach single-spin-flip times of $\approx$8 ps on a GeForce GTX 570 GPU. Of the prior implementations [38–42,51,52] of Ising and Edwards-Anderson model only the one by Weigel [41] is faster. But it uses multihit updates, which means each thread block updates for several flip attempts per spin, thus requiring the costly global synchronization less frequently. This is no problem if one is interested in equilibrium properties, but this changes the dynamics, as, e.g., visible by the reduced growth of correlations. Therefore, that is undesirable when one wants to actually study the nonequilibrium dynamic aging properties, as in the current work. The recent Ref. [53] was also brought to our attention, which reports faster single-spin-flip times of $\approx$3 ps. But this was done on a newer, more powerful GPU, so we cannot say how this directly compares to our implementation.

For the current work we only had access to two GPUs and consequently designed this approach for maximum efficiency per GPU. However, if one had access to a larger number, it would be preferable to be able to distribute samples better among GPUs. For this purpose one can simply reduce the number of samples to $M = 2^m$ ($m < 6$) by storing $64/M$ spins from different positions per sample in a multispin. A simple way is to split the system into $64/M$ equal parts along one dimension and assign spins at the same relative positions to the bits $\{i, i+M, i+2M, \ldots\}$ in the same multispin. This effectively makes it look and work like a smaller system, with the caveat that when coupling over the "periodic boundary" one has to rotate the bits of the multispin variable by M positions. The computational overhead for this change is negligible but it has two other problems. First this effectively shrinks our systems which might result in low occupancy and efficiency of the GPUs processors. But the bigger problem is that we cannot use the same random number for spins from the same sample. Thus we have to generate $64/M$ random numbers per kernel instead of just one. Because of this requirement, the synchronous multispin coding corresponding to $M = 1$ is very inefficient, and a more balanced choice like $M = 8$ is preferable.

IV. RESULTS

We simulated a total of 192 samples of randomly initialized 128^3 systems with two replicas each. The simulations were performed on two GeForce GTX 570. We took the parameters $T = 0.8$ and $p \in [0.5, 1]$ and performed 10^8 time steps, which takes about 63 h per batch of 64 multispin-coded samples. At the measure points the whole system configurations were simply stored to hard disk and the generated data was later postprocessed to gain access to the correlation functions.

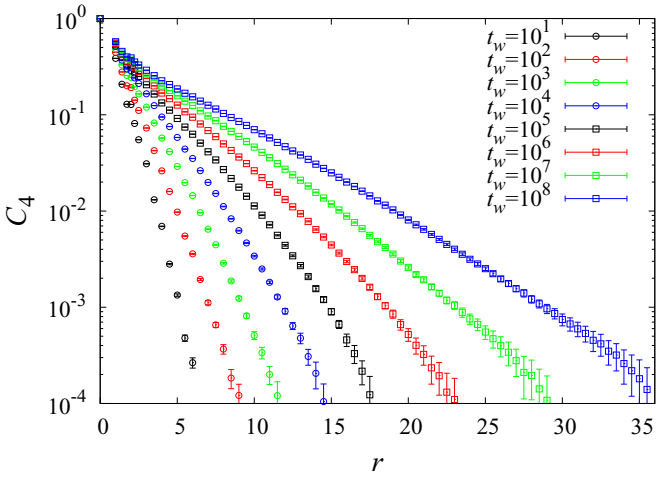

FIG. 2. (Color online) Spatial correlation C_4 over the distance r at different waiting times t_w for a 128^3 system at $p = 0.5$. Multiple close points were merged to give a clearer picture.

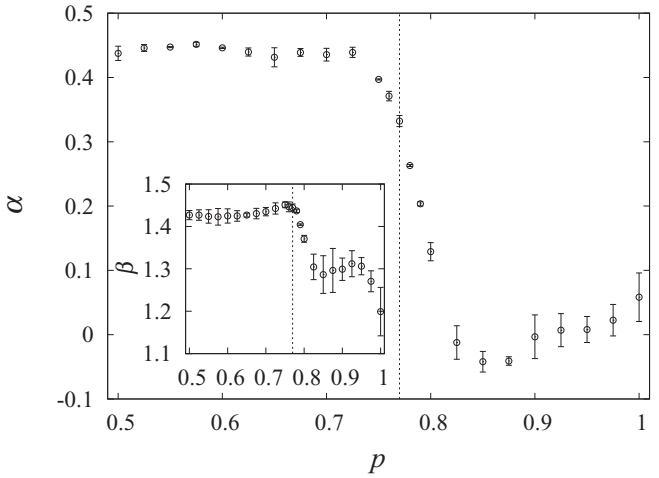

FIG. 3. Scaling exponent α of the spatial correlation over the bond probability p for a 128^3 system. The vertical line marks $p = p_c$. Inset: associated scaling exponent β.

An exemplary spatial correlation C_4 from (3) for $p = 0.5$ is shown in Fig. 2. The two replicas utilized for the calculation make correlations visible despite the model's inherent disorder. The curves show a seemingly exponential decay for larger distances. The steeper gradient for small distances is incorporated in the scaling form (4) with the power law $r^{-\alpha}$. As one would expect, the correlations spread to larger distances as time passes suggesting a growing length scale.

Our first approach to extracting this coherence length ξ is a fit of (4) with the stretched exponential form for g. However, the values $\alpha \approx 0.5$ and $\beta \approx 1.5$ which are suitable deep in the spin-glass phase ($p \approx 0.5$) might depend on the value of p. To get the most consistent values at a particular p we performed multifits of the curves for all different $t_w \geqslant 10$ at once, i.e., with universal values of α and β (independent of t_w), but individual values $\xi(t_w)$. Naturally the choice of points included in the fit can have an influence on the outcome. As such we generally restricted it to $r \geqslant 3$ and specifically found the cleanest results, when only using points at integer-valued distances r. Those are always located along the lattice axes. But as a reference we did the same fits also with all $r \geqslant 3$ and use these for calculating our error bars for α, β, and ξ. In detail we estimate our errors as the difference between the fit results for our restricted point set and the larger one plus both of the normal statistical errors from the two different fits. Still, this might underestimate the errors a bit, because for multifits statistical independence of the data is assumed, while in our case the measurements are from the same runs, just at different waiting times t_w; thus they are correlated.

The results for both exponents α and β as a function of the probability p are shown in Fig. 3. A strong change can be seen around the phase transition $p_c \approx 0.77$ from $0.4 < \alpha < 0.5$ and $1.4 < \beta < 1.45$ in the spin glass phase to $\alpha \approx 0$, which is to be expected in the ferromagnetic phase, and $\beta \approx 1.3$ for $p > p_c$. Thus the equilibrium phase disorder driven phase transition is well visible in the analysis of the nonequilibrium aging behavior. Note that when getting closer to $p = 1$, i.e., in the ferromagnetic phase, the system quickly develops a complicated behavior [54–57] not yet fully understood and not compatible with the scaling assumption in Eq. (4). Thus when addressing the range where the coherence length is small, we only have the first few times t_w available to work with, i.e., the fits according to (4) generally cannot be fitted as well at later times.

The second approach [29,33] for calculating ξ uses the integral estimation ξ_1 according to (6). Like in the original work we take the integrals up to the point where the value of C_4 first becomes smaller than three times its error, and approximate the rest of the integrals with our fitted function. Figure 4 shows results for the coherence length ξ as a function of the waiting time t_w for both different approaches, respectively. As is visible from the figure, both methods agree well for a large stretch of waiting times but disagree close to the beginning and the end. While the integral results look

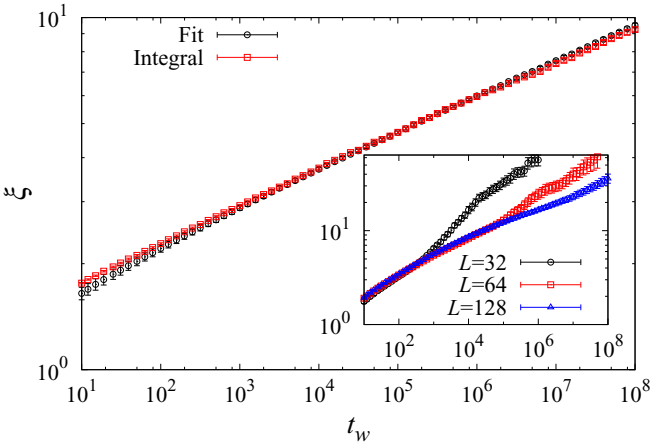

FIG. 4. (Color online) Coherence length ξ as a function of the waiting times t_w for a 128^3 system at $p = 0.5$. The values are calculated by fitting and integral estimation, respectively. Inset: results by fitting estimation for different side lengths $L \in \{32,64,128\}$ at $p = 0.9$.

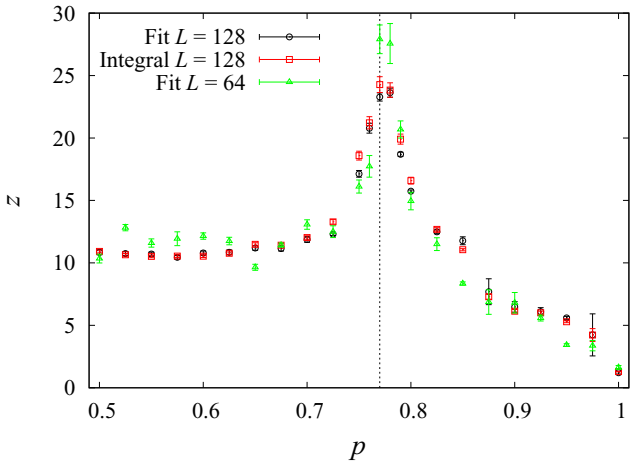

FIG. 5. (Color online) Power law exponent z of the coherence length as a function of the bond probability p for $L = 128^3$ and $L = 64$. For $L = 128$, the two curves correspond to the fitting and integral estimation of ξ, respectively, which agree pretty well. To give an impression of the finite-size effects, also the result from the fitting approach for $L = 64$ (exhibiting more systematic errors due to the small system size) is included. The vertical line marks $p = p_c$.

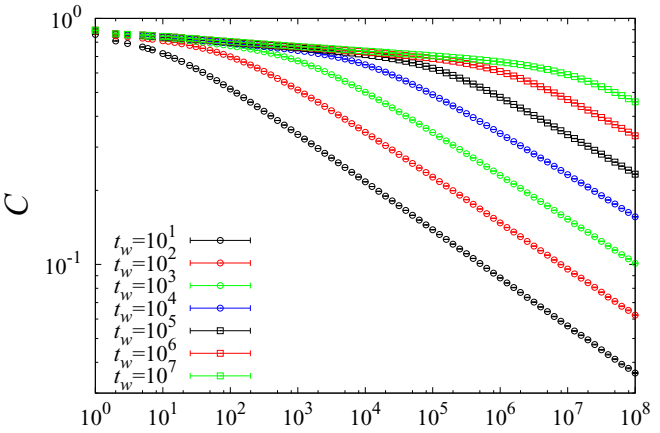

FIG. 6. (Color online) Autocorrelation C as a function of the time distance t at different waiting times t_w for a 128^3 system at $p = 0.5$.

closer to a power law, the fit results give a bit higher estimates for ξ at the end and bend down at short times. A grave problem arises with finite size effects in the ferromagnetic phase, as can be seen in the inset of Fig. 4. When ξ becomes comparable to the system's own length scale L, the values get overestimated and the systems start to actually equilibrate. This means we would need to go to even larger systems to get better results in the ferromagnetic phase at these temperatures.

Anyway, to study the dependence on the concentration p of ferromagnetic bonds, we fitted power laws of the form $\xi(t_w) \sim t_w^{1/z}$, which is the most-simple yet widespread approach. This power-law behavior however is subdued at the beginning. So for fitting purposes we found it is best to add a constant term, that then usually takes negative values. A crossover to logarithmic growth has been proposed and tested for the spin glass [25,26]. This is also compatible with our data, but we find no noticeable improvement over the simple power law, possibly because simulation times are still not long enough. Such a crossover was also found [58] for a strictly ferromagnetic random-bond model, which might behave similar to our case for $p > p_c$. Here we do not see a clear crossover either and the finite size effects add to the difficulty. The determined exponents z for different p are shown in Fig. 5 for both methods of extracting ξ. The phase transition can again be seen. Starting from a constant value between 10 and 11 in the spin glass phase z has a peak around the phase transition before decreasing in the ferromagnetic phase. Larger values of z correspond to slower growth of correlations and consequently overall slower behavior and equilibration. So it fits expectations that the spin glass phase has much higher values of z than the ferromagnetic phase. But interestingly we can see an even more pronounced slowdown in the critical region around the disorder-driven phase transition. Here the dynamics are so slow that one could expect even just a logarithmic growth of the coherence length with waiting time. Thus we also tried right at $p \approx p_c$ a fit to the functional form $\xi(t_w) \sim \log(t_w/t_0)^{\tilde{z}}$ with parameters t_0 and $\tilde{z}$. The fit worked as well as the power-law fit; thus we are not able to distinguish a very slow power-law growth from a logarithmic growth here. For values of p close to one, note that the finite-size effect in ξ affects the values of z, which causes a deviation from the expected $z = 2$ for $p = 1.0$. Furthermore, we tried to extrapolate the value of p_c from this data. For this purpose, we fitted a Gaussian to the peaks of the $L = 64$ and $L = 128$ data, while the data for $L = 32$ has a very pronounced peak at $p = 0.78$ (spacing $\Delta p = 0.01$). As a result, we obtain estimates $p_c(L = 32) = 0.78(1)$, $p_c(L = 64) = 0.777(2)$, and $p_c(L = 128) = 0.772(1)$. Thus no pronounced finite-size effect is visible; see also the peak of the $L = 64$ data in Fig. 5. This can be expected, since we analyzed nonequilibrium data in the time interval where the coherence length is much smaller than the system size. Thus it does not make much sense to try to extrapolate p_c for large system sizes, which appears anyway not necessary since the obtained values for p_c are very well compatible with the finite-size estimate from equilibrium studies.

In order to validate these results for ξ we will take a look at the autocorrelation C from (7). An example for $p = 0.5$ is shown in Fig. 6. Obviously, the use of the powerful GPU technique allowed us to obtain a very good statistics over large time scales such that we can observe very clean power-law behavior over many magnitudes of time. Transitions can be seen around $t = t_w$, respectively, from the equilibrium regime with slow algebraic decay to the aging regime with faster decay. This corresponds to the notion that up to time t_w the system is equilibrated on length scales of size $\xi(t_w)$ and it takes time $t > t_w$ to make a spin feel that a system is not equilibrated at longer length scales.

To obtain x, often called the equilibrium exponent in spin glass context and defined in Eq. (8), we perform fits inspired by Ref. [59] of the form $C_\infty(t) = A \cdot t^{-x} + M^2$. This is particularly difficult in the ferromagnetic phase $p > p_c$, since we performed the simulation at rather low temperatures relative to T_c. This means that the magnetization M is rather large, allowing only for a small drop of the correlation towards M^2 over time. Thus we first plotted $C_\infty(t) - M^2$ on a log-log

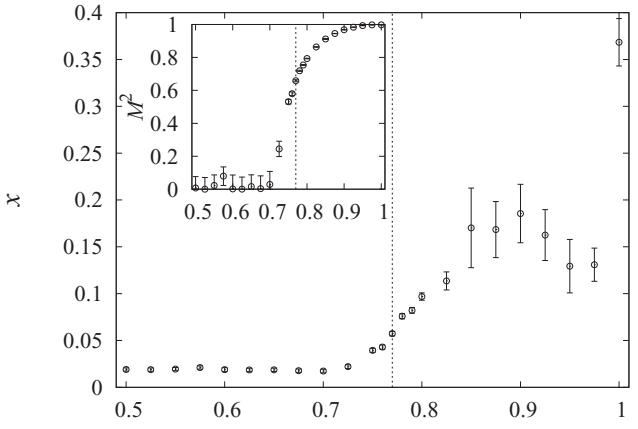

FIG. 7. Exponent x of the autocorrelation as a function of the bond probability p for a 128^3 system. The vertical line marks $p = p_c$. Inset: the square of the equilibrium magnetization estimated from a fit to the quasiequilibrated region of C.

scale and varied M^2 such that we observed a straight line on a time interval as large as possible. This time interval we used for the actual fit. Typically this resulted in an interval like $[10,10^4]$. We show in Fig. 7 the exponent x at different probabilities along with the squared magnetization in the inset. In the spin-glass phase M^2 naturally starts at basically zero but then makes its appearance slightly before the transition and grows, almost approaching 1 for $p \to 1$. The exponent begins growing at the same point from an initial value $x \approx 0.019$ for the spin glass $p = 0.5$ and reaching $x \approx 0.4$ for the ferromagnet $p = 1.0$. Deep in the ferromagnetic regime $M^2 \approx 1$; thus the correlation function drops only slightly. Thus the value of x is difficult to measure with our method in that regime.

Next we test the assumption that the aging part scales with $\xi(t_w + t)/\xi(t_w)$ in the spin-glass phase at $p = 0.5$ by trying a collapse in Fig. 8 using a multiplicative decomposition $C(t_w,t) = C_{\mathrm{eq}}(t) \cdot C_{\mathrm{age}}(t_w,t)$. We used the results for ξ previously obtained using the fit approach. This gave a better

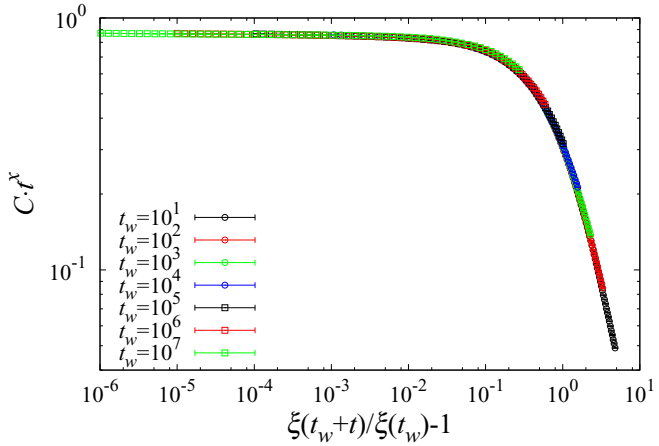

FIG. 8. (Color online) Collapse of the autocorrelation C at different waiting times t_w for a 128^3 system at $p = 0.5$. The values for ξ are interpolated from the fitting results. x is adjusted for best collapse.

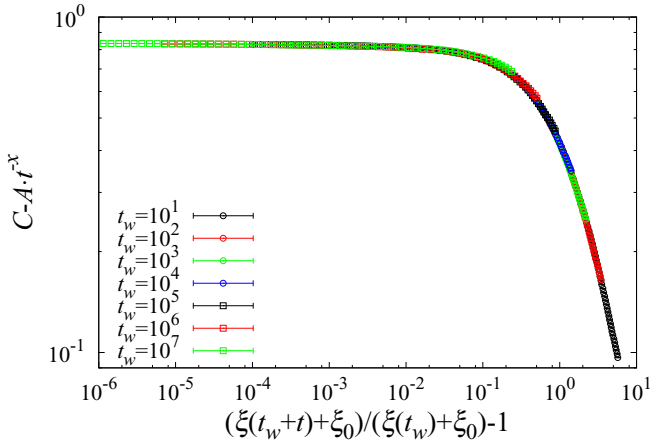

FIG. 9. (Color online) Collapse of the autocorrelation C at different waiting times t_w for a 128^3 system at $p = 0.8$. The values for ξ are interpolated from the integral results. x, A and an additive constant $\xi_0 \approx -0.66$ to ξ are adjusted for best collapse.

collapse than the integral results especially for small t_w. For the display in Fig. 8, we subtract 1 from our abscissa to make the collapse of values for $t \ll t_w$ and thus $\xi(t_w + t) \approx \xi(t_w)$ better visible. As can be seen from the figure, the quality of the collapse is very good.

The optimal value of $x \approx 0.016$ for the collapse is a bit smaller than the fit result shown in Fig. 7 and generally seems to be susceptible to the particular form of ξ. Note that we tried also collapses with an additive decomposition $C(t_w,t) = C_{\mathrm{eq}}(t) + C_{\mathrm{age}}(t_w,t)$ but this resulted everywhere in the spin-glass phase in a worse overlap of the curves even with its additional free parameter.

Starting in the region $p > 0.7$, when M^2 first becomes measurably larger than zero, the additive decomposition can give a good collapse as well and beyond the transition $p > 0.77$ the multiplicative collapse ceases working. Here $p = 0.8$ seems to be the furthest we can go into the ferromagnetic phase, before larger errors and finite-size effects prevent finding a good collapse at all. We show this collapse in Fig. 9 using the integral results for ξ this time, as they are slightly better suited here. In addition to the parameters x and A associated with C_{eq} we have to introduce an additive constant $\xi_0 \approx -0.66$ to ξ, see label of the x axis, in order for the collapse to work. This constant represents the corrections to scaling which arise at small correlation lengths, where the additive decomposition is apparently more susceptible to them. The exponent $x \approx 0.1$ is slightly larger than the fit result this time and again very sensitive. For the scaling of the x axis, we also tried different growth laws like power law, enhanced power law, or logarithm law [7,60] instead of the actually measured correlation lengths. This usually gave better collapses by adjusting the parameters, but the end results were unphysical, possibly due to overfitting. Thus we do not show these collapses here.

Finally, we also studied the aging behavior of the autocorrelation, i.e., at long times beyond t_w in the ferromagnetic phase. Here a power-law behavior $C(t) \sim t^{-\lambda}$ is visible. For intermediate waiting times, e.g., $t_w = 10^4$, the fit works well in the range $p_c < p < 0.95$ and the value of λ changes only

slowly with p, i.e., increases from $\lambda = 0.090(1)$ for $p = 0.8$ to $\lambda = 0.131(2)$ for $p = 0.95$. For smaller waiting times the dependence of λ on p is stronger, e.g., for $t_w = 100$ it grows from $\lambda = 0.116(1)$ at $p = 0.8$ to $\lambda = 0.209(1)$ at $p = 0.95$ and finally $\lambda = 0.257(2)$ at $p = 0.975$.

V. CONCLUSION

We were able to perform relatively long simulations of the random-bond Edwards-Anderson model at low temperatures for multiple different bond probabilities p from the spin-glass phase up to the ferromagnetic one. This was made possible largely through the use of general-purpose GPU (CUDA) computing, which has become feasible in recent years. Thus the CUDA-based approach allows for very fast simulations of spin glasses at much cheaper costs compared to standard CPU systems or even compared to specific FPGA-based hardware like the Janus computer [36].

The aging behavior of the spin-glass phase seen in previous studies, restricted to $p = 0.5$, was reproduced well. The main purpose of our work was to study the aging behavior of the system as a function of the variable fraction p of ferromagnetic bonds. Using this nonequilibrium study, we could easily detect the transition from the spin-glass phase to the ferromagnet, when altering the bond probability p. This was visible in all quantities we measured. Note that further into the ferromagnetic phase finite-size effects began to complicate matters considerably and we cannot give exact results there. Our entry point, the spatial correlation, contains information about the growing length scale of a system but, because we lack an explicit form for it, getting reliable values is difficult. The integral estimators taken from Refs. [29,33] have proven useful but do not seem to work as well at short times. A multifit of the assumed form (4) can help out for these cases. It also turns out the matching form changes noticeably when crossing the phase transition line. While the exponent β of the stretched exponential only lowers slightly, the other exponent α basically vanishes in the ferromagnetic phase, thus arriving at a simpler form.

The autocorrelation exhibits a quasiequilibrated part with power law behavior in the spin-glass phase. This part approaches a plateau for the ferromagnet associated with the square of the equilibrium magnetization M^2. The exponent x grows when moving deeper into the ferromagnet phase. On the other hand, the range over which a power law can be fitted becomes smaller; thus x is becoming harder to measure. The aging part, on the other hand, can be scaled well with the quotient $\xi(t_w + t)/\xi(t_w)$ in the spin-glass and a bit into the ferromagnetic phase, also giving credence to their calculated values. The assumption of power law growth works for the coherence length albeit with a small correction for early values. However, we did not delve into testing a crossover to logarithmic growth. The power law exponent z is naturally higher for the slower dynamics of the spin glass. But it also shows additional slowdown in the critical region around the phase transition. In essence all findings agree with the general expectations that in the ferromagnet the equilibrium state is more uniform and stationary and systems can arrive there much faster. Nevertheless, we were quite surprised how well the equilibrium disorder-driven transition shows up when measuring the nonequilibrium aging properties.

However, our whole approach was focused on the long time simulation of spin glasses and as such we could not get as good results for the ferromagnet. Because of the faster evolution a different emphasis would have to be placed to fare better. Also it can be seen as a bit questionable to make use of the trick of recycling random numbers for different samples without a strong influence of the disorder. In any case, beyond the physical results, the developed implementation and analysis methods can be used to proceed further efficiently with investigations of the equilibrium and nonequilibrium behavior of the random bond model.

ACKNOWLEDGMENT

We thank Martin Weigel and A. Peter Young for many interesting discussions and helpful suggestions.

[1] K. Binder and A. P. Young, Rev. Mod. Phys. **58**, 801 (1986).
[2] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
[3] K. H. Fischer and J. A. Hertz, *Spin Glasses* (Cambridge University Press, Cambridge, UK, 1993).
[4] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press, Oxford, 2001).
[5] N. Kawashima and H. Rieger, in *Frustrated Spin Systems*, 2nd ed., edited by H. T. Diep (World Scientific, Singapore, 2013), pp. 509–614.
[6] J. A. Mydosh, *Spin Glasses: An Experimental Introduction* (Taylor and Francis, London, 1993).
[7] J.-P. Bouchaud, L. F. Cugliandolo, J. Kurchan, and M. Mézard, in *Spin Glasses and Random Fields*, edited by A. P. Young (World Scientific, Singapore, 1998), pp. 161–224.
[8] F. Corberi, L. F. Cugliandolo, and H. Yoshino, in *Dynamical Heterogeneities in Glasses, Colloids, and Granular Media*, edited by L. Berthier, G. Biroli, J.-P. Bouchaud, L. Cipelletti, and W. van Saarloos (Oxford University Press, Oxford, 2011), pp. 370–406.
[9] D. Sherrington and S. Kirkpatrick, Phys. Rev. Lett. **35**, 1792 (1975).
[10] S. F. Edwards and P. W. Anderson, J. Phys. F **5**, 965 (1975).
[11] G. Parisi, Phys. Rev. Lett. **43**, 1754 (1979).
[12] G. Parisi, Phys. Rev. Lett. **50**, 1946 (1983).
[13] D. S. Fisher and D. A. Huse, Phys. Rev. Lett. **56**, 1601 (1986).
[14] D. S. Fisher and D. A. Huse, Phys. Rev. B **38**, 373 (1988).
[15] W. L. McMillan, Phys. Rev. B **30**, 476 (1984).
[16] N. Kawashima and A. P. Young, Phys. Rev. B **53**, R484 (1996).
[17] 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).

[18] H. G. Katzgraber, M. Körner, and A. P. Young, Phys. Rev. B **73**, 224432 (2006).
[19] M. Hasenbusch, A. Pelissetto, and E. Vicari, Phys. Rev. B **78**, 214205 (2008).
[20] 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. Tripiccione, and D. Yllanes, Phys. Rev. B **88**, 224416 (2013).
[21] H. Rieger, Physica A **224**, 267 (1996).
[22] J. Kisker, L. Santen, M. Schreckenberg, and H. Rieger, Phys. Rev. B **53**, 6418 (1996).
[23] E. Marinari, G. Parisi, J. J. Ruiz-Lorenzo, and F. Ritort, Phys. Rev. Lett. **76**, 843 (1996).
[24] T. Komori, H. Yoshino, and H. Takayama, J. Phys. Soc. Jpn. **69**, 1192 (2000).
[25] J.-P. Bouchaud, V. Dupuis, J. Hammann, and E. Vincent, Phys. Rev. B **65**, 024439 (2001).
[26] L. Berthier and J.-P. Bouchaud, Phys. Rev. B **66**, 054404 (2002).
[27] H. Yoshino, K. Hukushima, and H. Takayama, Phys. Rev. B **66**, 064431 (2002).
[28] L. D. C. Jaubert, C. Chamon, L. F. Cugliandolo, and M. Picco, J. Stat. Mech. (2007) P05001.
[29] F. Belletti, A. Cruz, L. A. Fernandez, A. Gordillo-Guerrero, M. Guidetti, A. Maiorano, F. Mantovani, E. Marinari, V. Martin-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. Tripiccione, and D. Yllanes, J. Stat. Phys. **135**, 1121 (2009).
[30] C. Chamon, F. Corberi, and L. F. Cugliandolo, J. Stat. Mech. (2011) P08015.
[31] Y. Ozeki and H. Nishimori, J. Phys. Soc. Jpn. **56**, 1568 (1987).
[32] A. K. Hartmann, Phys. Rev. B **59**, 3617 (1999).
[33] F. Belletti, M. Cotallo, A. Cruz, L. A. Fernandez, A. Gordillo-Guerrero, M. Guidetti, A. Maiorano, F. Mantovani, E. Marinari, V. Martin-Mayor, A. M. Sudupe, D. Navarro, G. Parisi, S. Perez-Gaviro, J. J. Ruiz-Lorenzo, S. F. Schifano, D. Sciretti, A. Tarancon, R. Tripiccione, J. L. Velasco, and D. Yllanes, Phys. Rev. Lett. **101**, 157201 (2008).
[34] E. Lippiello, F. Corberi, A. Sarracino, and M. Zannetti, Phys. Rev. B **77**, 212201 (2008).
[35] E. Lippiello, F. Corberi, A. Sarracino, and M. Zannetti, Phys. Rev. E **78**, 041120 (2008).
[36] F. Belletti, M. Cotallo, A. Cruz, L. A. Fernandez, A. Gordillo-Guerrero, M. Guidetti, A. Maiorano, F. Mantovani, E. Marinari, V. Martin-Mayor, A. Muñoz-Sudupe, D. Navarro, G. Parisi, S. Perez-Gaviro, M. Rossi, J. J. Ruiz-Lorenzo, S. F. Schifano, D. Sciretti, A. Tarancon, R. Tripiccione, J. L. Velasco, D. Yllanes, and G. Zanier, Comput. Sci. Eng. **11**, 48 (2009).
[37] NVIDIA CUDA C Programming Guide, Version 4.0, NVIDIA Corporation, 2011.
[38] T. Preis, P. Virnau, W. Paul, and J. J. Schneider, J. Comput. Phys. **228**, 4468 (2009).
[39] K. A. Hawick, A. Leist, and D. P. Playne, Int. J. Parallel Prog. **39**, 183 (2011).
[40] B. Block, P. Virnau, and T. Preis, Comput. Phys. Commun. **181**, 1549 (2010).
[41] M. Weigel, J. Comput. Phys. **231**, 3064 (2012).
[42] M. Guidetti, A. Maiorano, F. Mantovani, M. Pivanti, S. F. Schifano, and R. Tripiccione, in *Applied Parallel and Scientific Computing*, Lecture Notes in Computer Science Vol. 7133 (Springer, New York, 2012), pp. 220–230.
[43] H. W. Blöte, L. Shchur, and A. L. Talapov, Int. J. Mod. Phys. C **10**, 1137 (1999).
[44] F. Corberi, E. Lippiello, and M. Zannetti, J. Stat. Mech.: Theory Exp. (2004) P12007.
[45] D. A. Huse and D. S. Fisher, Phys. Rev. B **35**, 6841 (1987).
[46] L. Golubović and S. Feng, Phys. Rev. B **43**, 972 (1991).
[47] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, Oxford, 1999).
[48] M. Manssen, Master's thesis, Carl-von-Ossietzky University Oldenburg, 2011.
[49] D. B. Kirk and W.-m. W. Hwu, *Programming Massively Parallel Processors: A Hands-on Approach*, 2nd ed. (Morgan Kaufmann, Burlington, MA, 2012).
[50] M. Manssen, M. Weigel, and A. K. Hartmann, Eur. Phys. J. Spec. Top. **210**, 53 (2012).
[51] M. Weigel, Comput. Phys. Commun. **182**, 1833 (2011).
[52] M. Weigel and T. Yavors'kii, Phys. Proc. **15**, 92 (2011).
[53] M. Lulli, M. Bernaschi, and G. Parisi, arXiv:1411.0127.
[54] K. Barros, P. L. Krapivsky, and S. Redner, Phys. Rev. E **80**, 040101 (2009).
[55] J. Olejarz, P. L. Krapivsky, and S. Redner, Phys. Rev. Lett. **109**, 195702 (2012).
[56] T. Blanchard and M. Picco, Phys. Rev. E **88**, 032131 (2013).
[57] T. Blanchard, F. Corberi, L. F. Cugliandolo, and M. Picco, EPL (Europhys. Lett.) **106**, 66001 (2014).
[58] F. Corberi, E. Lippiello, A. Mukherjee, S. Puri, and M. Zannetti, J. Stat. Mech. (2011) P03016.
[59] R. Paul, G. Schehr, and H. Rieger, Phys. Rev. E **75**, 030104 (2007).
[60] L. F. Cugliandolo, Out of equilibrium dynamics of complex systems, lecture notes King's College, London, 2013 (unpublished).