

Optimal Control of Bit Erasure in Stochastic Random Access Memory

Songela W. Chen^{1,2} and David T. Limmer^{1,2,3,4,*}

¹Department of Chemistry, *University of California, Berkeley, California 94720, USA*

²Chemical Sciences Division, *Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

³Materials Sciences Division, *Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

⁴*Kavli Energy NanoScience Institute, Berkeley, California 94720, USA*

 (Received 27 January 2026; revised 14 April 2026; accepted 14 May 2026; published 16 June 2026)

Energy costs of information processing are growing exponentially. Bit erasure is a key problem in this energy-information nexus, and a number of seminal relationships have been deduced regarding the relationship between thermodynamic costs and memory storage. To continue making progress in the modern era, however, requires confronting thermodynamic costs in realistic physical systems, which operate away from equilibrium. Here, we explore the thermodynamic costs of bit erasure in a complementary metal-oxide-semiconductor model of two types of random-access memory. We find dynamic random-access memory dissipates the least amount of energy when operated in the quasistatic limit, where errors are also minimized. By contrast, static random-access memory is most efficiently operated in finite time due to the energy required to maintain the state of the bit. We demonstrate a numerically robust optimization scheme using mean field theory and automatic differentiation, finding optimal protocols compatible with electrical engineering insights. These results provide a framework for operating realistic circuits in thermodynamically advantageous ways.

DOI: [10.1103/bmsv-mlq5](https://doi.org/10.1103/bmsv-mlq5)

I. INTRODUCTION

With the explosive growth of information technology, data-center electricity consumption is projected to double in the next five years [1]. In the context of the global energy crisis, it is crucial to confront physical costs of computation. Bit erasure is a well-studied problem in equilibrium thermodynamics, dating back to Landauer's insight that at least $k_B T \ln 2$ of energy is required to erase one bit of information [2]. As such, it offers a benchmark for understanding the energetic costs of computation. Landauer's formulation and most subsequent studies, however, ignore the practical requirements of finite-time computation and the effect of the physical substrate implementing erasure that is often operated within a nonequilibrium steady state [3–6]. Although some studies have confronted finite-time bit erasure or considered nonequilibrium initial states, the majority of these results have limited applicability as they employ abstract confining potentials to model circuitry [7–12]. Here, we find optimal trade-offs between dissipated

heat and accuracy of bit erasure in a realistic model of a logical circuit, kept out-of-equilibrium by voltage differences across transistors. Our results leverage contemporary machine-learning-based numerical optimization techniques and formal insights from stochastic thermodynamics to expand our understanding of the interplay between thermodynamics and information processing. More generally, our approach illustrates the application of optimal control techniques to nonequilibrium steady states.

The vast majority of information processing devices today use complementary metal-oxide-semiconductor (CMOS) technology [13]. Until now, energy efficiency has improved by physically scaling down transistors on chips [14]. However, we are approaching the nanoscale regime where component sizes are comparable to the mean free path of electrons, and thermal noise effects must be considered. This caveat renders traditional macroscopic models used for circuit design inappropriate. Several CMOS-inspired models have been proposed in recent years to account for energetic costs and noise effects in a way that is thermodynamically consistent [15–17], more realistically than the Gaussian noise assumptions of standard SPICE simulations in electrical engineering [18]. While some prior studies probed reliability of information storage in these models [19–21], little work had investigated active erasure until recently [22,23]. One outstanding question is the discrepancy between the observation in

*Contact author: dlimmer@berkeley.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

statistical physics that dissipation is minimized in the quasistatic limit, versus the electrical engineering conventional knowledge of an unavoidable accumulating power dissipation [13]. We aim to bridge this gap in an architecture-specific manner using a stochastic thermodynamic framework.

In this study, we used one of these microscopic models to quantify the thermodynamic costs of bit erasure in a CMOS architecture of two different memory storage circuits, dynamic random-access memory (DRAM) and static random-access memory (SRAM). We use numerical optimization techniques to compare optimal erasure protocols, highlighting the contribution from the energy required to maintain the nonequilibrium steady states, as well as the reversible charging and discharging of the relevant capacitors. For CMOS circuits, bit erasure operates far from Landauer's bound, largely due to the sizable reversible heat from discharging the capacitors on which the bit is stored. Our results demonstrate that the trade-offs between information processing speed and dissipation depend on circuit architecture. For SRAM, there is a thermodynamically optimal time for completing a bit erasure. By contrast, DRAM dissipates the least amount of energy and is most accurate when operated in the quasistatic limit.

II. CIRCUIT MODEL OPTIMIZATION

We model CMOS implementations of memory storage devices using physical noise models consistent with circuits operated at low power [15,19]. Here we study random-access memory (RAM), which is used to store working data rather than long-term data. In Fig. 1(a), we show the general process. The circuits that store bits of information are connected to so-called wordline $V_w(t)$ and bitline $V_b(t)$ voltages, which act jointly as a time-dependent controller that manipulates an output voltage V_o and sets it to 0 irrespective of the initial state. Individual bits are connected to these time-dependent controls through access transistors, which function as on-off switches. Macroscopically these are deterministic processes, but for systems operating close to thermal energy scales, fluctuations can result in errors in bit erasure and different amounts of heat dissipated in different realizations of erasure. We will consider averages over independent erasure attempts.

In a typical CMOS circuit, like DRAM shown in Fig. 1(b) or SRAM in Fig. 1(c), V_w controls the flow of current through the access transistor, and V_b controls the state being written. The basic unit of time is set by the residence time of an electron in a transistor, τ_t , while changes to V_w or V_b affecting bit erasure occur over times $\tau_{op} \gg \tau_t$, illustrated in Fig. 1(d). The stochastic nature of the circuit is illustrated by the probability distribution before and after bit erasure in Fig. 1(e), whereby errors can occur

both spontaneously as well as the result of inaccurate control. The driving voltage V_d sets the scale of fluctuations and bias strength, and the threshold voltage V^* defines the logical 0 state. The rate of spontaneous errors is denoted as $1/\tau_{err}$, and as presented in Fig. 1(f) we will consider operation timescales faster than it $\tau_{op} \ll \tau_{err}$ throughout.

A. Stochastic circuit model

The components involved in either RAM device include electrodes, capacitors, and transistors. Additional components such as inductors could also be accommodated for more complex devices. We treat electrodes as ideal electron reservoirs at fixed potentials V_j . The capacitors are nonideal electron reservoirs, with voltages across the capacitor V_j that evolve dynamically. The charge on the capacitor is $qm_j = CV_j$, where q is the fundamental unit of charge, $m_j = \{-\infty, \dots, 0, 1, \dots, \infty\}$ is the number of electrons on the capacitor, and C is capacitance. In contrast to the electron reservoirs, we model transistors as a single Fermionic state, with occupancy $n_i = \{0, 1\}$ and onsite energy E_i representing its band gap. The transistor energy is modulated electrostatically by the voltage sources to which it is connected. The state space of a given device is given by the joint occupancies of the capacitors and transistors, $\mathbf{x} = \{n_i\}^{N_t} \otimes \{m_j\}^{N_c}$ where N_t is the number of transistors and N_c the number of capacitors. Throughout, we use i subscripts to reference transistor states and j subscripts to reference electrode or capacitor states. We report voltages with reference to a thermal voltage $V_T = 26$ meV at room temperature and take capacitance $C = 10q/V_T$.

The systems evolve stochastically according to the Markovian master equation

$$\partial_t \mathbf{p}(\mathbf{x}, t) = \mathbf{W} \mathbf{p}(\mathbf{x}, t), \quad (1)$$

where $\mathbf{p}(\mathbf{x})$ is the probability vector associated with the transistor and capacitor states and $\mathbf{W}$ is the stochastic generator, with matrix elements W_{ij} specifying the rates of transition from states j to i . The rates are defined with reference to the Fermi-Dirac distribution, with the transition rate of an electron from an electrode j to a transistor i given by

$$W_{ij}(V_j) = \Gamma (e^{\beta(E_i - qV_j)} + 1)^{-1} \quad (2)$$

and the reverse given by

$$W_{ji}(V_j) = \Gamma - \Gamma (e^{\beta(E_i - qV_j)} + 1)^{-1}, \quad (3)$$

where $\beta = 1/k_B T$ is the inverse temperature defined with Boltzmann's constant k_B and temperature T , Γ specifies the timescale for transitions and is physically set by the resistance of the transistor-electrode interface, E_i is a transistor energy, and V_j is a fixed voltage measured on an electrode




FIG. 1. Schematic of bit erasure and related timescales. (a) Controlled bit erasure, driven by time-dependent $V_w(t), V_b(t)$ to set a 0 state. (b) Dynamic random-access memory (DRAM) circuit. (c) Static random-access memory (SRAM) circuit. (d) Idealized trajectory showing evolution of bit state $V_o(t)$ (blue) in response to control $V_b(t)$ (orange). (e) Probability distribution of bit state before (orange) and after (blue) erasure, with threshold voltage V^* defining the 0 state. (f) Relevant timescales for electron residence time in a transistor τ_t , operation time τ_{op} , and stochastic error time τ_{err} .

or the instantaneous voltage on a capacitor measured using a midpoint rule [24]. We require $\Gamma \ll 1/\beta\hbar$ to ensure weak coupling such that electron hops can be treated as discrete memoryless events, and simulation times are expressed in units of $\tau_t = 1/\Gamma$ at room temperature. The forward and reverse transition rates chosen in this way ensure local detailed balance and thus thermodynamic consistency [25]. This model is able to reproduce shot noise characteristics of CMOS devices [26]. We use kinetic Monte Carlo simulations to propagate the stochastic dynamics numerically [27,28].

B. Thermodynamic characteristics and optimization

To quantify thermodynamic costs of the bit erasure, we measure mean heat dissipation as the product of the electron current and its conjugate affinity from all channels in the circuit,

$$\frac{dQ}{dt} = \sum_{j,i} J_{ji} \ln \frac{W_{ji}}{W_{ij}}, \quad (4)$$

where the sum is over the various pairs of electrodes and capacitors connected by individual transistors [29]. For electrode or capacitor j connected to transistor i , the current into electrode j is given by

$$J_{ji}(n_i) = W_{ji}n_i - W_{ij}(1 - n_i), \quad (5)$$

and fluctuates due to the occupancy of the transistor and electron reservoir. To quantify the performance of the bit erasure, we define the error ϵ as the probability of being

erroneously above a threshold voltage, V^* ,

$$\epsilon = \int_{V^*}^{\infty} dV_j p(V_j, \tau_{op}), \quad (6)$$

where $p(V_j, \tau_{op})$ denotes the probability density for V_j of the measured output voltage at the final time τ_{op} . We will take the threshold to be $0.1V_d$, which is slightly looser than current practice at $V_d = 40V_T$ [15].

To obtain optimal control protocols, we use automatic differentiation for nonequilibrium systems by Engel *et al.* and the Adam optimization algorithm using the JAX library [30–33]. We define a loss function combining the averaged dissipated heat and error,

$$\mathcal{L} = \langle Q \rangle + \lambda \epsilon, \quad (7)$$

where $\langle \dots \rangle$ denotes an average over independent realizations and λ is a scaling factor for relative weights of the error and heat. As the error is bounded, we choose λ large enough to ensure the minimal error at each control time. Details of the optimization may be found in Appendix B. Briefly, we use a nonlinear ansatz for the time dependences of both the wordline and bitline voltages and ensure that they are cyclical, $V_w(0) = V_w(\tau_{op})$ and $V_b(0) = V_b(\tau_{op})$ for operation time τ_{op} . Further, we restrict $V_w = [-3V_d, V_d]$ and $V_b = [0, V_d]$, as values outside of this range do not meaningfully change system dynamics and result in numerical instability. The coefficients in the control ansatz are initialized randomly and iterated stochastically in the direction of steepest descent. Rather than evaluating the loss and its gradient by propagating the master equation directly or for realizations of

the stochastic dynamics, we use a surrogate model based on a mean-field approximation to the stochastic dynamics described in Appendix A. This surrogate model simplifies the gradient calculation greatly, making it comparable to previous continuous space genetic algorithms [34]. As will be illustrated, the deterministic and fully differentiable mean-field dynamics are a very accurate representation of the typical trajectory of the system. Errors and mean dissipated heats are evaluated with the optimized protocols using kinetic Monte Carlo simulations.

III. DYNAMIC RANDOM-ACCESS MEMORY

The two main types of RAM are dynamic (DRAM) and static (SRAM). DRAM, shown in Fig. 2(a), consists of an N -type access transistor connected to a bitline electrode at potential V_b and a storage capacitor with capacitance C . The transistor is also connected to a wordline electrode at potential V_w , which controls the current through this transistor, functionally making it a switch. The onsite energy of the access transistor is given by $E_A = q(V_b - V_w)$. The voltage associated with charge on the capacitor serves as the output voltage, V_o , which stores the state of the bit. The total energy of the system is a sum of the energy of the access transistor and the output capacitor

$$E = n_A q(V_b - V_w) + \frac{CV_o^2}{2}, \quad (8)$$

where n_A is the occupancy of the access transistor and the output voltage $V_o = qm_o/C$ is given by the occupancy of the capacitor. In between writing or erasing the bit, the bitline is held at an intermediate voltage, so that charge gradually leaks away from the capacitor over time. The leakage current depends on the energy of the transistor through both V_w and V_b , as shown in Fig. 2(b). As a consequence of the leakage current, DRAM must be refreshed on timescales τ_{err} longer than the bit erasure protocols, τ_{op} . Given this separation of timescales, we do not consider the energetics of refreshing the memory in DRAM.

Representative bit erasure trajectories are shown in Fig. 2(c) starting from both 0 and 1 states. The two logical states are written into the capacitor through application of a bitline voltage V_b and setting the wordline voltage $V_w = 0$ to allow for current to flow. Writing the logical 1 state corresponds to setting $V_b = V_d > 0$, while for the logical 0 state, $V_b = 0$. Maintenance of the bit occurs by setting $V_b = V_d/2$ and $V_w = -2V_d$. As an example of a naive bit erasure protocol, analogous to that used in real CMOS device operation, at time $t = 0$ the bitline voltage is set instantaneously to $V_b = 0$, and the wordline voltage is turned “on” to $V_w = 0$ to allow current through the access transistor. The capacitor discharges with a characteristic RC time that is $\tau_{\text{RC}} \sim 10^3 \tau_t$, while dissipated heat accumulates in a manner that mirrors the voltage on the capacitor. Exact values of output voltages (top) and dissipated heat




FIG. 2. Dynamic random-access memory (DRAM) phenomenology. (a) The DRAM circuit consisting of an ideal bitline electrode at voltage V_b , an access transistor controlled by the wordline voltage V_w , and a capacitor with output voltage V_o indicating the bit state. (b) Average current across the transistor at constant values of V_w and different driving voltages V_d . (c) Example trajectories of bit erasure at $V_d/V_T = 3$ starting in logical 1 (orange) or 0 (blue) compared to the estimate from mean-field theory (black), for the output voltage (top) and heat (bottom).

Q (bottom) vary among stochastic trajectories, but their average is well described by a mean-field approximation. We leverage the quantitative accuracy of this mean-field theory to facilitate optimization. Starting in the 0 state, the average heat is 0, while starting in the 1 state, the average heat is given by the expected change in energy, $Q = -\Delta E \approx CV_d^2/2$. Averaged over a uniform distribution of bits, the mean heat from this naive protocol would be $\langle Q \rangle = -\Delta E \approx CV_d^2/4$.

To understand the trade-offs between dissipation and error, we have solved for optimal erasure protocols averaged over a uniform distribution of initial states of the bit 0 and 1. We consider the effects of both the operation time as well as the magnitude of the driving voltage, ensuring that the protocol is cyclic by enforcing $V_w(0) = V_w(\tau_{\text{op}})$ and $V_b(0) = V_b(\tau_{\text{op}})$. As a function of τ_{op} we observe in Fig. 3(a) that at short operation times, $\tau_{\text{op}} < \tau_{\text{RC}}$, the error remains high, around 0.5. The source of the error is incoming bits that start in the 1 state and do not have time to relax fully to the logical 0 state. Beyond this time, the error rapidly decreases to a plateau value. The plateau error is inversely correlated with driving voltage V_d , as larger V_d

results in larger separation between the two equilibrium distributions for the two bit states $V_b = V_d$ and $V_b = 0$, so errors are determined by the persistence of thermal fluctuations. The average heat dissipation in Fig. 3(b) increases up to this minimum time required to erase the bit. Beyond this time, dissipation decreases, consistent with expectations from previous equilibrium models of finite-time bit erasure. This is only possible by controlling both $V_w(t)$ and $V_b(t)$ with time. If we set $V_b = 0$ instantaneously as in standard CMOS operation, the minimum dissipation to achieve bit erasure is determined by the reversible charging energy of the capacitor $E = CV_d^2/2$, and dissipation plateaus to 1/2 of this value in the long-time limit rather than to zero (see Appendix C). Conversely, most prior studies on thermal noise in DRAM have manipulated the bitline voltage V_b while keeping the wordline voltage V_w constant [22,35,36], which is also slightly sub-optimal as will be shown.

To understand the dependence of the heat and error on operation times, we consider the behavior of the circuit under limiting optimal protocols. In the short time regime, $\tau_{op} < \tau_{RC}$, optimal protocols resemble Fig. 3(c), where V_b and V_w are set to zero abruptly and return to their initial condition similarly sharply near τ_{op} . This protocol effectively turns on the access transistor by bringing the wordline voltage near 0 for the duration of the protocol, then turns it off at the end—just as typical devices do in practice. Physically, this corresponds to the energy of the access transistors being equal to the Fermi energy of their associated bitline electrodes. Optimal protocols have V_w slightly below 0, pulling the energy of the access transistor closer to that of the electron exiting the capacitor to facilitate more downhill transitions than having the access transistor resonant in energy with the bitline electrode only. Meanwhile, the bitline voltage V_b is set to the desired 0 state to maximize the thermodynamic force to discharge the logical 1 or preserve the logical 0. This is reflected in representative trajectories of V_o under this protocol, which either stay fluctuating near 0 from the start or immediately relax toward 0 once V_w reaches near 0. Given the constraint of the short τ_{op} , the trajectories starting in 1 lack sufficient time to relax fully to 0, hence their end value is slightly above the threshold voltage $V^* = 0.1V_d$. However, we minimize the error by erasing the bit to the extent possible within this time and by ensuring incoming 0 states stay 0. Mirroring the output voltage, the dissipated heat either fluctuates around 0 for the initial logical 0 state or grows as the capacitor discharges when starting in the logical 1. For simulations starting in a thermal distribution, the transistor is occupied roughly half the time, and the energy lost from this electron accounts for the small finite dissipation observed for some trajectories starting from the logical 0 state.

In the long-time limit, $\tau_{op} > \tau_{RC}$, optimal protocols resemble Fig. 3(d), where the access transistor is turned




FIG. 3. Heat-accuracy trade-offs observed in DRAM, averaged over starting states 0 and 1. (a) Error and (b) dissipated heat for two driving voltages $V_d/V_T = 3$ (green circles) and $V_d/V_T = 6$ (purple squares) across operation times. Error bars indicate variation from degenerate protocols with optimization losses within 10% of the minimum for that operation time τ_{op} . Representative trajectories and optimal protocols with $V_d/V_T = 3$ for operation times (c) $\tau_{op}/\tau_t = 1000$ and (d) $\tau_{op}/\tau_t = 16000$, respectively, where stochastic dynamics starting in logical 1 (orange) or 0 (blue) state are consistent with mean-field dynamics (black). Red dashed line (top) indicates voltage threshold of $0.1V_d$. Green shaded region (bottom) indicates range of V_w where transistor is active.

“on” only for a brief amount of time while the bitline is adjusted quasistatically. From $t/\tau_t \approx 1000$ to $t/\tau_t \approx 6000$ in the example, V_w is near 0 so current can flow, and V_b decreases approximately linearly from $V_d/2$ to 0. During

this active part of the protocol, the capacitor stays in equilibrium with the bitline voltage, minimizing the voltage difference across the transistor and therefore the dissipated heat. During the remainder of the protocol, the capacitor is in equilibrium with the bitline and so there are no additional channels for dissipation on average. Absent the initial logical 0 state, such a protocol would result in zero heat dissipated in the long time limit. However, we use a single protocol to erase from either 0 or 1 states, and a protocol that minimizes dissipation starting from the 1 state is necessarily non-optimal for the 0 state. The finite dissipation at long times reflects dissipation incurred from the 0 state moving away from 0 towards equilibrium with the bitline V_b , and then returning back to 0 at the end of the protocol. Representative trajectories of V_o under this protocol starting in the logical 1 state exhibit a delay in their discharging before exhibiting an elongated relaxation to 0 due to the bitline being smoothly set to zero. In contrast, representative trajectories of V_o under this protocol starting in the logical 0 state exhibit a spike in values coincident with the turning on of the access transistor and bitline voltage at intermediate values, for example around $t/\tau_t = 2000$, before relaxing back to 0. The average heat dissipated for both initial bits mirrors the output voltage, and for long operation times obtains a net contribution from the logical 0 state. An extremized version of this behavior is described in Appendix D, and distributions of output voltages are shown in Appendix E.

In both time limits, the basic phenomenology we find is independent of the V_d value we use to set the logical 1 state, though quantitatively the asymptotic error is strongly dependent on its value and the heat dissipated at finite time grows as V_d^2 . In the long time limit, the heat dissipated will approach $CV_d^2/4$ or half the reversible charging energy of the capacitor, far greater than Landauer's bound. In principle there are contributions to the dissipated heat from the manipulation of the transistor energy, upon which work is being done. However, the characteristic relaxation time of the transistor, $\tau_t \ll \tau_{op}$, is much less than the protocol times we consider here. Thus, manipulation of the transistor maintains adiabaticity with respect to the protocol.

IV. STATIC RANDOM-ACCESS MEMORY

In contrast to the 1-transistor, 1-electrode, 1-capacitor design of the DRAM, the 6-transistor, 6-electrode, 2-capacitor SRAM we consider next is more complex. This device consists of two coupled NOT gates, the kinetic network for which is shown in Fig. 4(a). Each NOT gate has a source voltage, $V_s = 0$, a drain voltage, $V_d > 0$, a control voltage electrode held at $V_b^{(k)}$, and a capacitor whose voltage serves as the readout of the bit, $V_o^{(k)}$, where k labels the two NOT gates, $k = 1, 2$. Each NOT gate consists of an N -type and P -type transistor, which we model as a single

state. The onsite energy of each transistor is determined by the output voltage of the other NOT gate, for example

$$E_P^{(1)} = qV_o^{(2)}, \quad E_N^{(1)} = \frac{3}{2}qV_d - qV_o^{(2)}, \quad (9)$$

with occupation variables $n_{P,1}$ and $n_{N,1}$, and equivalent energies for the transistors on the second NOT gate, dependent on the output voltage of the first. This model of a NOT gate is able to reproduce expected voltage transfer curves in typical CMOS circuits [15]. The interdependency of the two NOT gates forms a positive feedback loop that maintains the state of the bit through a pitchfork bifurcation, exhibited in Fig. 4(b). The bistability depends on the size of the driving voltage, V_d , with an onset of bistability at $V_d/V_T \approx 2.5$, and a time to spontaneously flip a bit, τ_{err} , that grows exponentially with V_d [19]. The combined output of the two capacitors determines the state of the bit, and we use the convention that $(V_o^{(1)} > 0.9V_d) \wedge (V_o^{(2)} < 0.1V_d)$ defines the 0 state.

To control the state of the bit, each NOT-gate capacitor is connected to an access transistor whose energies depend on a joint wordline voltage $V_w^{(1)} = V_w^{(2)} = V_w$ and complementary bitline voltages $V_b^{(1)} = V_b$ and $V_b^{(2)} = V_d - V_b$, with onsite energies $E_A^{(k)} = q(V_b^{(k)} - V_w)$. The current behavior shown in Fig. 2(b) for DRAM also holds for the access transistors in SRAM. The resulting energy for the SRAM device is

$$E = \sum_{k=1,2} n_{P,k} E_P^{(k)} + n_{N,k} E_N^{(k)} + n_{A,k} E_A^{(k)} + \frac{C}{2} (V_o^{(k)})^2, \quad (10)$$

where the output voltages are dynamic variables, determined by occupation variables $m_o^{(k)}$, as $V_o^{(k)} = qm_o^{(k)}/C$.

SRAM is an intrinsically nonequilibrium system because the source and drain electrodes for each NOT gate are held at different voltages. Representative erasure trajectories are shown in Fig. 4(c) starting from both 0 and 1 states. As with DRAM, two logical states are written into the capacitor by applying a bitline voltage and turning on the wordline voltage $V_w = 0$ to allow current to flow across the access transistor. Writing the logical 1 state corresponds to setting $V_b = V_d > 0$, while for the logical 0, $V_b = 0$. Maintenance of the bit occurs by setting $V_b = V_d/2$ and $V_w = -2V_d$. Like DRAM, bit erasure for SRAM occurs with a naive protocol and with an analogous characteristic time τ_{RC} . However, the total dissipation in SRAM continues to increase after bit erasure has occurred. This so-called housekeeping heat [37] contribution derives from the constant voltage drop across each NOT gate and changes the heat-accuracy trade-offs for SRAM as compared to DRAM. The mean-field theory captures this housekeeping heat contribution alongside the overall average behavior with quantitative accuracy, justifying its use in optimization.




FIG. 4. Static random-access memory (SRAM) phenomenology. (a) One of two coupled NOT gates in the SRAM circuit. It contains an ideal bitline electrode at V_b , an access transistor controlled by the wordline V_w , and a capacitor with output voltage V_o indicating the bit state like in DRAM, but the capacitor is additionally connected to ideal source V_s and drain V_d electrodes with a nonequilibrium bias across them. (b) Pitchfork bifurcation behavior as a function of driving voltage V_d . Memory is stored when bistability emerges above a critical voltage $V_d/V_T \approx 2.5$. (c) Example trajectories of bit erasure at $V_d/V_T = 3$ starting in logical 1 (orange) or 0 (blue) compared to the estimate from mean-field theory (black) for coupled output voltages (top, middle) and dissipated heat (bottom).

To understand the dependence of the heat and error on operation time, we again consider the behavior of the circuit under limiting optimal protocols. In Fig. 5(a), we show error ϵ as a function of operation time τ_{op} at two driving voltages $V_d/V_T = 3$ and 6. For operation times shorter than a threshold time, the error remains high, while at long times the error drops to a plateau value. Both the minimum time to erase and the plateau error rate depend on the driving voltage, with smaller V_d requiring shorter times but yielding larger errors asymptotically. The origin of the error in the long time limit is analogous to DRAM where smaller driving voltages produce larger overlap between the stable bit states, allowing thermal fluctuations to weakly degrade memory. The V_d dependence to the threshold time to reach this plateau error is a manifestation of the increase in stability of the two logical states inherited from the bifurcation in SRAM. Larger V_d results in deeper metastability of both the logical 0 and logical 1, and thus a longer time to transition between them [19].

In Fig. 5(b), we show the operation time dependence of the dissipated heat Q . Below the threshold time required to erase the bit, the average heat dissipated grows weakly because little activity occurs in the circuit. Dissipation peaks near the threshold time, then slowly decreases as operation time increases like in DRAM. However, for SRAM dissipation increases in the long-time limit due to the housekeeping heat, which results from stochastic transitions generating a continuous electron current from source to drain in each NOT gate. The housekeeping heat is more significant for lower driving voltages because these stochastic transitions are more likely across smaller potential differences, generating larger leakage currents. Thus, the SRAM heat-accuracy trade-off becomes more dramatic

for low driving voltage $V_d/V_T = 3$ in the long-time limit, as housekeeping heat is a bigger contribution. As a consequence, optimal driving voltages depend on the desired operation time.

Representative optimal protocols are shown for two different τ_{op} in Figs. 5(c) and 5(d). For short times in Fig. 5(c), current is maximized through the access transistors by setting the wordline voltage V_w near 0, strongly driving the system toward the zero state induced by setting the bitline voltage to $V_b = 0$. Error and dissipation are therefore minimized within the time constraint, the latter because energy is directed efficiently toward bit erasure rather than through other parts of the circuit. By contrast, a long-time protocol in Fig. 5(d) maintains the access transistors in “off” states for most of the protocol. During the “active” portion of the protocol where $V_w \approx 0$, denoted by the shaded regions, V_b follows the energy of the discharging capacitor within the circuit. In this case, it is advantageous to minimize activity in the circuit during the “off” phase to minimize dissipation, while still ensuring that the bit is erased during the subsequent “on” phase. We found that the duration of the “on” part of the protocol could vary somewhat without changing the results dramatically, and also that the exact protocol during the “on” part could vary as long as V_w was near 0. Accordingly, the behavior of V_b while V_w is away from zero does not affect the state of the bit. Various optimal protocols are nearly degenerate, determined by judicious switching between the “off” and “on” phases of the control protocol.

In both limits of τ_{op} , the optimal protocol resembles that of DRAM because the same control parameters $V_w(t)$ and $V_b(t)$ are tuned on timescales much slower than the characteristic relaxation time of the transistor, $\tau_{op} \gg \tau_t$. Work




FIG. 5. Heat-accuracy trade-offs observed in SRAM, averaged over starting states 0 and 1. (a) Error and (b) dissipated heat for two driving voltages $V_d/V_T = 3$ (green circles) and $V_d/V_T = 6$ (purple squares) across operation times. Error bars indicate variation from degenerate protocols with optimization losses within 10% of the minimum for that operation time τ_{op} . Representative trajectories and optimal protocols with $V_d/V_T = 3$ for operation times (c) $\tau_{\text{op}}/\tau_t = 1000$ and (d) $\tau_{\text{op}}/\tau_t = 16000$, respectively, where stochastic dynamics starting in logical 1 (orange) or 0 (blue) state are consistent with mean-field dynamics (black). Red dashed line (top) indicates voltage threshold of $0.1V_d$. Green shaded region (bottom) indicates range of V_w where transistor is active.

done to manipulate the energy of the occupied transistor is negligible as compared to the reversible charging energy of the capacitor and especially to the housekeeping heat. Unlike DRAM, however, the dissipation is not minimized in the long time limit and grows unbounded due to the housekeeping heat. This time-extensive contribution was also observed in a similar study using an alternative thermodynamically consistent model of SRAM [23]. The optimal time to perform erasure depends on V_d and is coincident with the onset of the plateau in the error. For large V_d there is a range of optimal operation times $\tau_{\text{op}} > \tau_{\text{RC}}$, as the timescale to plateau the error and that associated with the growth of housekeeping heat are distinct, yielding an optimal dissipated heat on average of CV_d^2 . For small V_d there is less of a separation of timescales between the plateau of the error and the growth of housekeeping heat, so there is a narrow range of operation times able to reach this limiting heat. The basic phenomenology is reproduced with simple linear protocols motivated from the study of DRAM, illustrated in Appendix D.

V. DISCUSSION

We have shown that physical details are important when accounting for thermodynamic costs of bit erasure. For DRAM operated on timescales shorter than the memory refresh time, we find that a quasistatic protocol is optimal, in line with prior studies that assume equilibrium thermodynamics [3,8,38]. For SRAM, we find an optimal trade-off between error and dissipated heat at intermediate operation times due to the accumulation of housekeeping heat. This contribution from housekeeping heat has been recently pointed out in Ref. [23], but it was neglected in prior studies that work with simple potentials, even though background power consumption is well studied in electrical engineering [13]. As a consequence, we have laid a foundation for significant further study on the thermodynamics of modern computation.

By understanding the behavior of our time-dependent control protocols, we obtain a roadmap for future design principles in nanoscale devices using similar architectures. The approach we have developed is distinct from previous calculations that optimize erasure using the master equation directly [23], which while providing easy access to gradients is limited in system size. By using the kinetic Monte Carlo method in conjunction with gradients approximated from mean field theory, we have a methodology that can scale to large systems but is still optimizable. Given this numerically robust and procedurally flexible nonequilibrium optimization scheme, we hope to extend this framework to more complex circuits with larger-scale implications. Future work will help to elucidate other energetically costly operations used in contemporary CMOS computers, and hopefully help to attenuate the increasing energy demands of information processing technology.

ACKNOWLEDGMENTS

We thank Gavin Crooks, Sam Oaks-Leaf, Kyle Ray, and Michael DeWeese for helpful discussions. This work was supported by NSF Grant CHE-1954580 and by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, Materials Sciences and Engineering Division, under Contract No. DEAC02-05-CH11231 within the Fundamentals of Semiconductor Nanowire Program (KCPY23). This research used the Savio computational cluster resource provided by the Berkeley Research Computing program at the University of California, Berkeley (supported by the UC Berkeley Chancellor, Vice Chancellor for Research, and Chief Information Officer).

DATA AVAILABILITY

The data that support the findings of this article are openly available [39].

APPENDIX A: MEAN-FIELD THEORY

To enable the use of automatic differentiation and to accelerate numerical optimization, we use a mean-field theory to propagate the various components. We assume that transition rates depend on the mean value of each fluctuating quantity, such that we propagate coupled ordinary differential equations for the equations of motion of each component,

$$\frac{d\langle n_i \rangle}{dt} = \sum_j J_{ij}(\langle m_j \rangle, \langle n_i \rangle), \quad (\text{A1})$$

$$\frac{d\langle m_j \rangle}{dt} = \sum_i J_{ji}(\langle m_j \rangle, \langle n_i \rangle), \quad (\text{A2})$$

where the sum goes over the connected components. Consequently, we obtain two coupled ordinary differential equations for DRAM, with access transistor occupancy n_A and capacitor state m_o , and eight coupled ODEs for SRAM, including N -type, P -type, and access transistor occupancies $n_i^{(k)}$, and output capacitor state $m_o^{(k)}$ for each of $k = 1, 2$ NOT gates.

This mean-field theory is accurate in the macroscopic limit, where individual electronic transitions become uncorrelated relative to the overall scale of the system. In practice we find that working capacitances of $C = 10$ and driving voltages $V_d/V_T = 3, 6$ are sufficient to yield quantitatively accurate results for our system as shown in the main text. For this system, fluctuations are roughly symmetry about the mean and therefore the theory works well. If the distribution were more skewed or we were interested in information about the tails of the distribution, a more




FIG. 6. Loss function starting from random seed.

sophisticated theory capturing the full trajectory dynamics would be necessary. Using this framework, we propagate mean-field dynamics within the optimization process for numerical facility, then confirm the results using kinetic Monte Carlo simulations with the optimized parameters.

APPENDIX B: OPTIMIZATION DETAILS

In a high-dimensional and nonconvex parameter space subject to nonequilibrium driving, finding an optimal solution is difficult. Engel *et al.* recently developed a method to find optimal protocols for nonequilibrium systems using automatic differentiation [30]. We use this method with our fully differentiable mean-field theory to propagate dynamics. For our bit erasure setup, we found that a learning rate of 10^{-3} achieved sufficient convergence within 1000 iterations of the Adam optimizer with otherwise standard parameters. We used $\lambda = 10$ to scale the relative contributions in our loss function, chosen empirically to ensure complete bit erasure while still allowing a dynamic range of heat dissipation to influence the optimization. We used several different initializations to generate optimal protocols, then used the best results. First, for each length of protocol, we optimized ten sets of randomly initialized seed coefficients and took the best of these ten as the optimal protocol. An example of the convergence of the optimization loss is shown in Fig. 6, and a successful erasure protocol was obtained. However, the roughness of the optimization landscape created inconsistencies in optimal protocols across parameter spaces.

Noticing two classes of solutions from the random seeds at short and long times, we used the shortest time optimal protocol to seed the next shortest time optimization iteratively, and the same from the longest time protocol. There was significant hysteresis in the optimal protocols obtained from these two sequential seeds. Finally, we took the long-time-seeded protocols and refined the optimization for each time point by seeding from time points slightly shorter and longer than the target operation time. The refined protocols are used for the data in the main text.




FIG. 7. Heat-accuracy trade-offs observed in DRAM with only wordline V_w control, averaged over starting states 0 and 1. (a) Error and (b) dissipated heat for two driving voltages $V_d/V_T = 3$ (green) and $V_d/V_T = 6$ (purple) across operation times. (c) Representative optimal protocol, similar for all operation times. Solid lines denote data from automatic differentiation, and dashed lines are the same from the linear response theory. Dotted horizontal lines show $CV_d^2/4$, half the energy of a capacitor.

APPENDIX C: RESTRICTED OPTIMIZATION

In a standard CMOS device undergoing bit erasure, the bitline is precharged to $V_b = 0$ and the wordline V_w is turned on or off. We performed a single parameter optimization of the wordline only to probe the effects of modulating current through the access transistor. The overall takeaways hold from the two-dimensional optimization: that there is a minimum time required for erasure, that error plateaus for times longer than the minimum, and that there is a trade-off between error and dissipation. In Fig. 7(a), we show for DRAM the same plateau of error that was observed in the optimization with both V_b and V_w , and that error is inversely related to driving voltage. However, given a fixed bitline voltage $V_b = 0$, dissipation also plateaus to the reversible charging energy of the capacitor, $E = CV_d^2/2$, rather than decreasing in the long-time limit in Fig. 7(b).

Optimal protocols for this constrained case shown in Fig. 7(c) are remarkably simple. There is no opportunity to adjust the bitline voltage adiabatically, so any electron exiting the system goes to a bath at potential $V_b = 0$. The only adjustable behavior is flux through the access transistor, and since the system is relaxing to an equilibrium




FIG. 8. Dissipated heat for DRAM under a linearly varying bitline voltage $V_b(t)$.

state, we need only to turn on the access transistor to allow the relaxation process to occur. Hence, the protocol stays near $V_w = 0$ for most of the time, still respecting the boundary conditions where it is in the “off” state at $V_w(0) = V_w(\tau_{\text{op}}) = -2V_d$.

Tradeoffs and optimal protocols were also computed for SRAM in this restricted framework. At fixed bitline voltage $V_b = 0$, the minimum dissipation is the reversible charging energy of the two output capacitors, $E = CV_d^2$. As a nonequilibrium system, SRAM incurs predominantly housekeeping heat at long times. Optimal protocols here are also controlling only flux through the access transistor. We cannot simply turn on the access transistor in SRAM as we did in DRAM because it is not simply an equilibrium relaxation process. However, the behavior is similar to that of the optimization with both V_b and V_w for SRAM, where short operation times minimize error by maximizing the discharging flux through the access transistor, and long operation times minimize dissipation by minimizing overall activity.

APPENDIX D: LINEAR RESPONSE

Some insight into the optimal protocols can be gained using linear response theory [40], applied in the limit of long operation times where erasure is assumed and the dissipated heat minimized. Such a theory is applicable to equilibrium systems, so we apply it to DRAM in which maintenance of the bit is not considered, only the erasure, and leakage current is ignored. Under these assumptions for a protocol that controls both wordline and bitline voltages, $\mathbf{V}(t) = \{V_w(t), V_b(t)\}$, the dissipation is given by

$$\langle Q \rangle = \int_0^{\tau_{\text{op}}} dt \frac{d\mathbf{V}}{dt} \cdot \boldsymbol{\zeta} \cdot \frac{d\mathbf{V}}{dt}, \quad (\text{D1})$$

where $\boldsymbol{\zeta}$ is a friction matrix, given by a matrix of integrated time correlation functions. In the case of DRAM, the relevant matrix has entries that only depend on the




FIG. 9. Dissipated heat for SRAM, linearly varying bitline voltage $V_b(t)$ for driving voltages (a) $V_d/V_T = 3$ and (b) $V_d/V_T = 6$.

autocorrelation function of the occupancy of the access transistor, as the driving force is solely a function of n_A . It follows from

$$-\frac{\partial E}{\partial \mathbf{V}} = \{n_A, -n_A\} \quad (D2)$$

that the friction matrix ζ is

$$\zeta = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \zeta, \quad (D3)$$

where

$$\zeta = \beta \int_0^\infty dt \langle \delta n_A(0) \delta n_A(t) \rangle = \frac{\beta}{4\Gamma(1 + \cosh(\beta q V_w))} \quad (D4)$$

is a correlation function that can be evaluated exactly for DRAM by solving for the average occupancy in equilibrium since $\langle \delta n_A^2 \rangle = \langle n_A \rangle (1 - \langle n_A \rangle)$.

The friction matrix is thus a function of V_w but independent of V_b . As a consequence, optimal protocols for V_b are piecewise linear in time, while those for V_w satisfy an

Euler-Lagrange equation of the form

$$\frac{d^2 V_w}{dt^2} = \frac{\beta q \sinh(\beta q V_w)}{2 + 2 \cosh(\beta q V_w)} \left(\frac{dV_w}{dt} \right)^2 \quad (D5)$$

with boundary conditions $V_w(0) = -2V_d$ and $V_w(\tau_{op}/2) = 0$, corresponding to the “off” and “on” states of the access transistor, respectively. As a controlled bit erasure requires the same start and end states, we symmetrize the optimal protocol to return to $V_w(\tau_{op}) = -2V_d$.

Considering first only controlling V_w , although this analytical form cannot predict the optimal length of protocol τ_{op} , nor can it tell us how long we should keep the transistor “on,” it informs us about the shape of the optimal protocol, shown in the dashed line in Fig. 7(c). We observe the same rapid increase in the control parameter as obtained by automatic differentiation, followed by a slower transition towards $V_w = 0$. The corresponding thermodynamic characteristics in dashed lines in Figs. 7(a) and 7(b) show that this protocol performs comparably to the optimal protocol from automatic differentiation, only slightly worse quantitatively.

For the bitline voltage V_b , the friction tensor is uniform and thus the optimal protocol is a linear ramp between the desired parameters. In the long time limit, manipulating the bitline voltage is analogous to manipulating the equilibrium position of a harmonic potential, and this operation is well within the linear response regime. We measured dissipation for a protocol smoothly varying from $V_b(0^+) = V_d$ to $V_b(\tau_{op}^-) = 0$ in Fig. 8. At short times below the characteristic relaxation time of the system, $\tau_{op} < \tau_{RC}$, dissipation grows as the capacitor relaxes to equilibrium with the bitline electrode. Starting in the logical 1 state, there is a maximum dissipated heat near the characteristic relaxation time $\tau_{op} \approx \tau_{RC} \sim 10^3$ before reaching a quasistatic, reversible limit. Our choice of protocol going from $V_b(0^+) = V_d$ to $V_b(\tau_{op}^-) = 0$ allows the capacitor starting in the logical 1 state to remain in equilibrium with the bitline electrode throughout the protocol, minimizing dissipation in the long-time limit. However, this protocol is necessarily suboptimal starting in the logical 0 state, as the capacitor must first be charged to reach equilibrium with the bitline electrode before relaxing back to the 0 state at long times. This charging cost is reflected in the asymptotic behavior as $Q \rightarrow CV_d^2/2$. Because we require a single protocol for either starting state, the average behavior reflects both contributions. It is important to note that we instantaneously perturb the bitline voltage from its boundary conditions at $V_b(0) = V_d/2$ to $V_b(0^+) = V_d$, and the same for $V_b(\tau_{op}^-) = 0$ to $V_b(\tau_{op}) = V_d/2$. The original formulation in Ref. [40] does not allow for such a “bang-bang” protocol, but we may still observe the intervening behavior in the linear response regime.

We have also explored analogous protocols for SRAM, whose behavior changes as a function of driving voltage.




FIG. 10. Output distributions for two driving voltages $V_d = 3, 6$ and operation times $\tau_{\text{op}} = 1000, 16000$. Colors correspond to starting states 0 (blue) and 1 (orange). Histograms show data from replicates of KMC simulations, and solid lines denote the mean field estimate of the same condition. Dashed red lines show the threshold voltage $0.1 V_d$.

For a small driving voltage $V_d/V_T = 3$ in Fig. 9(a), housekeeping heat dominates over variations due to the control protocol as thermal noise is close to the scale of driven currents in the system. For a sufficiently high driving voltage $V_d/V_T = 6$ in Fig. 9(b), we observe similar behavior as DRAM at intermediate operation times up to $\tau_{\text{op}} \sim 10^4$, where the control is slow enough that the system can respond to it semireversibly. As the driving voltage increases, SRAM reverts to an equilibriumlike behavior when activity is low, as transitions are sufficiently high-energy that they are unlikely. However, housekeeping heat still dominates in the long-time limit.

APPENDIX E: OUTPUT DISTRIBUTIONS

To clarify the accuracy of the outputs, we show in Fig. 10 the distributions corresponding to the data in Fig. 3 for DRAM. Similar distributions are observed for SRAM. As discussed in the main text, short operation times $\tau_{\text{op}} = 1000$ (left column) fail to erase the bit starting from the 1 state. Accuracy improves with higher driving voltage $V_d/V_T = 6$ (bottom row).

- [3] T. Sagawa and M. Ueda, Minimal energy cost for thermodynamic information processing: Measurement and information erasure, *Phys. Rev. Lett.* **102**, 250602 (2009).
- [4] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz, Experimental verification of Landauer's principle linking information and thermodynamics, *Nature* **483**, 187 (2012).
- [5] G. W. Wimsatt, A. B. Boyd, P. M. Riechers, and J. P. Crutchfield, Refining Landauer's stack: Balancing error and dissipation when erasing information, *J. Stat. Phys.* **183**, 16 (2021).
- [6] S. Dago, J. Pereda, N. Barros, S. Ciliberto, and L. Bellon, Information and thermodynamics: Fast and precise approach to Landauer's bound in an underdamped micromechanical oscillator, *Phys. Rev. Lett.* **126**, 170601 (2021).
- [7] K. Proesmans, J. Ehrich, and J. Bechhoefer, Optimal finite-time bit erasure under full control, *Phys. Rev. E* **102**, 032105 (2020).
- [8] P. R. Zulkowski and M. R. DeWeese, Optimal finite-time erasure of a classical bit, *Phys. Rev. E* **89**, 052140 (2014).
- [9] L. T. Giorgini, R. Eichhorn, M. Das, W. Moon, and J. S. Wettlaufer, Thermodynamic cost of erasing information in finite time, *Phys. Rev. Res.* **5**, 023084 (2023).
- [10] A. B. Boyd, A. Patra, C. Jarzynski, and J. P. Crutchfield, Shortcuts to thermodynamic computing: The cost of fast and faithful information processing, *J. Stat. Phys.* **187**, 17 (2022).
- [11] M. A. Ciampini, T. Wenzl, M. Konopik, G. Thalhammer-Thurner, M. Aspelmeyer, E. Lutz, and N. Kiesel, Erasure of a nonequilibrium memory beyond Landauer's bound using levitated optomechanics with spatio-temporal optical control, *Phys. Rev. Res.* **7**, 043321 (2025).
- [12] S. Dago, S. Ciliberto, and L. Bellon, Reliability and operation cost of underdamped memories during cyclic erasures, *Adv. Phys. Res.* **3**, 2300074 (2024).
- [13] Y. Taur and T. H. Ning, *Fundamentals of Modern VLSI Devices* (Cambridge University Press, Cambridge, United Kingdom; New York, NY, 1998).
- [14] J. Koomey, S. Berard, M. Sanchez, and H. Wong, Implications of historical trends in the electrical efficiency of computing, *IEEE Ann. Hist. Comput.* **33**, 46 (2011).
- [15] C. Y. Gao and D. T. Limmer, Principles of low dissipation computing from a stochastic circuit model, *Phys. Rev. Res.* **3**, 033169 (2021).
- [16] N. Freitas, J.-C. Delvenne, and M. Esposito, Stochastic thermodynamics of nonlinear electronic circuits: A realistic framework for computing around kT , *Phys. Rev. X* **11**, 031064 (2021).
- [17] J. Gu and P. Gaspard, Counting statistics and microreversibility in stochastic models of transistors, *J. Stat. Mech: Theory Exp.* **2020**, 103206 (2020).
- [18] L. W. Nagel and D. Pederson, SPICE (Simulation Program with Integrated Circuit Emphasis), Tech. Rep. UCB/ERL M382 (University of California, Berkeley, 1973).
- [19] P. Helms, S. W. Chen, and D. T. Limmer, Stochastic thermodynamic bounds on logical circuit operation, *Phys. Rev. E* **111**, 034110 (2025).
- [20] N. Freitas, K. Proesmans, and M. Esposito, Reliability and entropy production in nonequilibrium electronic memories, *Phys. Rev. E* **105**, 034107 (2022).

- [1] International Energy Agency, Energy and AI, Tech. Rep. (International Energy Agency, 2025).
- [2] R. Landauer, Irreversibility and heat generation in the computing process, *IBM J. Res. Dev.* **5**, 183 (1961).

- [21] C. Murphy, S. Nicholson, N. Freitas, E. Penocchio, and T. Gingrich, Dissipation-reliability tradeoff for stochastic CMOS bits in series, [arXiv:2603.04658](#) [cond-mat].
- [22] T. Shimizu, K. Chida, G. Yamahata, and K. Nishiguchi, Thermodynamic constraints in dynamic random-access memory cells: Experimental verification of energy efficiency limits in information erasure, *Phys. Rev. Lett.* **136**, 117103 (2026).
- [23] T. Basile and K. Proesmans, Optimal control of static RAM erasure: Arbitrarily fast operation with finite dissipation, *New J. Phys.* **27**, 104601 (2025).
- [24] L. Brillouin, Can the rectifier become a thermodynamical demon? *Phys. Rev.* **78**, 627 (1950).
- [25] D. T. Limmer, *Statistical Mechanics and Stochastic Thermodynamics* (Oxford University Press, Oxford, United Kingdom; New York, NY, 2024).
- [26] R. Sarpeshkar, T. Delbruck, and C. A. Mead, White noise in MOS transistors and resistors, *IEEE Circuits Devices Mag.* **9**, 23 (1993).
- [27] D. T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, *J. Comput. Phys.* **22**, 403 (1976).
- [28] D. T. Gillespie, Stochastic simulation of chemical kinetics, *Annu. Rev. Phys. Chem.* **58**, 35 (2007).
- [29] M. Esposito and C. Van den Broeck, Three detailed fluctuation theorems, *Phys. Rev. Lett.* **104**, 090601 (2010).
- [30] M. C. Engel, J. A. Smith, and M. P. Brenner, Optimal control of nonequilibrium systems through automatic differentiation, *Phys. Rev. X* **13**, 041032 (2023).
- [31] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](#) [cs].
- [32] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: composable transformations of Python + NumPy programs, 2018, <http://github.com/jax-ml/jax>
- [33] P. Kidger, On neural differential equations, PhD thesis, University of Oxford, 2021.
- [34] S. Whitelam, Nonequilibrium formulation of varying-temperature bit erasure, *J. Stat. Mech: Theory Exp.* **2024**, 063203 (2024).
- [35] L. Van Brandt and J.-C. Delvenne, in *2023 IEEE 23rd International Conference on Nanotechnology (NANO)* (IEEE, Jeju City, Korea, Republic of, 2023), pp. 726–731.
- [36] L. Van Brandt and J.-C. Delvenne, Noise–dissipation relation for nonlinear electronic circuits, *Appl. Phys. Lett.* **122**, 263507 (2023).
- [37] U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines, *Rep. Prog. Phys.* **75**, 126001 (2012).
- [38] K. Proesmans, J. Ehrich, and J. Bechhoefer, Finite-time Landauer principle, *Phys. Rev. Lett.* **125**, 100602 (2020).
- [39] S. W. Chen and D. T. Limmer, Code and data for “Optimal control of bit erasure in stochastic random access memory” (2026), <https://doi.org/10.5281/zenodo.19582802>
- [40] D. A. Sivak and G. E. Crooks, Thermodynamic metrics and optimal paths, *Phys. Rev. Lett.* **108**, 190602 (2012).