q-state Potts model metastability study using optimized GPU-based Monte Carlo algorithms

Ezequiel E. Ferrero[1,3,‡], Juan P. De Francesco[2], Nicolás Wolovick[2], Sergio A. Cannas[2,3]

[1]*Centro Atómico Bariloche (CAB-CNEA), 8400 San Carlos de Bariloche, Argentina.*
[2]*Facultad de Matemática, Astronomía y Física; Universidad Nacional de Córdoba, Ciudad Universitaria, 5000 Córdoba, Argentina*
[3]*Consejo Nacional de Investigaciones Científicas y Técnicas* (CONICET)

‡ferrero@cab.cnea.gov.ar

Journal Ref.:
Comp. Phys. Commun. 183 (2012) 1578–1587

We implemented a GPU-based parallel code to perform Monte Carlo simulations of the two-dimensional q-state Potts model. The algorithm is based on a checkerboard update scheme and assigns independent random number generators to each thread. The implementation allows to simulate systems up to ~10^9 spins with an average time per spin flip of 0.147 ns on the fastest GPU card tested, representing a speedup up to 155×, compared with an optimized serial code running on a high-end CPU. The possibility of performing high speed simulations at large enough system sizes allowed us to provide a positive numerical evidence about the existence of metastability on very large systems based on Binder's criterion [1], namely, on the existence or not of specific heat singularities at spinodal temperatures different of the transition one.

The ferromagnetic q-states Potts model

$$\mathcal{H}_P = -J\sum_{(i,j)}\delta_K(s_i,s_j) \qquad s_i = 1,2,\ldots,q \qquad J_P > 0$$

$\delta_K(n_i,n_j)$ is the Kronecker delta

Order-disorder phase transition at exactly known temperature [2] in 2D: $\frac{k_B T_c}{J_P} = \frac{1}{\ln(1+\sqrt{q})}$

The transition is **2nd** order for **q=2,3,4** and **1st** order for ***q*>4**.

We choose $k_B/J = 1$

Considering the energy per spin $e = \langle H\rangle/N$ in the thermodynamic limit, the latent heat for $q > 4$ is [3]

$$e_d - e_o = 2\left(1+\frac{1}{\sqrt{2}}\right)\tanh\frac{\Theta}{2}\prod_{n=1}^{\infty}(\tanh n\Theta)^2 \qquad \text{where } \Theta = \arccos\sqrt{q}/2 \text{ and}$$

$$e_d = \lim_{N\to\infty}\frac{1}{N}\lim_{T\to T_c^+}\langle H\rangle, \qquad e_o = \lim_{N\to\infty}\frac{1}{N}\lim_{T\to T_c^-}\langle H\rangle.$$

Also $e_d + e_o = -2(1+1/\sqrt{q})$ from which the individual values of e_d and e_o can be obtained [2].

The order parameter is defined as $m = \frac{q(N_{max}/N-1)}{q-1}$ $\qquad N_{max} = \max(N_1, N_2, \ldots, N_q)$

At the transition the jump in the order parameter (for $q > 4$) is given by [4]

$$\Delta m = 1 - q^{-1} - 3q^{-2} - 9q^{-3} - 27q^{-4} - \cdots.$$

Chequerdboard update scheme GPU implementation

The parallelization strategy is based on a checkerboard scheme (since interactions are limited to nearest neighbors). Blacks (Whites) cells are independent and can be updated in parallel. The code's main features are [5]:

- Implementing Multiply With Carry RNG with FRAMESxFRAMES/2 independent generators.
- Computing multiple outputs per *thread*. Two consecutive kernels (black/white) of typically 512×512/2 threads are launched.
- Comprising the remapping of a two-dimensional stencil of four points in order to save memory transfers.
- Encoding each spin in a byte, allowing simulations with q<256 and L^2 < *available RAM.*
- Using registers for RNG states.
- Implementing parallel sums (butterfly-like algorithm) for energy and magnetization calculations, uses *shared memory*.

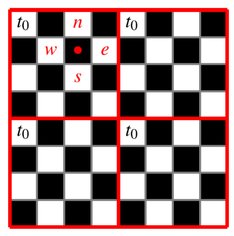

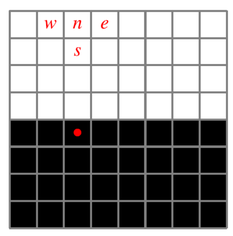

On the left: an 8×8 checkerboard framed in 4×4 (red marking), the cells updated by thread t_0 are singled out, we also marked the *north, east, south* and *west* neighbors of cell •. **On the right:** packed checkerboard showing first half of whites, where the neighboring cells *n*, *e*, *s*, *w* are marked, also in the second half of black cells • is singled out.

Performance

We compare our CUDA code running on NVIDIA GTX 280, 470 and 480 platforms, against a decent serial C code running on a high-end CPU core.

- GTX 280: spin flip times in the range [0.48 ns, 0.54 ns], which are 47× to 42× faster than those of the CPU code.
- GTX 470: [0.21 ns, 0.30 ns], speedup from 108× to 76×.
- GTX 480: [0.18 ns, 0.24 ns] speedups from 126× to 95×.
- GTX 480 (tuned version, CC 2.0) obtains 155× (0.147 ns).

There are two competing factors in the loop of the update kernel. One decreasing with *L* and is related to the amount of global memory movements per cell (RNG load/store global memory latency is distributed into more cells). The second factor is increasing with *L* and given by the inherent overhead incurred by a loop (comparison and branching).

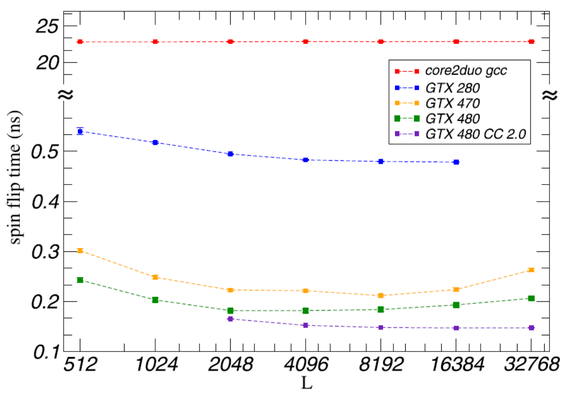

Spin flip time in nanoseconds vs. lattice size running on an Intel Core 2 Duo E8400@3.0 GHz CPU, and running on GTX 280, GTX 470 and GTX 480 NVIDIA GPUs. Averages are performed over 400 runs for the GPUs and 60 runs for the CPU. Error bars are smaller than symbol sizes when not showed.

Algorithm checking

In order to validate our CUDA code we run some typical simulations to measure well established results. We can see a very good agreement between data and exact results. It's worth noting that the data from the Table below is not the result of extrapolations of some finite size analysis, but the values from curves in the Figure on the right at the transition itself.

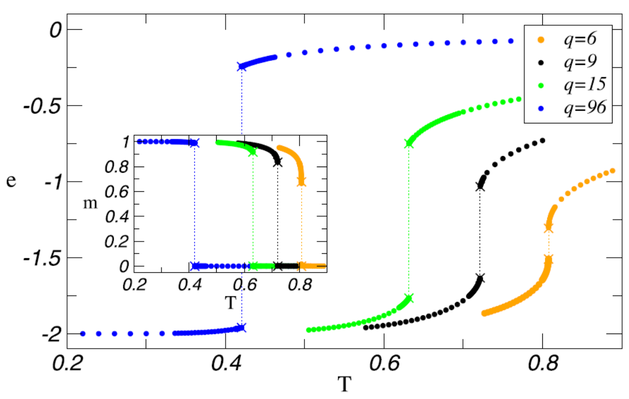

Equilibrium energy per spin *e* and magnetization *m* (inset) versus temperature for q = 9, 12, 15, 96. Exact values at the transition point are marked as crosses. Data comes from averages over 10 samples of linear system size L = 2048. Error bars are smaller than the symbol size.

Table. Comparison between calculated and known exact values of e_o , e_d, and Δm at the transition for different values of q. Results were obtained from averages over 10 samples of linear system size L = 2048 and equilibration and measurements times of at least 5×10^5 MCS each one.

q	$-e_o$		$-e_d$		Δm	
	Exact	Calculated	Exact	Calculated	Exact	Calculated
6	1.508980…	1.51(2)	1.307516…	1.306(1)	0.677083…	0.674(2)
9	1.633167…	1.6332(5)	1.033499…	1.0334(5)	0.834019…	0.8338(4)
15	1.765905…	1.7659(2)	0.750492…	0.7509(4)	0.916693…	0.9167(3)
96	1.960306…	1.96030(3)	0.243817…	0.24382(4)	0.989247…	0.98924(2)

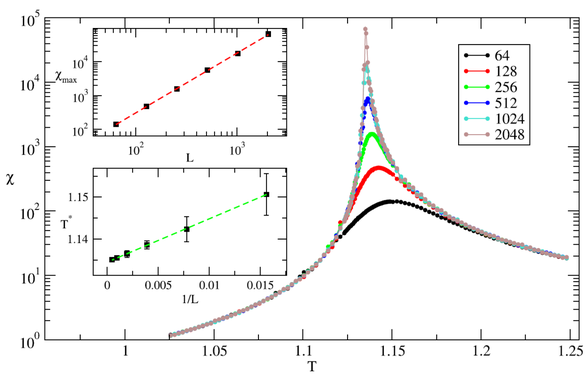

In addition, we test our code for the q = 2 (Ising) case. The figure shows the susceptibility of the order parameter.

The extrapolated value of the pseudo critical temperature $T^*(L)$ (defined as the location of the susceptibility maximum) for $L\to\infty$, 1.1345±0.0001, agrees with the exact value $T_c(q=2)$ = 1.134592... within an accuracy of the 0.009%.

The finite size scaling $\chi_{max} \sim L^{\gamma/\nu}$ is also satisfied.

Finite size scaling of the susceptibility for q=2. Main plot: χ as a function of temperature for different system linear sizes. Averages were taken over several samples ranging from 300 for small system sizes down to 50 and 15 for L = 1024 and L = 2048, respectively. We have used equally equilibration and measurement times of 2×10^5*MCS*, measuring quantities each 10*MCS*.

Metastability in the q-state Potts model and the Binder's criterion

Metastability is a verified fact in a finite system. It is known that below but close to T_c the system quickly relaxes to a disordered (paramagnetic) metastable state, with a life time that diverges as the quench temperature T approaches T_c.

A typical single-sample energy per spin vs time plot after a quench from infinite temperature to T=0.99T_t for q=24 and L=200. Snapshots at selected times are also shown with color coding for spin values.

Per spin energy histeresis cycles with temperature for a system size 2048×2048 and different temperature cooling/heating rates R.

Nevertheless, the existence of metastability in the thermodynamic limit is still an open problem. In Ref. [1] Prof. Binder studied static and dynamic critical behavior of the q-state Potts model using standard Monte Carlo procedures. When analyzing the q = 5 and 6 cases he realized that the transition is, in fact, a very weak first order transition, where pronounced "pseudocritical" phenomena occur. He studied system sizes from N=16×16 up to N=200×200, and observation times up to 10^4*MCS*. Within his analysis he was unable to distinguish between two different scenarios for the transition at q≥5 due to finite size effects taking place at the simulations. He proposed two self-avoiding possible scenarios for the transition.

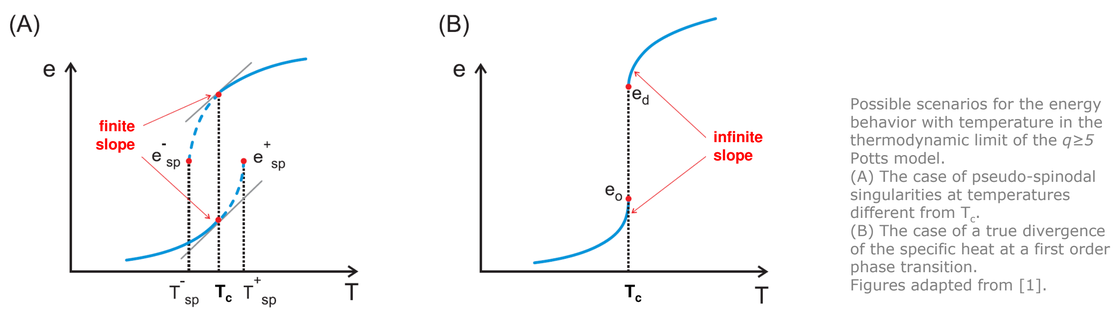

Possible scenarios for the energy behavior with temperature in the thermodynamic limit of the q≥5 Potts model.
(A) The case of pseudo-spinodal singularities at temperatures different from T_c.
(B) The case of a true divergence of the specific heat at a first order phase transition.
Figures adapted from [1].

If well defined metastable states occur, the energy could be represented in terms of a specific heat diverging at pseudospinodal temperatures T_{sp}^+ , T_{sp}^-

$$e_{T<T_c} = e_{sp}^- - A^-\left(1-T/T_{sp}^+\right)^{1-\alpha_-},$$
$$e_{T>T_c} = e_{sp}^+ - A^+\left(1-T_{sp}^-/T\right)^{1-\alpha_+}.$$

On the other hand, a power law divergence of the specific heat at T_c would imply the following behavior

$$e_{T<T_c} = e_o - A^-(1-T/T_c)^{1-\alpha_-},$$
$$e_{T>T_c} = e_d - A^+(1-T_c/T)^{1-\alpha_+}$$

$$\alpha_-, \alpha_+ > 0.$$

As pointed out by Binder [1], to observe the crossover (if it exists at all) a temperature resolution at least $\Delta T = T_c - T_{sp}^-$ for the high energy branch (or $\Delta T = T_{sp}^- - T_c$ for the low energy branch) is needed, where $\Delta T \equiv |T - Tc|$.

A numerical estimation of the lower spinodal temperature predicted by Short Time Dynamics [6] is given by

$$\frac{T_c - T_{sp}^-}{T_c} \simeq 0.0007\left(\ln(1+q-4)\right)^{2.81}$$

To attain the desired temperature resolution the system size has to be large enough, since finite size rounding errors are expected to decay as $1/L$. A rough estimation of the minimum size required to reduce the error $L \approx 1/\Delta T$ predicts the table on the right.

q	$\Delta T/T_t$	L
96	0.05	50
15	0.009	200
9	0.0035	400
6	0.00090	1400
5	0.00025	3400

Using our GPU-based code we measure equilibrium curves for $e_{T<Tc}$ ($e_{T>Tc}$) starting from an ordered (disordered) initial state and performing a cooling (heating) procedure approaching T_c *for large system sizes.*

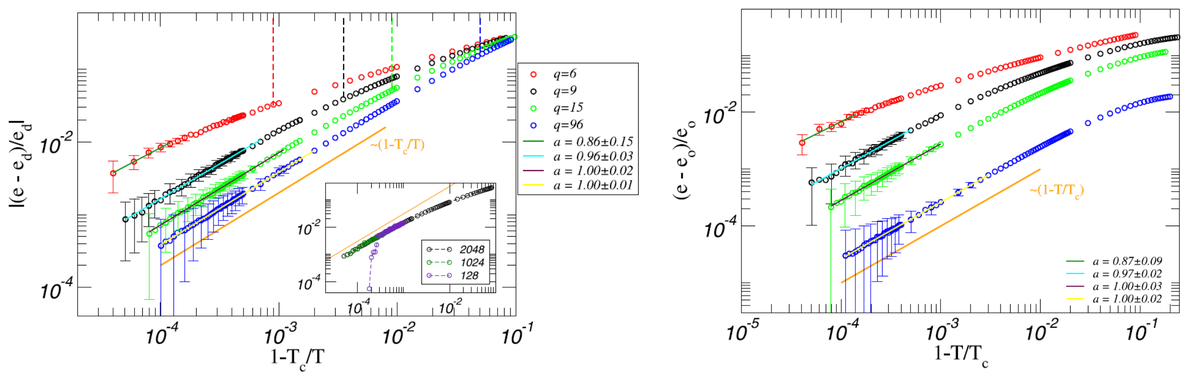

Log–log plot of energy differences versus temperatures $T>T_c$ ($T<T_c$), left (right) panel, for various q. Data correspond to averages over 20 samples of systems size L = 2048, equilibration times ranging from 5×10^4[*MCS*] to 2×10^5[*MCS*] and measurement times of 5 × 10^4[*MCS*], with sampling every 100[*MCS*]. Full color lines are power-law fits of the form $|(e - e_o)/e_o| = A(1-T_c/T)^a$ (left panel) and $(e-e_o)/e_o = A(1-T/T_c)^a$ (right panel). Resulting exponents a are showed in the labels. On the left panel, dashed vertical lines of different colors correspond to $T = T_c + \Delta T(q)$, with $\Delta T = T_c - T_{sp}^-$ estimated from the STD, the inset shows q=9 curves for different system sizes, the full orange curve indicates the slope 1.

In both figures a crossover of the curve's slope as we approach T_c can be observed for all values of q. Close enough to T_c, the curves for q=9,15,96 show exponents which are indistinguishable from 1, consistently with the existence of metastability and divergences at spinodal temperatures different from T_c, at least for q≥9.

Summary

- We implemented a CUDA-based parallel Monte Carlo algorithm to simulate the Statistical Mechanics of the q-state Potts model. The present algorithm allows the simulation of very large systems in very short times, namely ~10^9 spins with an average time per *MCS* of 0.15s.
- The key factors of the code are the per-thread independent RNG that is fast and takes only a few registers, the framing scheme that increases the amount of computation done by each thread and at the same time it bounds the number of independent RNG needed, and finally the cell-packing mapping that orders the memory access.
- We studied the metastability problem in the two-dimensional system based on Binder's criterion, namely, on the existence or not of specific heat singularities at spinodal temperatures different from the transition one (but very close to).
- Our results provide a positive numerical evidence about the existence of metastability on very large systems.
- Our CUDA implementation of the 2D q-state Potts model allow for feasible generalizations. In particular an extension of our code, developed for a recent study [7] in the q-state Potts glass, can be find in [8] .

References:

[1] K. Binder, *J. Stat. Phys*. **24**, 69 (1981).
[2] T. Kihara, Y. Midzuno, T. Shizume, *J. Phys. Soc. Japan.* **9**, 681 (1954).
[3] R.J. Baxter, *J. Phys. C* **6**, L445 (1973).
[4] R.J. Baxter, *J. Phys. A* **15**, 3329 (1982).
[5] *potts3* is available for download and use GNU GPL 3.0 at https://bitbucket.org/ezeferrero/potts
[6] E.S. Loscar, E.E. Ferrero, T.S. Grigera, S.A. Cannas, *J. Chem. Phys.* **131**, 024120 (2009).
[7] E.E. Ferrero, F. Romá, S. Bustingorry, P. M. Gleiser, *Phys. Rev. E* **86**, 031121 (2012).
[8] https://bitbucket.org/ezeferrero/potts-glass