Monte Carlo Simulations of the Edwards-Anderson Model using Graphics Processing Units

Ka-Ming Tam

Louisiana State University
Department of Physics and Astronomy
Center of Computation and Technology

Source code available for download
http://www.institute.loni.org/lasigma/package/ising/

Our Team

- Implementing algorithms for Graphics Processing Unit (GPU)

Ye Fang, grad. student
(LSU, ECE and CCT)

Zhifeng Yun
(LSU, CCT)

Sheng Feng, grad. student
(LSU, Phys. and Astronomy)

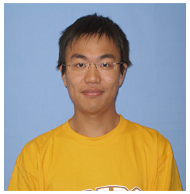

Bhupender Thakur
(LSU, HPC)

Ariana Papke, undergrad. student
(Gottingen, Maths)

J. Ramanujam
(LSU, ECE and CCT)

Karen Tomko
(Ohio Supercomputer center)

Juana Moreno
(LSU, Phys. and Astronomy)

Mark Jarrell
(LSU, Phys. and Astronomy)

nvidia GPU architecture

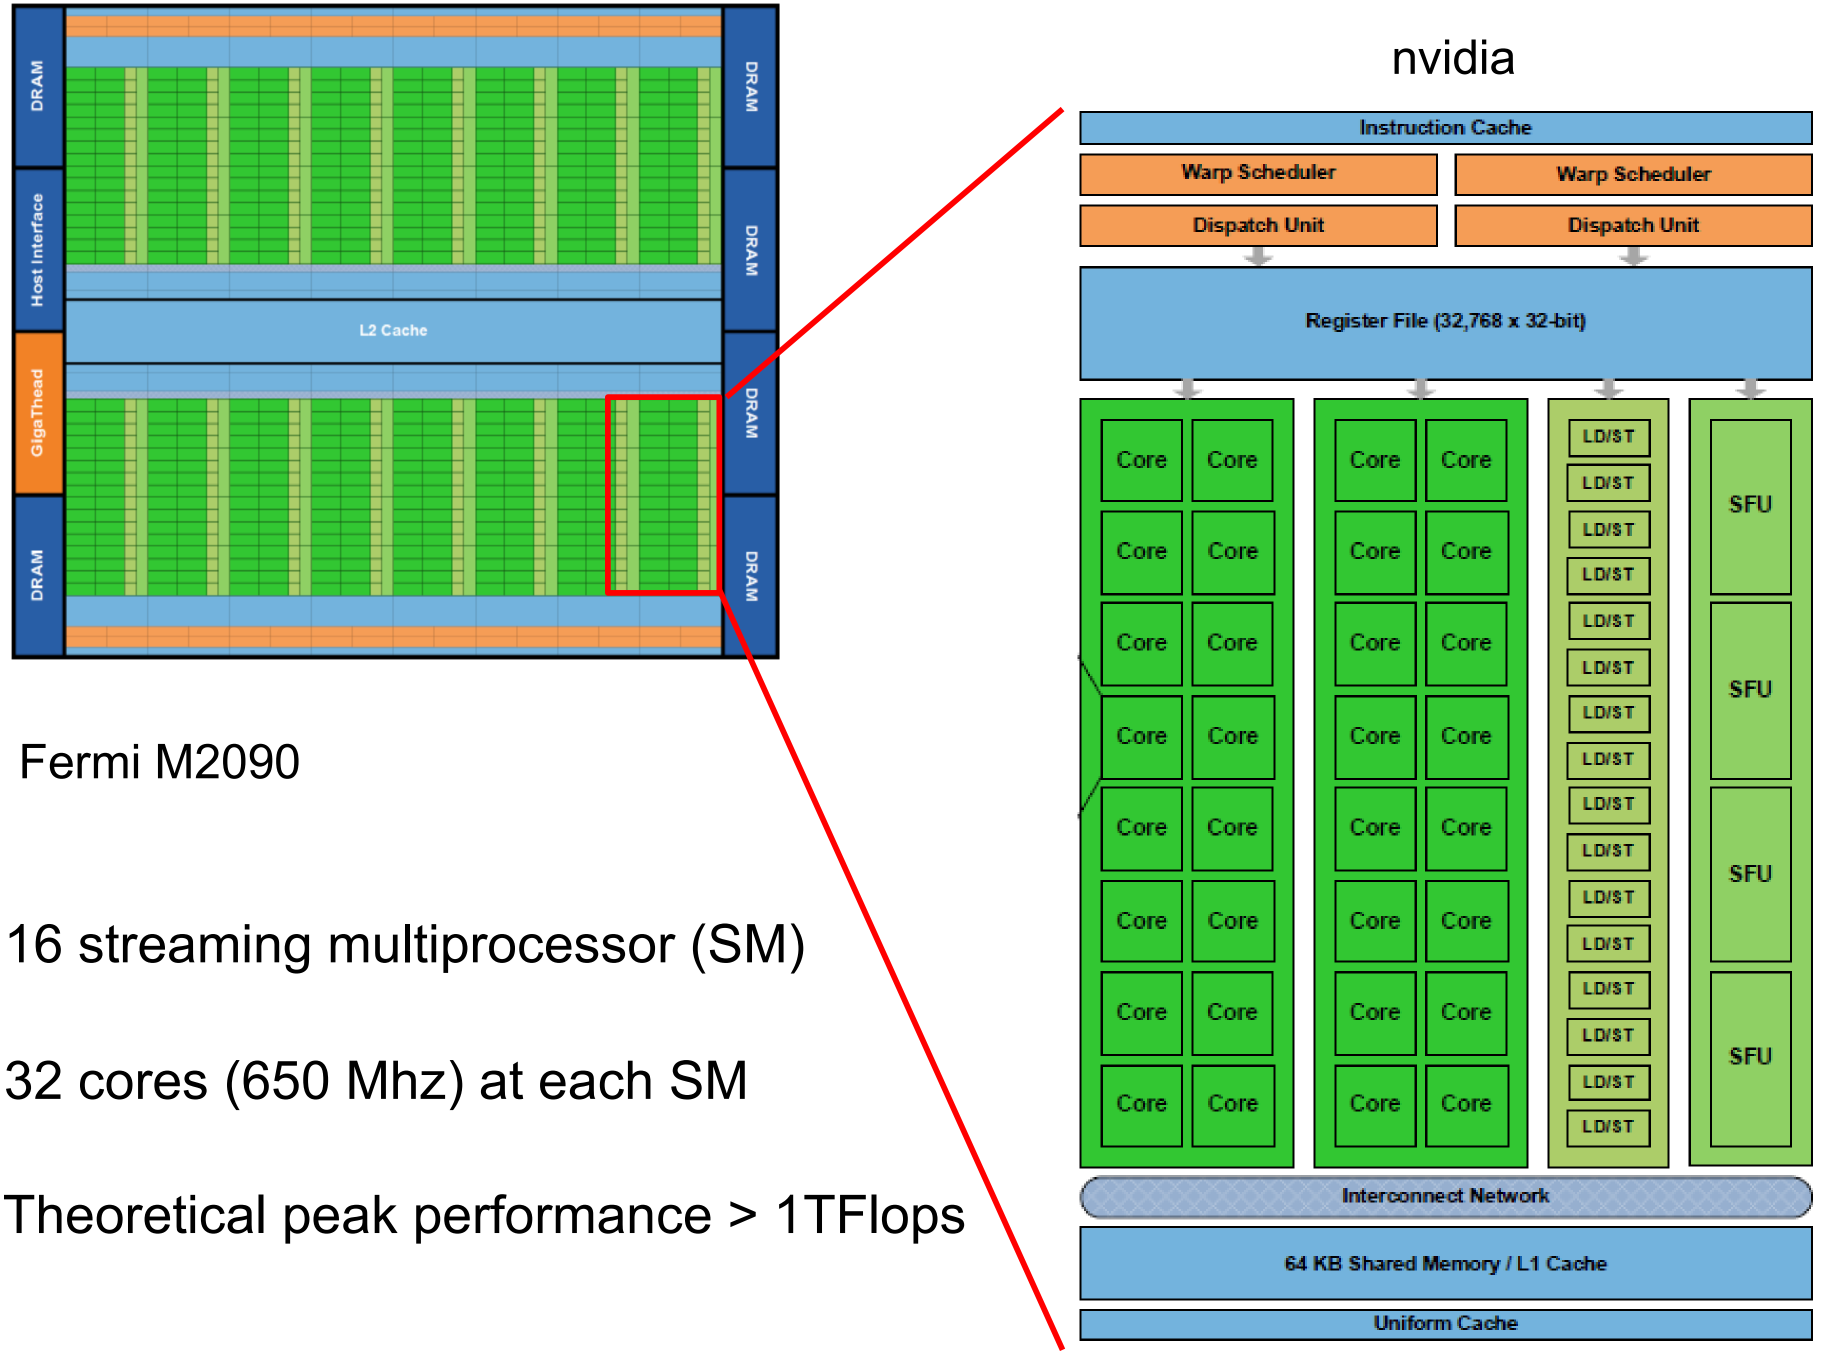

Fermi M2090

- 16 streaming multiprocessor (SM)
- 32 cores (650 Mhz) at each SM
- Theoretical peak performance > 1TFlops

GPU CUDA Memory Architecture

- Device memory:
  Large but slow (global memory)

- Shared memory:
  Fast but small (48KB)
  Cannot shared among thread blocks

- Registers
  Fast but small
  Cannot shared between threads

- Constant memory
  Read only, accessible for all threads

- Texture memory
  Read only,
  hardware filtering (interpolation)

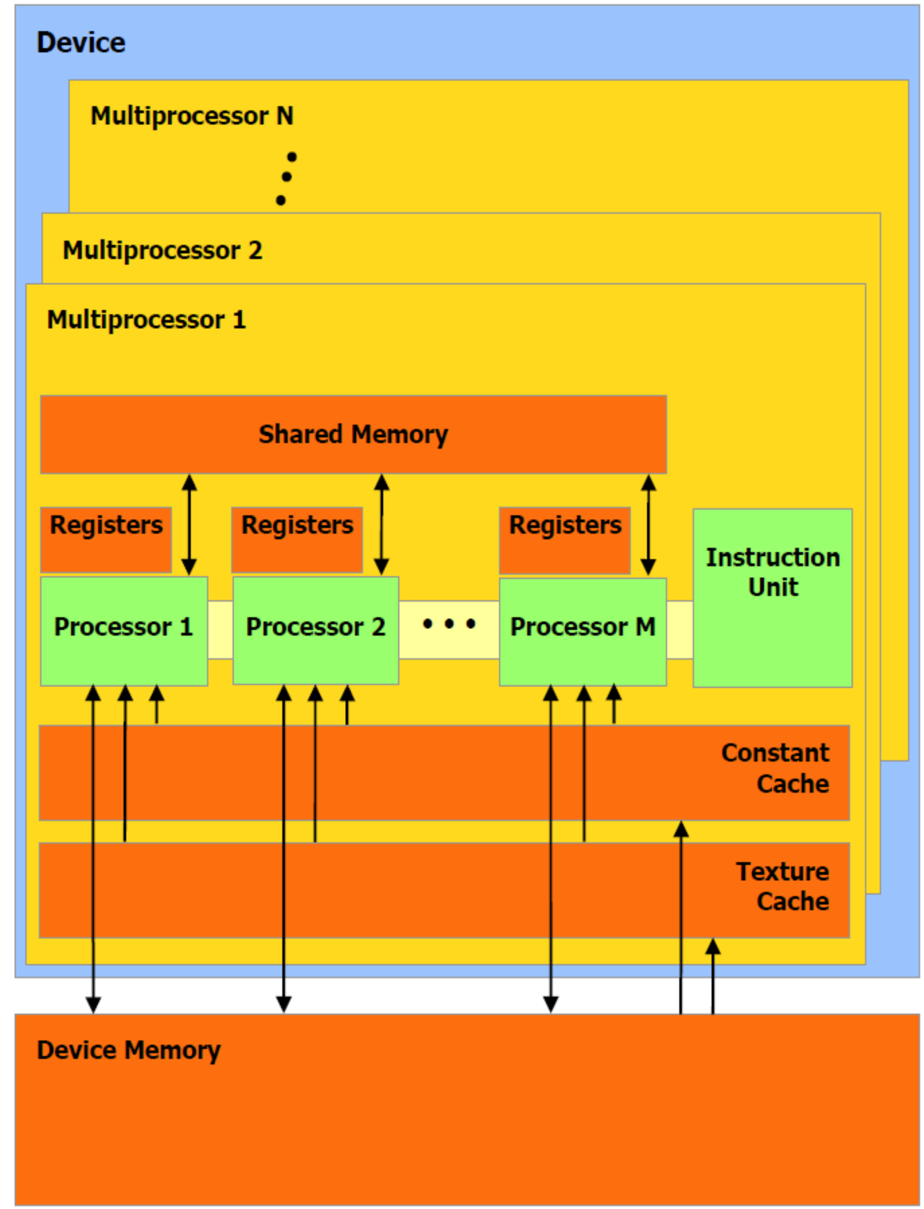

nvidia

nvidia CUDA programming paradigm

- A kernel is executed by a grid of thread blocks
- Each streaming multiprocessor can contains multiple threads
- 32 threads (1 wrap) are running in parallel
- Typical CUDA program

1. Allocate memory in the device
2. Set up the kernel
3. Launch the kernel
4. Transfer data from the device to the host

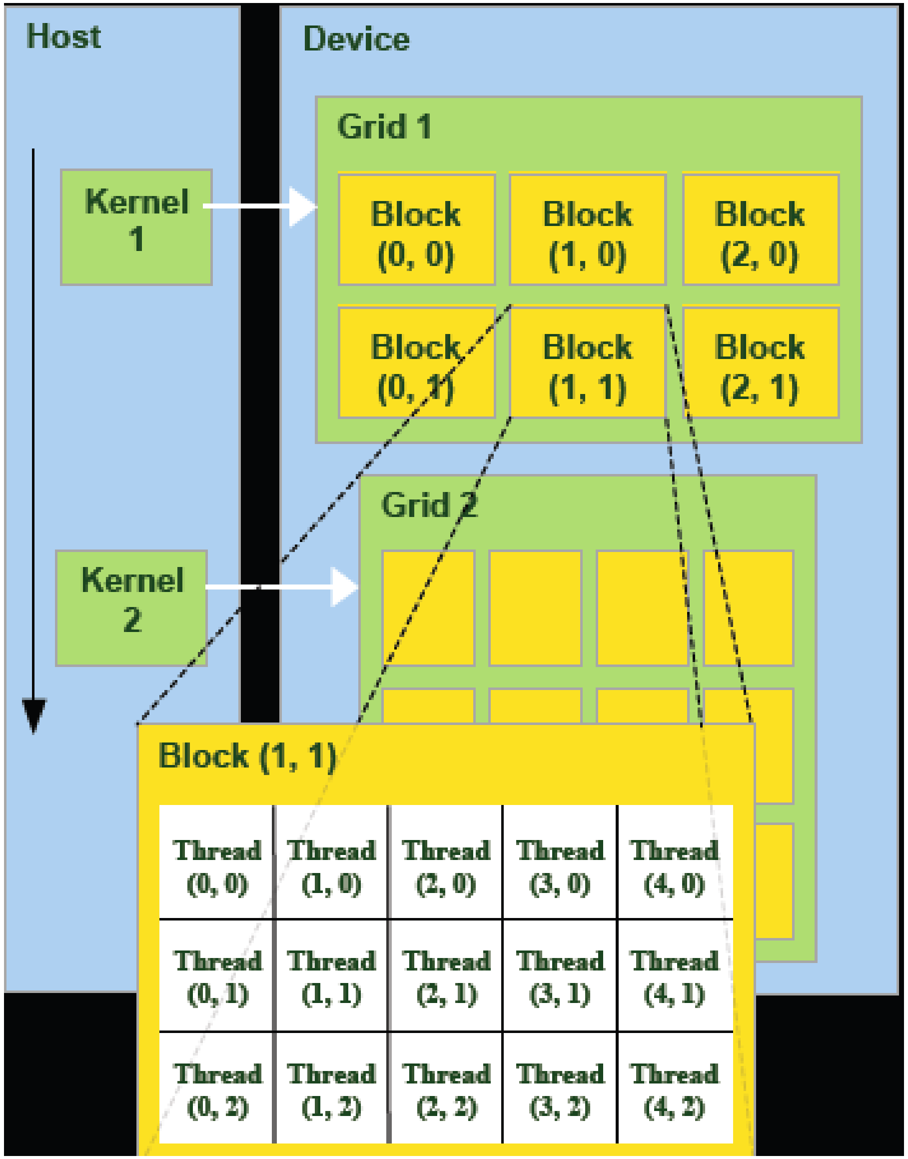

nvidia

Magnetic Transition

- Ferromagnetic Ising model $J_{i,j} = -1$.

$$H = \sum_{ij} J_{ij} s_i s_j - H \sum_i s_i$$

- Finite critical temperature transition from paramagnetic phase to ferromagnetic phase.

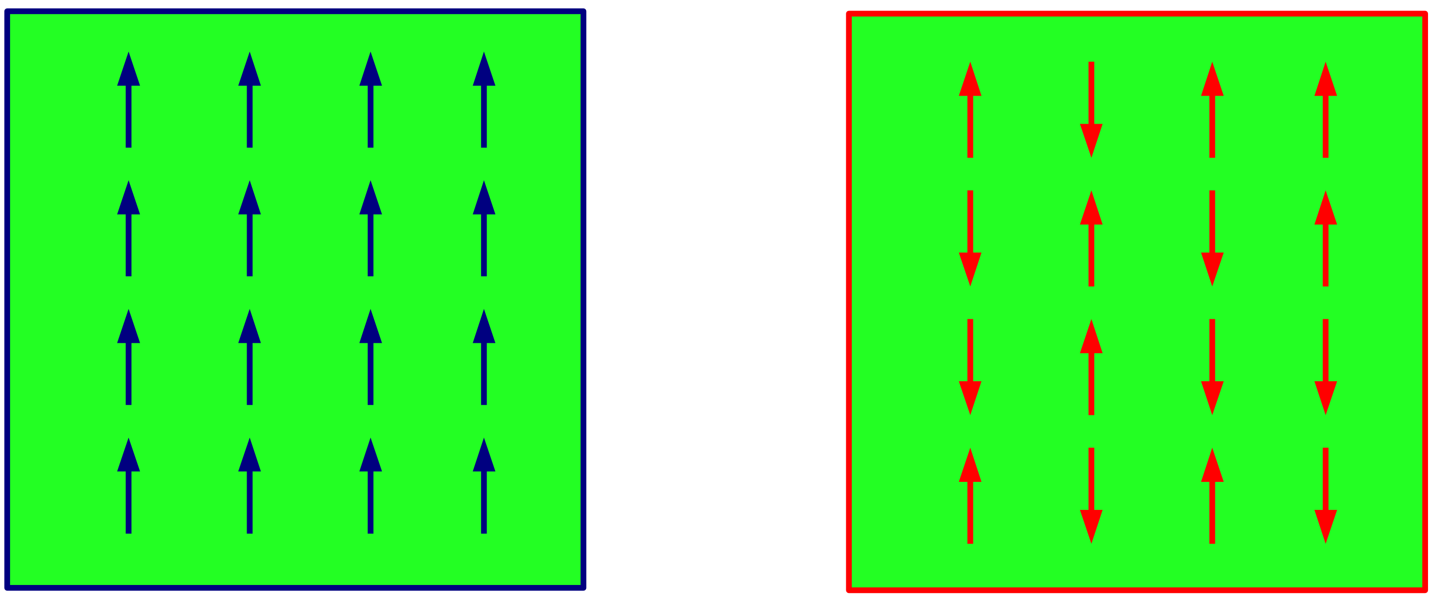

$T < T_c$ $\quad\quad$ $T > T_c$

Frustration

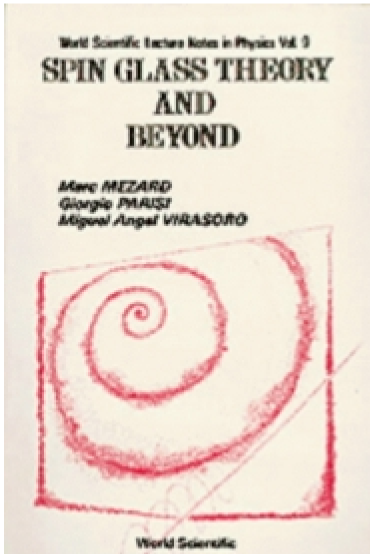

Mezard, Parisi, and Virasoro

Chapter 0

Often in life we find out that our goals are mutually incompatible: we have to renounce some of them and we feel frustrated. For example, I may want to be a friend of both Mr. White and Mr. Smith. Unfortunately, they hate each other: it is then rather difficult to be a good friend of both of them (a very frustrating situation).

The situation is more complex when many individuals are present. In a classical tragedy the scenario may be the following: there is a fight between two groups and the various characters on the scene have to choose sides. In addition they all have strong personal feelings, positive or negative, towards each other (it is a tragedy!) Some of them are friends and some are enemies.

Frustration

- Consider an Anti-Ferromagnetic Ising model $J_{i,j}=1$.

$$H=\sum_{ij} J_{ij} s_i s_j - H\sum_i s_i$$

- Cannot minimize the energy of every bond.

e.g. 3 spins

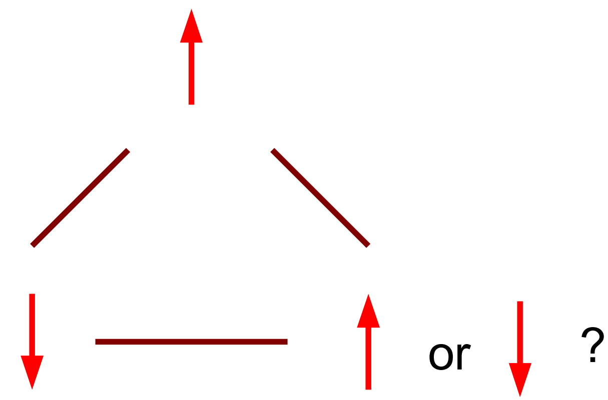

- No unique spin configuration which can minimize the energy.

- New concepts:
  - Classical “spin liquid”
  - Degenerate ground states
  - Finite zero-temperature entropy
  - May lead to “fractionalized” excitations, dual descriptions

Frustration from Randomness

- Same Ising model, but the coupling is random $J_{i,j}=1$ or $J_{i,j}=-1$

$$H=\sum_{ij} J_{ij} s_i s_j - H\sum_i s_i$$

- Example: a square parquet with one ferromagnetic bond.

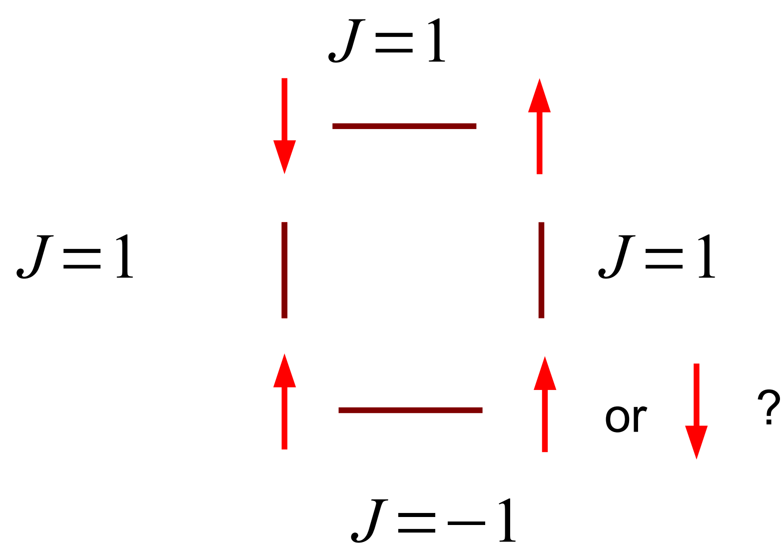

- Frustration due to randomness, instead of purely from geometry.

Other Examples of Geometrical Frustration

- 2D:Kagome lattice (AFM vector spin model)

$$H=\sum_{ij} J\vec{s}_i\cdot\vec{s}_j$$

$$=J\sum_{triangle}(\vec{s}_{i,1}+\vec{s}_{i,2}+\vec{s}_{i,3})^2+const.$$

Any configuration which satisfies the constraint that the total moment in each triangle is zero can be a valid ground state.

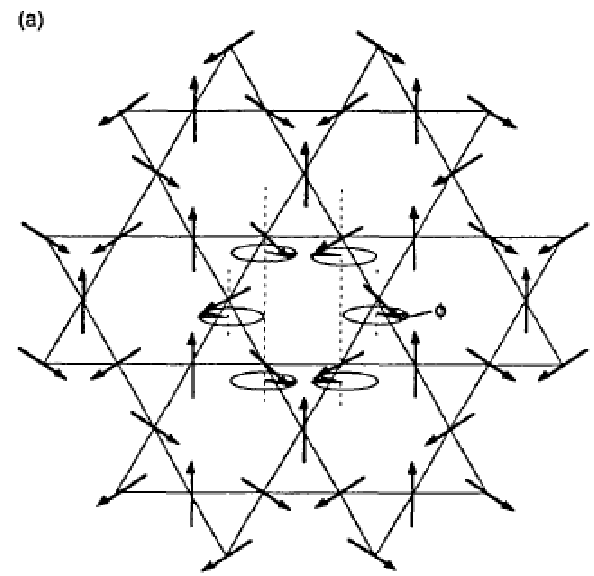

Non-abelian vortices from anisotropic spin exchange.

$$H=\sum_{ij} J(\vec{s}_i\cdot\vec{s}_j-\epsilon\, s_i^z s_j^z)$$

Topological spin glass due to the freezing of non-abelian vortices.

P. Chandra, P. Coleman, and I. Ritchey, PRB (1993)

Other Examples of Geometrical Frustration

- 3D:Pyrochlore lattice (AFM vector spin model)

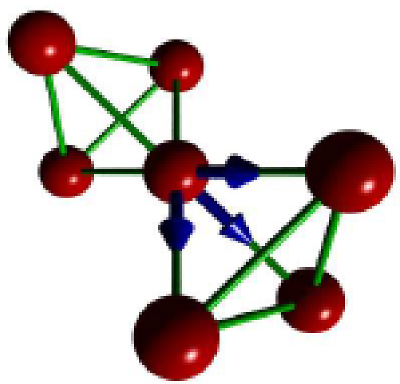

$$H=\sum_{ij} J\,\vec{s}_i\cdot\vec{s}_j$$
$$=J\sum_{tetrahedron}\left(\vec{s}_{i,1}+\vec{s}_{i,2}+\vec{s}_{i,3}+\vec{s}_{i,4}\right)^2+const.$$

- Spin Glass phase from
  Geometrical frustration + very small randomness

$H=\sum_{ij}(J+\delta_{ij})\,\vec{s}_i\cdot\vec{s}_j$ where the randomness $\delta_{ij}\ll J$

- Spin glass without explicitly competing ferromagnetic and antiferromagnetic coupling

Greedan, et al. Solid State Commun (1986)
Gingras, et al. PRL (1996)
Saunders and Chalker, PRL (2007)
KMT, et al. arxiv:1009.1272

Outstanding challenges of spin glass simulations

- Limitations of spin glass simulations.

1. Long equilibration time for equilibrium simulations.

2. Cluster algorithm usually won't help (there are exceptions).

3. A shortlist of the methods proposed.

   Simulated annealing (Kirkpatrick, Gelatt, Vecchi, ...)

   Multicanonical (Berg,...)

   Multi-variate Multicanonical (Hatano, Gubernatis,...)

   Parallel tempering (Geyer, Swendsen, Wang, Hukushima, Nemoto, Marinari, Parisi,...)

4. Parallel tempering seems to be most efficient for most glassy systems? Entropic barrier?

Parallel tempering

- Schematic representation of a rough "landscape".

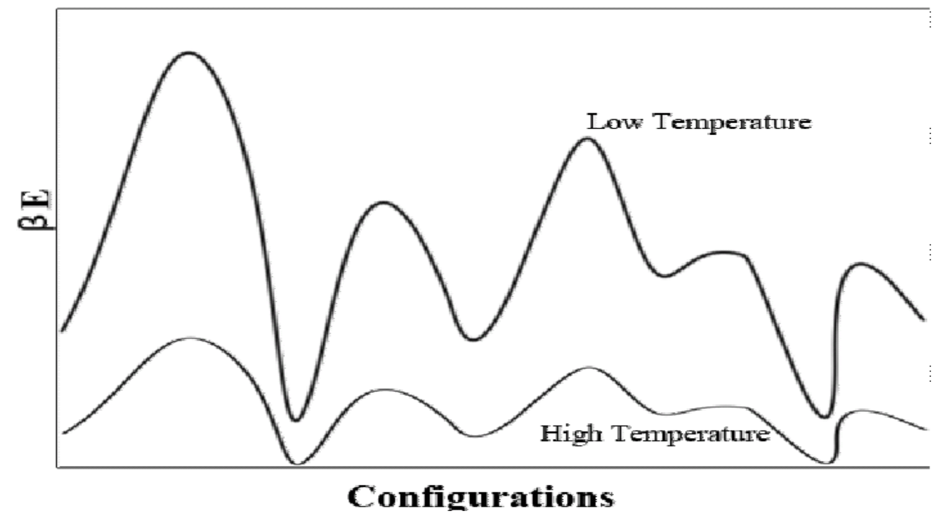

- Periodically swap samples at different temperatures.

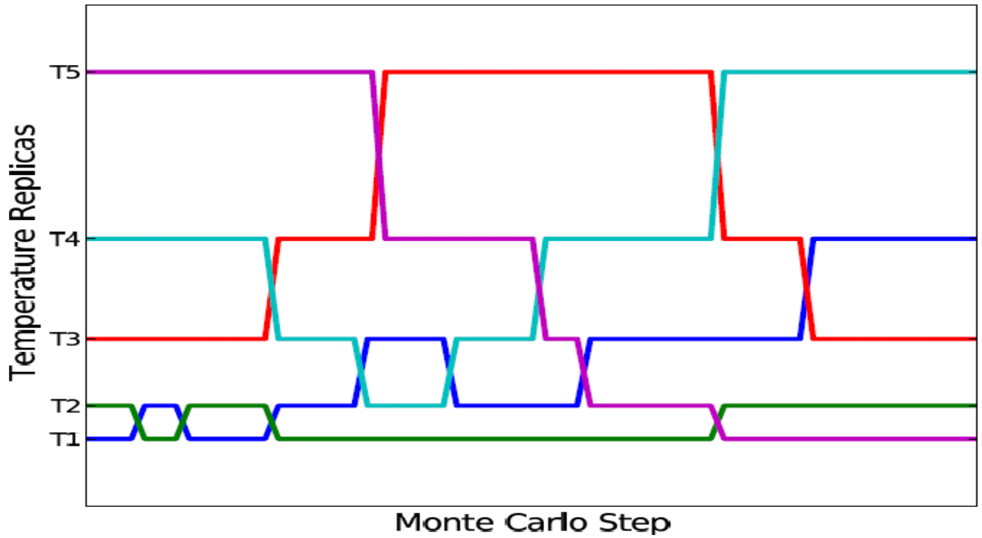

Computation aspects of Spin Glass Simulations

- The system sizes for equilibrium studies are relatively small, usually ~ 10000 sites. **(small memory)**

- Number of different disorder realizations is large, at least about >1000 and can be over 10^5 realizations. **(large number of Independent simulations)**

- Parallel tempering requires a set of temperatures for the simulations, usually around 20-30. **(again, large number of almost independent simulations)**

- For Ising spin glass, the operations on the spins are just bit-manipulation. **(simple instruction)**

Small memory allocation

Simple operations

GPU computing

Some Related Works

Discrete Spin

- FPGA (Janus collaborations)
- M. Weigel, J. Comput. Phys. (2012)
- T. Levy, G. Cohen, and E. Rabani, J. Chem Theory Comput. (2012)

Continuous Spin

- M. Bernaschi, G.Paris, L. Parisi, arXiv (2010)
- T. Yavors'kii and M. Weigel, Eur. Phys. J. Special Topics (2012)

Lattice structure: 3D Stencil

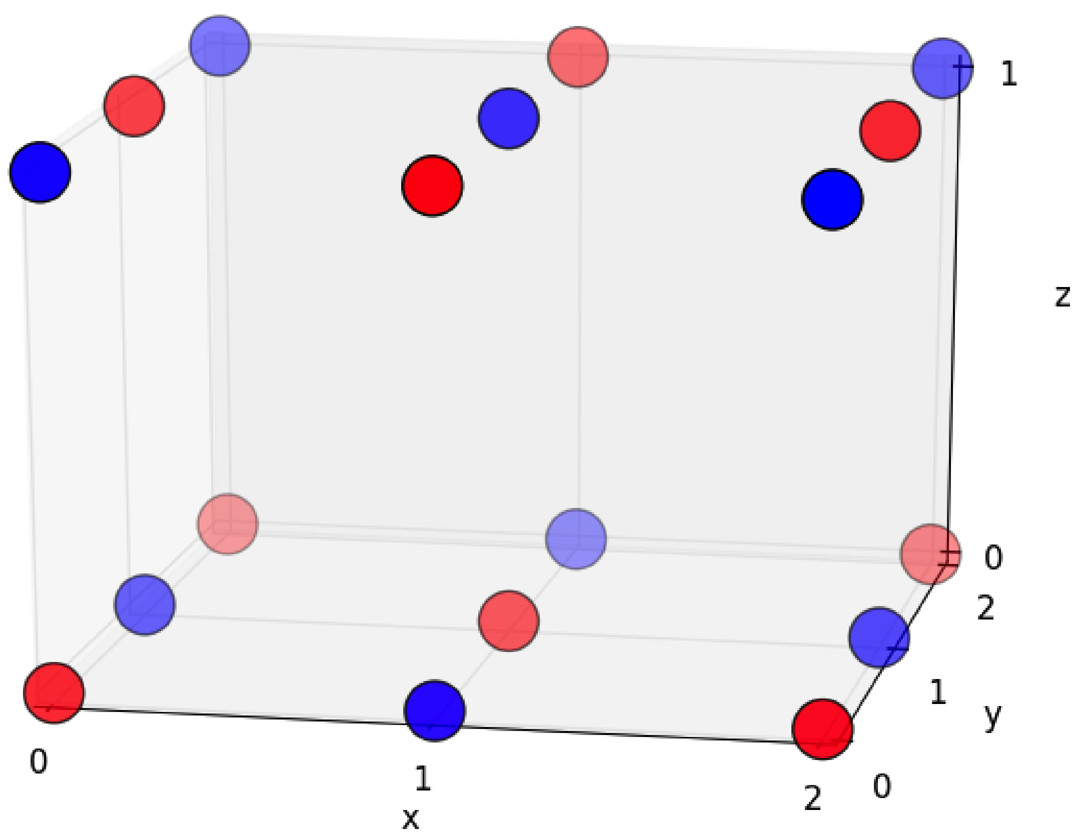

$$H = \sum_{ij} J_{ij} s_i s_j - h \sum_i s_i$$

- Nearest neighbors coupling in a cubic lattice

Lattice structure: 3D Stencil

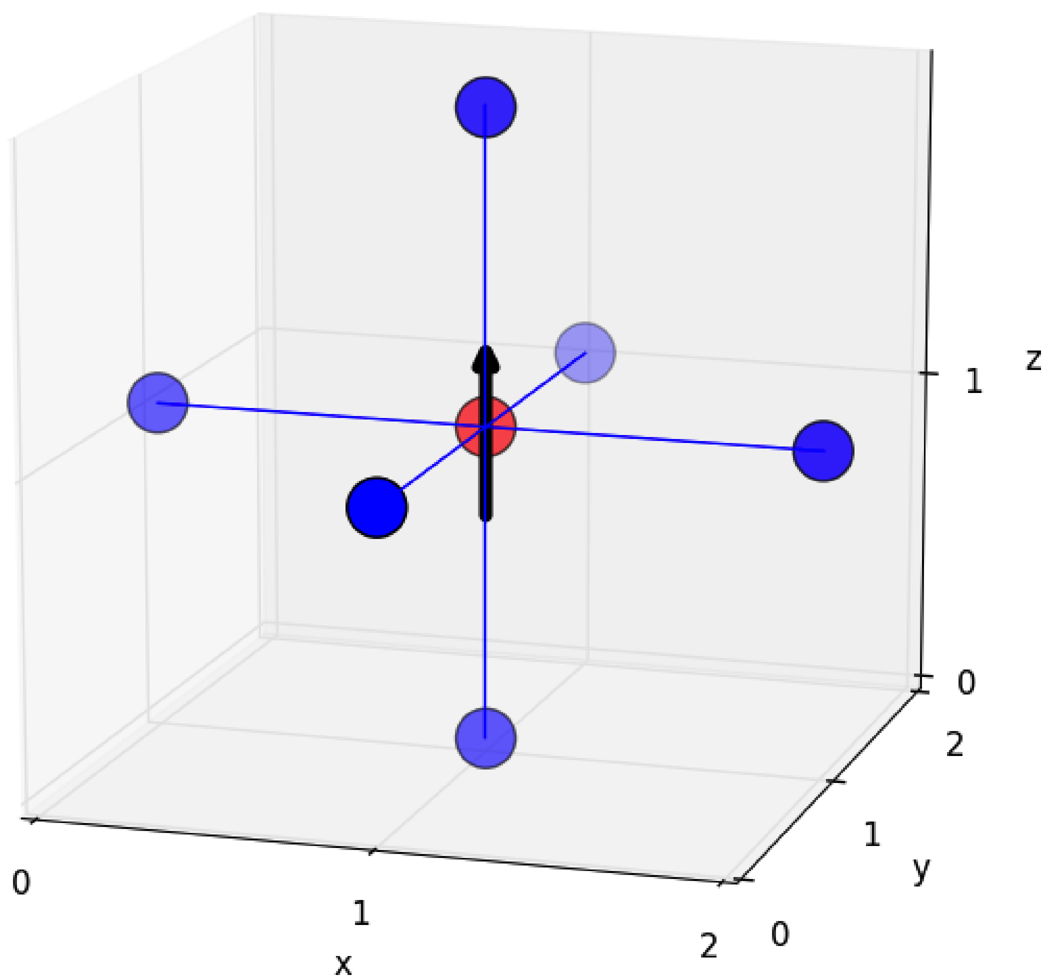

• Six nearest neighbors for each spin

Multi-spin coding

• Discrete spins: e.g. Ising model, Potts model, lattice gas,...

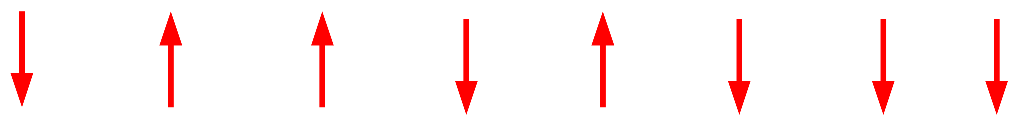

• General idea: **Store many spins at one word**

Rebbi, Creutz, … 80's

e.g. a 8-bit word

0 1 1 0 1 0 0 0

• Advantages: **'Simplify' the instruction as bit manipulation**

Important for GPU implementation

Saving memory allocation

Critical for GPU implementation

Multi-spin coding

- Calculating local energy
  (with bimodal random coupling)

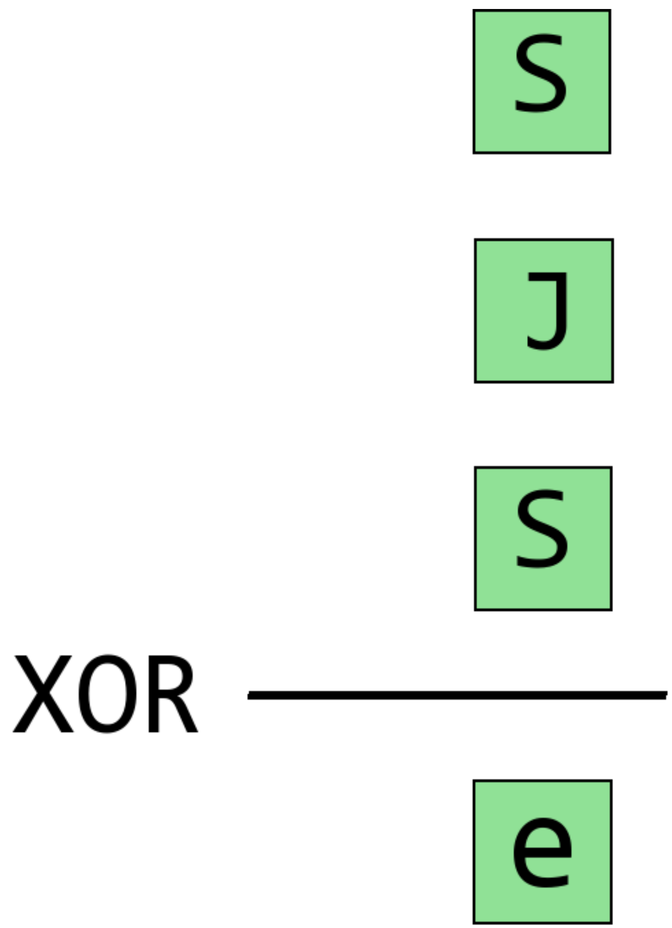

Multi-spin coding

- Accumulating the local energy for six nearest neighbors

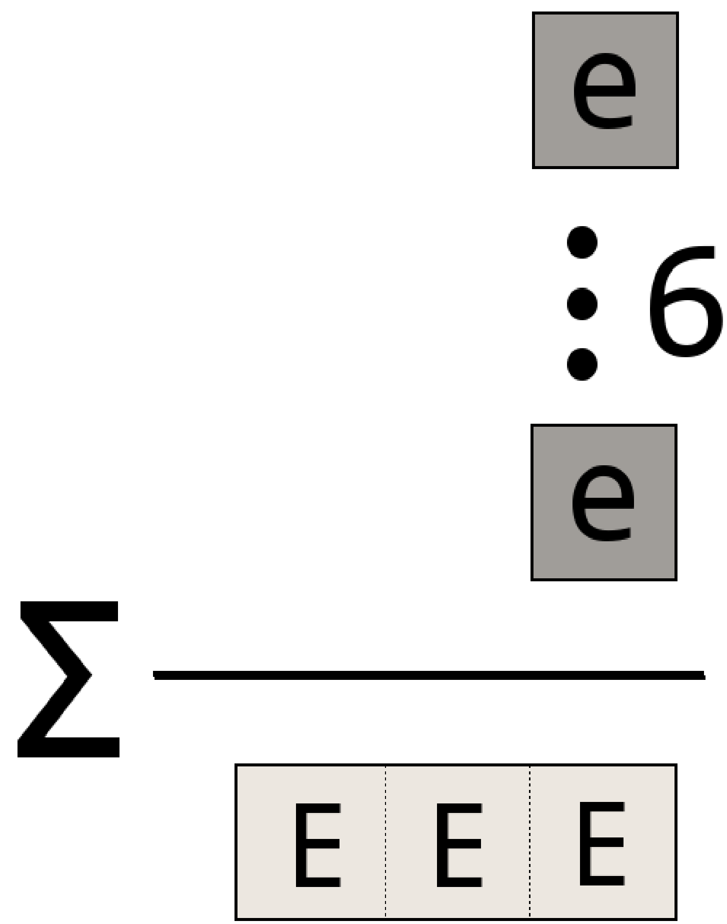

- Three bits are required

Multi-spin coding (single bit scheme)

- One bit per spin for updating local energy

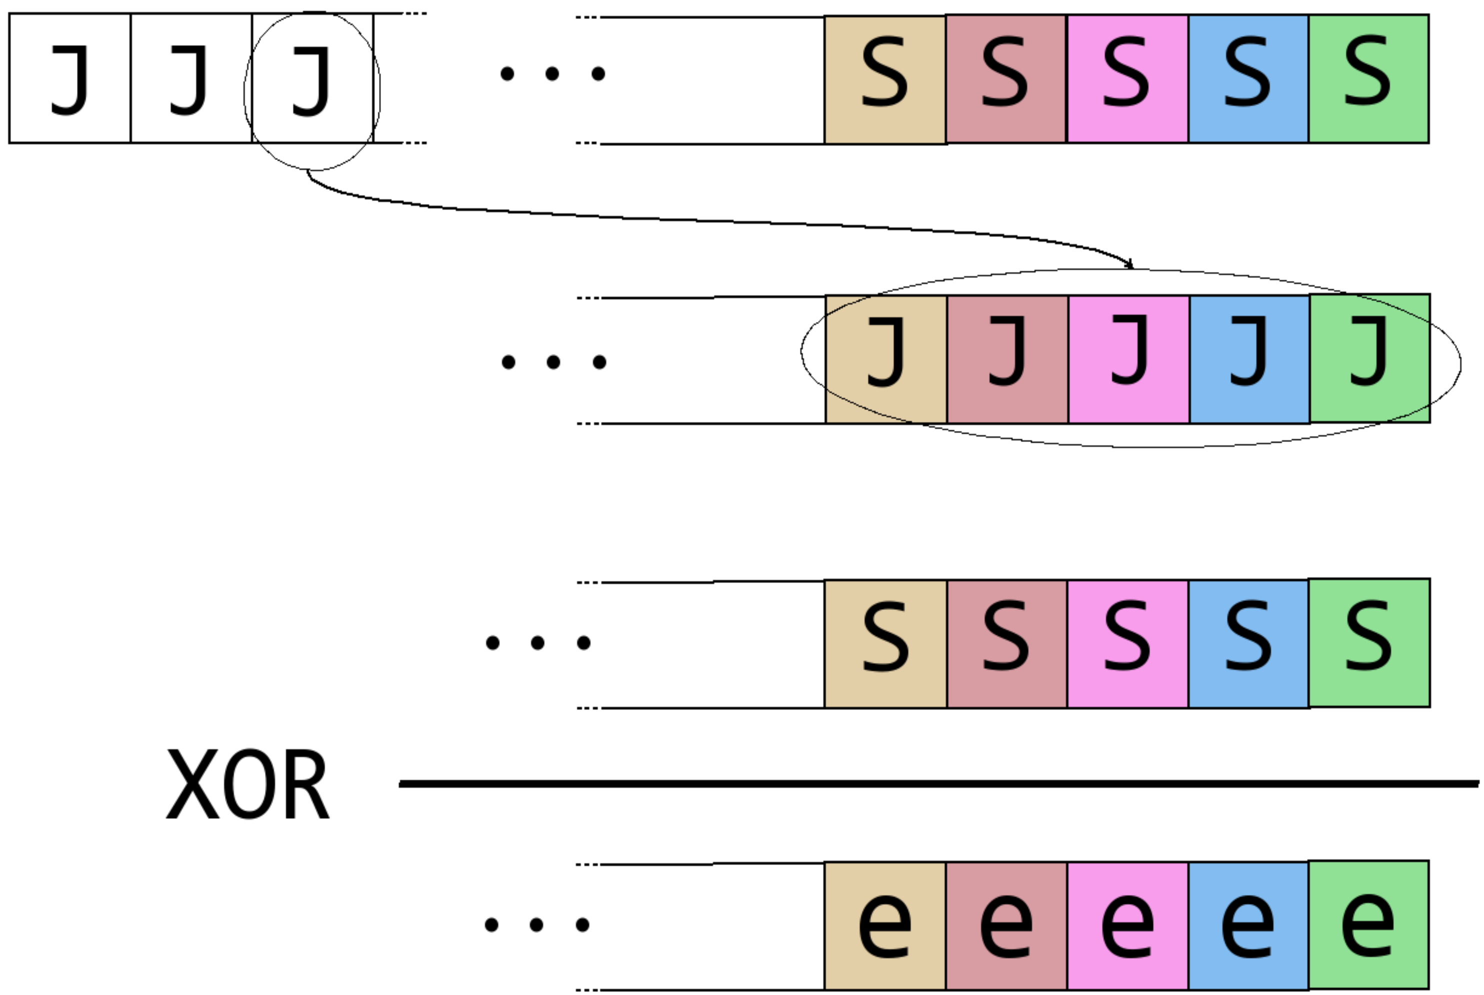

Multi-spin coding

- One bit per spin for calculating local energy

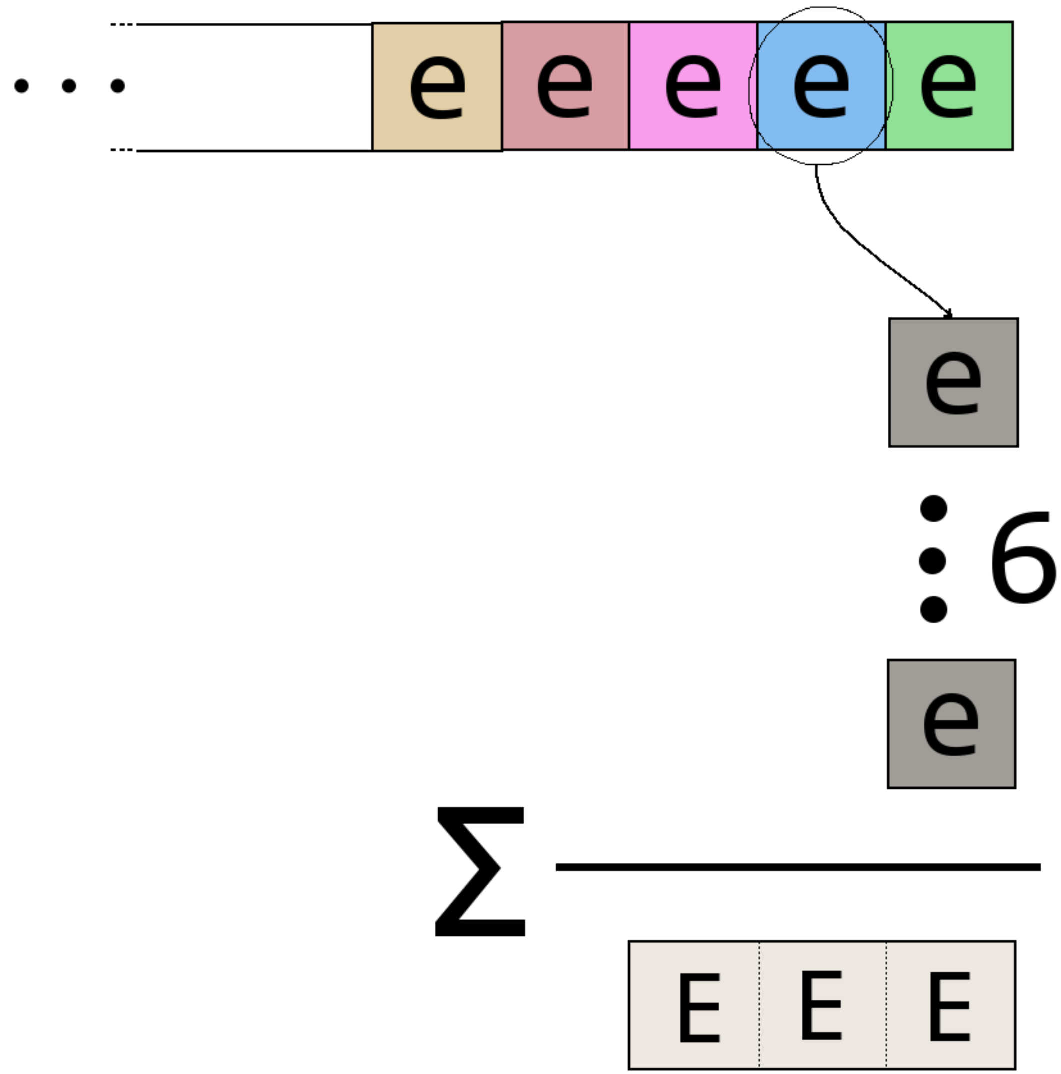

Multi-spin coding (three bits scheme)

- Three bits per spin for calculating energy

- Shifting the bit for the random coupling

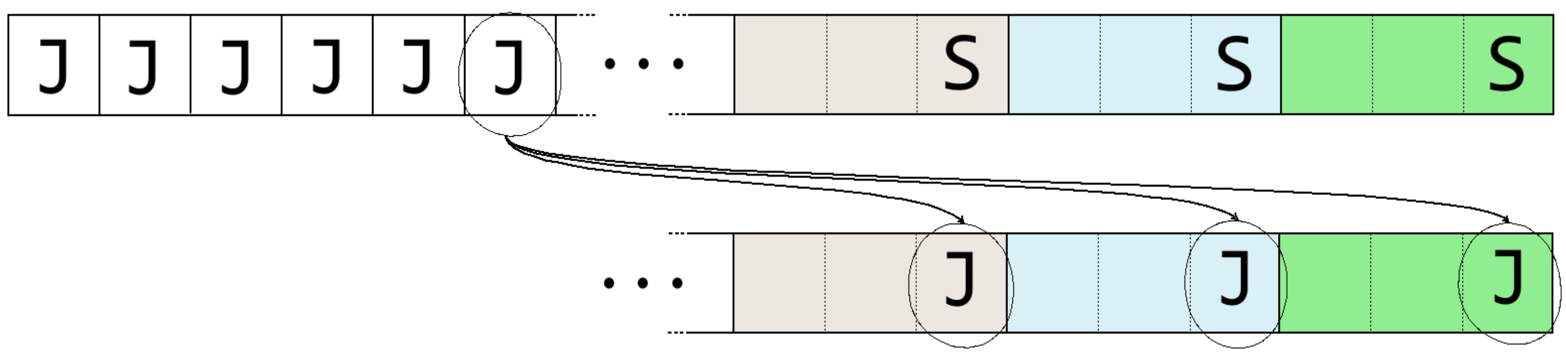

- Calculating the energy at each bond

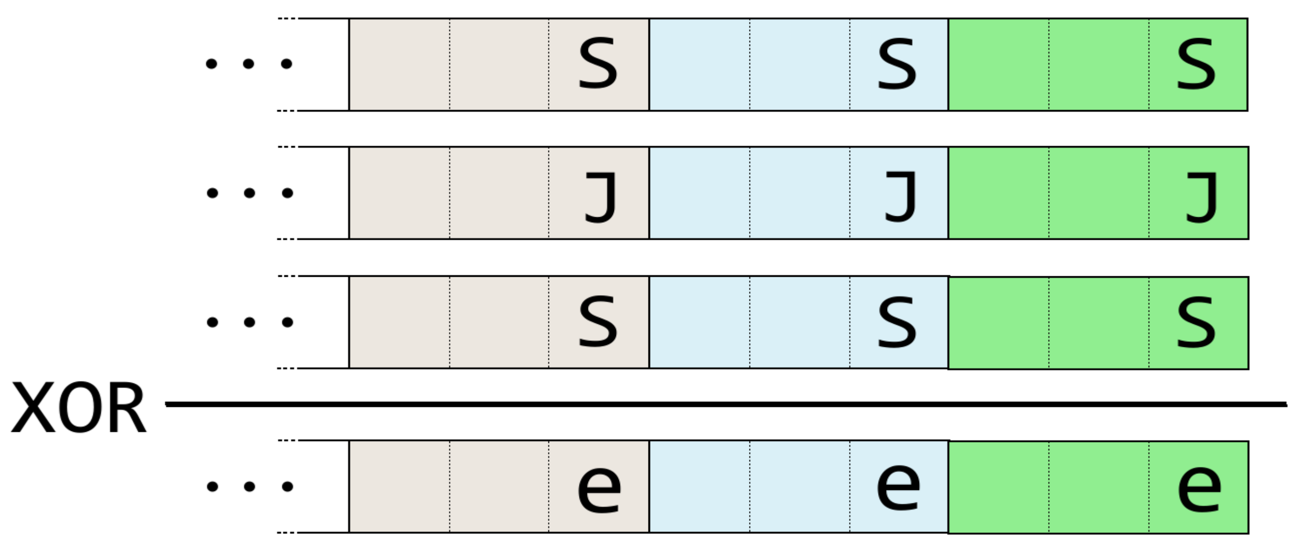

Multi-spin coding

- Sum over all six nearest neighbors.

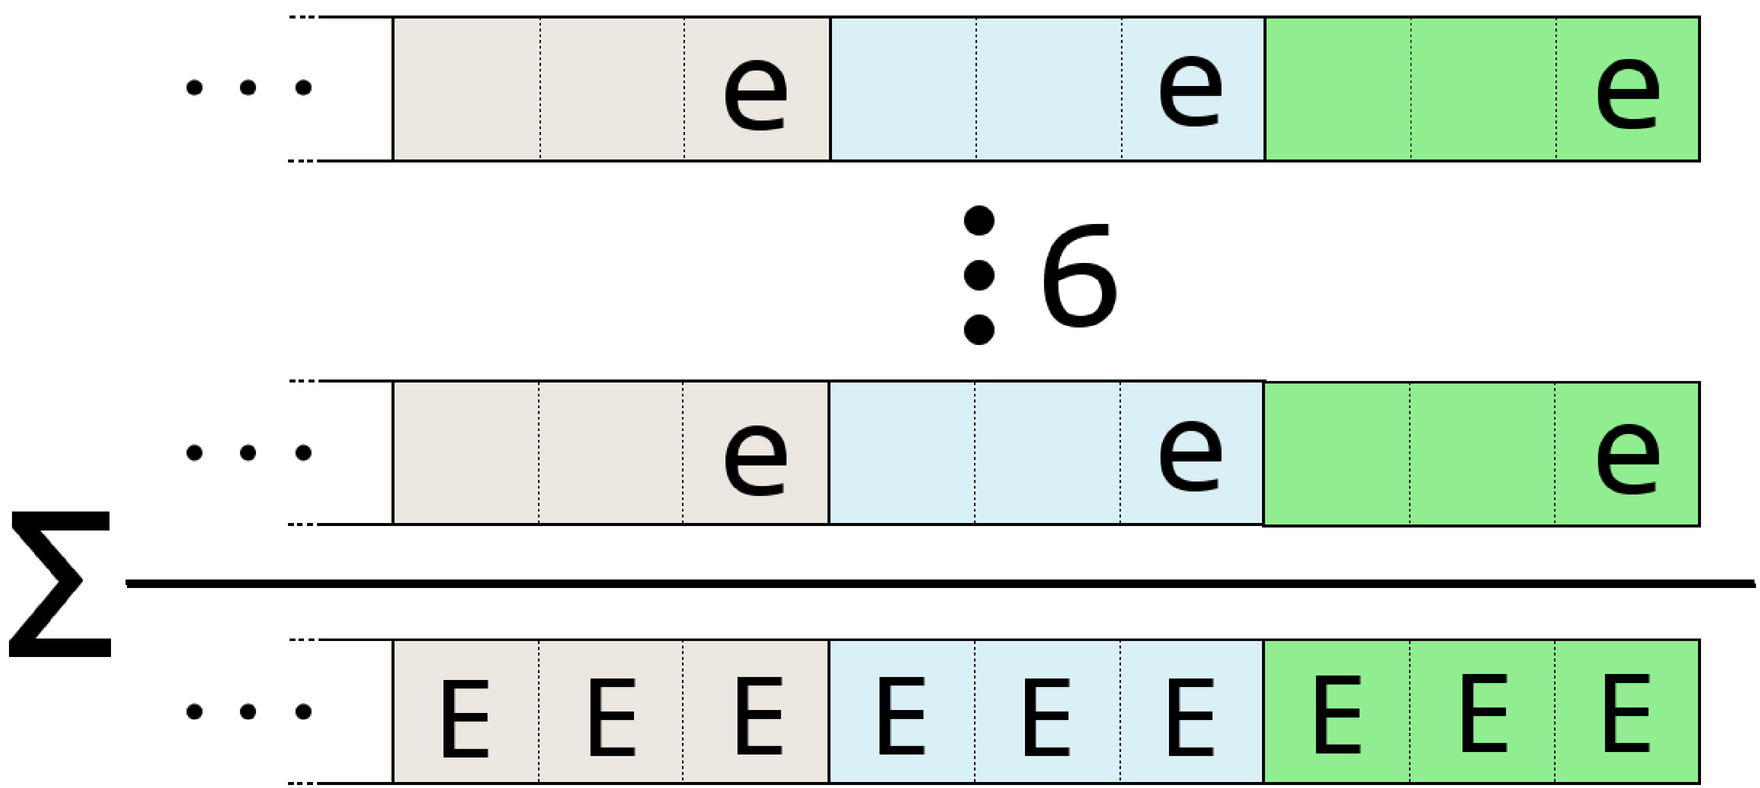

Multi-spin coding (4 bits scheme)

- Combined the energy bits with the spin bits

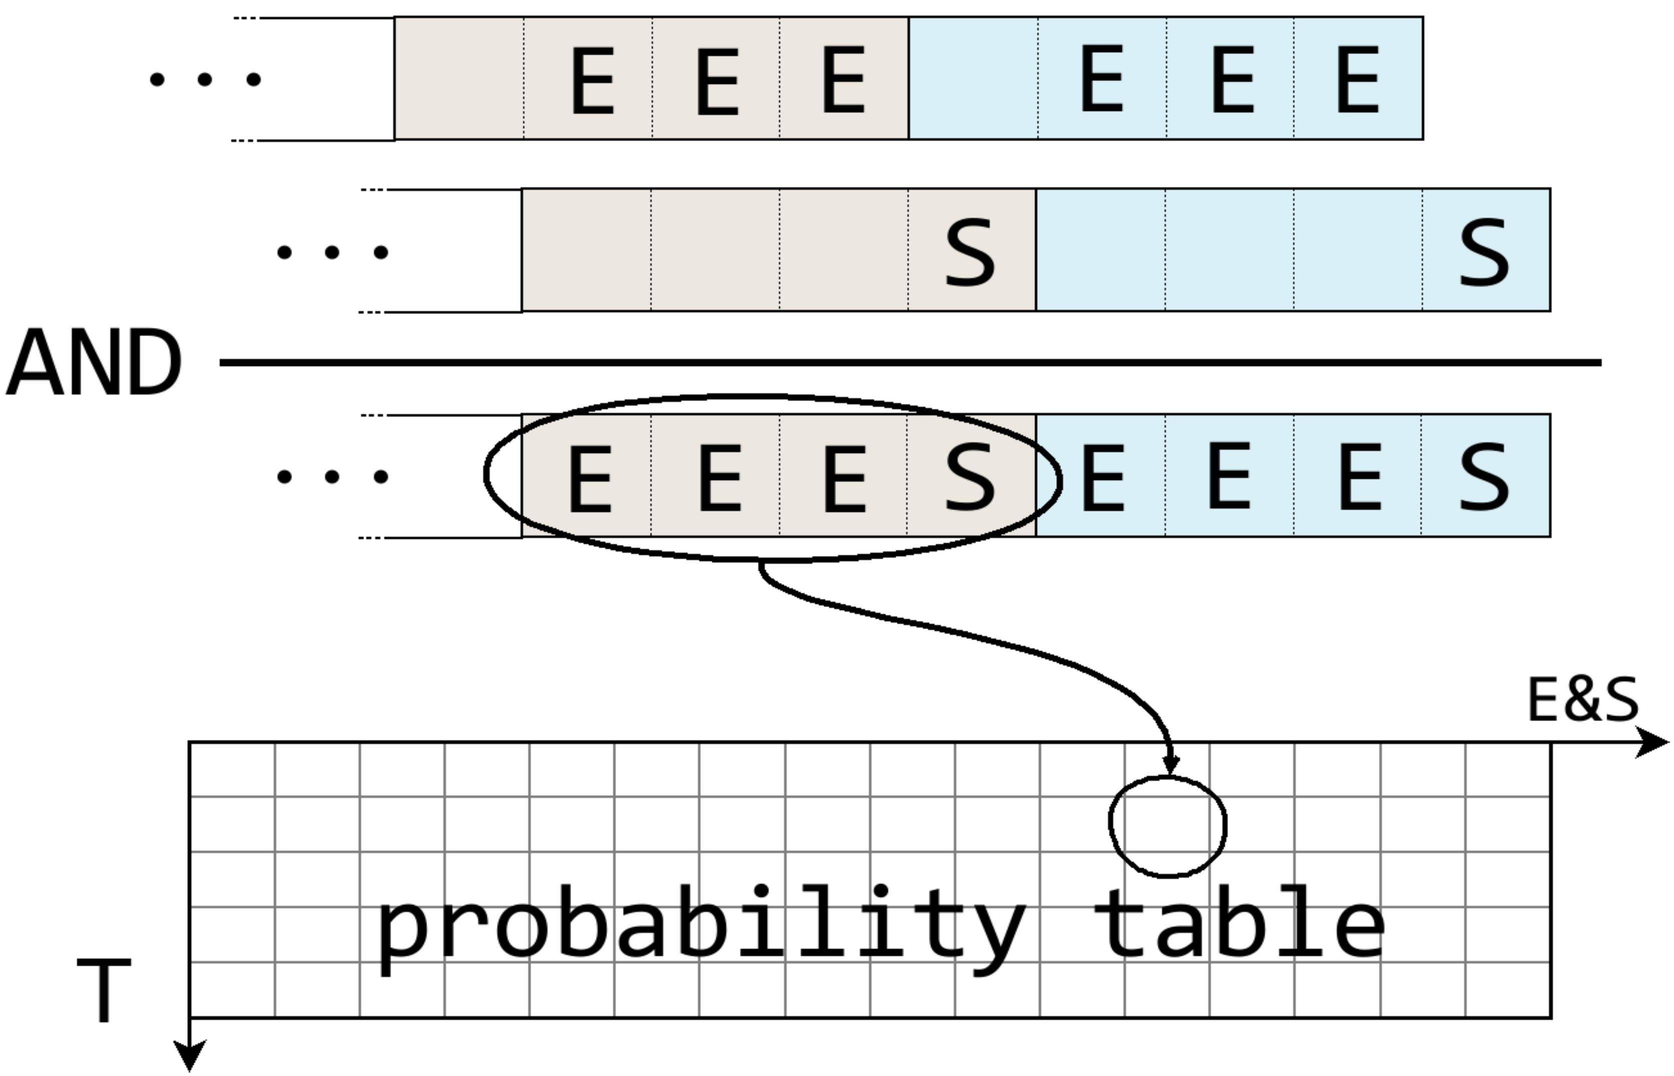

Comparison of Multi-spin coding schemes

•4 bits scheme required extra memory space

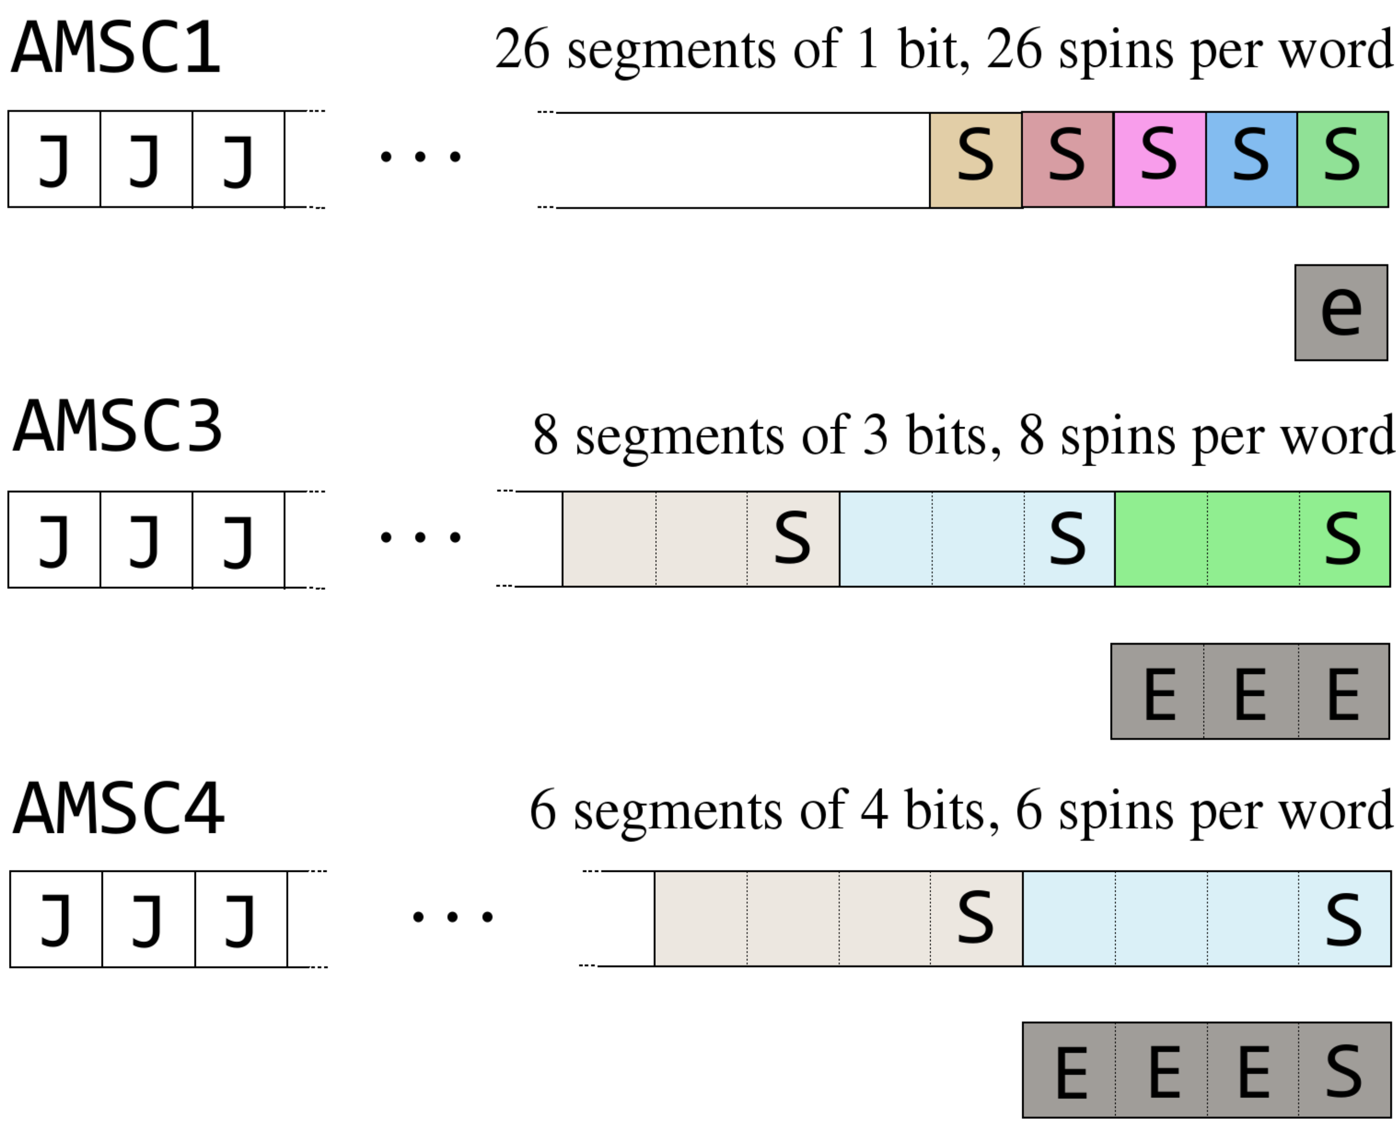

Multi-spin coding (4 bits scheme)

- Use 1-bit to store one spin.

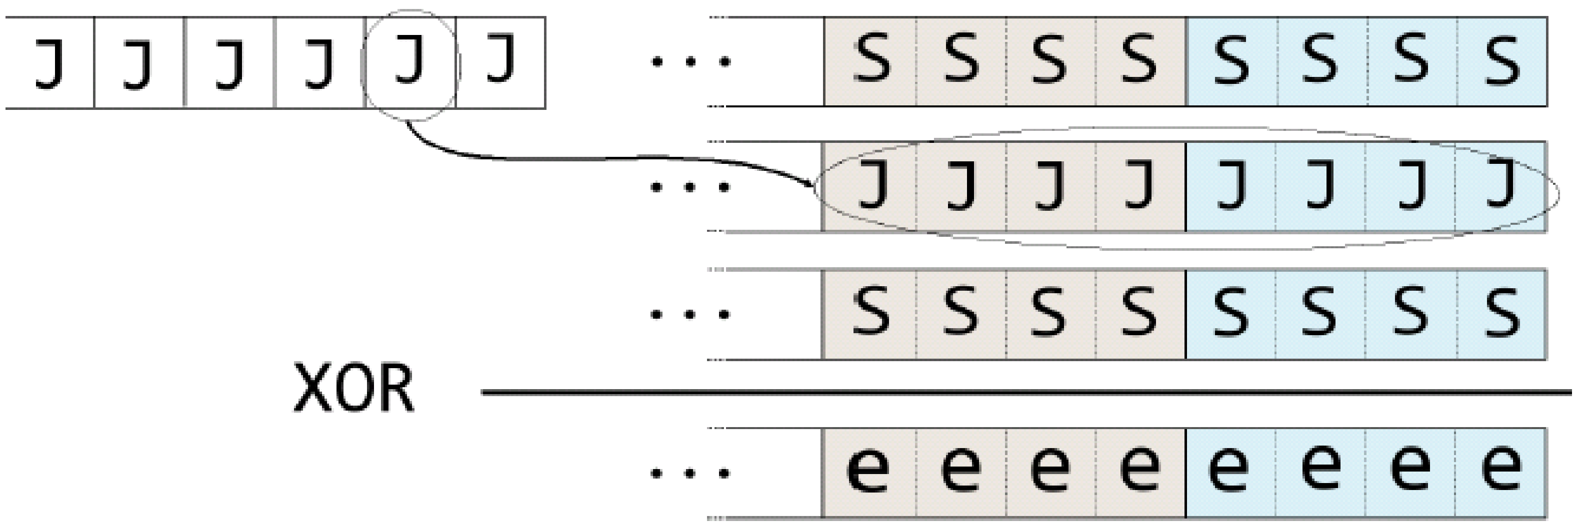

Multi-spin coding (4 bits scheme)

- Bit mask to expand the local energy 'e' into 4-bit segment

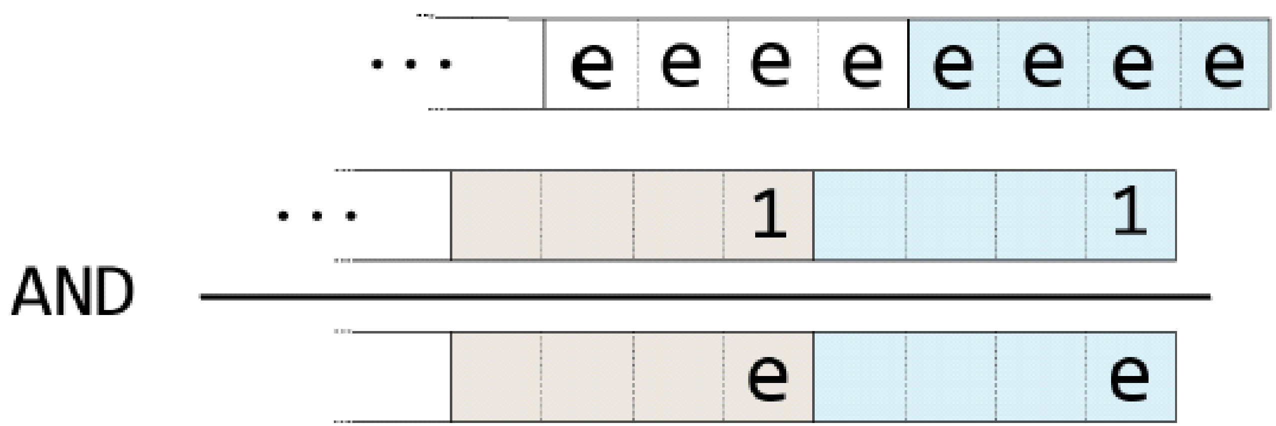

Multi-spin coding (4 bits scheme)

- Sum over the local energy for all 6 neighbors in a 4-bit format

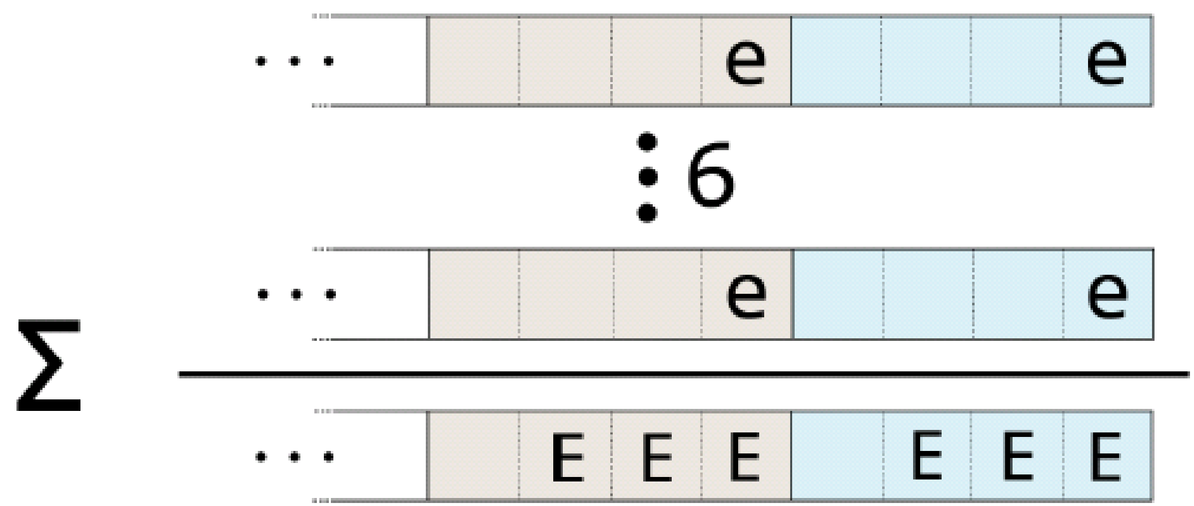

Multi-spin coding (4 bits scheme)

- Concatenate the 3-bit for energy and the 1-bit for spin state.

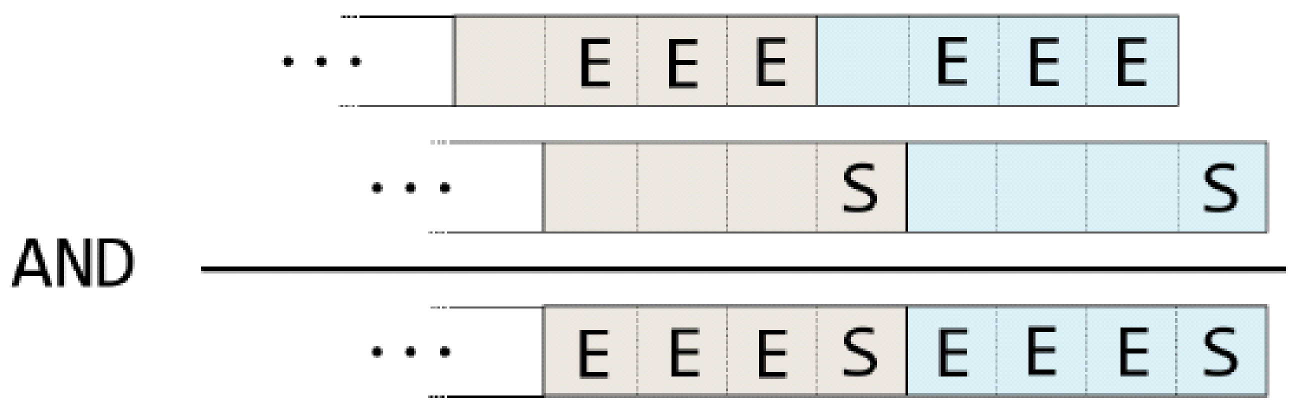

Comparison of multi-spin coding scheme

Parallel tempering move

- All replicas of the same realization are in the same streaming processor
- Overhead of calculating the total energy
- Swap the rows of the table for tempering move

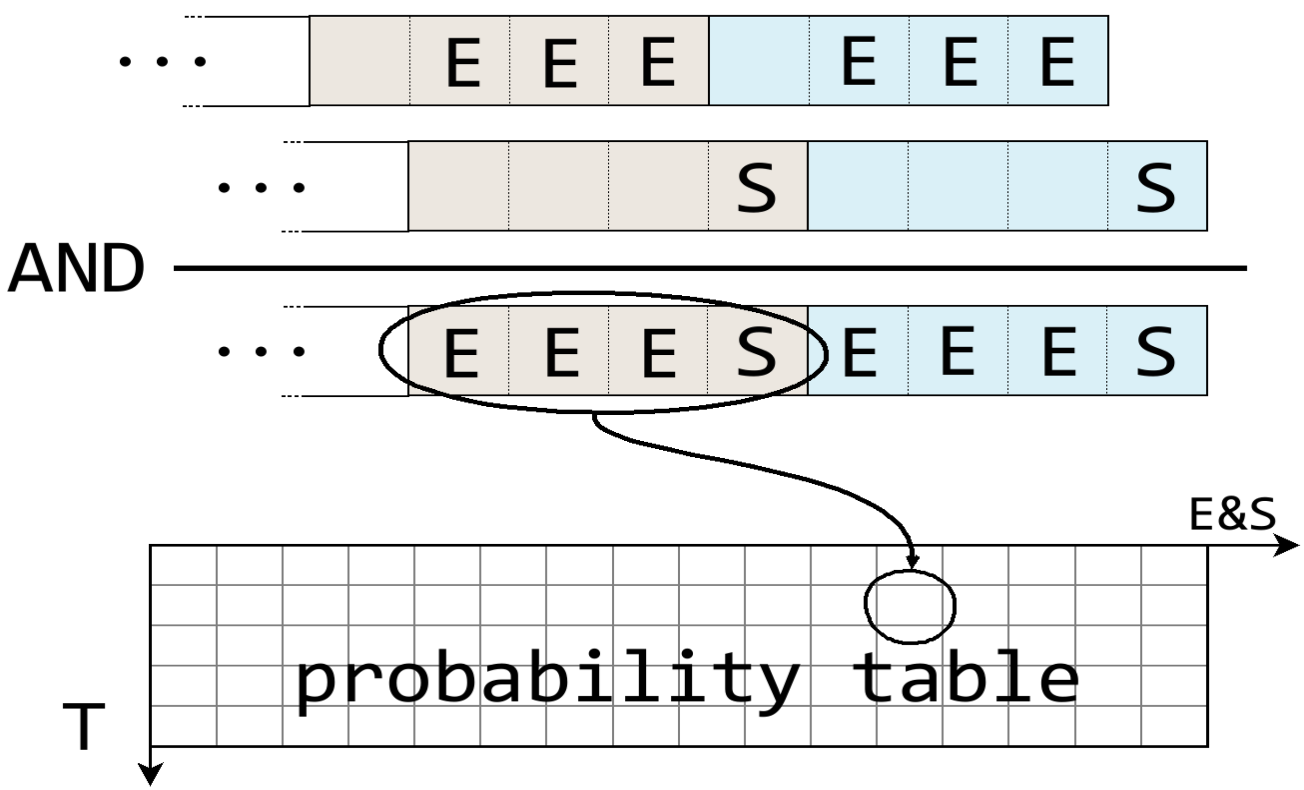

Outstanding challenges of spin glass simulations

- Interpretation of the data

1. What quantities for capturing phase transition.
   (overlap order parameter?)

$$q=(1/N)\sum_i s_i^{\mu} s_i^{\nu}$$

2. Scale invariance at the critical point.

3. Crossing of dimensionless quantities, e.g. Binder ratio.

4. Correlation length (assuming exponential decaying correlation).

5. Avoiding k=0 susceptibility?

6. Ratio of susceptibilities at two different finite momenta?

Quantities for identifying phase transition

Measured quantities

Binder ratio

$$B=\frac{1}{2}\left\{3-\frac{[\langle(q-[\langle q\rangle])^4\rangle]}{[\langle(q-[\langle q\rangle])^2\rangle]}\right\}$$

Spin glass susceptibility

$$\chi(\boldsymbol{k})=[\langle(q(\boldsymbol{k})-[\langle q(\boldsymbol{k})\rangle])^2\rangle]$$

Correlation length

$$\xi=\frac{1}{[\sin(2\pi/L)]}\sqrt{\frac{(\chi(\boldsymbol{k}=(0,0,0)))}{(\chi(\boldsymbol{k}=(2\pi/L,0,0)))}-1}$$

Susceptibility ratio

$$R_{12}=\frac{\chi(\boldsymbol{k}=(2\pi/L,0,0))}{\chi(\boldsymbol{k}=(2\pi/L,2\pi/L,0))}$$

Spin glass at different dimensions

- Infinite dimension → Exact mean field solution with RSB.
- 6 dimension → Non-Gaussian terms are irrelevant, upper critical dimension.
- 3 dimension → Solid numerical evidence that spin glass exists at finite temperature.
- 2.5 dimension → Could be the lower critical dimension?

- **Whether the spin glass phase below six dimension has some resemblance of the mean field solution?**
- **Non-trivial distribution of spin overlap?**
- **Finite critical field?**

Some recent studies

- No crossing in correlation length for 3D under field

Young and Katzgraber 2004

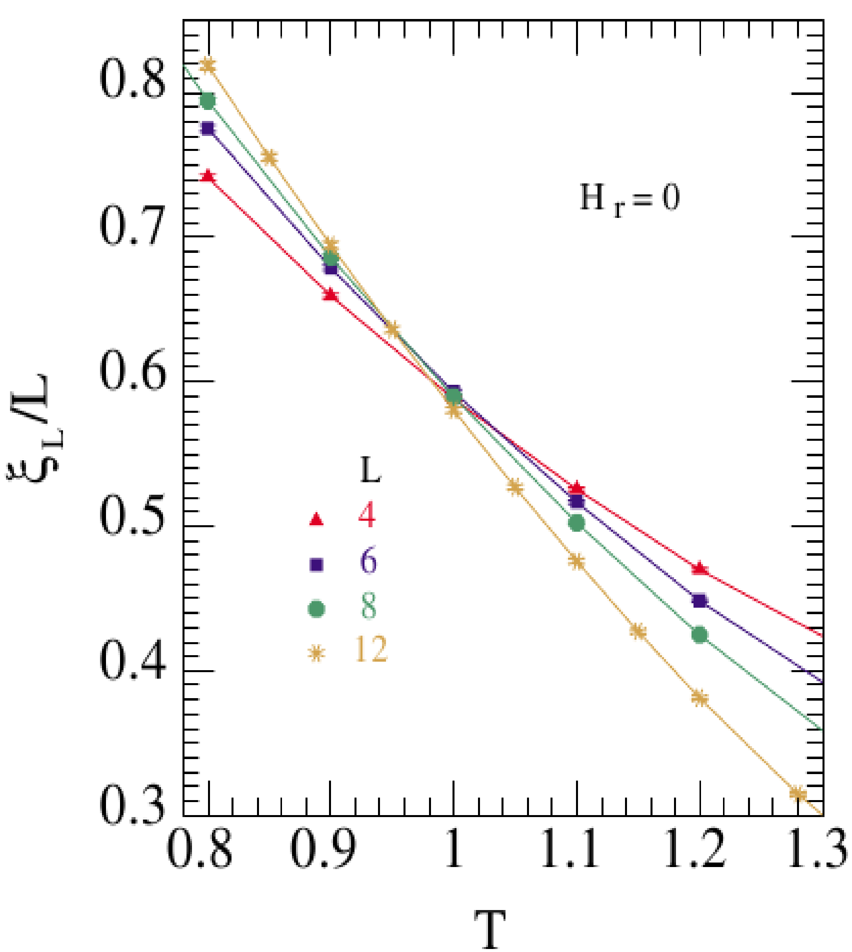

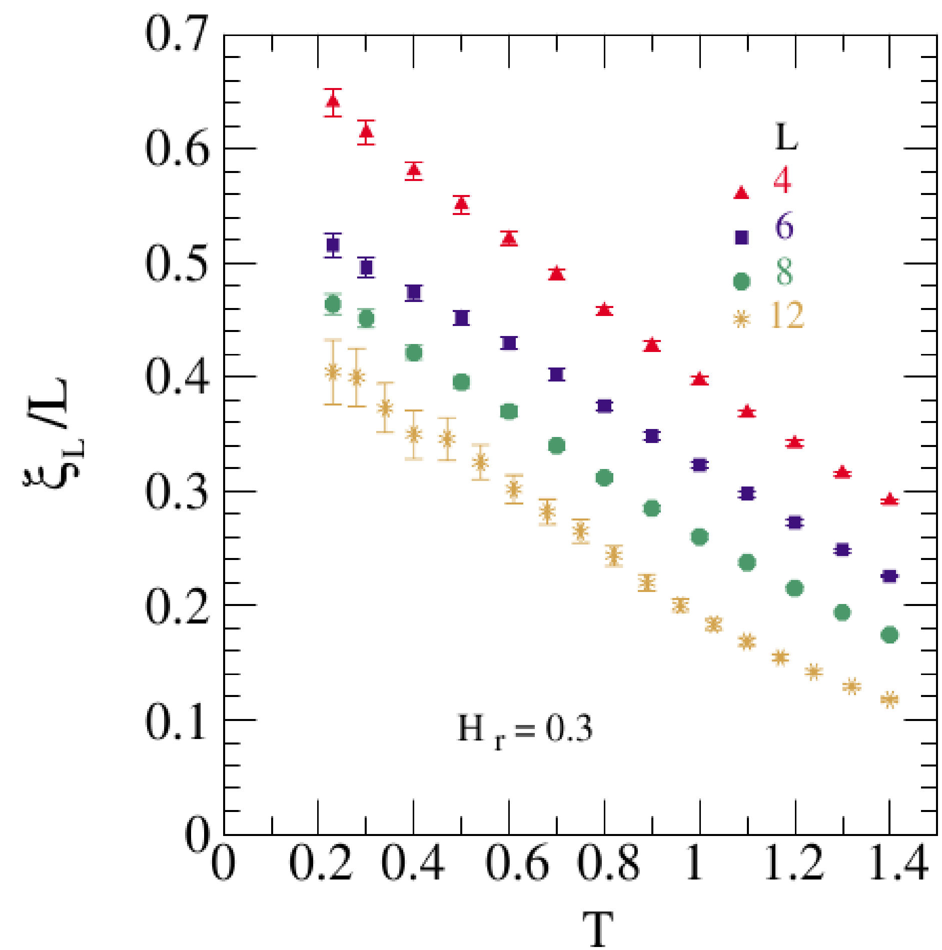

Some recent studies

- No ultrametricity in 3D

Hed, Young, Domany 2004

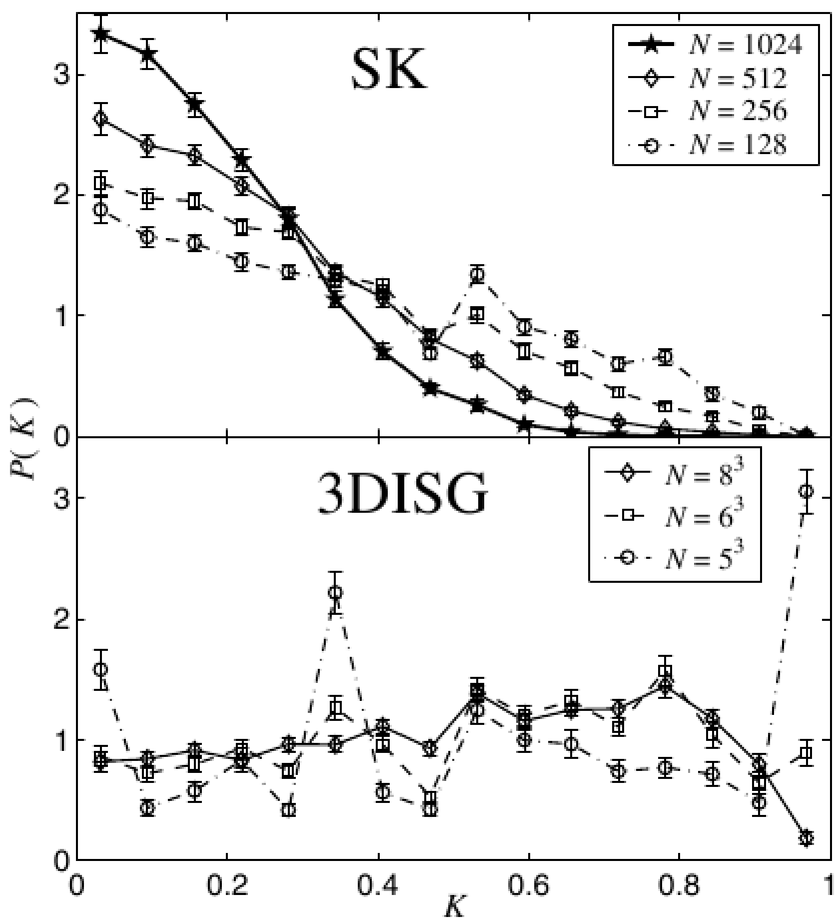

- Measure for the ultrametricity

$$K_{\mu,\nu,\rho} = (d_{\mu,\rho} - d_{\nu,\rho}) / d_{\mu,\nu}$$

$$d_{\mu,\nu} = (1 - q_{\mu,\nu}) / 2$$

- Ultrametricity in the state space implies $K_{\mu,\nu,\rho} = 0$

Some recent studies

- Effective 1D long range model, no AT line in 3D, 4D data give conflicting results.

Larson, Katzgraber, Moore, Young 2012

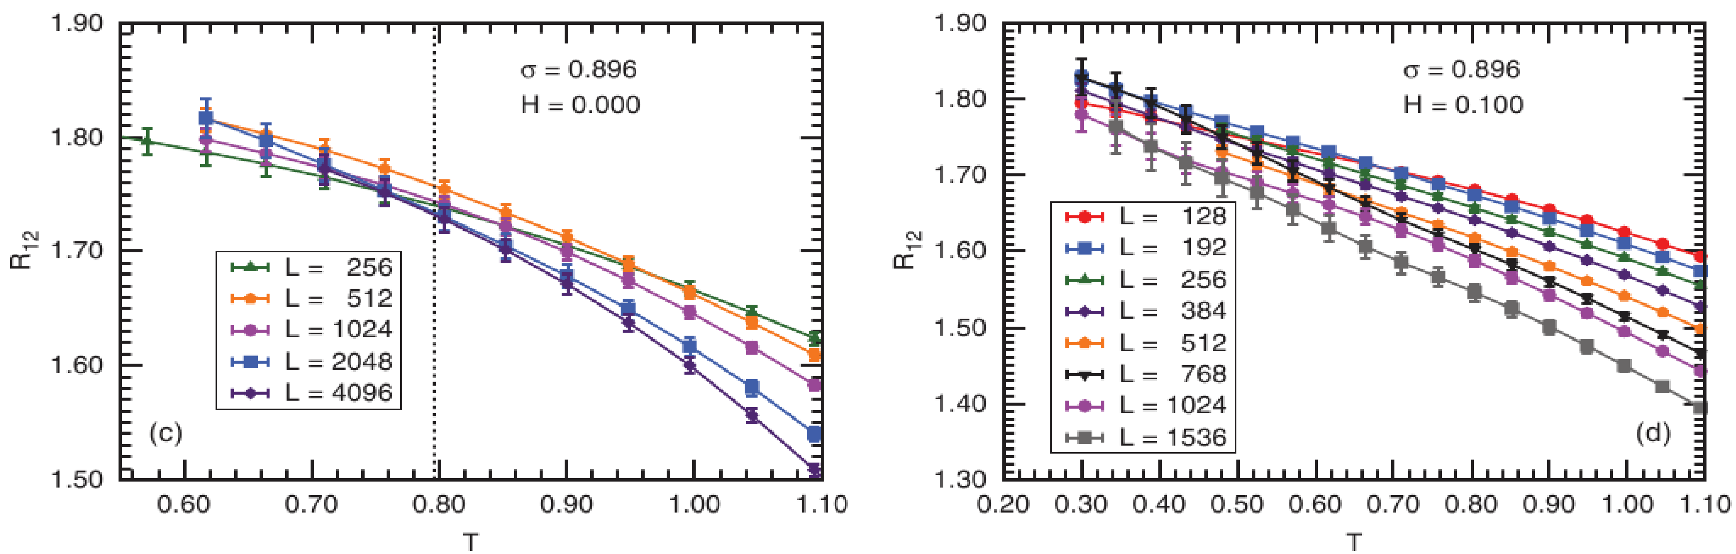

$$H = -\sum_{i,j} \epsilon_{i,j} J_{i,j} S_i S_j - \sum_i h_i S_i$$

$\epsilon_{i,j} = 0, 1$ with probability $P_{i,j} \sim r_{i,j}^{-2\sigma}$

Some recent studies

- Crossing of susceptibility ratio for 4D under field
  Banos et al., Jauns collaboration 2012

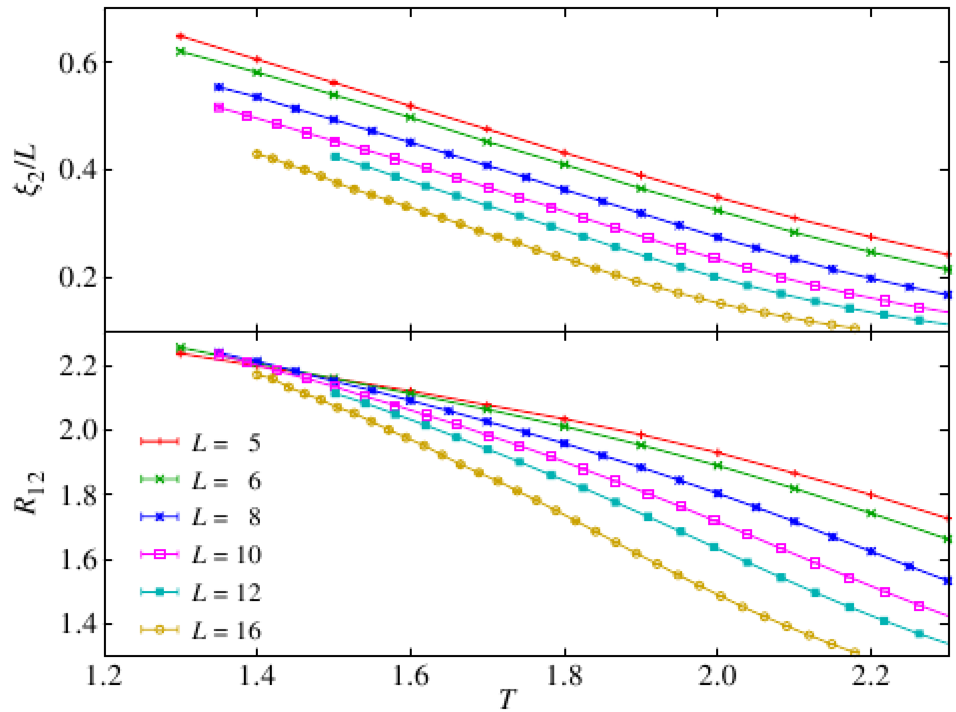

Benchmarking GPU Results

- Sanity Checks for the GPU code

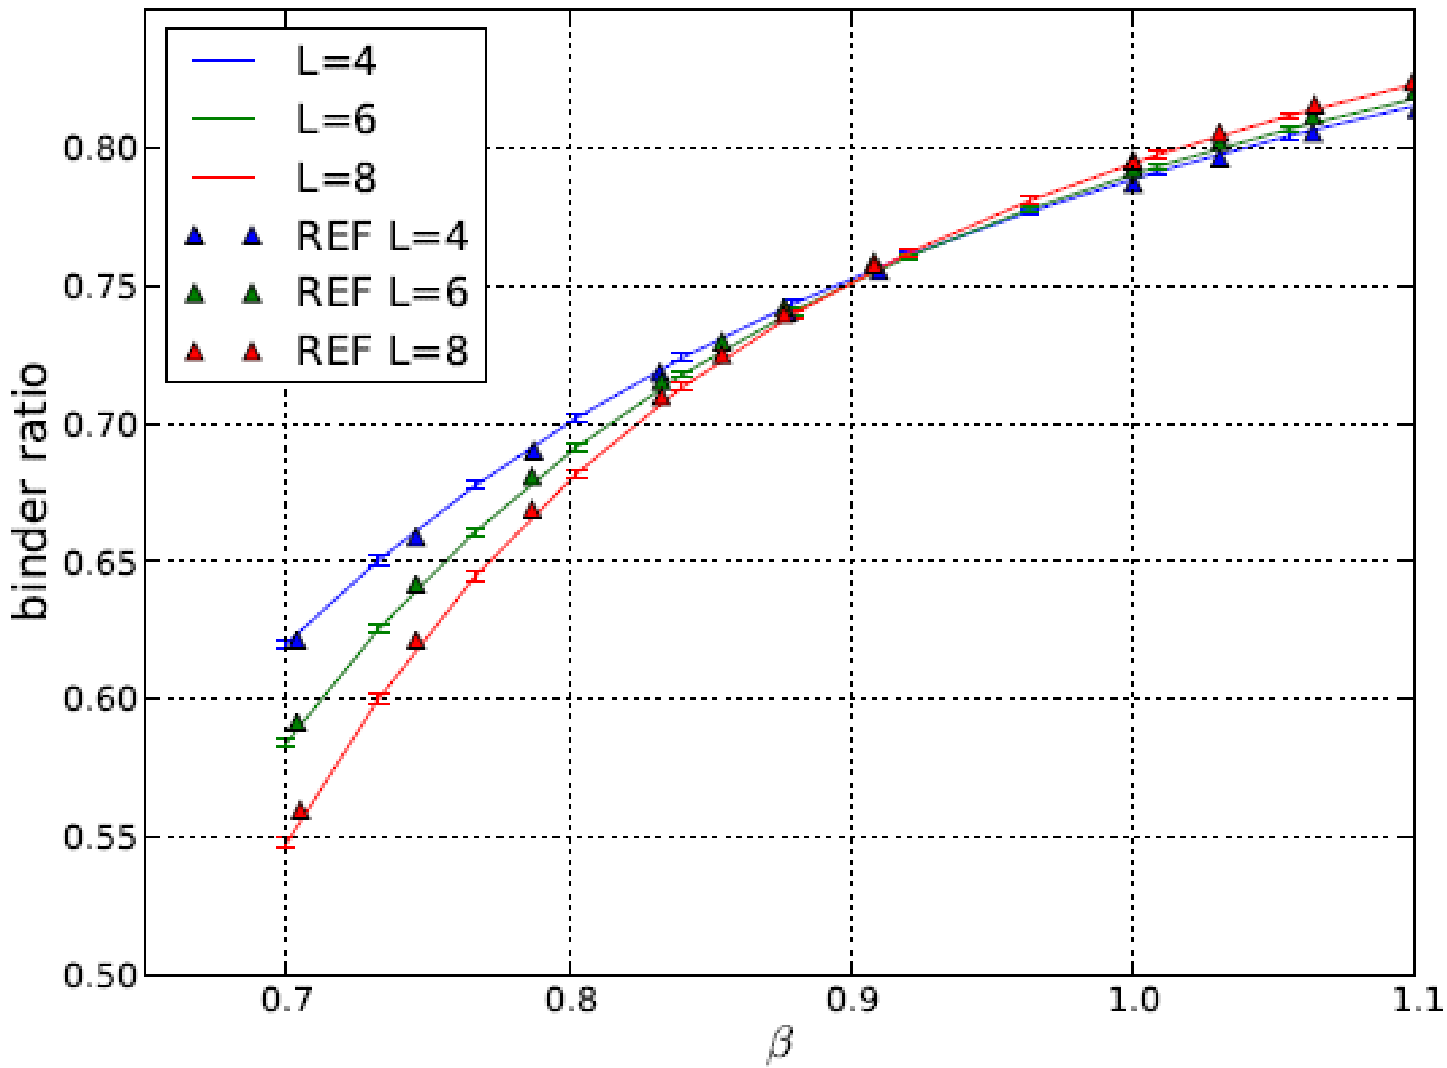

REF data extracted from Katzgraber, Koerner, and Young PRB (2006)

Binder Ratio at a Finite Field

- Binder ratio at h=0.1

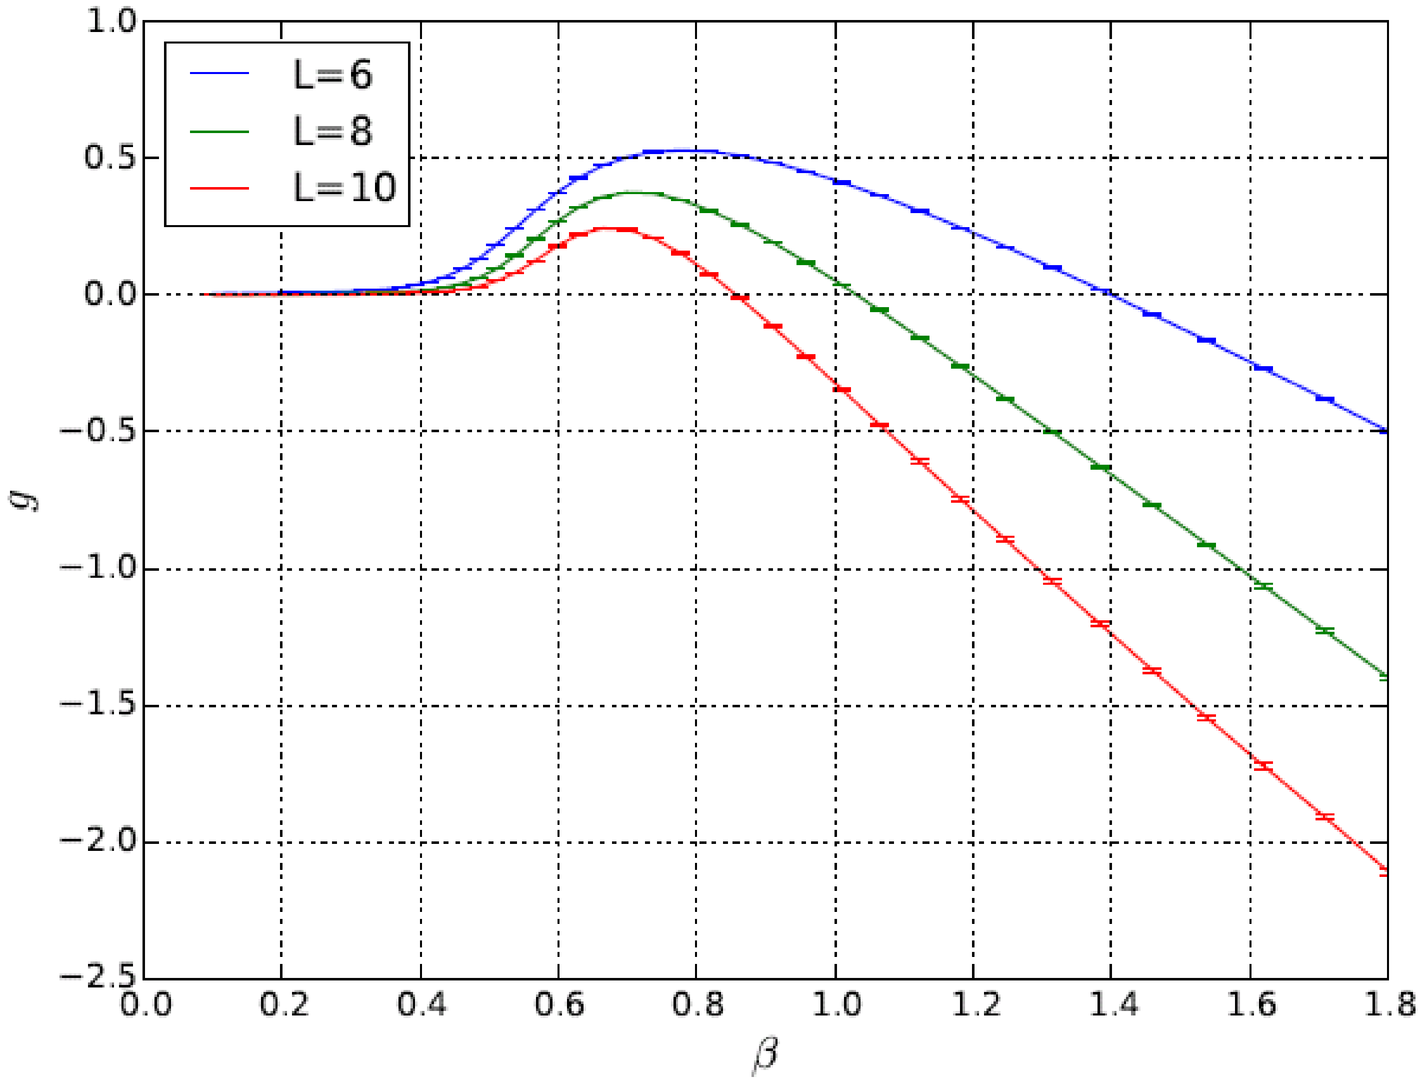

Correlation Length at a Finite Field

- Correlation length at h=0.1

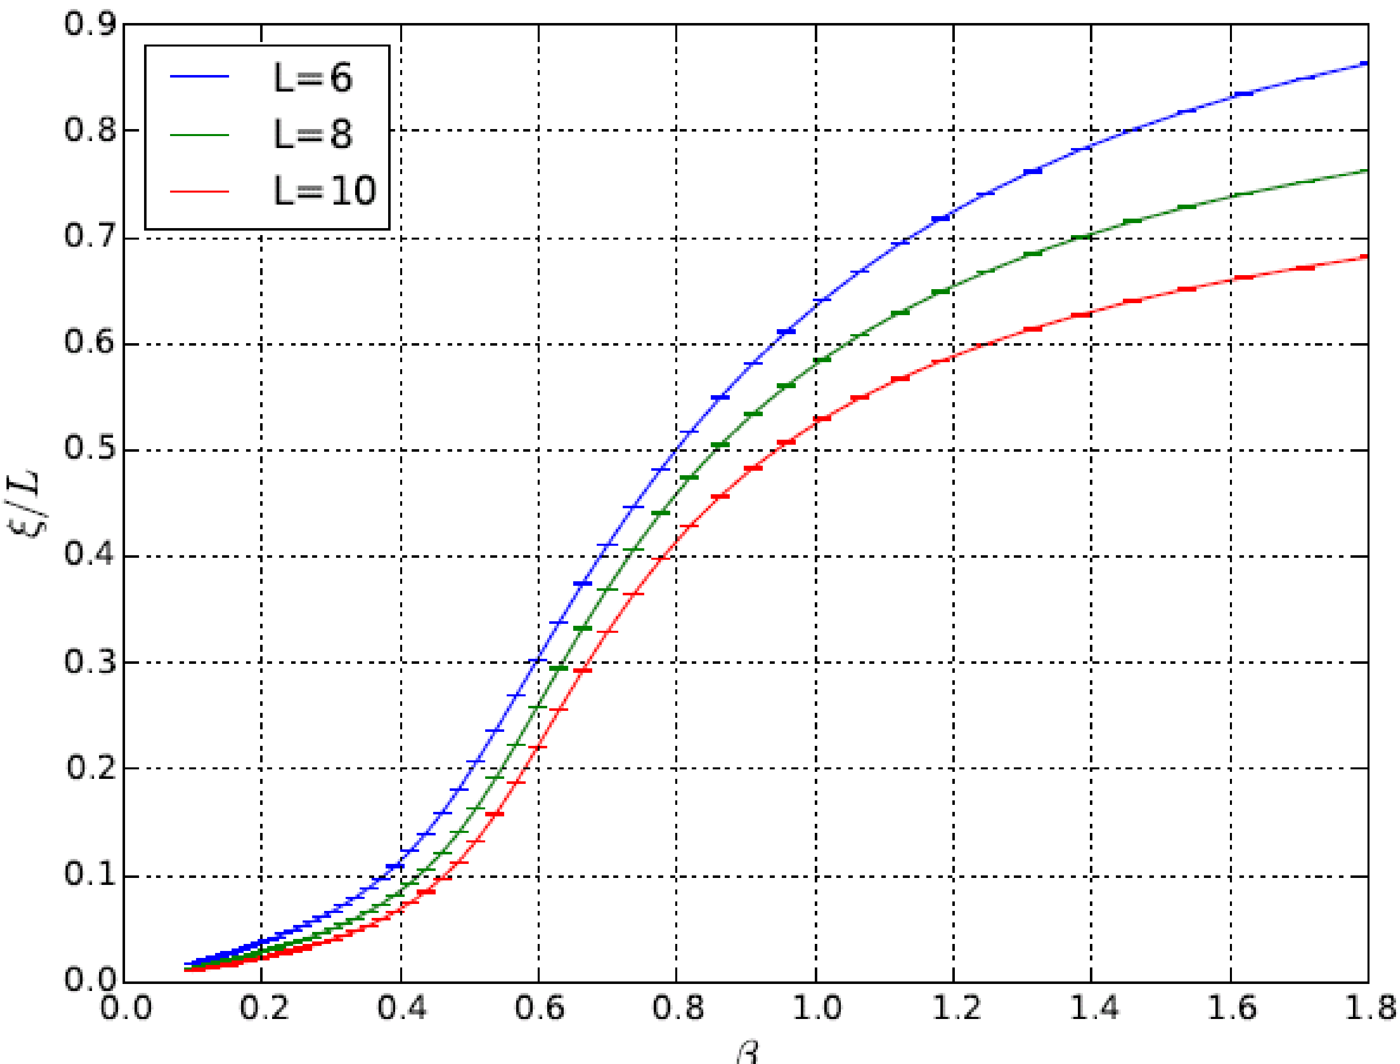

No crossing for correlation length,
consistent with Young and Katzgraber PRL 2004.

Susceptibility Ratio at a Finite Field

- R_{12} at h=0.1

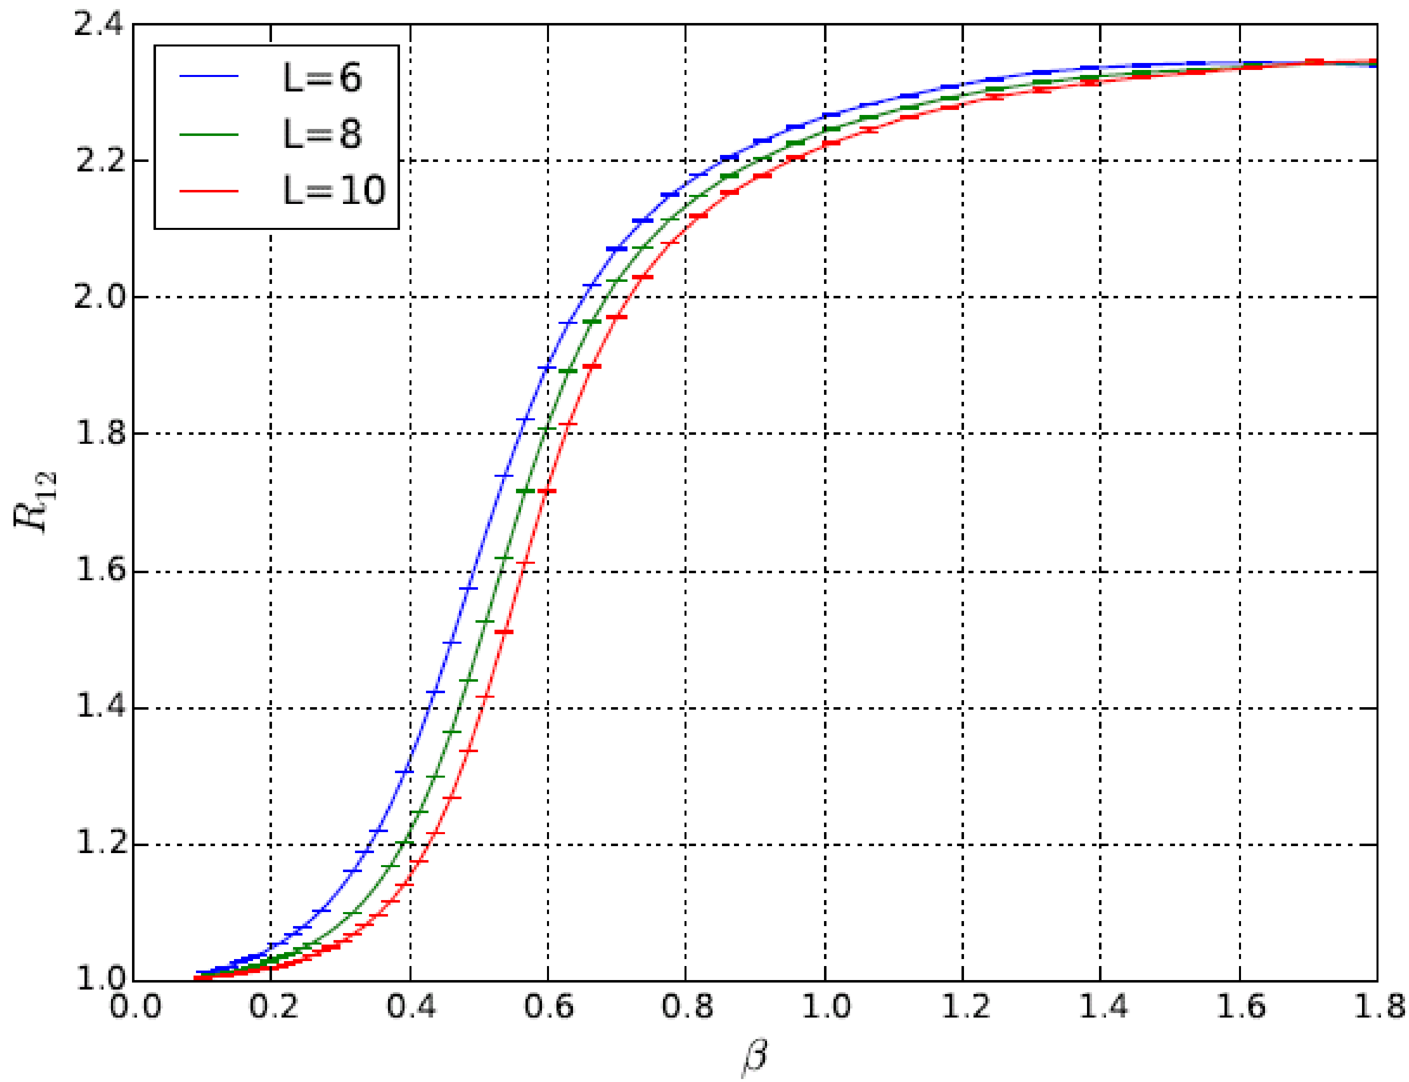

Tests for equilibration

- R_{12} at h=0.1, L=10

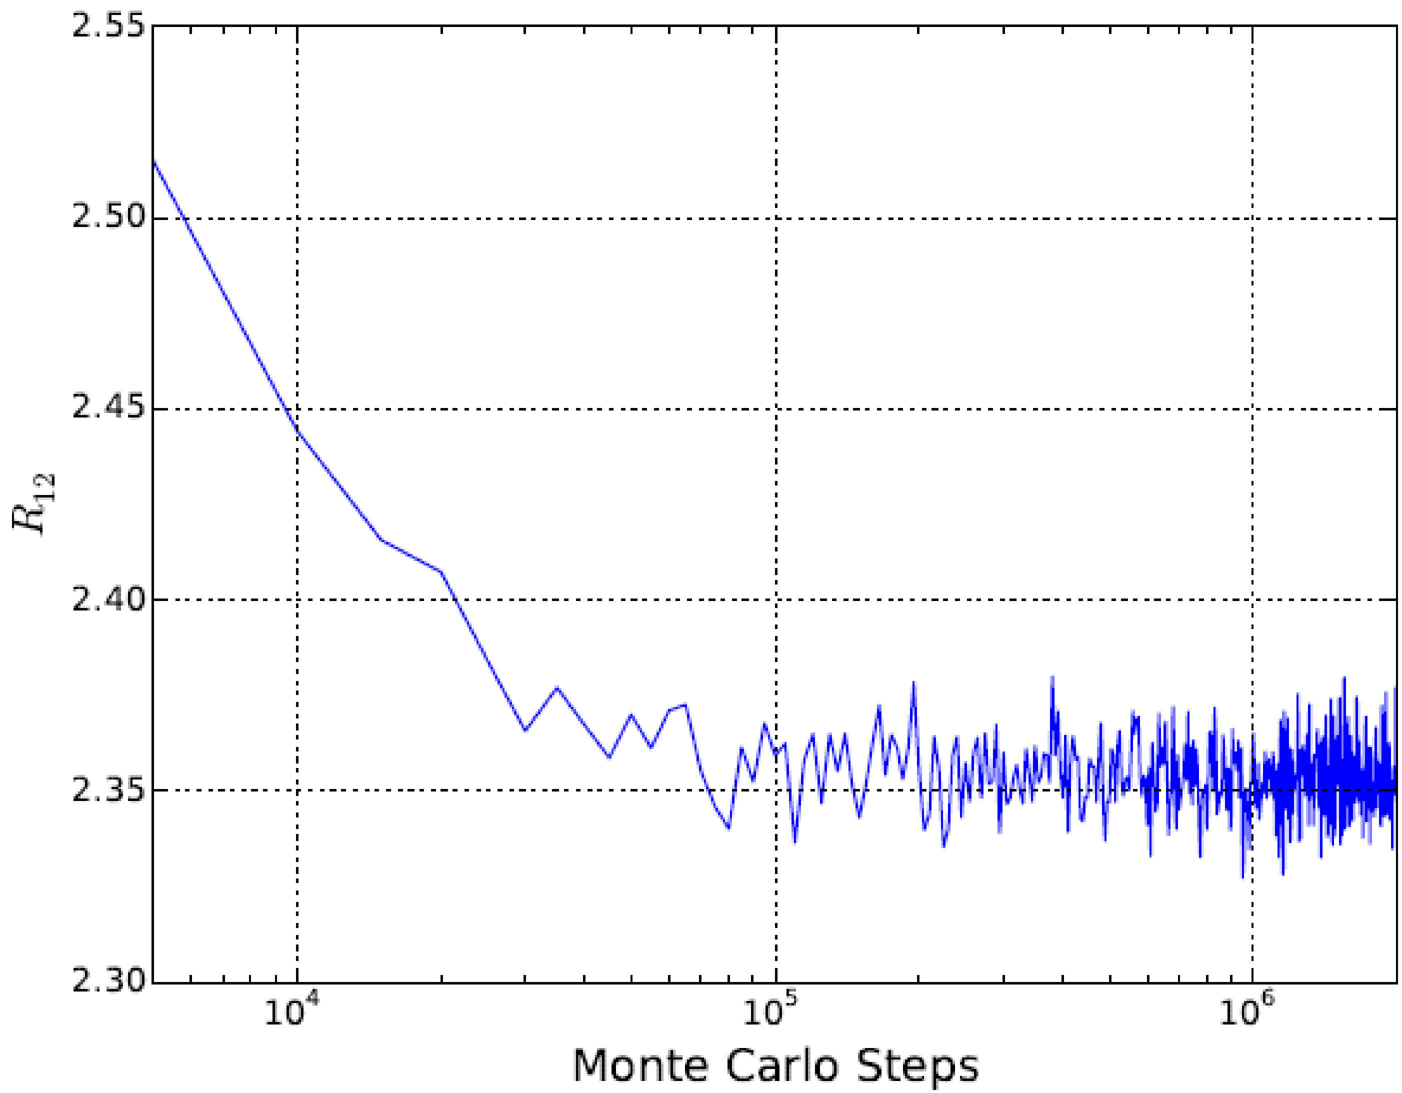

Effects from the number of realizations

- R_{12} at h=0.1

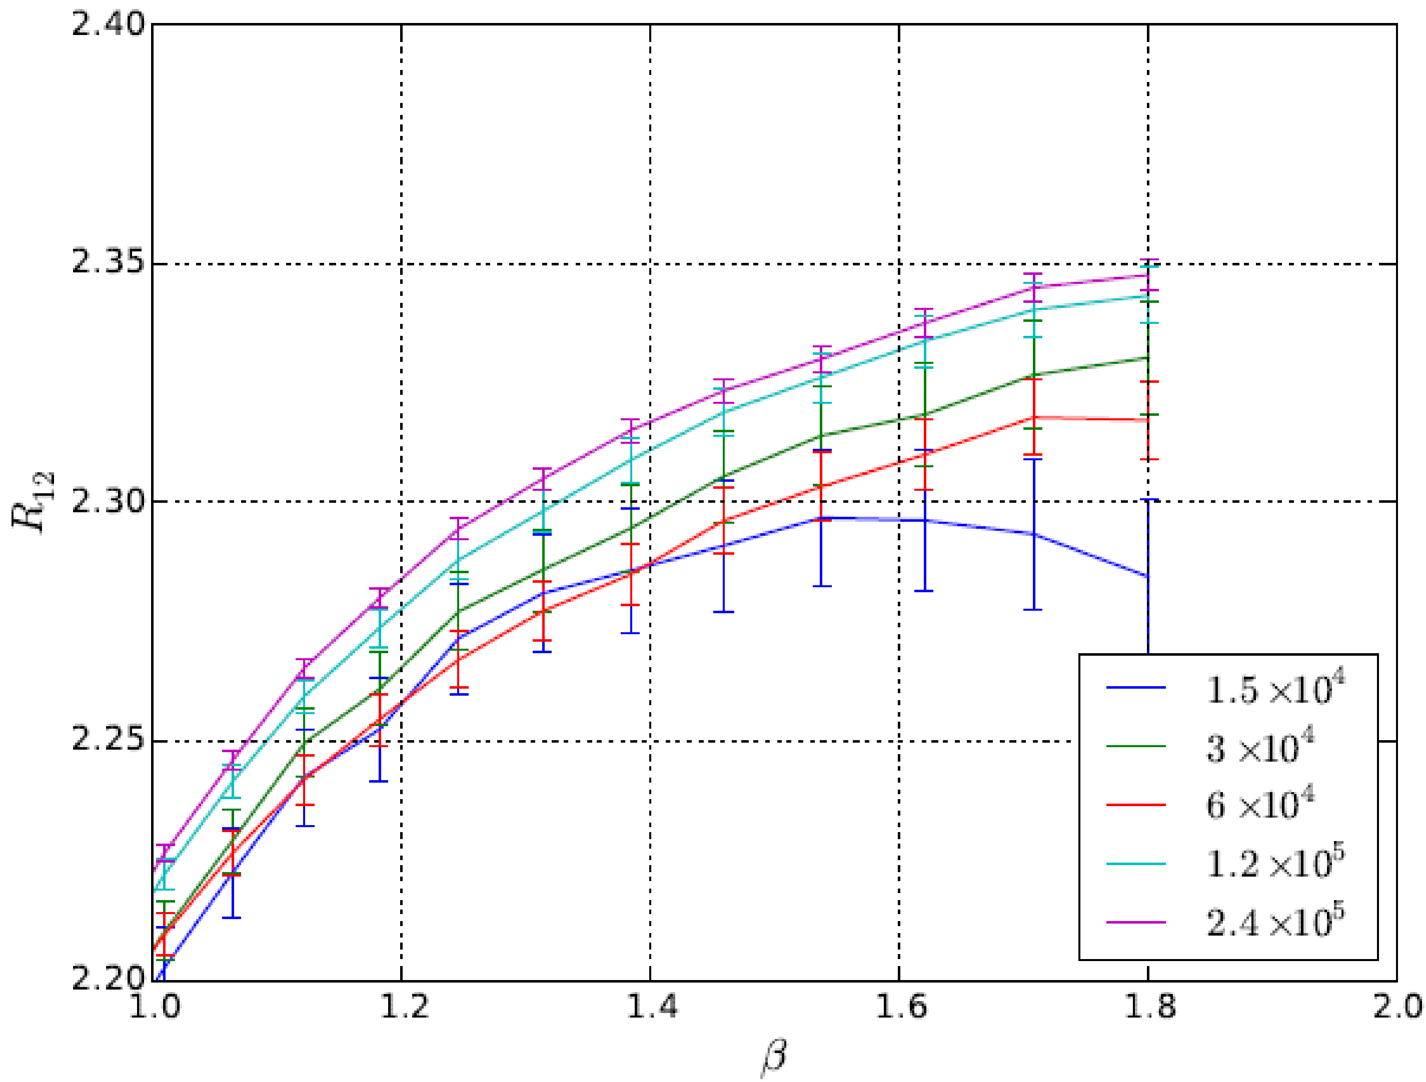

Susceptibility Ratio at a Finite Field

- R_{12} at h=0.1

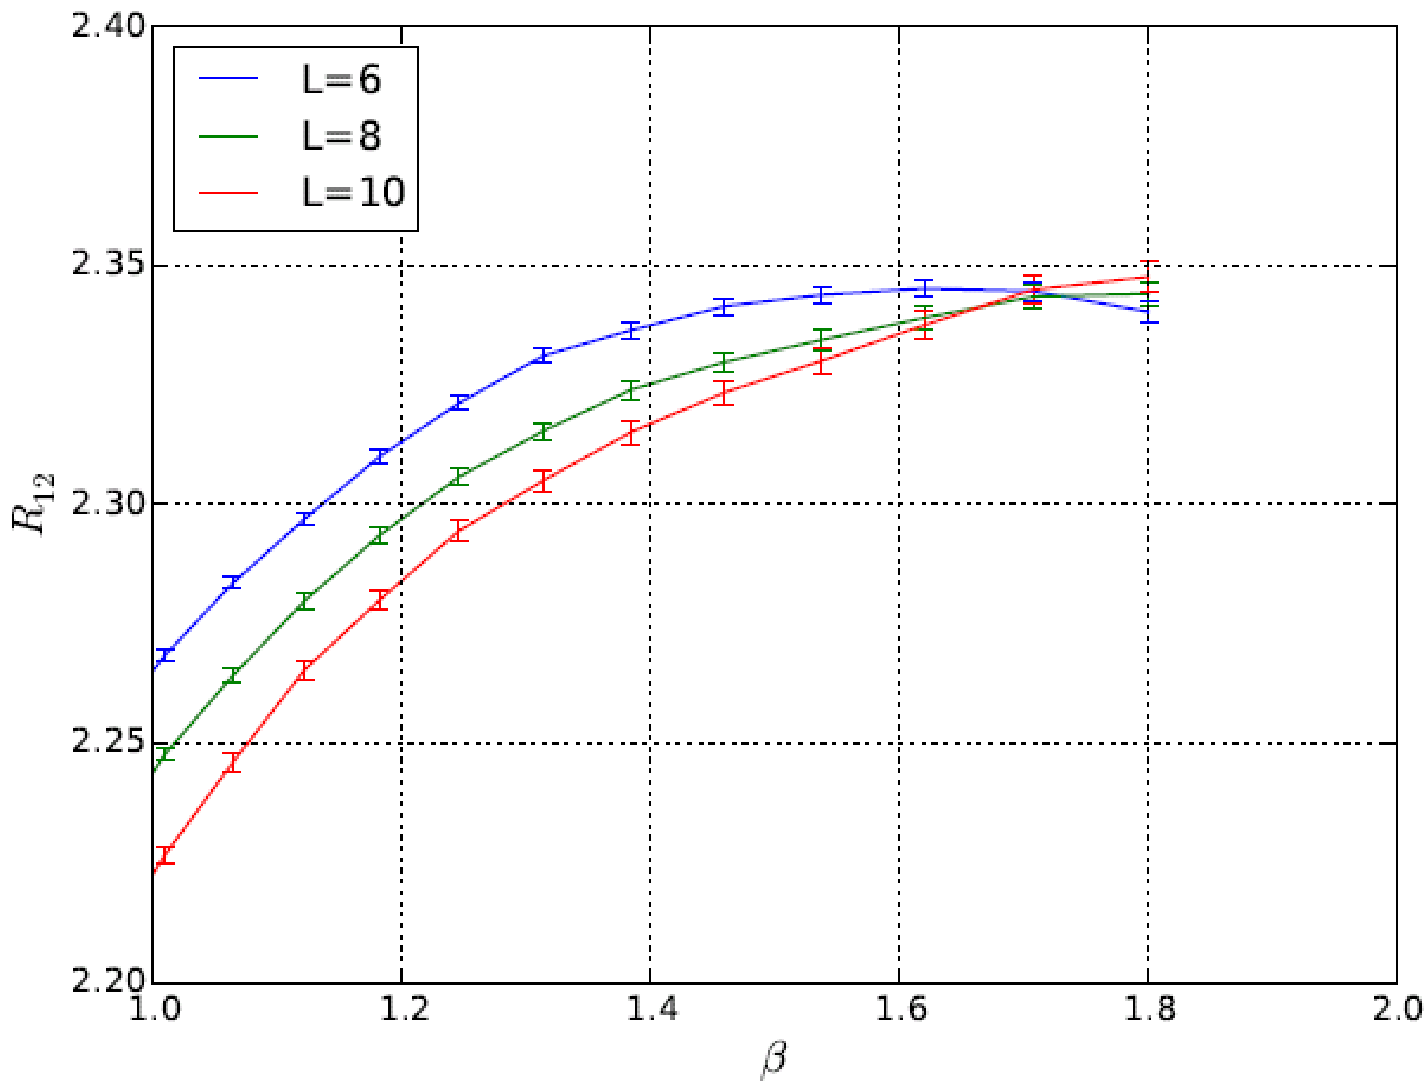

Conclusion

- GPU computing has a good potential for simulating glassy systems.
- You can buy a GPU card at Bestbuy, Amazon... for ~$500.
- We implemented parallel tempering and multispin coding algorithm for the simulation of EA model on GPU, 34ps / spin flip with 1 PT move for every 10 sweeps, non-shared random number using GTX580 card.
- We reproduce the results for 3D EA model without external field.
- The ratio of susceptibility is noise, more involved simulations are currently being done.

LA-SiGMA
Louisiana Alliance for Simulation-Guided Materials Applications

http://lasigma.loni.org/