Aging in the three-dimensional Random Field Ising Model

Sebastian von Ohr,[*] Markus Manssen, and Alexander K. Hartmann

Institute of Physics, Carl von Ossietzky University, 26111 Oldenburg, Germany

(Dated: January 2, 2017)

We studied the nonequilibrium aging behavior of the Random Field Ising Model in three dimensions for various values of the disorder strength. This allowed us to investigate how the aging behavior changes across the ferromagnetic-paramagnetic phase transition. We investigated a large system size of $N = 256^3$ spins and up to 10^8 Monte Carlo sweeps. To reach these necessary long simulation times we employed an implementation running on Intel Xeon Phi coprocessors, reaching single spin flip times as short as 6 ps. We measured typical correlation functions in space and time to extract a growing length scale and corresponding exponents.

I. INTRODUCTION

Studying the relationship between equilibrium (static) and non-equilibrium (dynamic) behavior is one of the main aims of statistical physics. A prominent example are structural glasses [1], where it is not even clear whether the drastic changes in dynamic behavior are accompanied by an equilibrium phase transition. On the other hand, lattice spin models like spin glasses and random field systems [2] exhibit for sure phase transitions in high enough spatial dimensions, but the dynamics are still not fully understood.

The Random Field Ising Model [3] (RFIM) has been studied extensively in theory [4] and in experiments [5], realized using dilute antiferromagnets. The equilibrium properties are reasonably well understood. Here it is particularly convenient that in any dimension exact ground states can be calculated numerically for large system sizes using mapping to the maximum-flow problem [6], along the full disorder parameter line. Nevertheless, the dynamics is still under discussion. Of particular interest is the aging behavior, i.e., the dynamics resulting from starting in an equilibrium state followed by a rapid parameter change to a target point in phase space resulting in an non-equilibrium situation [7, 8]. Here we study the quench from a random configuration at infinite temperature to a low temperature with suitably small random-field disorder, which results in domains of parallel spins forming and growing over time. The disorder introduced by the random field pins the domain wall and slows down the growth of domains.

Typically, literature on aging concentrates on isolated target points in phase diagrams. Here we are interested in correlating the aging behavior with a disorder-driven phase transition. So far, aging in the RFIM has been analyzed mostly in the ferromagnetic phase [9–12]. Some recent articles [13, 14] cover the two-dimensional (2D) RFIM in the disordered phase, but no studies of the 3D RFIM in a larger space of the phase diagram are known to the authors. This is probably be due to the fact that such studies require many-parameter long-run simulations of large systems, going beyond the amount of computational resources usually available. With the introduction of comparable low-cost cards like GPU and Intel-Phi, such studies become feasible. Similar to previous work [15] analyzing the aging across the spin glass-ferromagnet transition in the Edwards-Anderson model, we here present results for the dynamic behavior for the 3D RFIM across the disorder-driven ferromagnet-paramagnet transition. We will be looking at the spatial correlation and discuss different methods of extracting the coherence length and corresponding exponents from it. We will also look at the autocorrelation and try to collapse it by rescaling the time in units of the coherence length, also yielding suitably defined exponents.

To reach sufficient long simulation times we implemented the model on the Intel Xeon Phi coprocessor. These cards offer performance comparable to Graphic Processing Units (GPUs), but the architecture is more similar to current CPUs, just with more cores. Parallelization does not require learning a new programming extension as for GPUs, but can be done using well-known techniques, e.g., OpenMP [16], MPI [17] or just creating threads manually. Porting a simulation from CPU to Xeon Phi cards is straightforward; however, to fully utilize the performance a lot of knowledge about the architecture and careful optimization is necessary. Using our optimized implementation of the model we were able to simulate 10^8 sweeps for 64 disorder samples of a large $N = 256^3$ system at many different values of the disorder strength.

The remainder of this article is structured as follows. In Sec. II we describe the RFIM and the observables used to characterize the aging of the system. Section III describes details of the implementation on Xeon Phi cards. Results of the simulation are presented in Sec. IV. We close with our conclusions in Sec. V.

II. MODEL

The Random Field Ising Model describes a D-dimensional cubic system of side length L containing $N = L^D$ Ising spins $S_i = \pm 1$. The Hamiltonian is given by

$$H(S) = -J \sum_{\langle i,j \rangle} S_i S_j - \sum_i h_i S_i \qquad (1)$$

where the sum runs over nearest neighbors $\langle i, j \rangle$ and the field $h_i = h_0 \varepsilon_i$ with $\varepsilon_i = \pm 1$ being a quenched random variable. Here, we apply a symmetric bimodal distribution $P(\varepsilon_i) =$
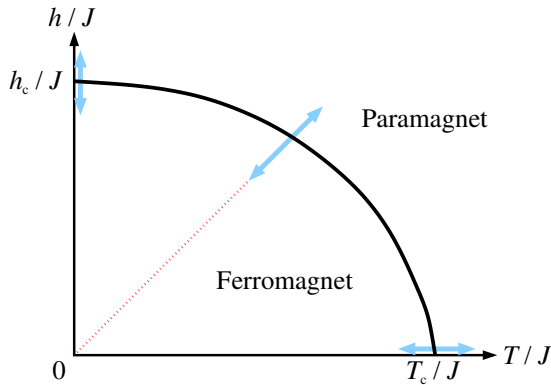
---

[*] sebastian.von.ohr@uni-oldenburg.de

FIG. 1. (Color online) Sketch of the phase diagram: for low temperatures and low disorder, the system is in a ferromagnetic phase, otherwise in a paramagnetic phase.

$[\delta(\varepsilon_i - 1) + \delta(\varepsilon_i + 1)]/2$. The boundary conditions are periodic in all directions. The parameter h_0 controls the strength of the random field. For $h_0 = 0$ the well known pure Ising Model is reproduced with a paramagnetic phase at high temperatures $T > T_c$ and a ferromagnetic phase at low temperatures. With $h_0 > 0$ additional disorder is introduced, which lowers the transition temperature to the paramagnetic phase. In Fig. 1 a phase diagram for $D \geq 3$ dimensions is shown [18]. Even at $T = 0$ the system is in a paramagnetic phase if the field strength becomes too large $h_0 > h_c$. In the remainder of the article will only be concerned with the case $D = 3$, which has a zero-disordered critical temperature $T_c \approx 4.5115$ [19] and a zero-temperature critical field strength $h_c \approx 2.20$ [20, 21].

Simulations start with random initial configurations, emulating a quench from infinite temperature. A single-spin flip Monte Carlo simulation is performed, see below for details. We then examine the system at different waiting times t_w (measured in Monte Carlo sweeps) after the beginning of the simulation. The order parameter is the magnetization

$$m = \frac{1}{N} \sum_i S_i. \tag{2}$$

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

$$C_2(r, t_w) = \frac{1}{N} \sum_i S_i(t_w) S_{i+r}(t_w) \tag{3}$$

between two points. 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 spatial correlation function, for a recent comparison in the case of the three-dimensional random-bond (spin-glass) model see, e.g., Ref. 15. Most approaches are based on the assumption that C_2 follows the functional form

$$C_2(r, t_w) \propto r^{-\alpha} g\left(\frac{r}{\xi(t_w)}\right), \tag{4}$$

where the function g is approximately a stretched exponential $g(x) \approx \exp(-x^\beta)$. Extraction of ξ and corresponding exponents works by fitting (4) to the data of C_2.

Alternatively one can use integral estimators, first used for spin-glasses [22], which work without assuming the functional form of g. This is done by calculating the integral

$$I_k(t_w) = \int_0^{L/2} r^k C_2(r, t_w) \, dr \tag{5}$$

which allows the calculation of the coherence length using

$$\xi_{k,k+1}(t_w) = \frac{I_{k+1}(t_w)}{I_k(t_w)} \propto \xi(t_w). \tag{6}$$

A value of $k = 1$ is recommended in Ref. 22 as a tradeoff between systematic errors for low k values and statistical errors for larger k. This method also allows to determine the exponent α since $I_1 \propto \xi_{1,2}^{2-\alpha}$.

A different, but very simple, method of extracting the coherence length uses the inverse density of defects [12]. A defect is a spin with at least one antiparallel neighbor. With the number of defects $D(t_w)$ the coherence length is given by $\xi(t_w) = N/D(t_w)$.

Another observable of interest is the autocorrelation

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

comparing the same system at different times. It is expected to split into two parts. The first quasi-equilibrated part for $t \ll t_w$ takes the form [23–27] of a power law

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

with a characteristic exponent x. The later aging part $C_{age}(t, t_w) = f(\xi(t_w + t)/\xi(t_w))$ is expected [7, 8] to depend only of the ratio of the coherence lengths at the two times. For long waiting times $\lim_{t_w \to \infty} C(t, t_w) = C_{eq}(t) + B^2$ a plateau is expected [28, 29], with B equal to the equilibrium magnetization M in the ferromagnetic phase.

III. IMPLEMENTATION

We implemented a standard Metropolis Monte Carlo simulation [30] of the model for the Intel Xeon Phi 3120P coprocessor. It is based on an earlier implementation for GPUs, but in the remainder of the article we cover only the implementation for Xeon Phi cards. For details of a similar GPU implementation see Ref. 15. As a reference for the Xeon Phi architecture we refer to the official documentation [31, 32]. It's main points are the 57 Pentium based cores with 4-way Hyper-Threading and the 512 bit wide vector processing unit.

In each sweep every spin is updated as follows. First, the energy change

$$\Delta E_i = 2JS_i \sum_{j \in N(i)} S_j + 2h_i S_i \tag{9}$$
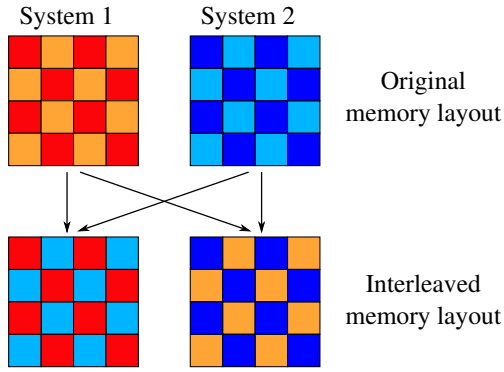
FIG. 2. (Color online) Example memory layout of a 2D system, before and after exchanging the odd (light color) spins. Each spin in one interleaved memory block can be updated independently because all neighbors are stored in the other memory block.

for a flip of spin i is calculated. The sum runs over the neighboring sites $N(i)$ of spin i. Next, a random number is generated and the spin i is flipped with probability

$$p_{\text{accept}} = \min[1, \exp(-\Delta E_i/T)] . \qquad (10)$$

To keep all cores of the Xeon Phi card busy, spins need to be updated simultaneously. If each thread would update a cubic subset of the system in a linear fashion then spins on the border of each subset have neighbors in a different subset and thus depend on data owned by a different thread. If threads are not synchronized these neighboring spins may change in a non-deterministic way, generating irreproducible results. Instead of synchronizing threads on the border, we chose to update spins in a checkerboard pattern. First, all even (sum of the x, y, z coordinates are even) spins are updated and the odd spins stay the same. Next, threads are synchronized and the roles are switched and all odd spins are updated. However, this update scheme results in an unfavorable access pattern to the memory, since we want to use SIMD instructions and load large chunks of continuous data at a time. Instead, we simulate two systems simultaneously and interleave them in memory, so that the even spins from the first system are in same memory block as the odd spins from the second system. Fig. 2 shows how spins would be exchanged in a 2D system. With this technique one can update all spins in one memory block because all the neighboring spins are stored in the other block. This also simplifies the spin update routine because spins can be updated in a linear fashion again, which effectively alternates between two samples. One sweep consists of updating both memory blocks.

We also implemented multispin coding, which means, that each spin is represented by a single bit and stored together with other spins in a single byte or some bigger data type. We used a 32 bit data type to store 32 spins of the same location but of different samples. Since the random field only takes the values $\pm h_0$ it can also be encoded using the same technique. Together with the memory interleaving this results in 64 samples for each run of the simulation. The calculation of the energy difference is then mapped to bitwise logic operations. In the case of the 3D RFIM there are a total of 14 possible

energy differences (spin aligned with $6, \ldots, 0$ neighbors, i.e., seven cases, times two cases for alignment or antialignment with field) and the spin flip probability can be precalculated for these cases and stored in a lookup table. The only thing that cannot easily be mapped to bitwise logic operations is the random number generation. So, we just use the same random number for 32 samples.

Most traditional pseudorandom number generators (PRNG) like the Mersenne Twister [33] are unsuitable for highly parallel architectures because each instance requires comparative large amounts of memory, i.e., about 2.5 kB for the Mersenne Twister. Depending on the PRNG it's also nontrivial to initialize multiple instances of the same generator so that the output streams do not overlap for a sufficient number of calls. Instead, we used a counter-based random number generator [34], namely the Philox PRNG. The difference compared to traditional generators is that they do not operate one some internal state, which is advanced with every number generated. Instead, they consist of a deterministic function, taking a key and a counter as parameters, and return a random number generated from those parameters. The key is chosen so that it's different in each thread and the counter is incremented after each function call. In our simulation we used the index $0 \ldots N-1$ of the spin as the key. This effectively reduces the state of the PRNG to a single counter, which is shared across all threads. Since the reference implementation by the authors of Ref. 34 doesn't support the Xeon Phi architecture, we implemented our own optimized version using SIMD instructions.

With the described optimizations we reached an effective single spin-flip attempt time of $\approx 6\,\text{ps}$ on the Xeon Phi card, corresponding to 57 (processors) $\times 32$ (spins per integer) $\times 16$ (vector processing unit) parallel flip attempts per $1.8 \cdot 10^{-7}\,\text{s}$, equal to roughly 180 cycles at the $10^{-9}\,\text{s}$ cycle time of the processors on the card. Compared to $\approx 9\,\text{ps}$ of our previous implementation on a GeForce GTX 570 GPU the Xeon Phi card implementation is faster. But this was expected since the GPU is older than the Xeon Phi 3120P card. Overall, the performance of the Xeon Phi card seems to be comparable to GPUs. One advantage of the Xeon Phi card was, however, that the reported spin flip time was archived with only 64 samples while the GPU implementation used 128 samples. This is because operations with 32 bit integers on the GPU are only marginal faster than operations with 64 bit integers, while on the Xeon Phi card 32 bit calculations are twice as fast as 64 bit operations. So, reducing the number of samples on the GPU to 64, increases the single spin flip time to $\approx 13\,\text{ps}$.

IV. RESULTS

For system size $N = 256^3$ we simulated 64 samples of randomly initialized realizations. Simulations were performed on 8 Intel Xeon Phi 3120P cards. The temperature was fixed to $T = 0.8$ and the field strength was chosen in the range $h_0 \in [1.2, 3.2]$, covering the ferromagnetic and paramagnetic phase. We do not know the precise critical field strength at $T = 0.8$, but analysis of the phase diagram [18] suggest that
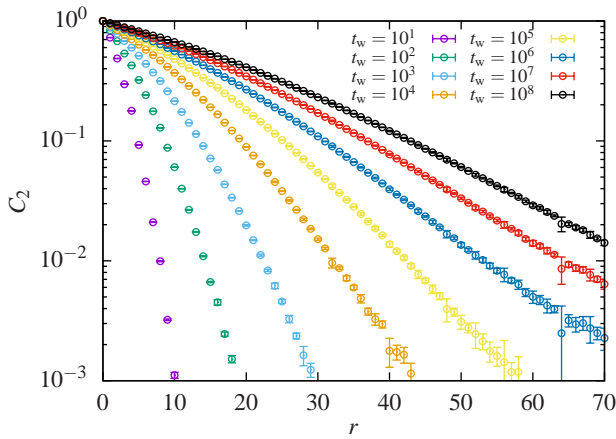
FIG. 3. (Color online) Spatial correlation C_2 over the distance r for different waiting times t_w of a 256^3 system with field strength set to $h_0 = 1.6$. Only data points at integer values of r are shown to give a clearer picture.

FIG. 4. (Color online) Coherence length $\xi_{1,2}$ as function of the waiting time t_w for different field strengths h_0. The coherence length was extracted from the spatial correlation using integral estimators. Inset: Section of the same data on a semi-log scale.

the $T = 0$ value h_c is still very close to the real critical field strength at $T = 0.8$. In most simulations 10^8 sweeps were performed, which took about a week on a single Xeon Phi card. At specific waiting times t_w we saved the whole spin configuration to disk, so that they can be post-processed to calculate the observables. For $h_0 \geq 2.4$ we performed only 10^7 sweeps, mainly because they are close to equilibration and would require a large amount of disk space since the disordered spin configurations can not be compressed very well.

An exemplary spatial correction C_2 for $h_0 = 1.6$ is shown in Fig. 3. The curves show an almost exponential decay with a slight bend, which is captured by the exponent β. The spatial correlation decreases more slowly for larger waiting times, suggesting a growing length scale.

Extracting the coherence length from the data proved to be complicated and different approaches were tried. In the end a combined approach of the integral method and fits to the tail of the spatial correlation gave the best results. Integrals (6) are calculated by numerically integrating the data until the value first becomes smaller than three times it's error. The remaining part of the integral is approximated by fitting (4) to the data and then integrating the fitted function. The choice of the fit range is not critical since only the tail of the fitted function is used and the contribution to the integral is small. The exponent α is expected to be 0 for the RFIM, but was added as a parameter nevertheless since previous works only analyzed the system in the ferromagnetic region.

The extracted coherence length is shown in Fig. 4 for different disorder values. Previous studies of the RFIM in one to three dimensions [9–12, 14, 36, 37] showed a crossover from a power law to logarithmic growth in the asymptotic regime. This behavior of the coherence length seems to be universal for disordered Ising systems [38–41]. Since comparatively high field strengths h_0 were chosen for our simulation, the power law regime is not really visible, but instead the crossover to logarithmic growth happens almost immediately. The curves for $h_0 \geq 2.8$ equilibrate within the simulated wait-
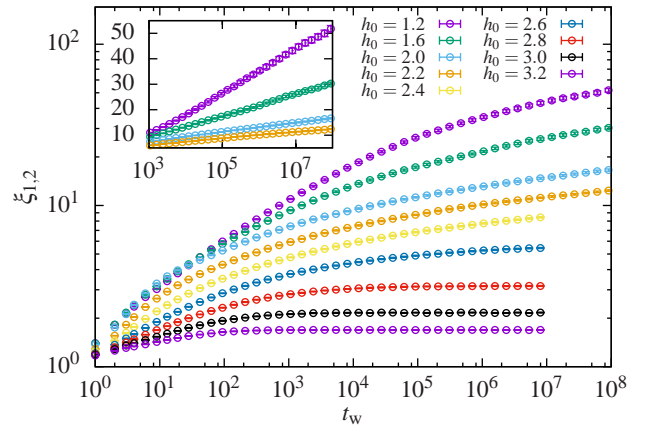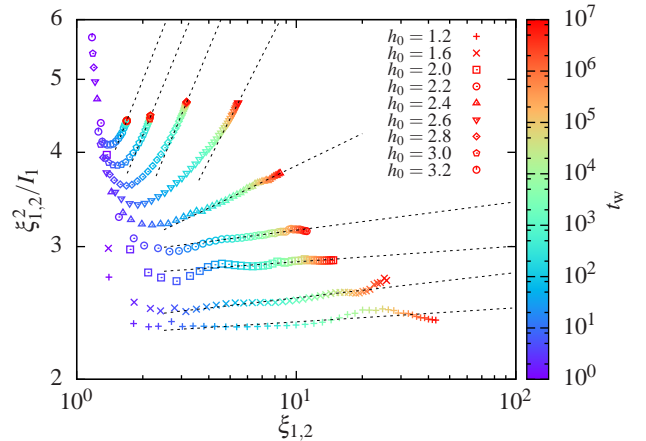

FIG. 5. (Color online) Power law fits (dashed lines) to extract the exponent α from the integral estimators. The color gradient shows from which waiting time the values of $\xi_{1,2}$ and I_1 were extracted.

ing time. In the beginning the curves for $h_0 = 1.6$ and $h_0 = 2.0$ seem to grow faster than the $h_0 = 1.2$ curve, but this is a result of the undetermined proportionality factor.

It was checked that the coherence length shows no finite size effects by simulating a $L = 512$ system for up to $t_w = 10^7$ sweeps and extracting the coherence length. It is essentially the same as for the $L = 256$ system (not shown). Therefore, we believe that the observables are free from finite size effects, even for $t_w = 10^8$.

The integral estimators also provides a means to extract the exponent α, using the relation $I_1 \propto \xi_{1,2}^{2-\alpha}$. By plotting $\xi_{1,2}^2/I_1$ over $\xi_{1,2}$, as depicted in Fig. 5, the exponent can be extracted using a power law fit. It is observed, that the curves do not show a clear power law, especially for small $\xi_{1,2}$. However, for larger $\xi_{1,2}$ all curves become more or less straight lines. Power laws were fitted to those parts, subjectively choosing the beginning of the fit range for every curve. It can be seen
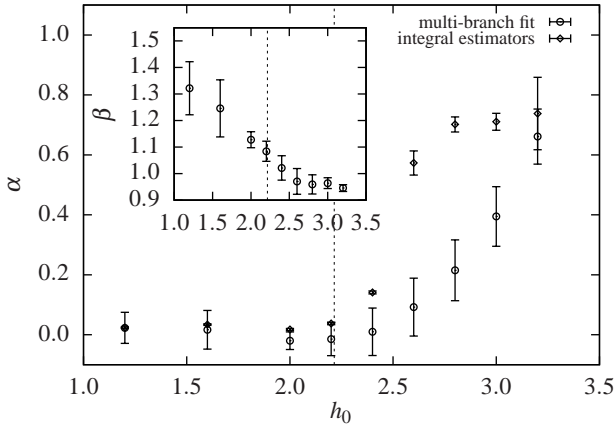
FIG. 6. Scaling exponent α of the spatial correlation over the field strength h_0 for a 256^3 system, extracted using different methods. The vertical line marks $h_0 = h_c$. Inset: Associated scaling exponent β.

FIG. 7. (Color online) Autocorrelation C as a function of time t for different waiting times t_w at field strength $h_0 = 1.2$.

FIG. 8. (Color online) Autocorrelation C as a function of time t for different waiting times t_w at field strength $h_0 = 2.8$. Note that the data for $t_w = 10^4$, 10^5 and 10^6 falls on top of each other.

that there is a clear crossover from a slope close to zero for $h_0 \leq 2.2$ to a larger slope for larger h_0. Note that for high disorder values the system equilibrates and the coherence length stops growing. For these curves only the short waiting times contribute to the power law, as it can be seen from the color gradient. The resulting exponent α for different disorder values is displayed in Fig. 6, together with the exponents α and β extracted using a multi-branch fit of (4) to the spatial correlation. By using a multi-branch fit [35] the exponents are forced to a common optimal value for all waiting times. The different methods to extract α mostly agree, except for the range $h_0 \in [2.4, 3]$ where the α from the integral estimators grows faster. Below the critical field strength the exponent α is zero within the error margin, just like expected. In the paramagnetic region α grows quickly with increasing h_0. We also tried to extract the exponent β again using a fit with (4), but with the parameters α and ξ fixed to the already extracted values from the integral estimator and with an additional adjustable prefactor. Because of the discrepancies for the exponent α, no reasonable fit was possible in the $h_0 \in [2.4, 3]$ range. Therefore, we only displayed the β results from the multi-branch fit. The exponent β also undergoes a change near $h \approx 2.6$, i.e., a bit beyond the phase transition.

Next, we look at the autocorrelation from (7). An exemplary curve for $h_0 = 1.2$ is shown in Fig. 7. The two parts of the autocorrelation with a transition around $t \approx t_w$ can be observed. The first quasi-equilibrated part looks like a constant function while the second aging part first decays as a power law which then slows down for short waiting times. In the paramagnetic region the system equilibrates and the autocorrelation relaxes to a plateau, as can be seen in Fig. 8 for the case $h_0 = 2.8$. This is because the spins are mostly aligned in the direction of the field and the autocorrelation shows only the fluctuation around this configuration. To extract the exponent x of the quasi-equilibrated part (8) a fit of the form $C_\infty(t) = A \cdot t^{-x} + B^2$ was performed. In the paramagnetic phase the value of B can be easily read from the height of the plateau. The fit range was chosen by plotting $C(t) - B^2$ on
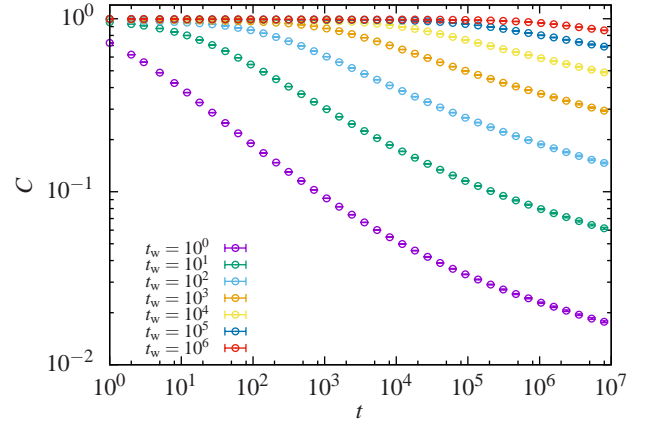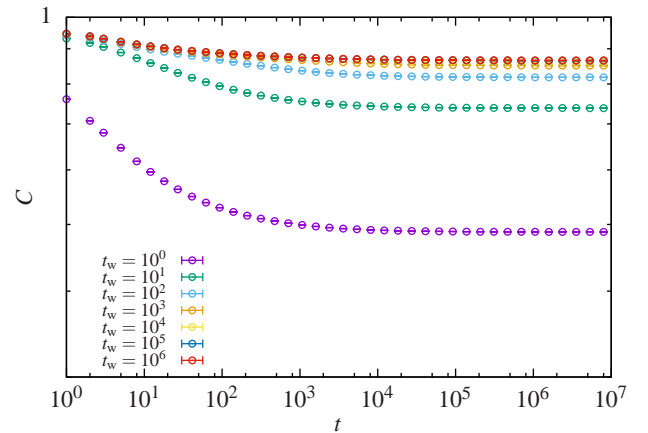
a log-log scale and restricting the fit to the straight part of this curve. In the ferromagnetic phase the exponent x is close to 0 and doesn't allow to determine B with reasonable accuracy. Therefore, we restricted the fit to the straight part of $C(t, t_w)$ and show only the results for the exponent x in Fig. 9. Here again the change in dynamics is visible slightly above the critical field strength.

Last we check the assumption that the aging part scales with $\xi(t_w + t)/\xi(t_w)$. In the ferromagnetic phase we can just ignore the quasi-equilibrated part, because it is close to a constant function, and do a collapse of C directly, as depicted in Fig. 10 for the case $h_0 = 1.2$. We subtract one from the abscissa to make the collapse for values $t \ll t_w$ better visible. It can be seen that the quality of the collapse is very good.

In the paramagnetic phase the quasi-equilibrated part is not constant and has to be accounted for in the collapse. We tested an additive decomposition $C(t, t_w) = C_{eq}(t) + C_{age}(t, t_w)$ and also a multiplicative decomposition $C(t, t_w) = C_{eq}(t) \cdot C_{age}(t, t_w)$. We found that the additive decomposition gave
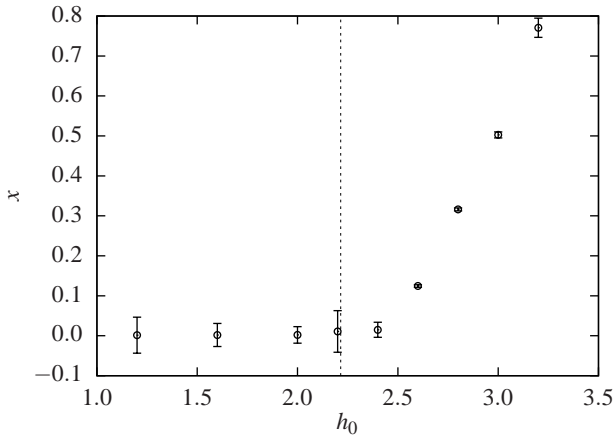
FIG. 9. Equilibrium exponent x for different values of disorder h_0. The vertical line marks $h_0 = h_c$.

FIG. 11. (Color online) Collapse of autocorrelation C with correction for the quasi-equilibrated part C_{eq} for different waiting times t_w at field strength $h_0 = 2.8$. An additional additive constant $\xi_0 \approx -1.333$ was added to the coherence length to give a better collapse.

V. CONCLUSION

For the Random Field Ising Magnet, we performed long-time Monte Carlo simulations for large systems ($N = 256^3$) at low temperature and different disorder strengths h_0, covering the ferromagnetic and paramagnetic phase. This was made possible by the usage of several Intel Xeon Phi cards and an optimized implementation of the model. Getting the first working simulation on the Xeon Phi card was relatively easy because of the similar architecture to CPUs. However, such a naive implementation was slower than the same naive implementation on a current CPU. The main reason for this seems to be the poor usage of the vector processing unit. Therefore, we rewrote most of the code and directly accessed the vector processing unit by using SIMD intrinsics. Implementing the same model on GPUs was more complicated in the beginning, but already the first version outperformed the CPU. We observed that the overall time spent optimizing for the Xeon Phi card and GPUs is roughly the same.

By using this implementation, we were able to study the aging behavior of the model. We analyzed the results by looking at the spatial correlation and the autocorrelation. The coherence length was extracted from the spatial correlation using integral estimators. Further exponents were extracted using multi-branch fits. Moving into the paramagnetic phase complicated the analysis since the spatial correlation is very short and the system equilibrates within the simulated timespan. At the ferromagnet-paramagnet phase transition we found a clear change of the aging behavior, as visible in the exponents α, β and x. However, the exponent β and x show a transition slightly above h_c, which might indicate a connection to a percolation transition like that found slightly above the phase transition for the similar RFIM with Gaussian disorder [42]. Note that in a previous study [15] of the Edwards-Anderson model a change in dynamics was found that coincides fully with the ferromagnet-spin glass transition, i.e., it was visible
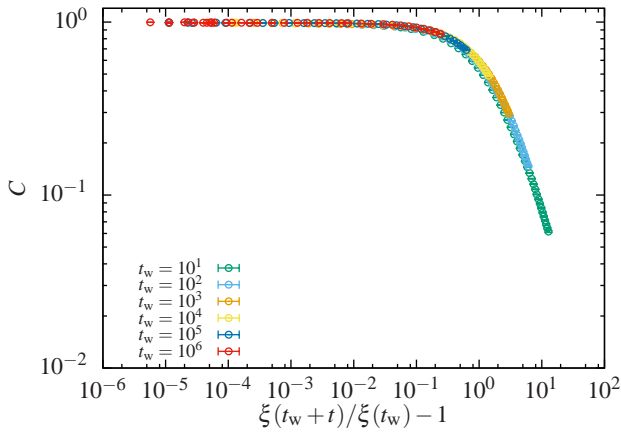

FIG. 10. (Color online) Collapse of autocorrelation C for different waiting times t_w at field strength $h_0 = 1.2$.

a better collapse in the $t \ll t_w$ range, as shown in Fig. 11. The parameters A, x and ξ_0 were adjusted for the best collapse. The parameter ξ_0 accounts for finite-size corrections to scaling, e.g., due to the discreteness of the lattice, which becomes notable at small scales, which is in particular true in the paramagnetic phase (see Fig. 4). The corrections are not known here, thus our approach is purely heuristic. Nevertheless, for $\xi \to \infty$ the used term $[\xi(t_w + t) + \xi_0]/[\xi(t_w) + \xi_0]$ crosses over to $\xi(t_w + t)/\xi(t_w)$ as required. The magnitude $\xi_0 \approx -1.333$ of the correction-to-scaling parameter is small and comparable to the lattice spacing. Nevertheless, without introducing the parameter ξ_0, the collapse in considerably worse (not shown here).

For this example the parameter x is close to the one extracted using a fit, as shown in Fig. 9, but in the ferromagnetic phase the collapse gives unrealistic large values for x. Except for the tails of the individual curves the collapse works reasonably well. The slight increase toward the end can probably be attributed to the system reaching equilibrium.
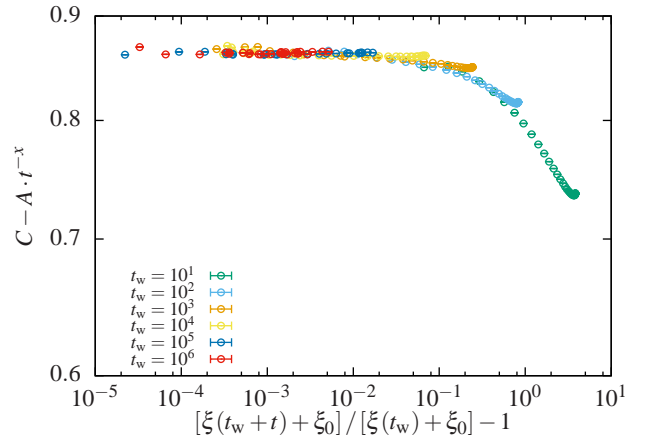
in all measured exponents simultaneously.

The autocorrelation splits into a quasi-equilibrated part, which follows a power law with the exponent x, and a later aging part. In the ferromagnetic phase x is close to 0 and starts growing in the paramagnetic phase. Thus, we can mostly ignore the quasi-equilibrated part in the ferromagnetic phase and do a collapse of the aging part by plotting the autocorrelation over the quotient $\xi(t_\mathrm{w}+t)/\xi(t_\mathrm{w})$. In the paramagnetic phase the collapse is becoming more difficult with increasing disorder strength. We found that an additive decomposition of the autocorrelation and an additive constant to the coherence length gave the best collapse.

Overall, we found that the equilibrium phase-transition behavior of the model and maybe a particular percolation transition are well reflected in the dynamic observables. Due to the availability of relatively inexpensive yet powerful GPU and Intel-Phi architectures, one could easily extend our studies to the RFIM in higher dimensions [43, 44], to RFIM systems with correlations [45]. Furthermore, it could be interesting to investigate whether the aging behavior can be understood in terms of non-trivial low lying excitations [46]. Certainly, it would be interesting to perform such studies for other spin models to investigate the relationship between the dynamic behavior and equilibrium phase transitions.

[1] K. Binder and W. Kob, *Glassy Materials and Diosordred Solids* (World Scientific, Singapore, 2011).

[2] A. P. Young, ed., *Spin glasses and random fields* (World Scientific, Singapore, 1998).

[3] Y. Imry and S.-k. Ma, Phys. Rev. Lett. **35**, 1399 (1975).

[4] T. Nattermann, in *Spin Glasses and Random Fields*, edited by A. P. Young (World Scientific, 1997), pp. 277–298.

[5] D. P. Belanger, in *Spin Glasses and Random Fields*, edited by A. P. Young (World Scientific, 1997), pp. 251–275.

[6] A. K. Hartmann and H. Rieger, *Optimization Algorithms in Physics* (Wiley-VCH, Weinheim, 2001).

[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, 1997), pp. 161–223.

[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, 2011), pp. 370–406.

[9] M. Rao and A. Chakrabarti, Phys. Rev. Lett. **71**, 3501 (1993).

[10] E. Oguz, J. Phys. A **27**, 2985 (1994).

[11] C. Aron, C. Chamon, L. F. Cugliandolo, and M. Picco, J. Stat. Mech. **2008**, P05016 (2008).

[12] F. Corberi, E. Lippiello, A. Mukherjee, S. Puri, and M. Zannetti, Phys. Rev. E **85**, 021141 (2012).

[13] S. Sinha and P. K. Mandal, Phys. Rev. E **87**, 022121 (2013).

[14] P. K. Mandal and S. Sinha, Phys. Rev. E **89**, 042144 (2014).

[15] M. Manssen and A. K. Hartmann, Phys. Rev. B **91**, 174433 (2015).

[16] *OpenMP Application Programming Interface, Version 4.5*, OpenMP Architecture Review Board (2015), URL http://www.openmp.org/mp-documents/openmp-4.5.pdf.

[17] *MPI: A Message-Passing Interface Standard, Version 3.1*, Message Passing Interface Forum (2015), URL http://www.mpi-forum.org/docs/mpi-3.1/mpi31-report.pdf.

[18] B. Ahrens, J. Xiao, A. K. Hartmann, and H. G. Katzgraber, Phys. Rev. B **88**, 174408 (2013).

[19] K. Binder and E. Luijten, Phys. Rep. **344**, 179 (2001), ISSN 0370-1573.

[20] A. K. Hartmann and U. Nowak, Eur. Phys. J. B **7**, 105 (1999).

[21] N. G. Fytas and A. Malakis, Eur. Phys. J. B **61**, 111 (2008), ISSN 1434-6028.

[22] F. Belletti, M. Cotallo, A. Cruz, L. A. Fernandez, A. Gordillo-Guerrero, M. Guidetti, A. Maiorano, F. Mantovani, E. Marinari, V. Martin-Mayor, et al., Phys. Rev. Lett. **101**, 157201 (2008).

[23] H. Rieger, Physica A **224**, 267 (1996).

[24] J. Kisker, L. Santen, M. Schreckenberg, and H. Rieger, Phys. Rev. B **53**, 6418 (1996).

[25] T. Komori, H. Yoshino, and H. Takayama, J. Phys. Soc. Jpn. **69**, 1192 (2000).

[26] L. Berthier and J.-P. Bouchaud, Phys. Rev. B **66**, 054404 (2002).

[27] L. D. C. Jaubert, C. Chamon, L. F. Cugliandolo, and M. Picco, J. Stat. Mech. **2007**, P05001 (2007).

[28] D. A. Huse and D. S. Fisher, Phys. Rev. B **35**, 6841 (1987).

[29] L. Golubović and S. Feng, Phys. Rev. B **43**, 972 (1991).

[30] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, 1999).

[31] *Intel Xeon Phi Coprocessor System Software Developers Guide*, Intel Corporation (2014), URL https://software.intel.com/sites/default/files/managed/

[32] *Intel Xeon Phi Coprocessor Vector Microarchitecture*, Intel Corporation (2012), URL https://software.intel.com/sites/default/files/article/

[33] M. Matsumoto and T. Nishimura, ACM Trans. Model. Comput. Simul. **8**, 3 (1998), ISSN 1049-3301.

[34] J. K. Salmon, M. A. Moraes, R. O. Dror, and D. E. Shaw, in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis* (ACM, New York, NY, USA, 2011), SC '11, pp. 16:1–16:12, ISBN 978-1-4503-0771-0.

[35] A. K. Hartmann, *Big Practical Guide to Computer Simulations* (World Scientific, Singapore, 2015).

[36] D. S. Fisher, P. Le Doussal, and C. Monthus, Phys. Rev. E **64**, 066107 (2001).

[37] F. Corberi, A. d. Candia, E. Lippiello, and M. Zannetti, Phys. Rev. E **65**, 046114 (2002).

[38] H. Park and M. Pleimling, Phys. Rev. B **82**, 144406 (2010).

[39] H. Park and M. Pleimling, Eur. Phys. J. B **85**, 300 (2012).

[40] F. Corberi, E. Lippiello, A. Mukherjee, S. Puri, and M. Zannetti, J. Stat. Mech. **2011**, P03016 (2011).

[41] F. Corberi, E. Lippiello, A. Mukherjee, S. Puri, and M. Zannetti, Phys. Rev. E **88**, 042129 (2013).

[42] E. T. Seppälä, A. M. Pulkkinen, and M. J. Alava, Phys. Rev. B

66, 144403 (2002).

[43] A. K. Hartmann, Phys. Rev. B **65**, 174427 (2002).

[44] B. Ahrens and A. K. Hartmann, Phys. Rev. B **83**, 014205 (2011).

[45] B. Ahrens and A. K. Hartmann, Phys. Rev. B **84**, 144202 (2011).

[46] M. Zumsande, M. J. Alava, and A. K. Hartmann, J. Stat. Mech. **2008**, P02012 (2008).