Name: Caleb Cook College: Harvard College Caleb Cook

Simulating Aging Behavior of the Random-Site Ising Model on GPUs

One goal of modern condensed matter physics is to understand the effect of disorder on the ordering properties of magnetic materials. Although nonequilibrium domain growth in nondisordered systems has been studied extensively [1], progress in uncovering the aging properties of their disordered counterparts has been stymied by the slow dynamics characteristic of such systems. Due to its inclusion of both frustration and disorder, the prominent spin glass model in particular exhibits infamously slow dynamics and thus presents real computational challenges in simulating out-of-equilibrium aging on a reasonable time scale.

One approach is to instead study simpler models containing disorder but not frustration, an example of which is the random-site Ising model (RSIM). Moreover, one can exploit the inherent parallelizability of lattice spin models through the use of graphics processing units (GPUs), which are hardware accelerators whose high efficiency stems from their ability to carry out massively parallel computations. This summer, I worked with Markus Manssen and Professor Alexander Hartmann at the University of Oldenburg to run GPU simulations of the RSIM using Nvidia's CUDA C platform and analyze the resulting dynamical correlations in space and time.

The RSIM is a side-length L, d-dimensional lattice (in this case: $d = 2, 3$) of spins $S_i = \pm 1$ with Hamiltonian given by

$$\mathcal{H} = -\sum_{\langle i,j \rangle} (\rho_i S_i) \cdot (\rho_j S_j)$$

where the sum is taken over nearest-neighbor spin pairs. The ρ_i here represent site occupancy and are random variables taken from the distribution $P(\rho) = p\delta_{\rho,1} + (1-p)\delta_{\rho,0}$. We ran simulations from an initially disordered state at infinite temperature and quenched to a second temperature $T < T_c(p)$, where $T_c(p)$ is the critical temperature for some site-dilution p. Subsequent time evolution of the system is determined by our modified Metropolis algorithm, designed to utilize the parallel capabilities of the GPU.

In the RSIM, a given spin interacts exclusively with neighboring spins, each of which evolve independently from one another. We therefore split the lattice in a black/white checkerboard fashion, and then evolve (via the standard Metropolis algorithm) each independent half in parallel. Moreover, since the state of a given spin is considered four times per time-step, we require each parallely-executing spin thread to preload neighboring spins into *shared memory*, a specialized memory cache that allows the GPU to minimize time fetching spin states.

We achieve further optimization through use of the novel parallel random number generator Philox [2], which has been shown to be both extremely fast and of high statistical quality (Crush-resistant), in obtaining the random numbers required to simulate probabilistic spin flipping. We also represent each spin as a binary bit (*multispin coding*), allowing us to quickly calculate spin interactions with the bitwise XOR operator as well as run 128 simultaneous simulations (2 checkerboards of 64-bit integers) for statistical averaging purposes. The net result of these and further optimizations is that our algorithm flips Ising spins at a rate of ~ 8 and ~ 14 picoseconds/flip in the 2D and 3D RSIM, respectively.

When studying nonequilibrium dynamics, two quantities of special importance are the *two-time correlation function* C and the *spatial correlation function* S, defined to be

$$C(t_w, t_w + t) = \frac{1}{L^d} \sum_i \overline{\langle S_i(t_w) \cdot S_i(t_w + t) \rangle} \qquad S(r, t) = \frac{1}{L^d} \sum_{\vec{x}} \left\{ \overline{\langle S_{\vec{x}}(t) \cdot S_{\vec{x}+\vec{r}}(t) \rangle} \right\}$$

where $\overline{\langle \cdots \rangle}$ indicates an average first over thermal noise and then the site disorder, and $\{\cdots\}$ indicates an

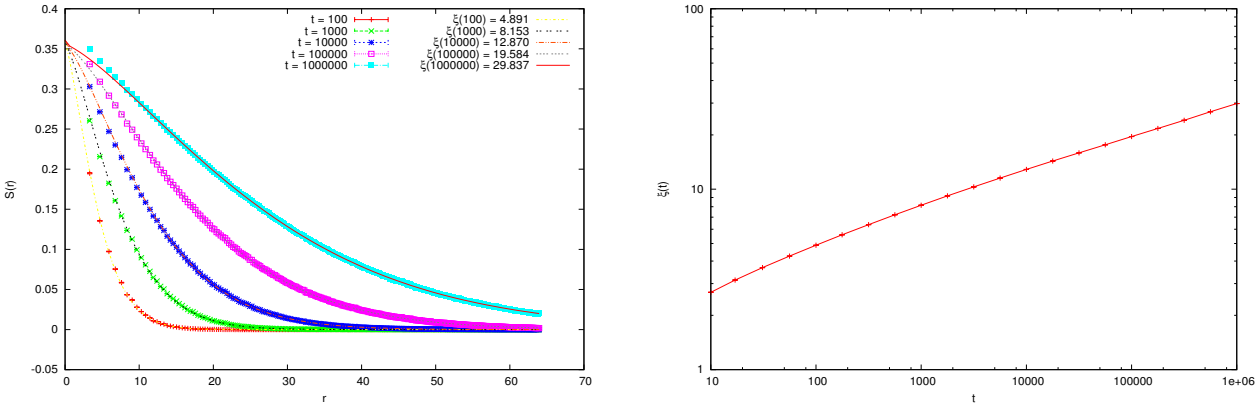

Figure 1: (Left) Fitting of $S_{\text{fit}}(r,t)$ to $S(r,t)$ for $r > 5$ at logarithmically-spaced times t. Not all times/fits shown. (Right) Growth of the fit-extracted dynamical correlation length $\xi(t)$ over time. ($d = 3$, $L = 256$, $T = 0.8$, $p = 0.75$).

average over spins a distance $|\vec{r}|$ away. From $S(r,t)$ one can extract the *dynamical correlation length* $\xi(t)$, which captures the growing length scale over which spins are highly correlated. We follow [3] in extracting $\xi(t)$ by fitting to the functional form

$$S_{\text{fit}}(r,t) = Ar^{-\alpha} \exp\left[-\left(\frac{r}{\xi(t)}\right)^{\beta}\right]$$

in which A, α, and β are time-independent parameters. Sample results obtained by fitting $S_{\text{fit}}(r,t)$ to our preliminary simulation data and extracting the corresponding growth behavior of $\xi(t)$ are shown above.

While the growth of the correlation length in the coarsening, nondisordered Ising model is known to follow $\xi(t) \sim t^{1/z}$, theoretical arguments made by Huse and Henly [4] predict an eventual slowing down in the presence of disorder, leading to an asymptotic regime in which $\xi(t) \sim (\log t)^{1/\psi}$, where $\psi_{2\text{D}} = 1/4$ and $\psi_{3\text{D}} \approx 0.55$. However, previous attempts at verifying the disordered growth law of Huse and Henly have been unsuccessful [5], often citing the long simulation times required even in the absence of frustration to reach the asymptotic regime.

Our preliminary simulation data exhibit deviations from the standard power law growth of $\xi(t)$, but attempts at fitting this data to the power of a logarithm have been inconclusive. However, we have only recently concluded the development phase of our algorithm and intend to start much longer simulations in the near future. With the speed benefits of the GPU, we are optimistic that we will be able to reach the asymptotic regime, where we hope to better characterize the slow growth of $\xi(t)$ and extend previous results [5] concerning the scaling behavior of the two-time correlation function in disordered systems.

References

[1] A. J. Bray. Theory of phase ordering kinetics. Adv. Phys. **43**, 357 (1994).

[2] J. K. Salmon, M. A. Moraes, R. O. Dror, and D. E. Shaw. Parallel random numbers: As easy as 1, 2, 3 in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, NY, USA (2011).

[3] E. Marinari, G. Parisi, and J. J. Ruiz-Lorenzo. Numerical Simulations of Spin Glass Systems in *Spin Glasses and Random Fields*, ed. A. P. Young, World Scientific (1997).

[4] D. A. Huse and C. L. Henley. Pinning and roughening of domain walls in Ising systems due to random impurities. Phys. Rev. Lett. **54**, 25 (1985).

[5] H. Park and M. Pleimling. Aging in coarsening diluted ferromagnets. Phys. Rev. B. **82**, 14 (2010).