Microcanonical simulated annealing: Massively parallel Monte Carlo simulations with sporadic random-number generation

Massimo Bernaschi[2], Claudio Chilin[4,5], Luis Antonio Fernandez[4], Isidoro González-Adalid Pemartin[2], Enzo Marinari[5], Víctor Martín-Mayor[4], Giorgio Parisi[4], Federico Ricci-Tersenghi[4], Juan Jesús Ruiz-Lorenzo[6], **David Yllanes**[1,3]

[1]Universidad de Zaragoza, Zaragoza, Spain, [2]Istituto per le Applicazioni del Calcolo, CNR, Rome, Italy, [3]Fundación ARAID, Zaragoza, Spain

Optimisation problems are ubiquitous —both in everyday life and in a plethora of applications in science and technology— and Monte Caralo methods are a precious tool to solve them. The task can often be expressed in terms of minimising a cost function, for example, by incrementally reducing the "temperature" control parameter, starting from a very high value and going down close to zero (where the "Boltzmann" measure concentrates at the minimum of the cost function). This is the well-known simulated annealing (SA) algorithm [1]. Spin glasses [2] play a central role in this realm as paradigmatic complex systems, since the computation of their ground state is an NP-complete problem.

High-perfomance computations of complex systems require massively parallel codes, which are are plagued by a major drawback: they are extremely greedy for (pseudo) random numbers. A possible cure is the use of microcanonical methods, which minimise random-number generation. One example is the Creutz algorithm [3], which uses an ensemble that includes a single auxiliary variable (a demon), but which leads to significant slowing down for spin glasses [4]. Alternatives that consider an extensive number of demons, such as the Lustig method [5], present different drawbacks: (i) unsuitability for parallel computations, (ii) loss of frugality in the use of random numbers and (iii) convergence to the microcanonical, rather than canonical, equilibrium ensemble.

We present a formalism, Microcanonical Simulated Annealing or MicSA, which cures these problems. First, we keep the demons as explicit dynamical variables, rather than integrating them out, which addresses problems (i) and (ii). Second, we generalise the SA algorithm to the microcanonical context. More precisely, we cool the system down (lower its energy) just by refreshing the probability distribution of the demons at a logarithmic time rate. This is the only part of the algorithm that involves random numbers and its computational cost is negligible. We hence obtain results indistinguishable (after a simple time rescaling) from those produced using a canonical simulation on the Janus II special-purpose computer [6]. This is shown in the figure, which plots the excess energy per spin over its equilibrium value for a large spin-glass sample as a function of time, comparing MicSA dynamics at time t with Metropolis dynamics at time kt.

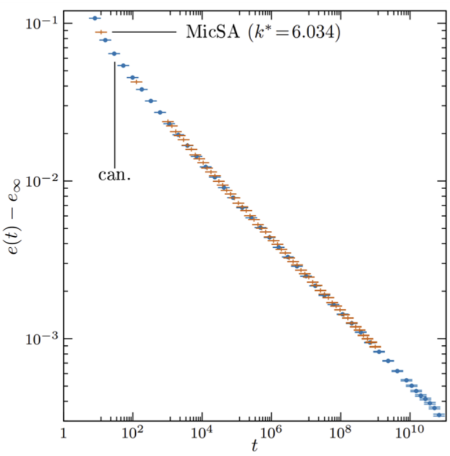

Our algorithm [7] is fully adapted to massively parallel computations and, in fact, outperforms conventional algorithms on latest-generation GPUs. Since, in canonical methods, the total fraction of computer time dedicated to random-number generation increases as the hardware grows more sophisticated, we expect that MicSA will be especially efficient for special-purpose computing platforms.

References:
[1] S. Kirkpatrick, C. D. Gelatt, M. P. Vecchi. Science 220 (4598) (1983) 671–680.
[2] E. D. Dahlberg et al., Rev. Mod. Phys. 97, 045005 (2025).
[3] M. Creutz, Phys. Rev. Lett. 50 (1983) 1411.
[4] J.J. Ruiz-Lorenzo, C. Ullod, Comp. Phys. Comm. 125 (1) (2000) 210–220.
[5] R. Lustig, J. Chem. Phys. 109 (1998) 8816.
[6] Janus Collabortion, Comp. Phys. Comm 185 (2014) 550–559.
[7] M. Bernaschi et al., arXiv:2506.16240.