

Thermodynamic Linear Algebra

Gavin E. Crooks

NORMAL Computing

arXiv:2308.05660

Thermodynamic Linear Algebra


Maxwell Aifer, Kaelan Donatella, Max Hunter Gordon,
Thomas Ahle, Daniel Simpson, Gavin Crooks, Patrick J. Coles
Normal Computing Corporation, New York, New York, USA

Linear algebraic primitives are at the core of many modern algorithms in engineering, science, and machine learning. Hence, accelerating these primitives with novel computing hardware would have


Thermodynamic Computing

Thermodynamic Computing 1911.01968 Tom Conte et al.

- Harness nature's innate computational capacity
 - Use the underlying physics to compute (compute closer to the hardware)
- Noise as a resource, not a curse




Transistor Energy Scaling



<https://ariaresearch.substack.com/p/spotlight-on-suraj-scaling-compute>

Improving Compute per kT



1) Algorithms

2) Quantum Computing

3) Novel (classical) Hardware

The Hardware Lottery

Sara Hooker

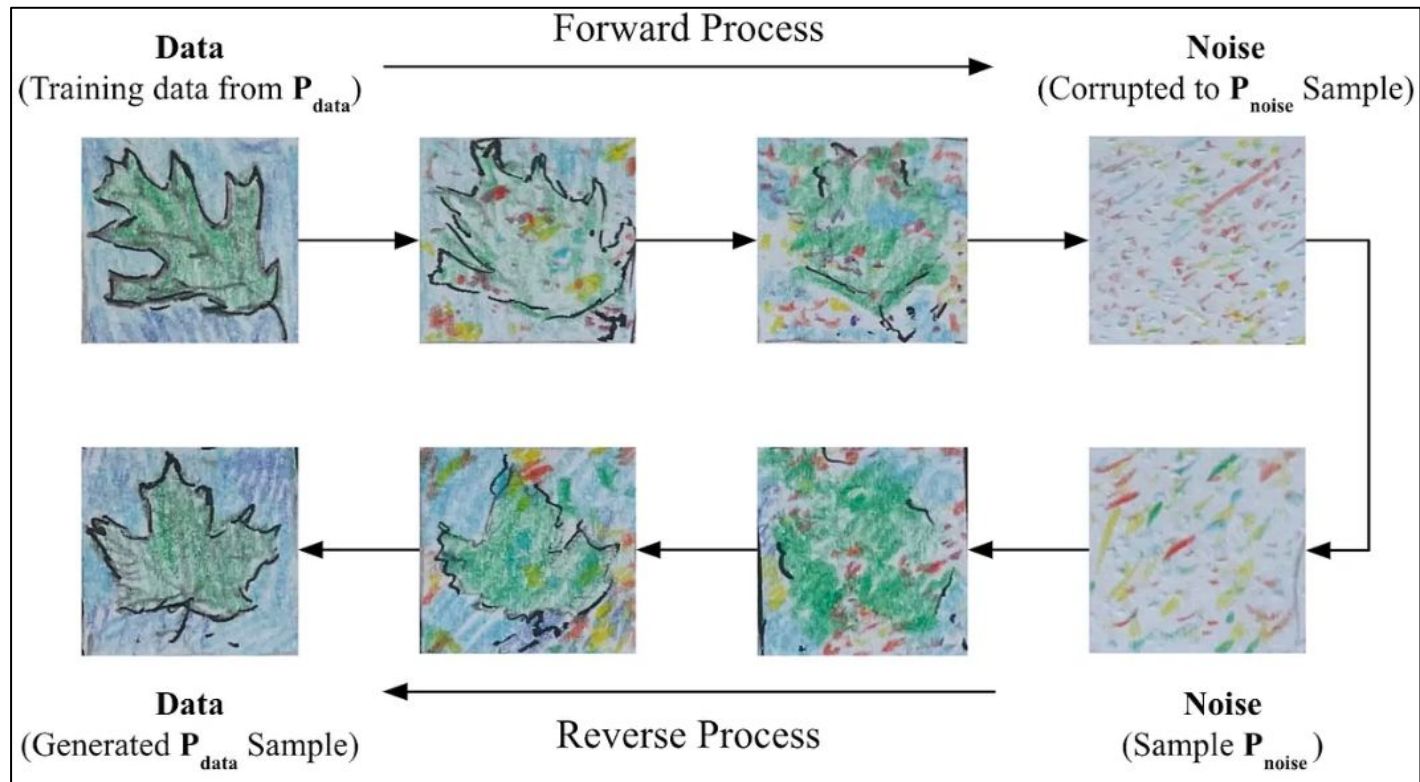
Google Research, Brain Team

shooker@google.com


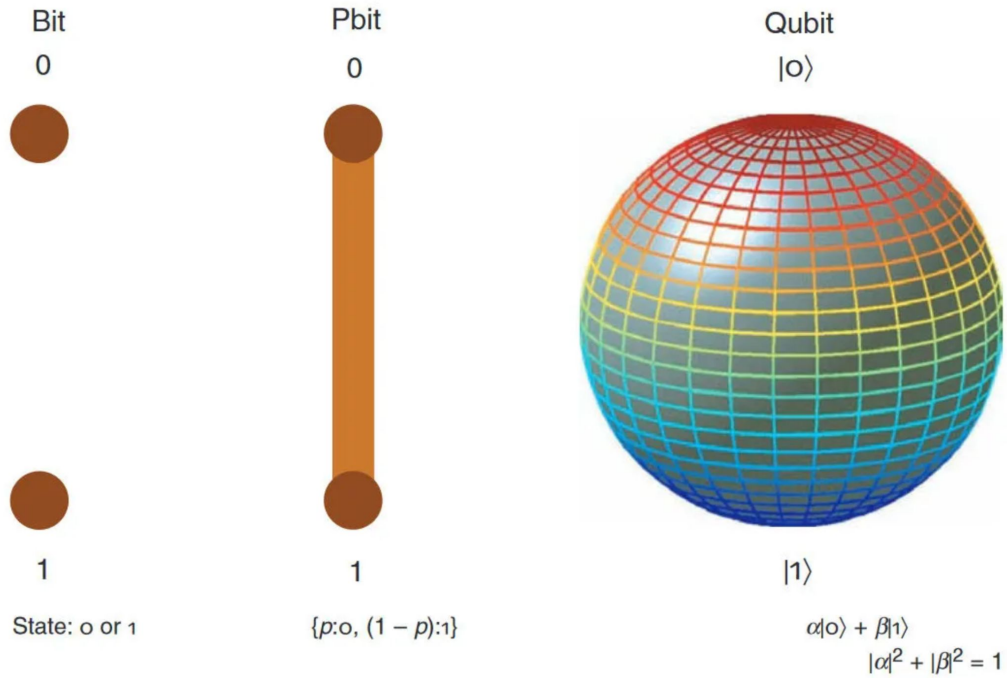
Abstract

Hardware, systems and algorithms research communities have historically had different incentive structures and fluctuating motivation to engage with each other explicitly. This historical treatment is odd given that hardware and software have frequently determined which research ideas succeed (and fail). This essay introduces the term hardware lottery to describe when a research idea wins because it is suited to the available software and hardware and *not* because the idea is superior to alternative research directions.


Noise and computation



Computational Building Blocks



RC Stochastic Unit



Stochastic noise sources

Thermal noise

$$v_{\text{tn}} = \sqrt{4k_B T R \Delta f}$$

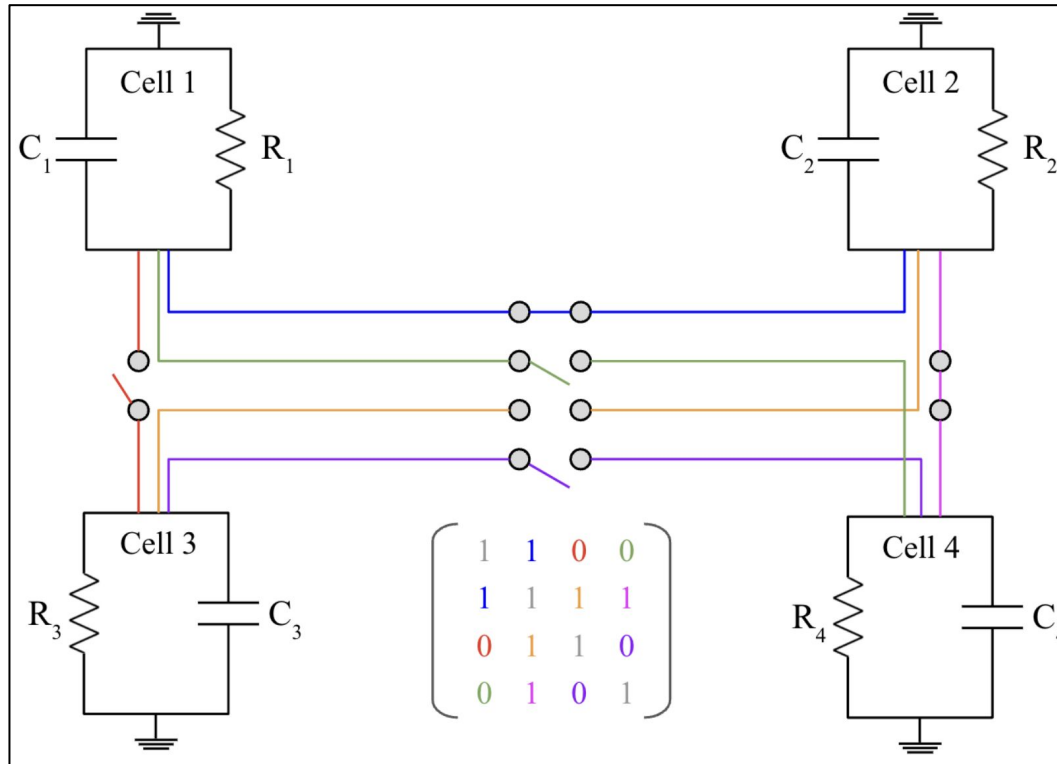
Shot noise

$$I_{\text{tn}} = \sqrt{2q|I|\Delta f}$$

Time evolution:

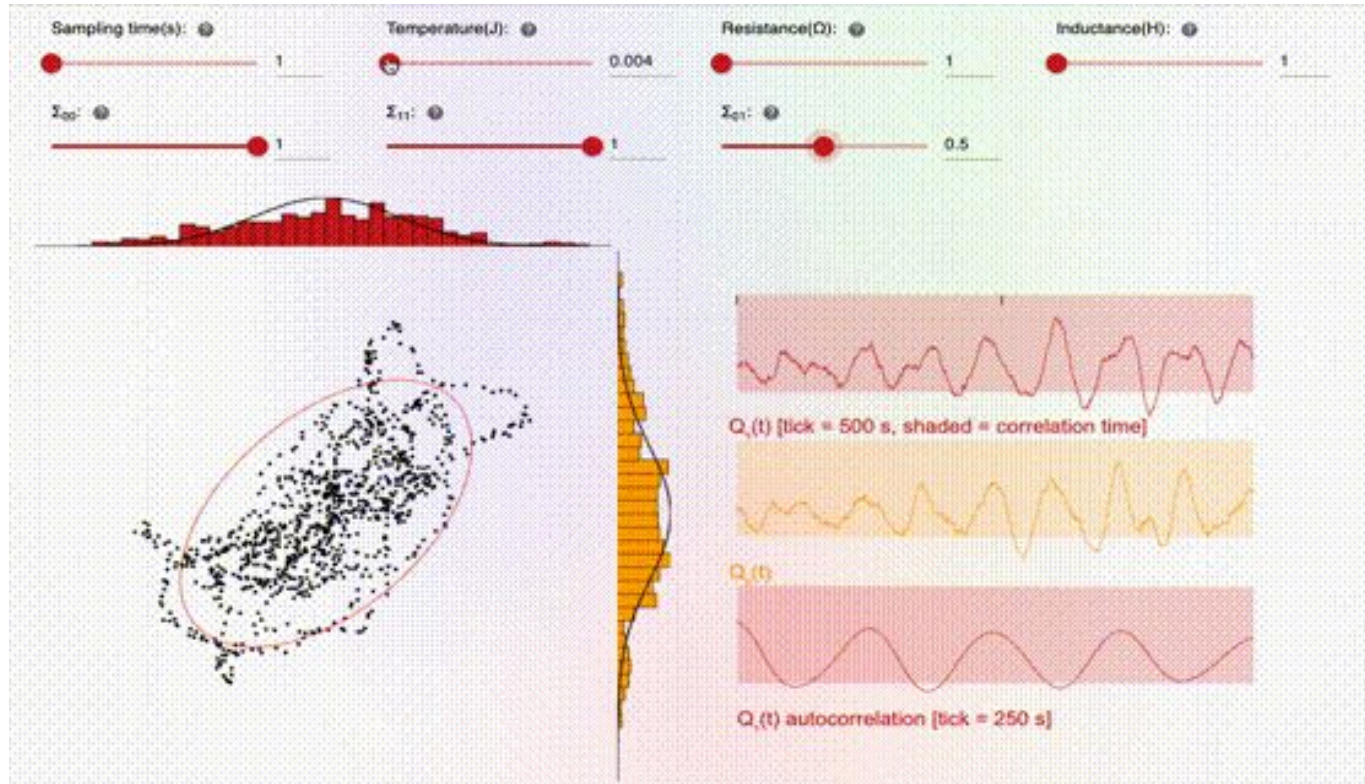
$$-\frac{dv(t)}{dt} = \frac{v(t) + \delta v(t)}{RC}$$

Stochastic Processing Unit (SPU)



Thermal Playground

<https://app.normalcomputing.ai/composer>



Stochastic Processing Unit Dynamics

Overdamped or Underdamped Langevin dynamics

$$d\mathbf{p} = [\mathbf{f} - BM^{-1}\mathbf{p}] dt + D d\mathbf{w}$$

$$d\mathbf{x} = M^{-1}\mathbf{p} dt$$

$$\mathbf{f} = -\nabla_{\mathbf{x}}U_{\theta}$$

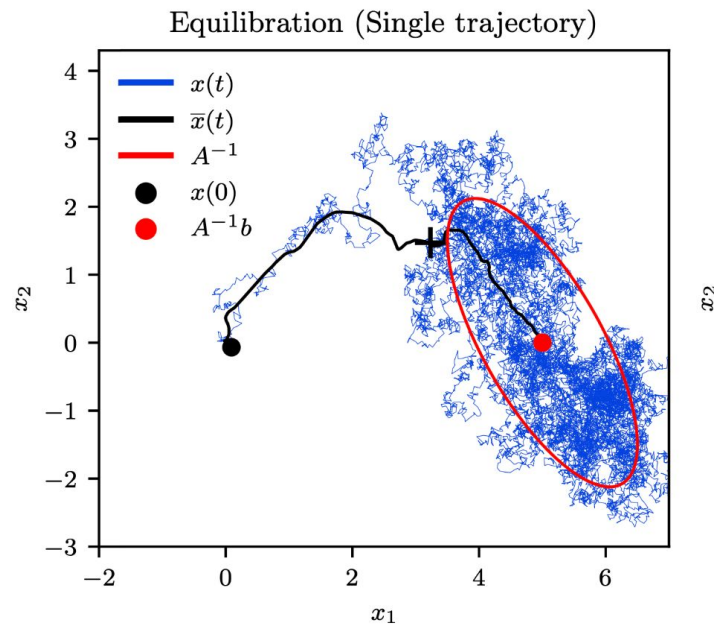
Gaussian Sampling with Stochastic Processing Unit

For harmonic oscillator system, at thermal equilibrium, x is Gaussian distributed:

$$V(x) = \frac{1}{2} x^\top A x - b^\top x$$


$$p(x) = \frac{1}{Z} e^{-\beta \frac{1}{2} x^\top A x - \beta b^\top x}$$

$$x \sim \mathcal{N}[A^{-1}b, \beta^{-1}A^{-1}]$$



Matrix Inversion with Stochastic Processing Unit



$$V(x) = \frac{1}{2} x^\top A x - b^\top x \quad \longrightarrow \quad x \sim \mathcal{N}[A^{-1}b, \beta^{-1} A^{-1}]$$



Thermodynamic Advantage

Problem	Digital SOTA	This work (Overdamped)	This work (Underdamped)
Linear System	$O(\min\{d^\omega, d^2\sqrt{\kappa}\})$	$O(d\kappa^2\varepsilon^{-2})$	$O(d\sqrt{\kappa}\varepsilon^{-2})$

$\omega=2.3$ matrix multiplication scaling
 κ Condition number
 ε accuracy



Matrix Determinant with Stochastic Processing Unit

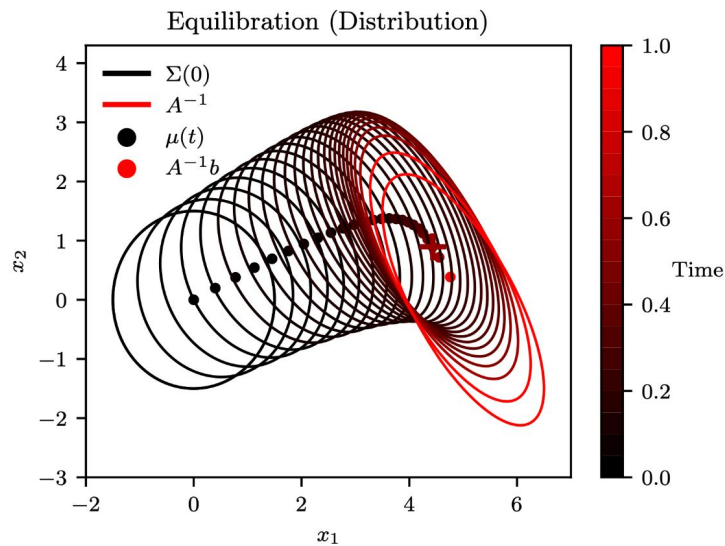
$$f_{\mu;\Sigma}(x) = (2\pi)^{-d/2} |\Sigma|^{-1/2} \exp\left(-\frac{1}{2}x^\top \Sigma^{-1}x\right),$$

$$\Delta F = F_2 - F_1 = -\beta^{-1} \ln \left(\frac{\int dx e^{-\beta V_2(x)}}{\int dx e^{-\beta V_1(x)}} \right).$$

$$\Delta F = -\beta^{-1} \ln \left(\sqrt{\frac{|A_2^{-1}|}{|A_1^{-1}|}} \right) = -\beta^{-1} \ln \left(\sqrt{\frac{|A_1|}{|A_2|}} \right).$$

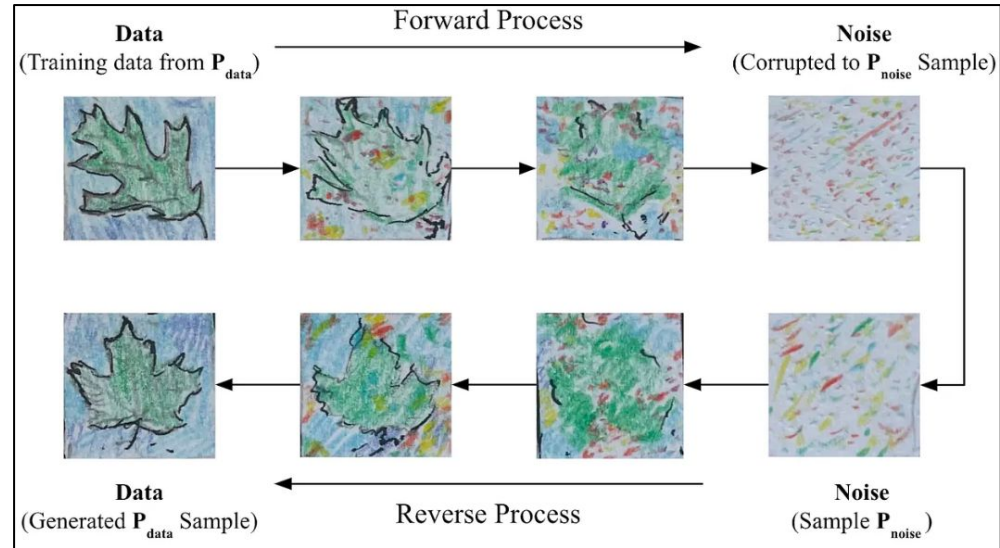

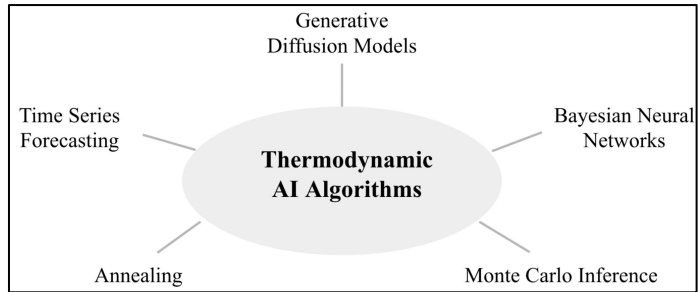
Matrix Determinant with Stochastic Processing Unit (2)

$$\Delta F = -\beta^{-1} \ln \left(\sqrt{\frac{|A_2^{-1}|}{|A_1^{-1}|}} \right) = -\beta^{-1} \ln \left(\sqrt{\frac{|A_1|}{|A_2|}} \right)$$



$$e^{-\beta \Delta F} = \langle e^{-\beta W} \rangle$$

Thermodynamic AI



NORMAL Computing



Thermodynamic Linear Algebra

Maxwell Aifer, Kaelan Donatella, Max Hunter Gordon,
Thomas Ahle, Daniel Simpson, Gavin Crooks, Patrick J. Coles
Normal Computing Corporation, New York, New York, USA

Linear algebraic primitives are at the core of many modern algorithms in engineering, science, and machine learning. Hence, accelerating these primitives with novel computing hardware would have tremendous economic impact. Quantum computing has been proposed for this purpose, although the resource requirements are far beyond current technological capabilities, so this approach remains long-term in timescale. Here we consider an alternative physics-based computing paradigm based on classical thermodynamics, to provide a near-term approach to accelerating linear algebra.

At first sight, thermodynamics and linear algebra seem to be unrelated fields. In this work, we connect solving linear algebra problems to sampling from the thermodynamic equilibrium distribution of a system of coupled harmonic oscillators. We present simple thermodynamic algorithms for (1) solving linear systems of equations, (2) computing matrix inverses, (3) computing matrix determinants, and (4) solving Lyapunov equations. Under reasonable assumptions, we rigorously establish asymptotic speedups for our algorithms, relative to digital methods, that scale linearly in matrix dimension. Our algorithms exploit thermodynamic principles like ergodicity, entropy, and equilibration, highlighting the deep connection between these two seemingly distinct fields, and opening up algebraic applications for thermodynamic computing hardware.

<https://normalcomputing.ai>

<https://normalcomputing.substack.com/>

arXiv:2308.05660

Upcoming: Berkeley Physics Colloquium, Mon Nov 13th