Digitally Optimized Initializations for Fast Thermodynamic Computing

Mattia Moroder[1], Felix Binder[1], John Goold[1]
[1]Trinity College Dublin, Dublin, Ireland

Thermodynamic computing is an emerging paradigm in which computations are performed by physical systems evolving under stochastic dynamics. A prototypical example is solving a linear system $Ax = b$, where one considers a collection of harmonic oscillators, encodes the matrix A in the all-to-all couplings and the vector b as a constant force, and obtains the solution $x = A^{-1}b$ from the mean of the equilibrium distribution of the oscillators' displacements. More generally, various linear algebra operations can be implemented by such physical systems whose equilibrium statistics encode the solution [1]. Importantly, this paradigm is not merely theoretical: thermodynamic computing hardware has recently been implemented using noisy LC circuits, establishing its feasibility and opening the door to potentially energy-efficient alternatives to digital computing [2].

A central challenge in this framework is that computation time is governed by thermalization dynamics. In practice, the time required to reach equilibrium, set by the slowest relaxation modes, constitutes a major bottleneck. This raises the question of whether anomalous thermalization phenomena, such as the Mpemba effect [3], where certain initial conditions lead to faster relaxation, can be exploited to accelerate computation. In this context, speeding up thermalization corresponds directly to accelerating the underlying algorithm.

In this talk, based on [4], I develop a hybrid digital-thermodynamic framework that leverages this idea. I consider thermodynamic computing schemes in which the thermalization dynamics are governed by overdamped Langevin equations. In standard implementations, the oscillators are typically initialized in their equilibrium position, which leads to slow convergence due to long-lived relaxation modes. Instead, I propose to digitally optimize the initialization by exploiting partial spectral information of the problem instance. Concretely, given a matrix, I compute a small number of low-lying eigenpairs and use them to construct initial conditions orthogonal to the slowest Fokker-Planck modes, thereby suppressing the dominant relaxation bottlenecks, as illustrated schematically in the accompanying figure.

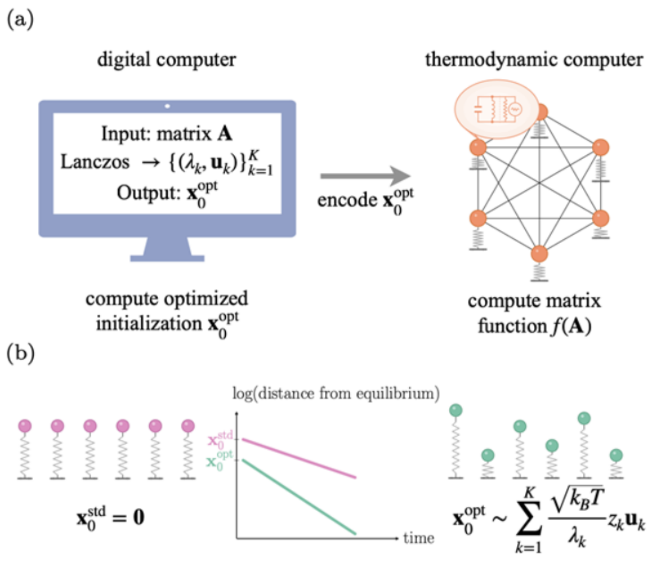

I derive analytical predictions linking the convergence rate of the dynamics to the spectral properties of the underlying matrix, thereby quantifying how initialization modifies the effective relaxation times. I test these predictions on two classes of random matrix ensembles and consider matrix inversion and determinant estimation as computational tasks. In all cases, I find good agreement between theory and numerical simulations, and observe substantial speedups compared to standard initialization strategies.

These results demonstrate that anomalous thermalization can be engineered as a computational resource, opening a new interface between nonequilibrium thermodynamics and algorithm design.

References:

[1] Maxwell Aifer, Kaelan Donatella, Max Hunter Gordon, Samuel Duffield, Thomas Ahle, Daniel Simpson, Gavin E. Crooks, Patrick J. Coles, npj Unconventional Computing volume 1, Article number: 13 (2024)

[2] Denis Melanson, Mohammad Abu Khater, Maxwell Aifer, Kaelan Donatella, Max Hunter Gordon, Thomas Ahle, Gavin Crooks, Antonio J. Martinez, Faris Sbahi, Patrick J. Coles, Nature Communications volume 16, Article number: 3757 (2025)

[3] Zhiyue Lu and Oren Raz, PNAS 114 (20) 5083-5088 (2017)

[4] Mattia Moroder, Felix C. Binder, John Goold, arXiv:2603.24183 (2026)