# Bit-Packed Damaged Lattice Potts Model Simulations with CUDA and GPUs

Mitchell G.B. Johnson and K.A. Hawick

Email: { m.johnson, k.a.hawick}@massey.ac.nz

Complex Systems and Simulations Group, Computer Science, Institute of Information and Mathematical Sciences, Massey University, Albany, North Shore 102-904, New Zealand

http://complexity.massey.ac.nz

## Introduction

New magnetic materials such as manganites exhibit exciting properties that are leading to new technological devices such as dense storage media for next generation hard disk drives. Understanding the properties of such materials continues to be a challenge and often computer simulation and modelling is the only practical means of investigation of complex microscopic phenomena. Numerical experiments based on a simulation model can form the basis for making macroscopic measurements which can be compared with measurements made of macroscopic properties of real materials.

The Griffiths phase is a regime exhibited by random spin systems and there is considerable recent research reporting the presence of this phenomena in real experoments with manganite materials. Comparisons can be made with random magnetic spin models and in particular much work has been done with the diluted or bond-damaged Ising model of a magnet including important theoretical studies and comparisons with disordered dielectric materials. The dilute Ising model has formed the basis for a very large body of research on phase transitions and critical phenomena. Reported research has covered the dynamics of dilute magnets and possible scaling laws explaining the relaxation behaviour as well as residual persistemce effects.

The Potts model extends the Ising model by allowing an arbitrary number of spin states – the Ising model only supports two – and which has a range of other applications including quantum gravity as well as practical magnetism in materials.

There are a number of useful quantities that can be measured from a simulated Potts or Ising system including the critical exponents that govern how properties such as the magnetization or heat capacity of the system change with temperature. Of particular interest to us in this present work is how the critical temperature itself changes as the magnetic system is damaged. Changing the fraction of bonds present is akin to simulating the faults and defects and polycrystalline structure of real magnetic materials. These can be characterised by a damage parameter that specifies how many of the spin-spin interaction bonds are actually present.

Work has been reported in the literature concerning how damage is formed and spreads in an Ising system and it is of great interest to study this in diluted Potts models also. Some work has been reported on the Potts model on graphs and on two dimensional systems, but not in damaged arbitrary-Q state and higher dimensional Potts systems.

It is necesary however to be able to simulate large simulated system sizes with in excess of $10^6$ spins in two dimensions or even $10^9$ spins in three dimensions. As we discuss in this paper, it is also necessary to be able to run simulations for a relatively large number of steps. This is firstly to properly thermally equilibrate systems so that they are truly representative of the simulated anneal temperature. Secondly, to make high quality measurements on simulated model the run time statistics need to be collected over long run time sequences of uncorrelated sample model configurations. Data parallel computing has long been a readily applicable technique to speed up Ising and Potts model simulations. Various approaches can be used to distribute the spin-updates across independent processing elements in a manner than does not affect the fundamental properties of the model. Graphical Processing Units (GPUs) are a power modern platform implementing the data parallel ideas that used to require dedicated supercomputers in the 1980s and 1990s. GPUs have been successfully used for simulating Ising and Potts models on regular undamaged attices. Work is also reported on simulating the Ising model on irregular graph structures such as small world networks. To our knowledge however no group has reported attempts to simulate damaged lattice Potts or Ising systems using GPUs.

## Method

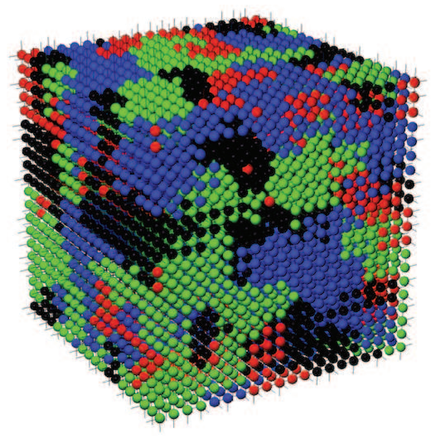

**FIGURE 1: A $20^3$ Potts model with $p = 0.6$ damaged links**

The Potts model is a generalisation of the well-known Ising model. Each spin takes on one of possible $Q$ values and the spin sites are coupled to their nearest neighbours so that for a ferromagnetic coupling, it is energetically favourable for the spin values to align. For systems like the Ising and Potts models there is a critical temperature $T_c$, above which the system remains random with no magnetic moment, but below which the system has aligned clusters of spins and exhibits a definite magnetic moment. These spin models can be simulated as described below but to accurately determine their properties requires simulation of quite large model systems and quite long simulation times are required to ensure the models are properly equilibrated and genuinely representative of the anneal temperature, before measurements are made on them.

We set up our d-dimensional model on a periodic mesh of $N = L^d$ spins. For simplicity in the work reported here we use the same lattice length $L$ in all dimensions $(x, y, z, ...)$. The Hamiltonian (in dimensionless units) of the Potts model is:

$$-\frac{1}{k_B T}\mathcal{H} = \sum_{<i,j>} K_{i,j}\delta\sigma_i, \sigma_j, K_{i,j} > 0 \qquad (1)$$

where $k_B$ is Boltzmann's constant, $T$ is the temperature and $K_{i,j}$ is the dimensionless coupling between nearest neighbouring spins denoted by convention as $< i, j >$. The summation is over all nearest-neighbour pairs, and is over $2dN$ possible bonds in a rectilinear hypercubic mesh and $\delta_{\sigma_i,\sigma_j}$ is the Kronecker delta function, which is zero unless the two spins $(\sigma_i, \sigma_j)$ are equal. We work with generalised spins $\sigma_i$ that take on any of the $Q$ possible values, labelled $\sigma_i = 0, 1, ..., Q - 1$ in a $Q$-state Potts model. We introduce damaged bonds by choosing the bond couplings from the bimodal distribution:

$$P(K_{i,j}) = p\delta(K_{i,j} - K) + (1 - p)\delta(K_{i,j} \qquad (2)$$

using the Dirac delta function notation and parameter $0 <= p <= 1$. If probability parameter $p$ is unity and all bonds are present, then this case corresponds to the undamaged Potts model. Decreasing $p$ thus introduces damage. The dynamics of the model is then to select spin sites and attempt to flip their spin to a different value and accept the change probabilistically using the Boltzmann distribution weighting factor $\exp(-\Delta\mathcal{H})$ determined from the change in energy Hamiltonian that would potentially occur from the flip.

## Bit-Packing

Simulating a large q-state Potts system requires a large amount of memory, Using a naive method of storing the spin value in a single integer for a $8192^2$ sized 2D system the memory used will be 262 megabytes excluding memory needed for random number generation. This is exacerbated in a damaged lattice system where each link must be tracked, to accomplish this we use an `unsigned short` (that occupies two bytes) for each site to store both the spin and the status of the neighbor links. Using dynamic bit manipulation to pack spins and neighbor links together allows various configurations of dimensions (D) and states to be stored in a single show. An example of this bit configuration in a 3D 8-state Potts model with one broken link is shown in Figure . This does limit the number of states that can be simulated and varies according to the number of dimensions.

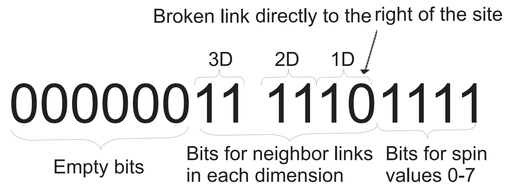


**FIGURE 2: The arrangement of bits for a single site of an 8 state, 3 dimensional Potts model with1 broken neighbor link.**

The algorithm for dynamic allocation of bits for each lattice point is outlined in Algorithm below.

**declare** $Q\_MASK \leftarrow (Q + (Q \mod 2))/2)$
**for** $k \leftarrow 0$ to $N$ **do**
  $site[k] \leftarrow site[k]$ **OR** $Q\_MASK$
  **for** $d \leftarrow 0$ to $D * 2$ **do**
    **declare** $random \leftarrow RandomInteger(0 - 1)$
    $site[k] \leftarrow site[k]$ **OR** ($random$ **SHIFT LEFT** ($d$+$Q\_MASK$))
  **end for**
**end for**

## GPU Implementation

The use of GPGPU in the simulation of regular square lattice spin models such at the Ising and Q-state Potts has been proven to produce significant speedups over CPU implementations. To exploit this massively parallel approach we utilize the CUDA extension to the C language developed by NVidia to leverage the SIMT (single instruction multiple thread) architecture of NVidia GPU's. On the current generation of hardware this approach allows us to run up to up to 512 concurrent threads on a single GPU each computing an independent spin calculation.

Global Memory sizes of modern GPUs that we have used to test our simulation on range from 800Mbytes to 6Gbytes. They provide enough memory that we can store large systems efficiently on device memory instead of swapping section in and out from host memory as may be required for really large systems in many dimensions.

The Random number generator used on this simulation is the Ran2 generator presented in The book Numerical Recipes. This has been proven to be a reliable source of random numbers and has passed all of the die hard tests for random numbers. we assign an integer in global memory for each random number instance for each thread to access based on its thread ID.

## Results and Discussion

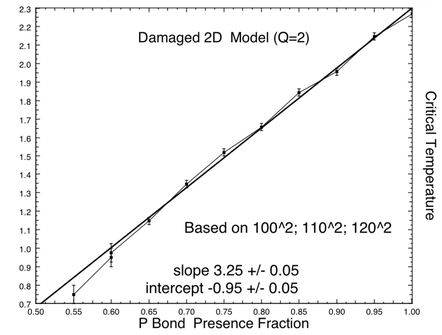


**FIGURE 3: Dependance of critical temperature of the damage parameter $p$ for a $Q = 2$ system**

Figure shows how the critical temperature varies when parameter $p$ is lowered to introduce damage into the $d = 2, Q = 2$ model (Ising case) in two dimensions. This data are preliminary, being based upon small system sized runs. The method described was applied using three different sized model simulations and recording the intercept of the three independent Binder cumulant plots. As can be seen, introducing damage lowers the critical temperature almost linearly. Near the bond-percolation threshold of $p = 0.5$ the system is no longer represented by a single connected cluster of spin sites and the phase transitional behaviour disappears.

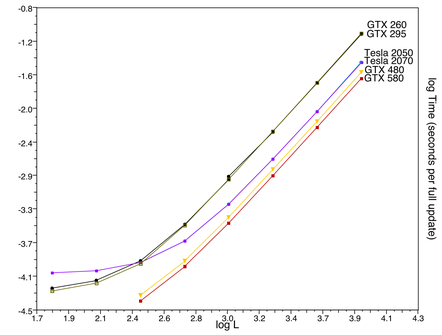


**FIGURE 4: Graph of Log L vs Log Time in seconds per full system update in two - dimensions when Q = 3 and D = 2 and no broken bonds.**

The time taken function for a full update of the Potts system is approximately linear in the number of spins providing $N$ is big enough and shows that the new architecture of the the GTX480, GTX580 and the Tesla cards produce a significant improvement over the older cards. This is attributable to the new memory formats including a L1 and L2 cache, larger multiprocessors and more cores. The GTX260 and the GTX 290 have the same performance and this is due to the higher clock speed of the 260 compensating for its smaller number of processors. The difference between the two generations of GPU are not as significant as has been seen in different types of simulation and we suspect that this is due to the highly efficient memory management with crinkling of the checkerboard update pattern. Our data for spin-flip times show little variation to once fully utilizing all cores this is in contrast to the results presented in other work.

For full paper and references please see the full text at http://complexity.massey.ac.nz/