Supercomputing, computational physics and phase transitions

Lev Shchur

Laboratory for Computational Physics
NRU Higher School of Economics
and
Landau Institute for Theoretical Physics RAS

The work was carried out as part of the fundamental research program
at the National Research University Higher School of Economics

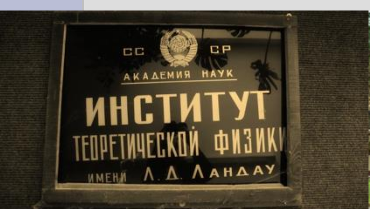

RSCDays2025, MSU, Moscow, 29.09.2025

Supercomputing, computational physics and phase transitions

- Modern computing systems are heterogeneous in terms of processor types
- They consist of a large number of nodes containing various types of computing elements

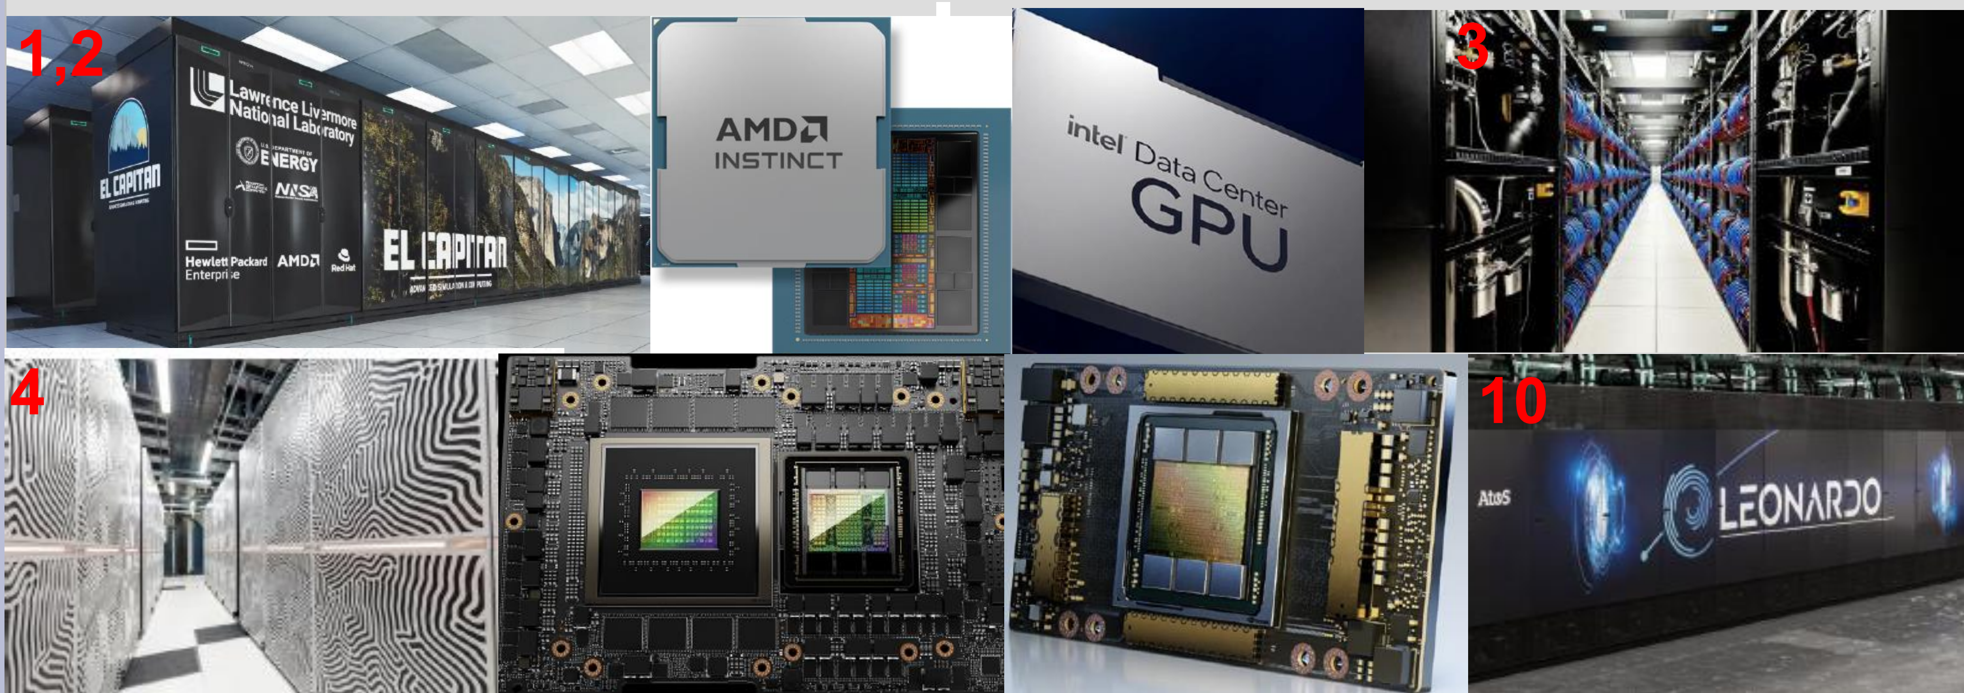

Population annealing methods using hybrid parallel computing architecture

- Most processor devices contain nested computing elements
- This is the path to further following Moore's Law

- **Developing software capable of loading all computing elements with useful tasks is a rather non-trivial task**

There are two options:

- Develop a universal, potentially fully scalable approach
- Invent an algorithm that will be effective for a specific, albeit broad, set of tasks

Population annealing algorithm

1. The population annealing (PA) algorithm may be suitable for study systems with rough free energy landscapes (spin glasses, molecular dynamics, polymer folding, optimization problems, …).
2. PA combines the power of well-known efficient algorithms - simulated annealing, Boltzmann weighted differential reproduction, and sequential Monte Carlo processes.
3. Capable of bringing the replica population to equilibrium even in low temperature regions.
4. It provides an ideal assessment of free energy and entropy.
5. Enables efficient parallel implementation – multiple replicas.

Hukushima & Iba, AIP Conf. Proc. (2003) - idea
Machta, Phys. Rev. E. (2010) – CPU
Barash, Weigel, Borovsky, Janke & LS, Comp. Phys. Comm. (2017) – GPU
Russkov, Chulkevich & LS, Comp. Phys. Comm. (2021) – GPU/MPI
Weigel, Barash, Janke & LS, Phys. Rev. E (2021) – detailed analysis of PopAnn

Population annealing algorithm

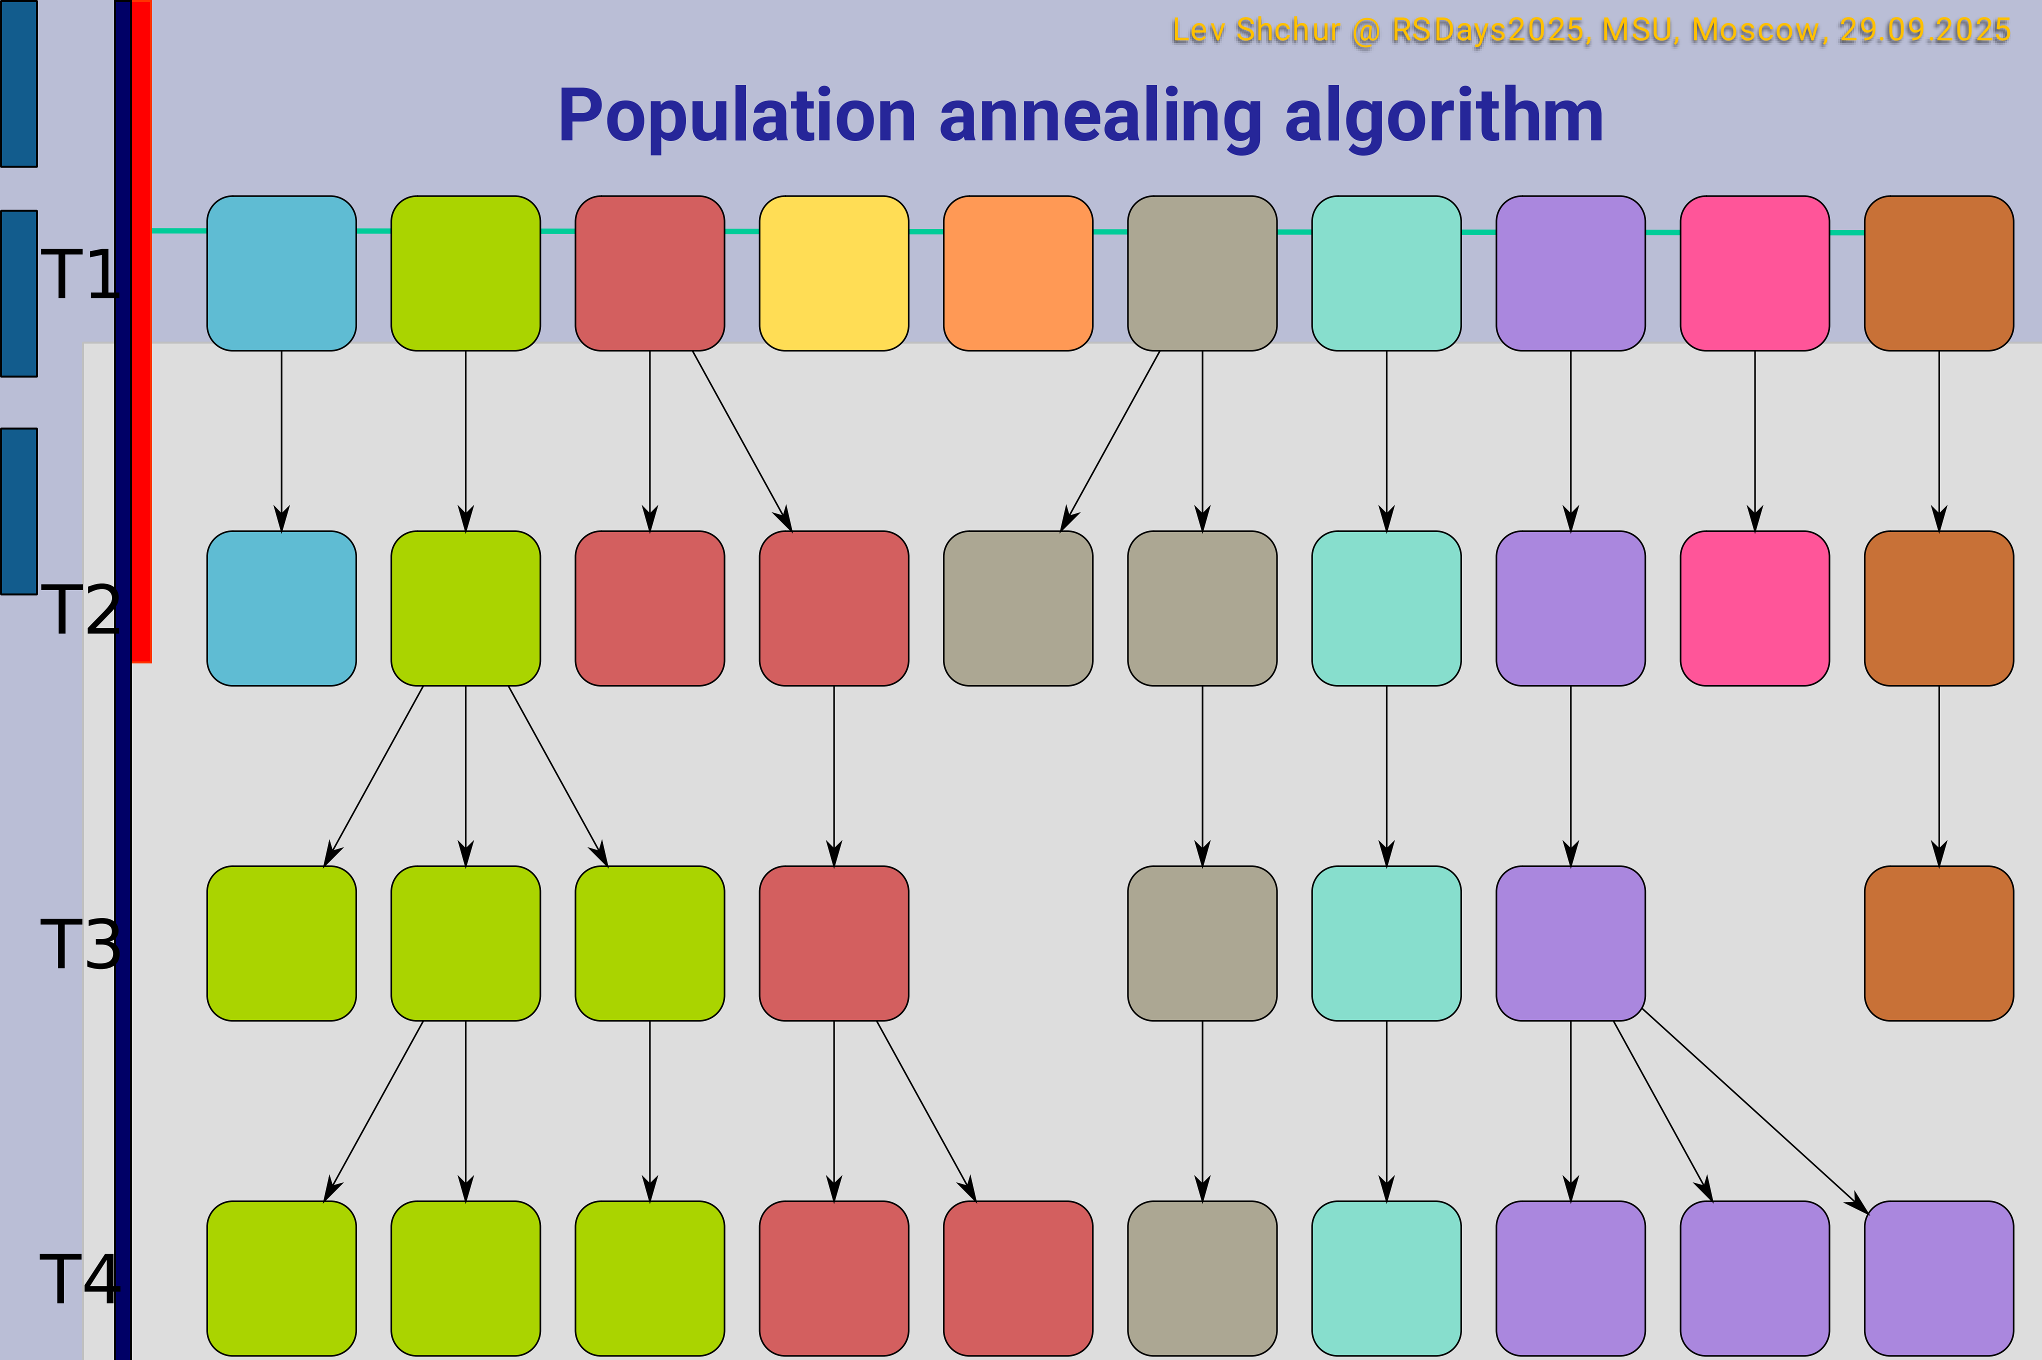

PopAnn, CUDA: frustrated Ising ferromagnet on the stacked triangular lattice

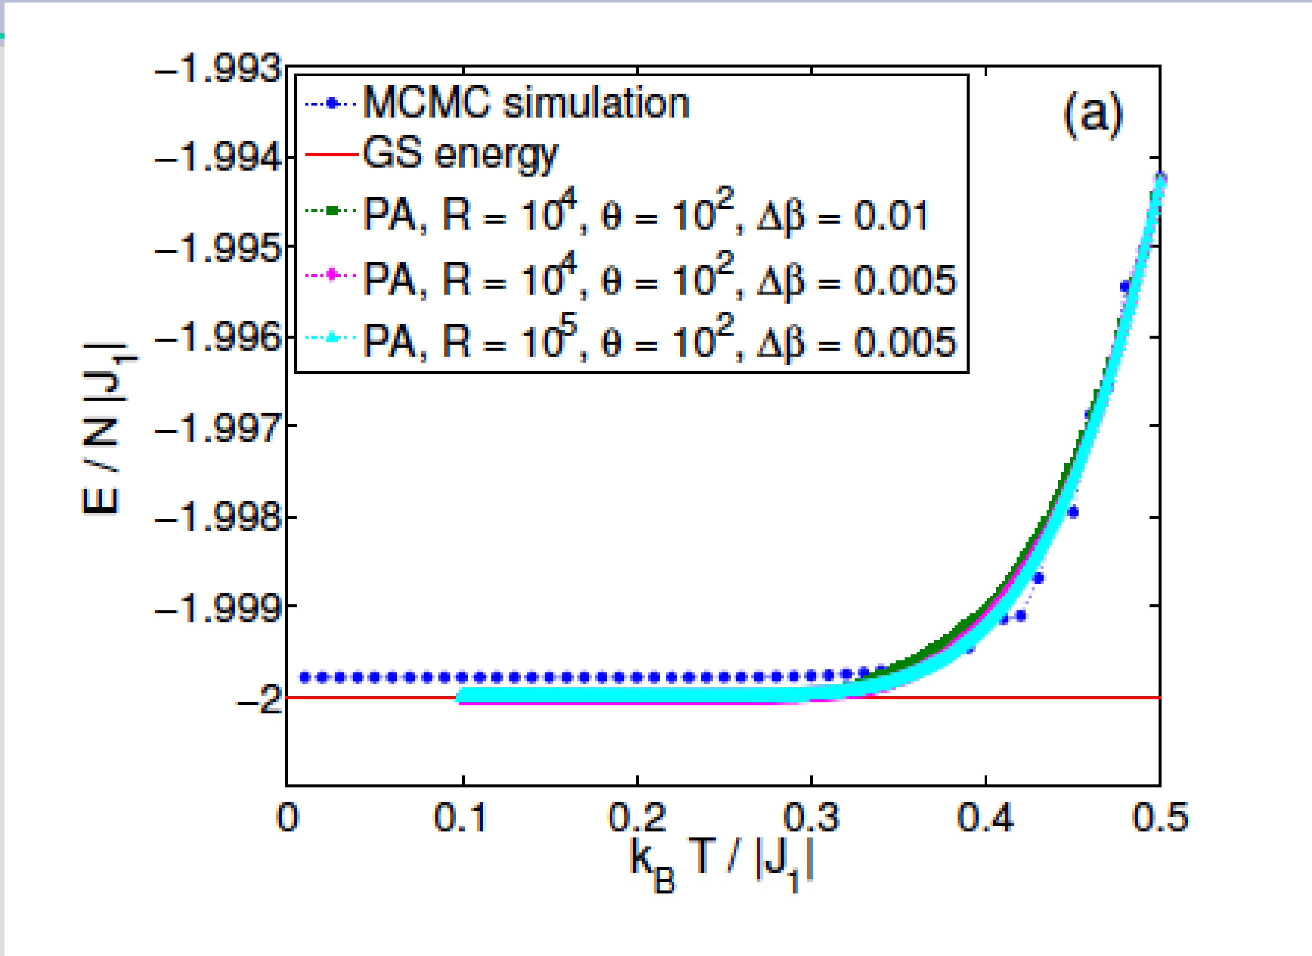

EPJ Web of Conferences **108**, 02016 (2016)

Potts model with Population Annealing

Cooling

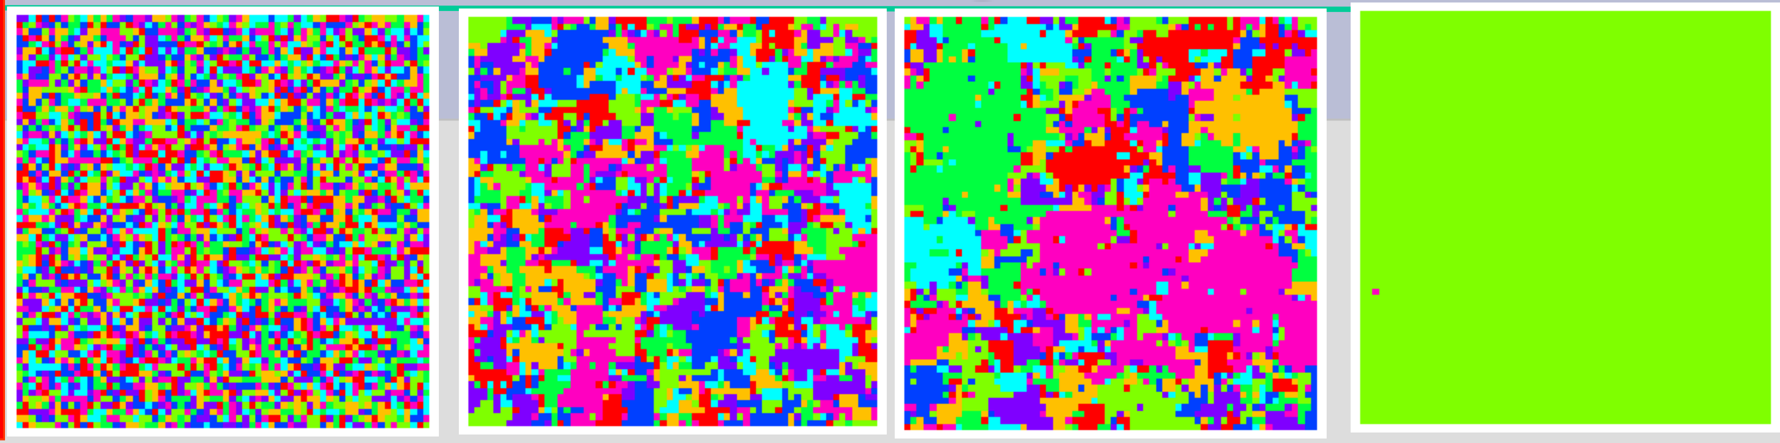

Heating

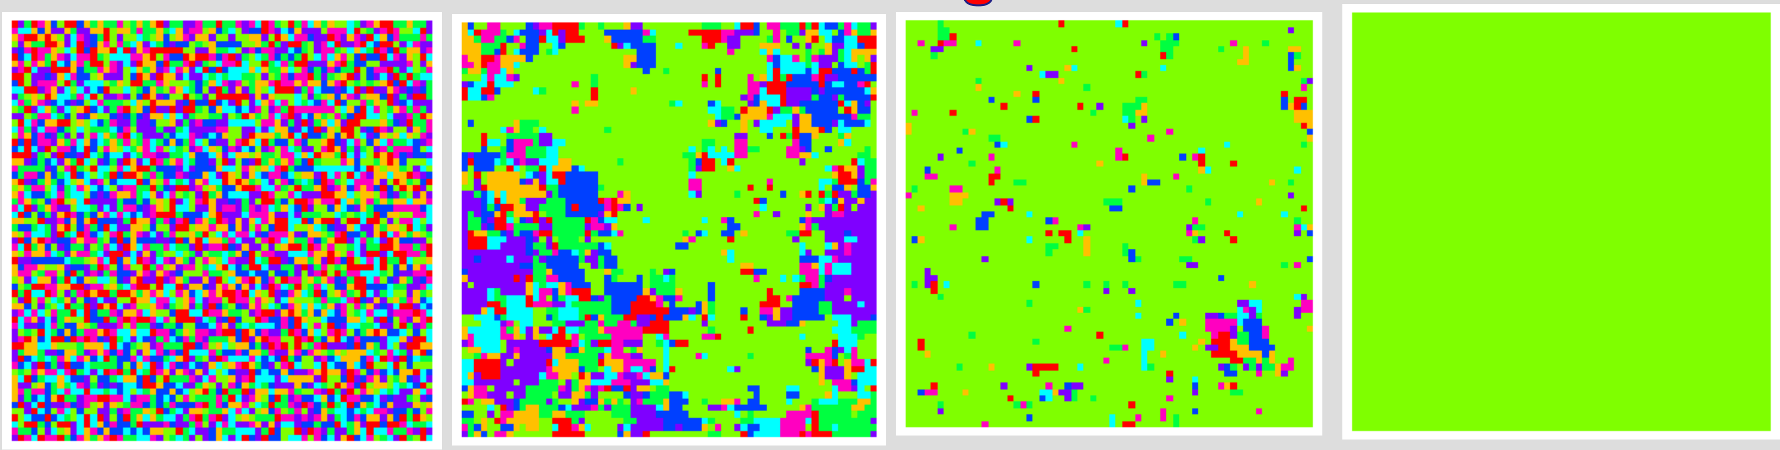

Barash, Weigel, Janke & LS, EPJ (2017)

Potts model with Population Annealing Hysteresis

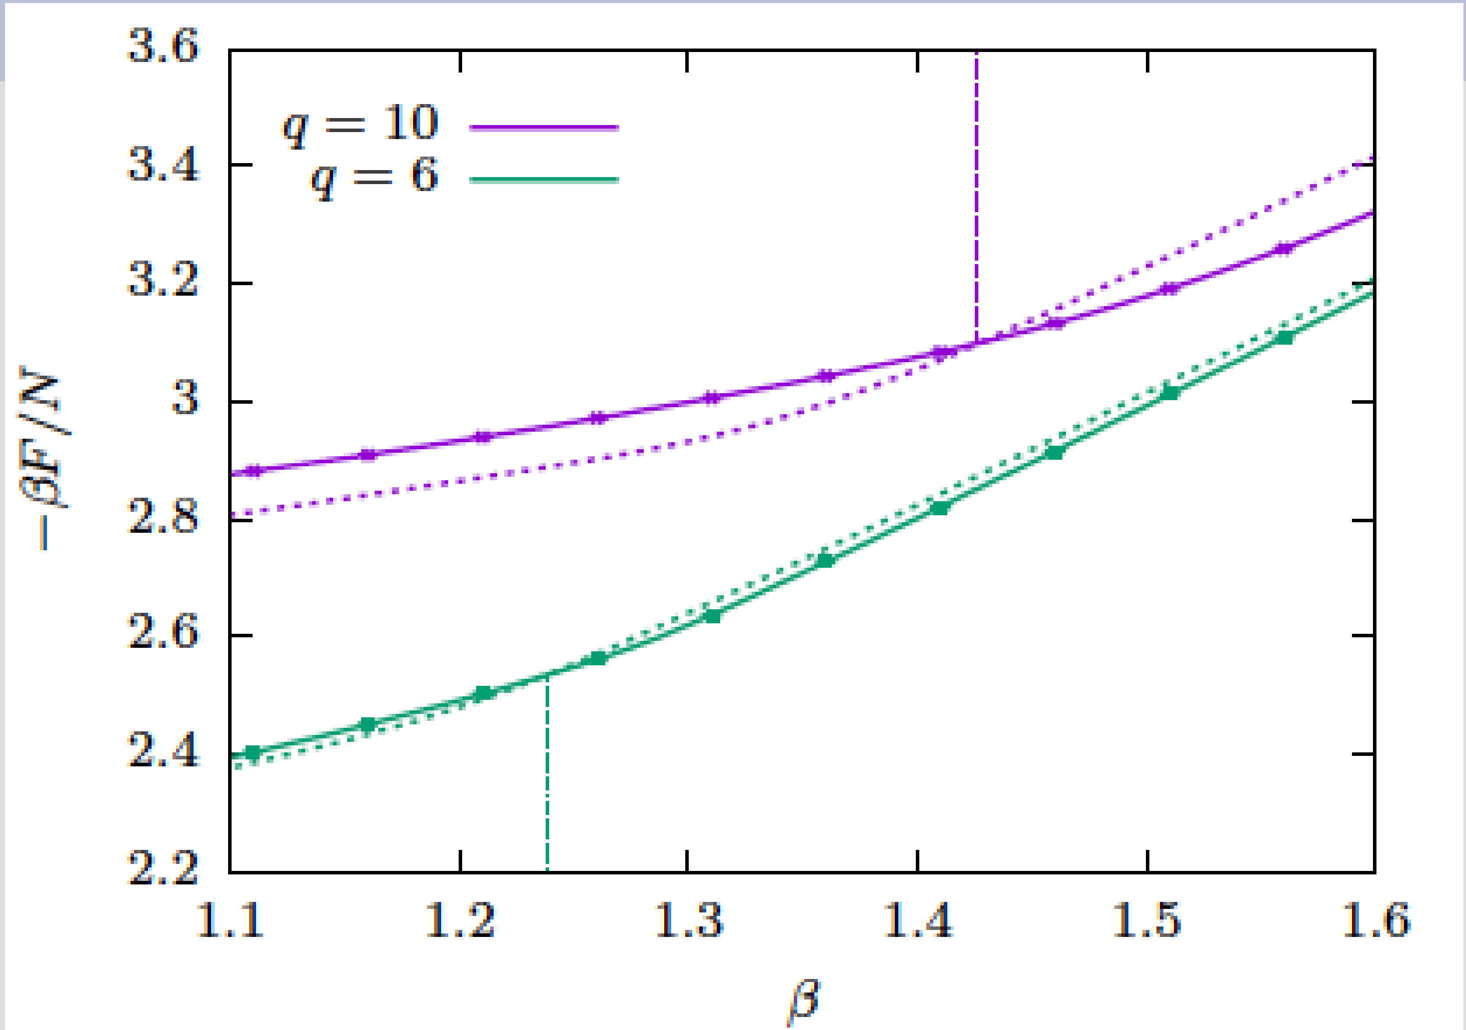

Population Annealing - Simulations of Biopolymers

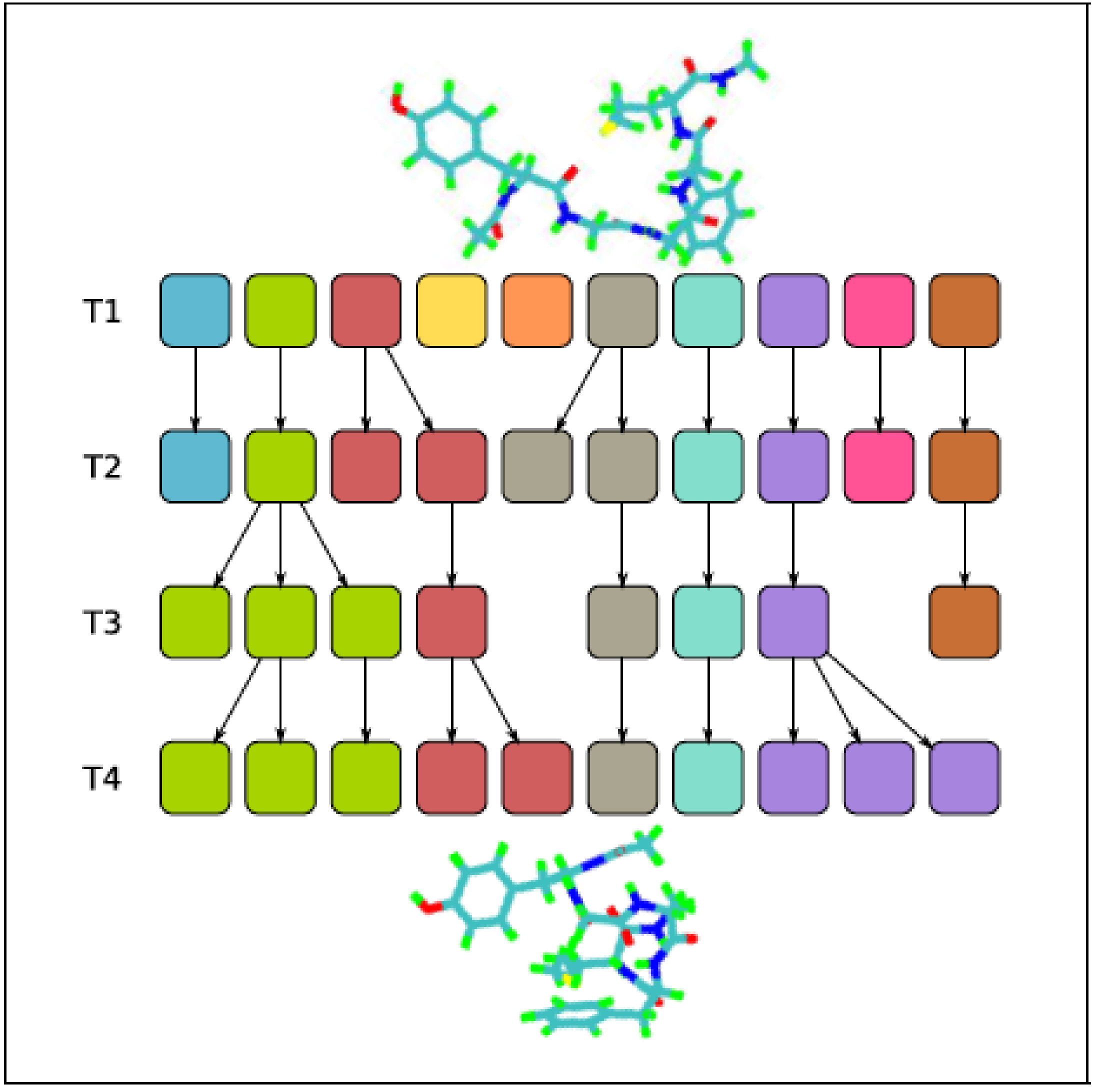

Population Annealing - Simulations of Biopolymers

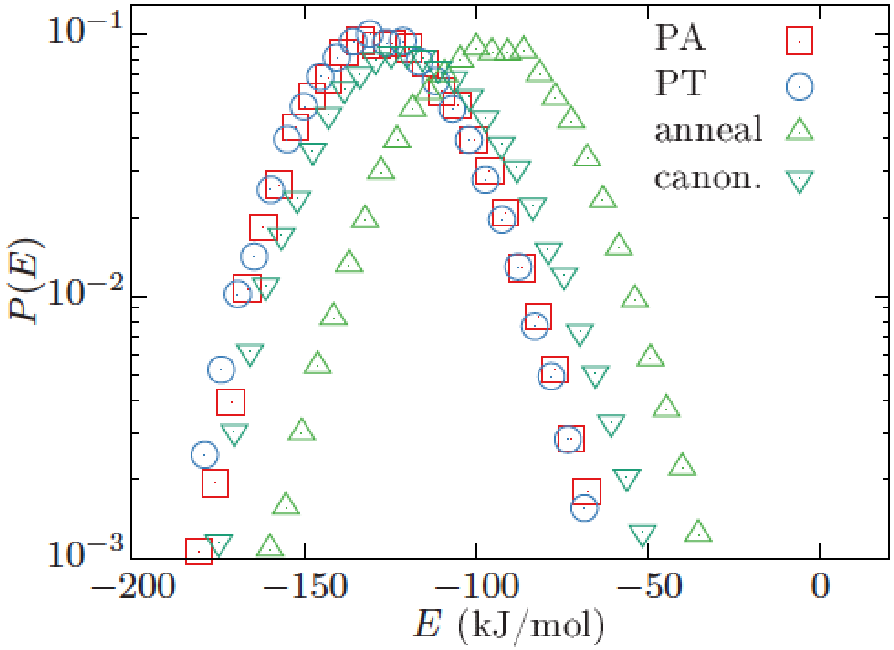

Figure 3: Energy histograms at the lowest temperature, $T = 200$ K, as obtained from the population annealing (PA), parallel tempering (PT), population annealing without resampling (“anneal”), and canonical (“canon.”) simulations, respectively.

Simulations
on the full-scale of cHARISMa (HSE)

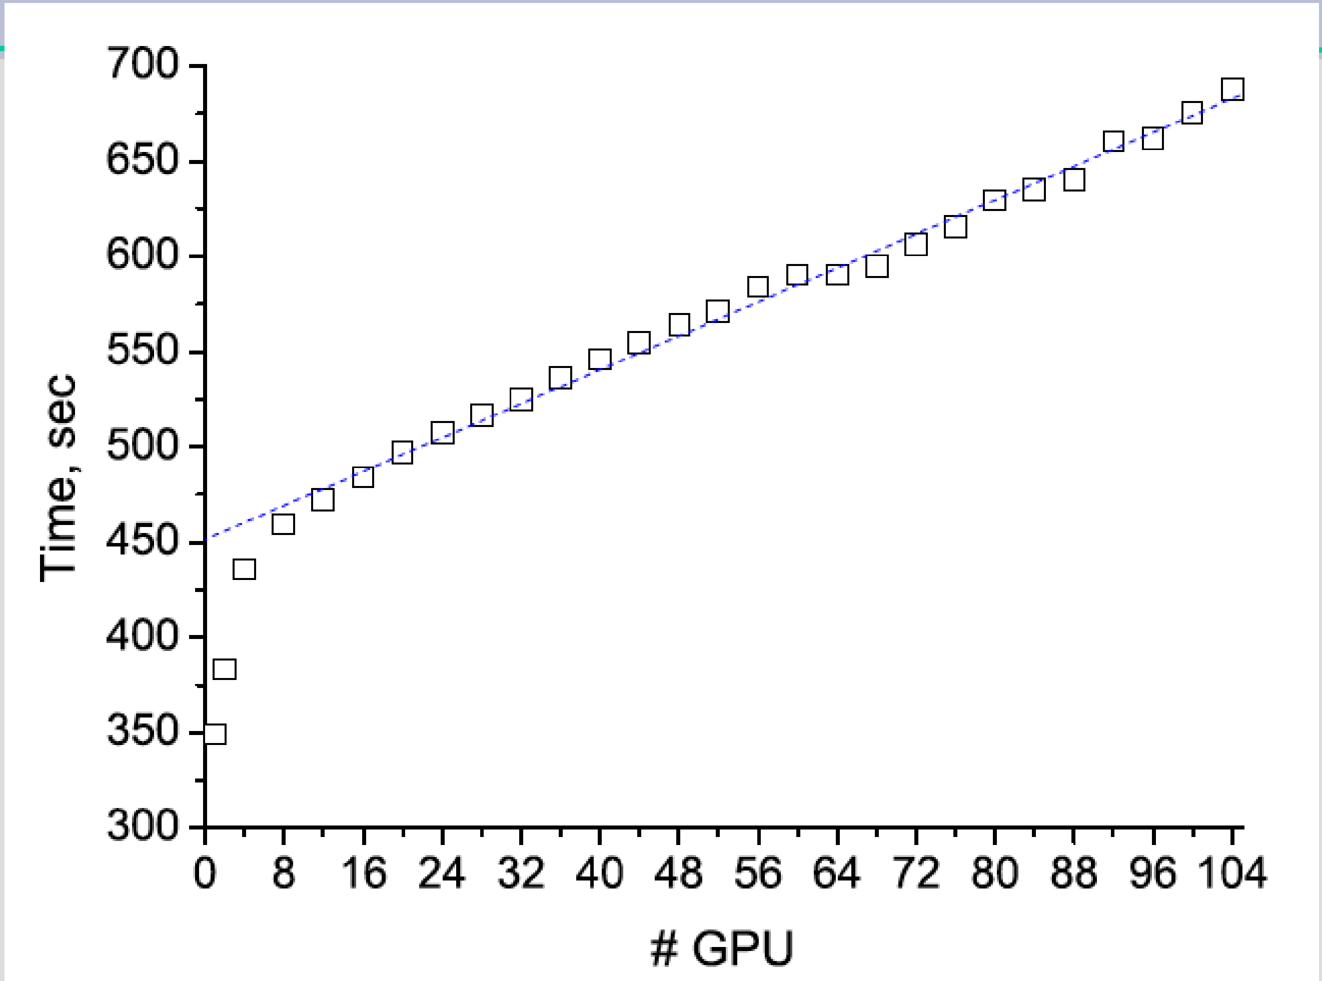

26 nodes: 2x Intel Xeon Gold 6152 + 4x V100

Simulations
on the full-scale of cHARISMa (HSE)

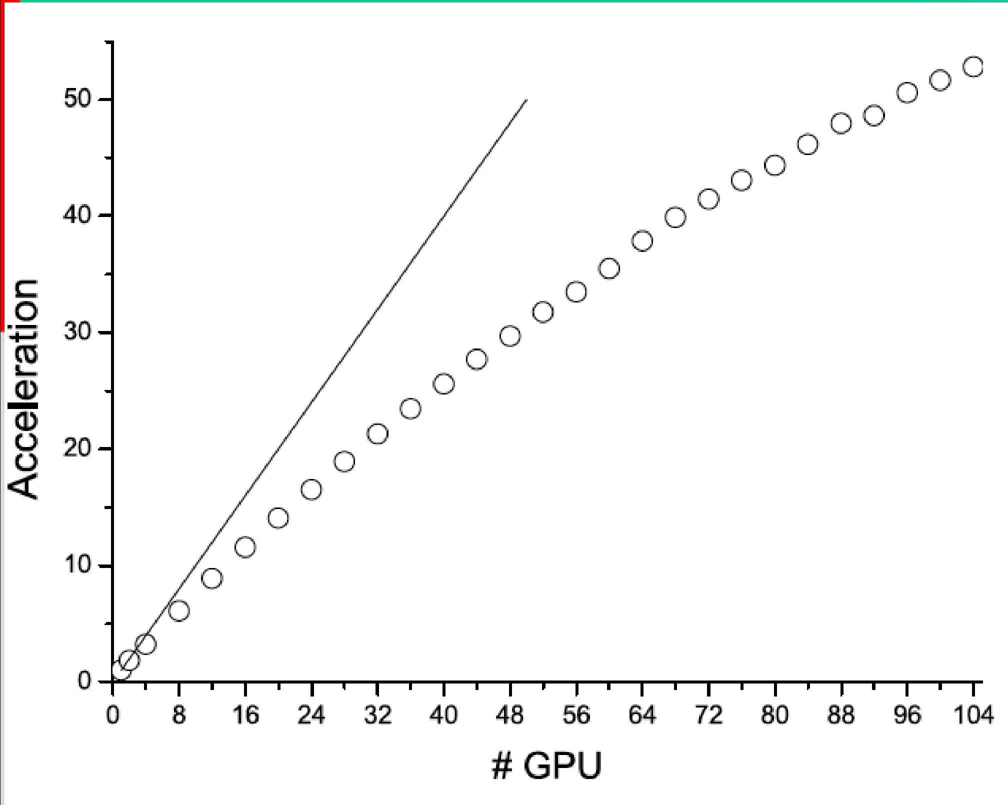

more than 2 million replicas in parallel

26 nodes: 2x Intel Xeon Gold 6152 + 4x V100

MPI/CUDA PA implementation

Technical problems solved:

- Load balancing of all GPU nodes with an approximately same number of replicas;
- Blocks of replicas = 2000;
- 10 blocks of replicas per GPU;

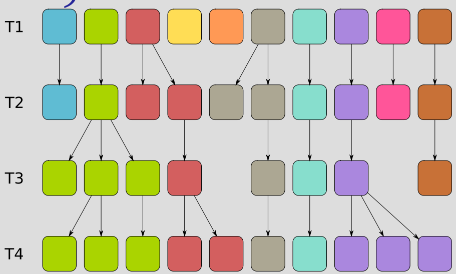

- Minimization of the extensive memory exchange between nodes;
- Approx. 2.5 million replicas in one run
  - technically impossible with the smaller clusters!

Microcanonical population annealing algorithm

- Initial configuration of a large number of replicas
- Setting the energy ceiling (floor)
- Transition from one energy level to another
- Monte Carlo at infinite temperature
- Calculating the proportion of replicas at the upper (lower) energy level
- Entropy calculation → Density of States → Functions of interest

Rose & Machta, Phys. Rev. E. (2019) – GPU on the base of PA-2017
Mozolenko & LS, Phys. Rev. E. (2024) – modified GPU on the base of Rose-Machta
Mozolenko, Fadeeva & LS, Phys. Rev. E. (2024) – comparison with Wang-Landau
Sukhoverkhova, Mozolenko & LS, Phys. Rev. E. (2025) – combined with ML

Microcanonical population annealing algorithm

$$S^c(E) = \ln(\epsilon(E)) + \sum_{E'>E} \ln(1-\epsilon(E')),$$

$$S^f(E) = \ln(\epsilon(E)) + \sum_{E'<E} \ln(1-\epsilon(E')).$$

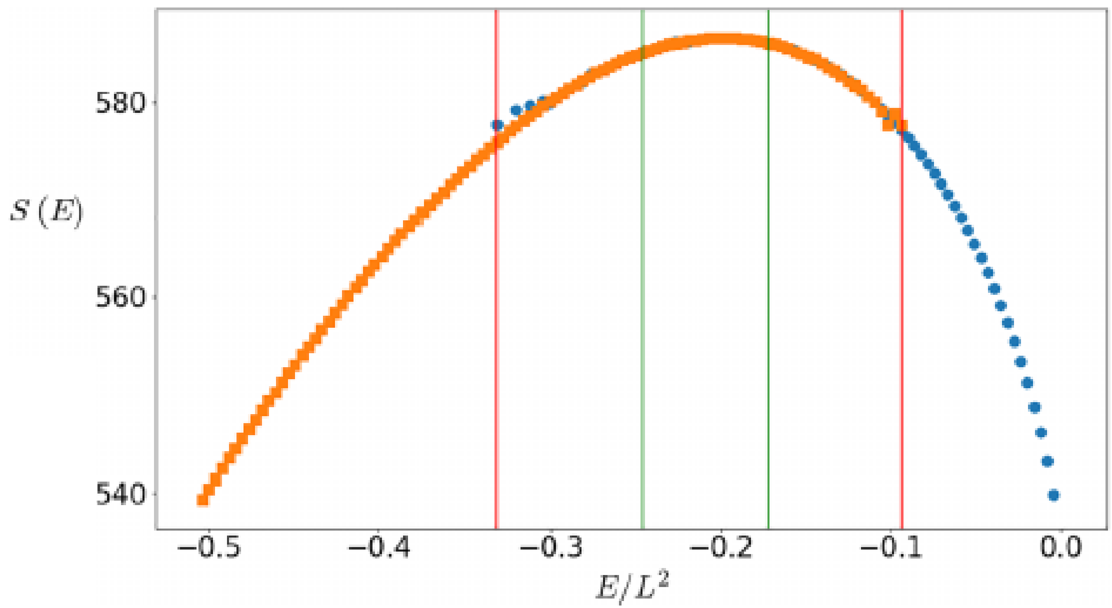

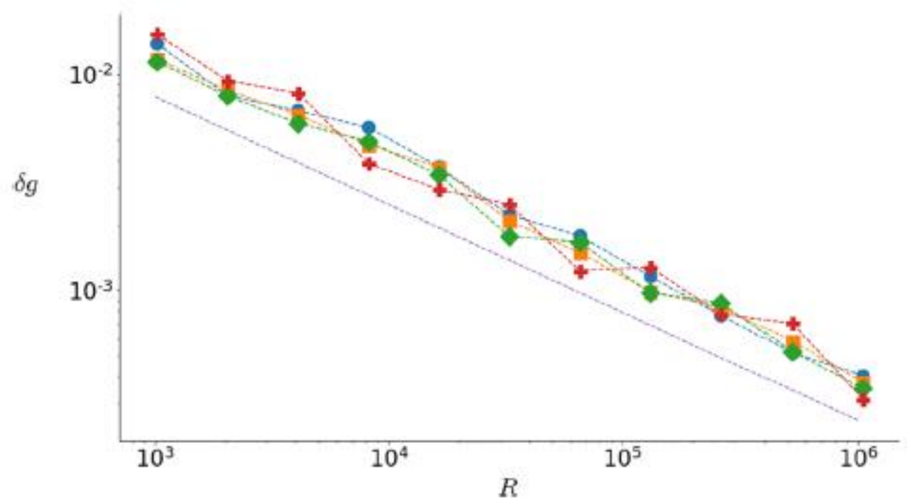

Mozolenko & LS, Phys. Rev. E. 109 (2024)

Microcanonical population annealing algorithm

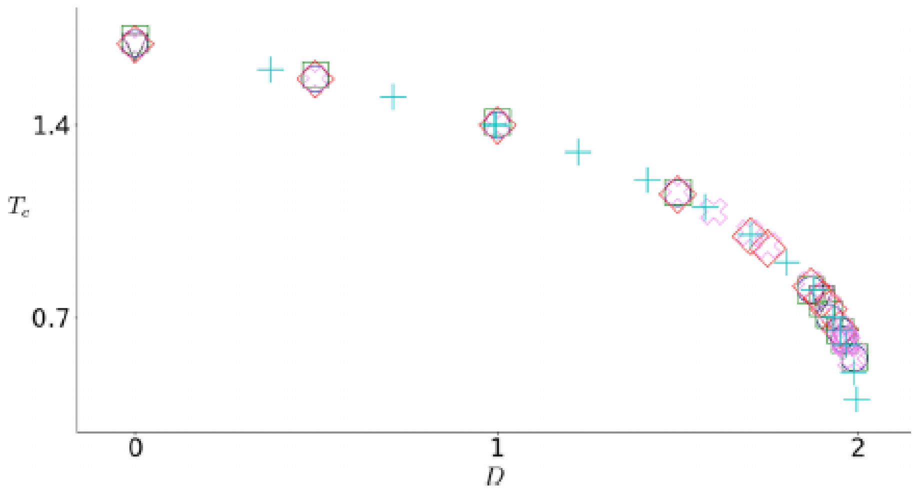

FIG. 1. Phase diagram obtained by different methods: transfer matrix [30] (blue circles), Monte Carlo [31] (black triangles), Wang-Landau [32] (green squares), high-energy and low-energy expansions [5,33] (red diamonds), microcanonical algorithm [34] (cyan pluses), and microcanonical population annealing (current work, violet crosses). The error bars are much smaller than the symbols.

The Blume-Capel model [1,2] in the absence of a magnetic field is described by a Hamiltonian,

$$H = -J \sum_{\langle i,j \rangle} \sigma_i \sigma_j + \Delta \sum_i \sigma_i^2, \quad (1)$$

Mozolenko & LS, Phys. Rev. E. (2024)

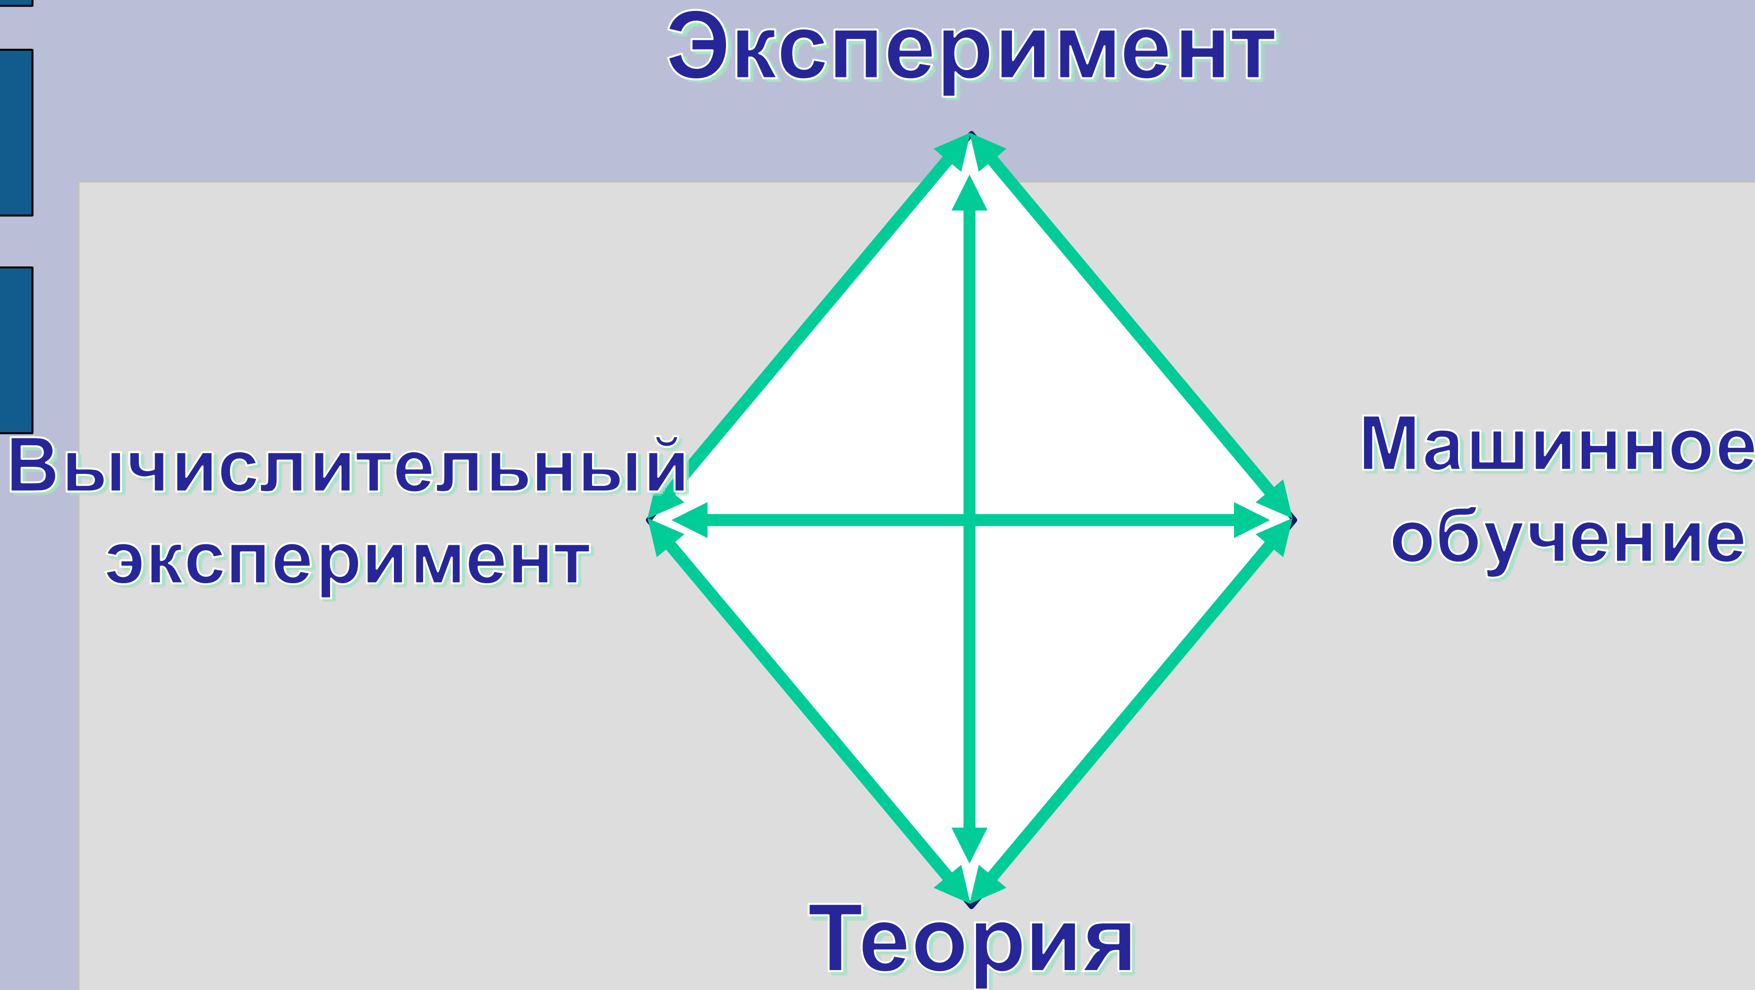
Эксперимент
Вычислительный эксперимент
Машинное обучение
Теория

Microcanonical population annealing algorithm and ternary classification

- 2^{17} replicas by microcanonical PA
- 2^{13} uncorrelated replicas at each energy level
- Ternary classification with the value of disordered energy and ordered energy by supervised learning
- Probability of ordered phase, mixed phase, and disordered phase

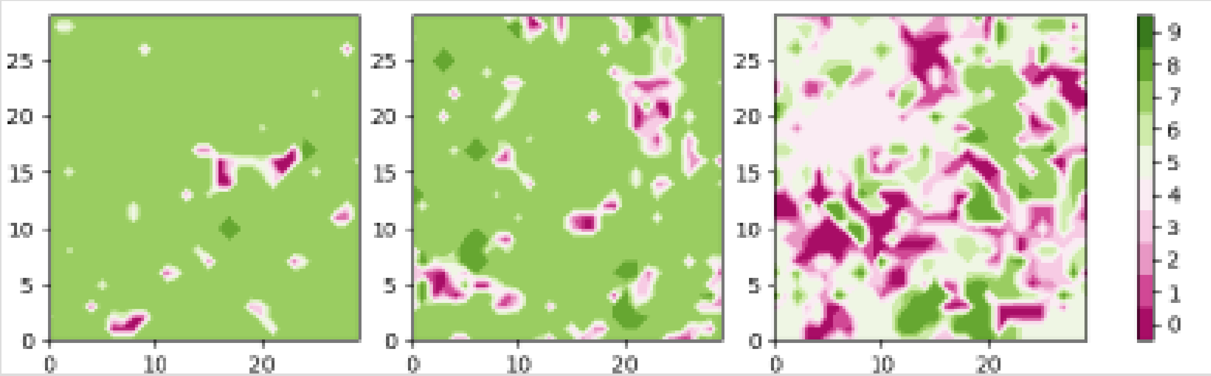

OS if $e < e_o$;

CS if $e_o < e < e_d$

DS if $e > e_d$.

Sukhoverkhova, Mozolenko & LS, Phase probabilities in first-order transitions using machine learning, Phys. Rev. E. (2025)

Microcanonical population annealing algorithm and ternary classification

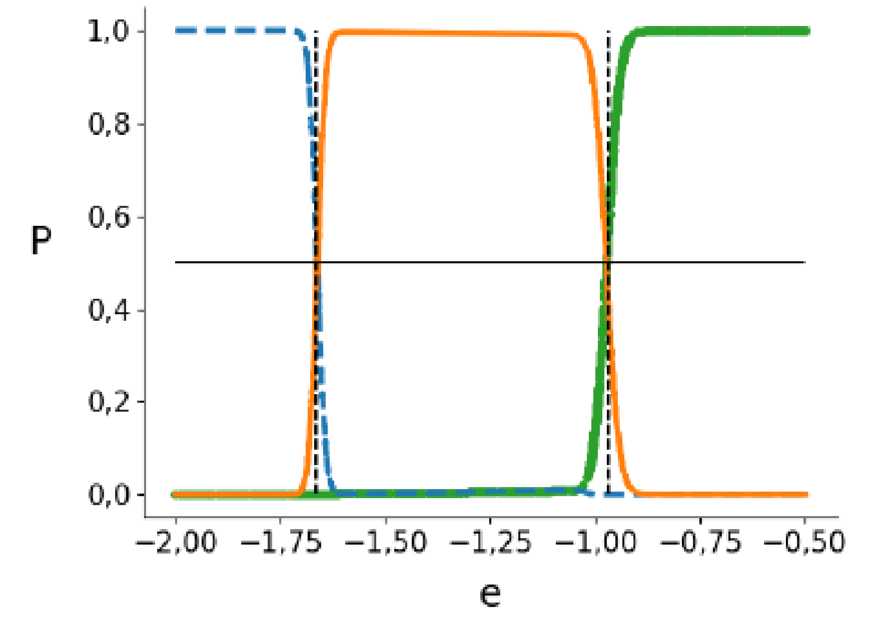

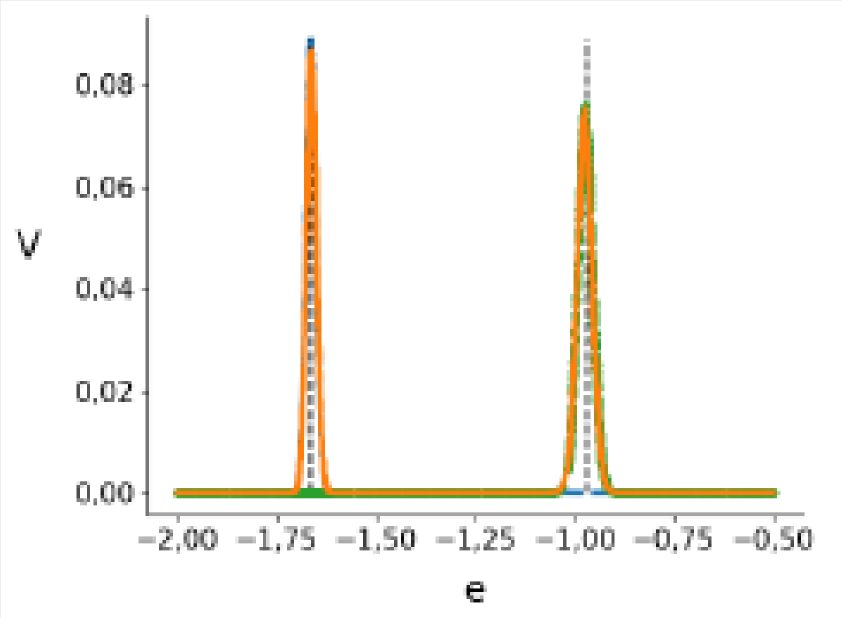

$$P_{xS}(e) = \frac{1}{N_{test}} \sum_{i=1}^{N_{test}} p^i_{xS}(e)$$

$$V_{xS} = \frac{1}{N_{test}} \sum_{i=1}^{N_{test}} \left[p^i_{xS}(e)\right]^2 - \left[\frac{1}{N_{test}} \sum_{i=1}^{N_{test}} p^i_{xS}(e)\right]^2$$

Sukhoverkhova, Mozolenko & LS, Phase probabilities in first-order transitions using machine learning, Phys. Rev. E. (2025)

Microcanonical population annealing algorithm CUDA+MPI

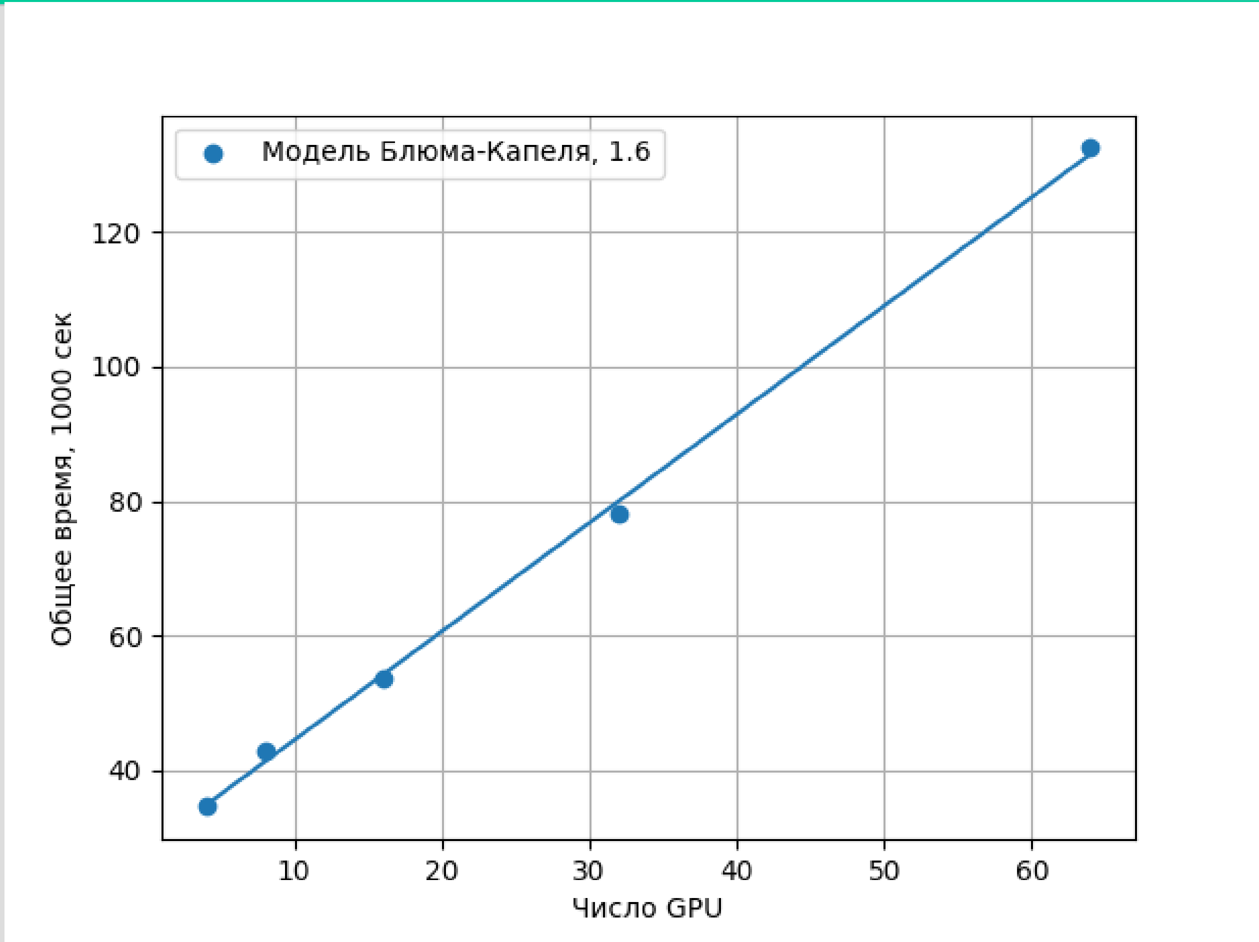

Russkov, XIX Int. Conf. Supercomputing and Mathematical Modeling, Sarov (2024)

Conclusion

- The *population annealing* approach is a promising tool for solving a number of problems:
    - "complex" ground state,
    - rough energy landscape,
    - optimization problems.
- The *Population annealing* algorithm and the *Microcanonical Population annealing* algorithm are natural candidates for large-scale and *fully scalable simulations* with *heterogeneous parallelism*.