

Metastable Dynamical Computing with Energy Landscapes: A Primer

Christian Z. Pratt,^{*} Kyle J. Ray,[†] and James P. Crutchfield^{*‡}
*Complexity Sciences Center and Department of Physics and Astronomy,
 University of California, Davis, One Shields Avenue, Davis, CA 95616*
 (Dated: May 4, 2026)

Smartphones, laptops, and data centers are CMOS-based technologies that ushered our world into the information age of the 21st century. Despite notable advantages for scalable computing, CMOS implementations come with surprisingly large energetic costs. This challenge has revitalized scientific and engineering interest in energy-efficient information-processing designs. One current paradigm—dynamical computing—controls the location and shape of minima in potential energy landscapes that are connected to a thermal environment. The landscape supports distinguishable metastable energy minima that serve as a system’s mesoscopic memory states. Information is represented by microstate distributions. Dynamically manipulating the memory states then corresponds to information processing. This framing provides a natural description of the associated thermodynamic transformations and required resources. Appealing to bifurcation theory, a computational protocol in the metastable regime can be analyzed by tracking the evolution of fixed points in the state space. We illustrate the paradigm’s capabilities by performing 1-bit and 2-bit computations with double-well and quadruple-well potentials, respectively. These illustrate how dynamical computing can serve as a basis for designing universal logic gates and investigating their out-of-equilibrium thermodynamic performance.

Keywords: unconventional computing, thermodynamic computing, dynamical systems, bifurcation theory

We review a dynamical computing paradigm that manipulates microstate distributions in a metastable potential-energy landscape immersed in a thermal environment. A candidate landscape contains energy minima separated by sufficiently large barriers that stabilize mesoscopic memory states. Changing the landscape’s structure corresponds to information processing. Tracking fixed-point bifurcations then provides a path to designing energy-efficient computations.

I. INTRODUCTION

Over recent decades, integrated-circuit technologies contributed to a global abundance of computing—from laptops and smartphones to acre-sized data centers. While potentially offering boundless opportunities for humanity, these devices generate $\mathcal{O}(10^4)$ times more heat [1–4] than what is theoretically required [5]. Even before the recent rapid appearance of large language models [6], energy consumption for computational purposes was projected to reach 20% of global energy demand by 2030 [7]. It is clear that these CMOS-based technologies cannot serve global computational needs much longer.

To address the ever-growing demands, alternative computing paradigms are being explored in industry and academia. One approach, first explicated by Landauer in 1961 [5] and subsequently followed up by others [8–15],

focuses on controlling the minima in the potential energy landscape of a system coupled to a thermal environment. Coarse graining the continuous microstate phase space into a discrete set of high-probability regions surrounding the energy minima yields long-lived mesoscopic metastable memory states. If the landscape contains pairs of energy minima, then these information-bearing degrees of freedom correspond to the logical 0s and 1s in binary computing. Changing the memory states’ dynamics in the landscape corresponds to performing information processing—creating, destroying, and transforming the 0s and 1s. Equivalently, tracking the fixed points of the landscape’s associated dynamics permits the design of computational protocols.

We first detail the intuition underlying dynamical landscape computing in Section II. After this, we illustrate the design paradigm by reviewing two computations: Section III reviews the 1-bit information erasure operation [5]; then, Section IV outlines 2-bit *control erasure* [16], a generalization of Landauer’s erasure in a two-dimensional landscape. These examples demonstrate how dynamical landscapes can serve as a platform both for designing classical logic gates [16, 17] and for investigating their out-of-equilibrium thermodynamic performance [8–15].

II. FOUNDATIONS OF DYNAMICAL LANDSCAPE COMPUTING

How can a noisy dynamical system perform classical computations? We investigate this by building up the canonical energy landscape and its associated dynamics for 1-bit computing. First, let’s consider an arbitrary one-dimensional flat landscape $U(x)$ that is initially un-

^{*} czpratt@ucdavis.edu

[†] kjray@ucdavis.edu

[‡] Corresponding author: chaos@ucdavis.edu

connected to a thermal environment—see Fig. 1(a). We then introduce a particle into the potential to represent system information. In this setup, the particle freely travels in either direction along the x -axis according to its initial position and velocity. In other words, there is no notion of discrete information storage.

One way to implement discrete information storage in the landscape is to add energy wells, which are individually separated by energy barriers of equal height. Fig. 1(b) shows one example, where infinitely many minima are contained in the landscape. Now, coarse grain the continuum of microstates, grouping those into initial positions in each well with total energy less than the barrier height, into a discrete information-bearing memory state. Although these wells permit information storage, there is an infinite number of them. In contrast, today's complementary metal–oxide–semiconductor (CMOS) devices store bits—a unit of information with only two states, 0 or 1—to represent information.

With binary computational states in mind, we simplify the landscape to two wells separated by one energy barrier, as shown in Fig. 1(c). This provides binary information storage. The only force acting on the particle comes from the potential $-dU(x)/dx$. In other words, the system is conservative: If the particle has low enough energy, it will be confined to its current well; conversely, if it has enough energy to travel over the barrier from one minimum to the other, the particle will continue to oscillate back and forth without losing energy, being unable to settle into either well.

To mitigate this, we couple the system to a thermal environment, i.e., a sufficiently large ideal heat bath at equilibrium temperature T . As a result, a damping force $-\gamma dx/dt$ with coefficient γ , as well as a thermal noise term $\eta(t)$, now acts on the particle. A common choice for modeling $\eta(t)$ is Gaussian white noise, such that the particle evolves according to Langevin dynamics [18, 19]:

$$m \frac{d^2x}{dt^2} = -\gamma \frac{dx}{dt} - \frac{dU(x)}{dx} + \eta(t). \quad (1)$$

Alternatively said, the particle now evolves according to Newton's 2nd law due to a dissipative force, a conservative force, and a stochastic force. Despite the generality of Eq. (1), we restrict the system to follow overdamped Langevin dynamics, in which the damping coefficient $\gamma \gg 1$, rendering the particle's inertial effects to be negligibly small [20]:

$$\gamma \frac{dx}{dt} = -\frac{dU(x)}{dx} + \eta(t). \quad (2)$$

Consequently, the particle will settle into the phase space region surrounding a minimum, as seen in Fig. 1(d). In other words, the particle's velocity-relaxation timescale is $\tau_m = m/\gamma \ll \tau_{\text{well}}$, where τ_{well} characterizes the timescale over which the particle settles into its respective well. By coupling the system to the heat reservoir, its state is now better represented as a distribution of microstates. The particle now represents various stochastic

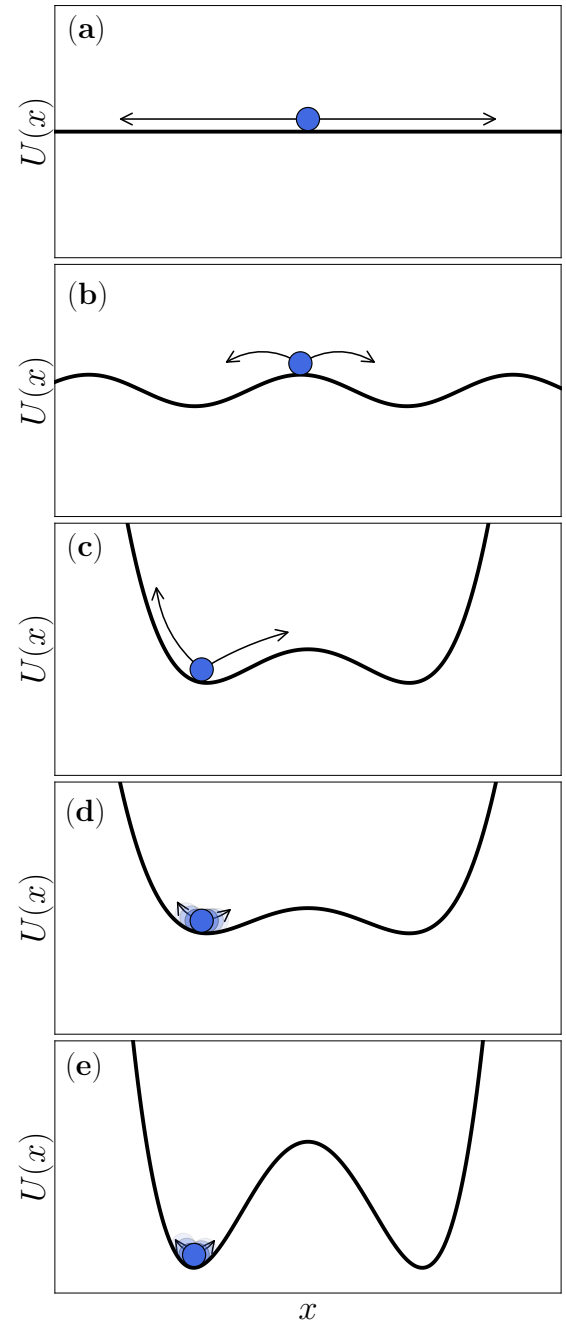


FIG. 1. Constructing a potential energy landscape $U(x)$ that stores and processes a single bit of information: (a) Flat landscape. (b) Landscape with an arbitrary number of energy minima. (c) Instantiating two minima in the landscape, but with no thermal environment connection. (d) Two wells but with damped and noise-perturbed particle motion. (e) Double-well potential with a sufficiently large barrier height for stable information storage.

realizations of its dynamics. Thermal noise randomly perturbs the particle within its respective well, mixing and merging trajectories with different histories into the same long-time behavior.

However, rare noise fluctuations can be large enough to cause transitions between the wells. This renders the landscape's mesoscopic memory states to not be purely stable, but *metastable* [21] as they balance noise and dissipation. For the one-dimensional potential in Fig. 1(d), the particle's average energy in a metastable state is on the scale of $k_B T$, where k_B is the Boltzmann constant [18, 21, 22]. If the energy barrier height is not much larger than $k_B T$, there is a high probability that the particle will fluctuate over the barrier, producing a thermally-activated information storage error [18, 22]. More precisely, the probability of an energy fluctuation of size ΔE is proportional to $\exp(-\Delta E/k_B T)$ [18, 21, 22], where ΔE represents the height of the energy barrier.

Reducing these errors involves raising the energy barrier height to be sufficiently higher than $k_B T$, as illustrated in Fig. 1(e). Since the distribution is only stable relative to a timescale that depends on the height of the energy barrier, this now ensures that the state can be stored over long time scales with an exponentially small probability of error. More quantitatively, $\tau_m \ll \tau_{\text{well}} \ll \tau_{\text{escape}}$, where τ_{escape} denotes the timescale of escapes [21]. We now deem the phase space regions surrounding minima to be the *metastable memory states* of the landscape. Together, the mesoscopic states on either side of the x -axis represent one bit of information. This landscape serves as our computational substrate for executing 1-bit computations. Sections III and IV, consider protocols such that $\tau_m \ll \tau_{\text{well}} \ll \tau_{\text{protocol}} \ll \tau_{\text{escape}}$ for which τ_{protocol} gives the duration of the landscape control protocol.

As a final note, although our focus is on the dynamical-systems nature of the operations, the framework readily supports explicit thermodynamic accounting. Following Sekimoto, the stochastic work cost $W = \int_0^{\tau_{\text{protocol}}} dt \partial_t U(x(t), t)$ of an operation is the change in system energy induced by driving the time-dependent potential, while the heat exchanged with the bath $Q = \Delta U - W$ is the system's response to that driving [19, 23].

The double-well potential in Fig. 1(e) can be described by a 4th-order polynomial:

$$U(x) = \frac{1}{4}ax^4 - \frac{1}{2}bx^2 + cx, \quad (3)$$

where a , b , and c are real-valued scalars that serve as the potential's control parameters. Their time dependence is the landscape's *control protocol*. Generally speaking, the potential can be continuously transformed into a single well depending on the values of a and b , while c tilts the potential to the positive (negative) side of the x -axis if its value is negative (positive).

To illustrate the role of the associated state-space geometry and the use of bifurcation theory, consider a concrete example of a fixed-point analysis for Eq. (3) which can produce Fig. 1(e). First, let $a = 1$ so that b provides complete control of the barrier height. Setting $c = 0$ enforces a potential symmetric about $x = 0$. To start the fixed point analysis, we find values of x that set the

derivative of $U(x)$ with respect to x to zero. These critical values $\{x^*\}$ are the potential's fixed points:

$$\left. \frac{dU(x)}{dx} \right|_{x=x^*} = x^{*3} - bx^* = 0, \quad (4)$$

with solutions $x^* = 0, \pm\sqrt{b}$. We have 3 real fixed points when b is positive. We can assess the stability of $U(x = x^*)$ to understand how microstate distributions interact with the potential by evaluating $-d^2U(x = x^*)/dx^2$:

$$-\frac{d^2U(x = x^*)}{dx^2} = 3x^{*2} - b. \quad (5)$$

Next, $x^* = \pm\sqrt{b}$ results in $d^2U(x^* = \pm\sqrt{b})/dx^2 = 2b > 0$. Now, both fixed points are stable. This validates our choice of the potential in Fig. 1(e).

III. SINGLE-BIT ERASURE

By changing the landscape's barrier height and the locations of its minima, particles can be controlled to travel between the system's memory states. Said differently, dynamically manipulating the landscape's metastable states corresponds to information processing.

We will demonstrate this with two information erasure protocols that are inspired by the the Restore-to-One (RT1) protocol [5], whose differences can be seen through a bifurcation theory perspective: A pitchfork bifurcation protocol versus a saddle-node bifurcation protocol.

A. Pitchfork protocol

We first introduce a particle into the landscape. If it is sampled in equilibrium, it has a 50% probability of being in one or the other well. Regardless of the particle's initial well, the goal of the computation is to bring it to the 1 state [5]. We show this in Fig. 2(a) at $t = t_0$ with each particle representing a potential starting location.

Assuming that a particle remains in local metastable equilibrium throughout the protocol (i.e., $\tau_{\text{well}} \ll \tau_{\text{protocol}}$), we can equivalently view the system response from a bifurcation theory perspective by tracking only the potential's fixed points. A particle then represents a sampling of a locally-stable distribution in the landscape. Its mean location corresponds to a respective landscape's stable fixed point. And, its higher moments are determined by the local curvature of the potential that guides the particle's motion in the region surrounding an energy minima. This perspective is illustrated in Fig. 2(b) at $t = t_0$, where the potential has two (one) stable (unstable) fixed points that are colored as blue (yellow) diamonds.

First, we lower the energy barrier separating the two wells. This transforms the double-well potential into a single well and causes the two possible particle locations to converge to a single minimum, as displayed in

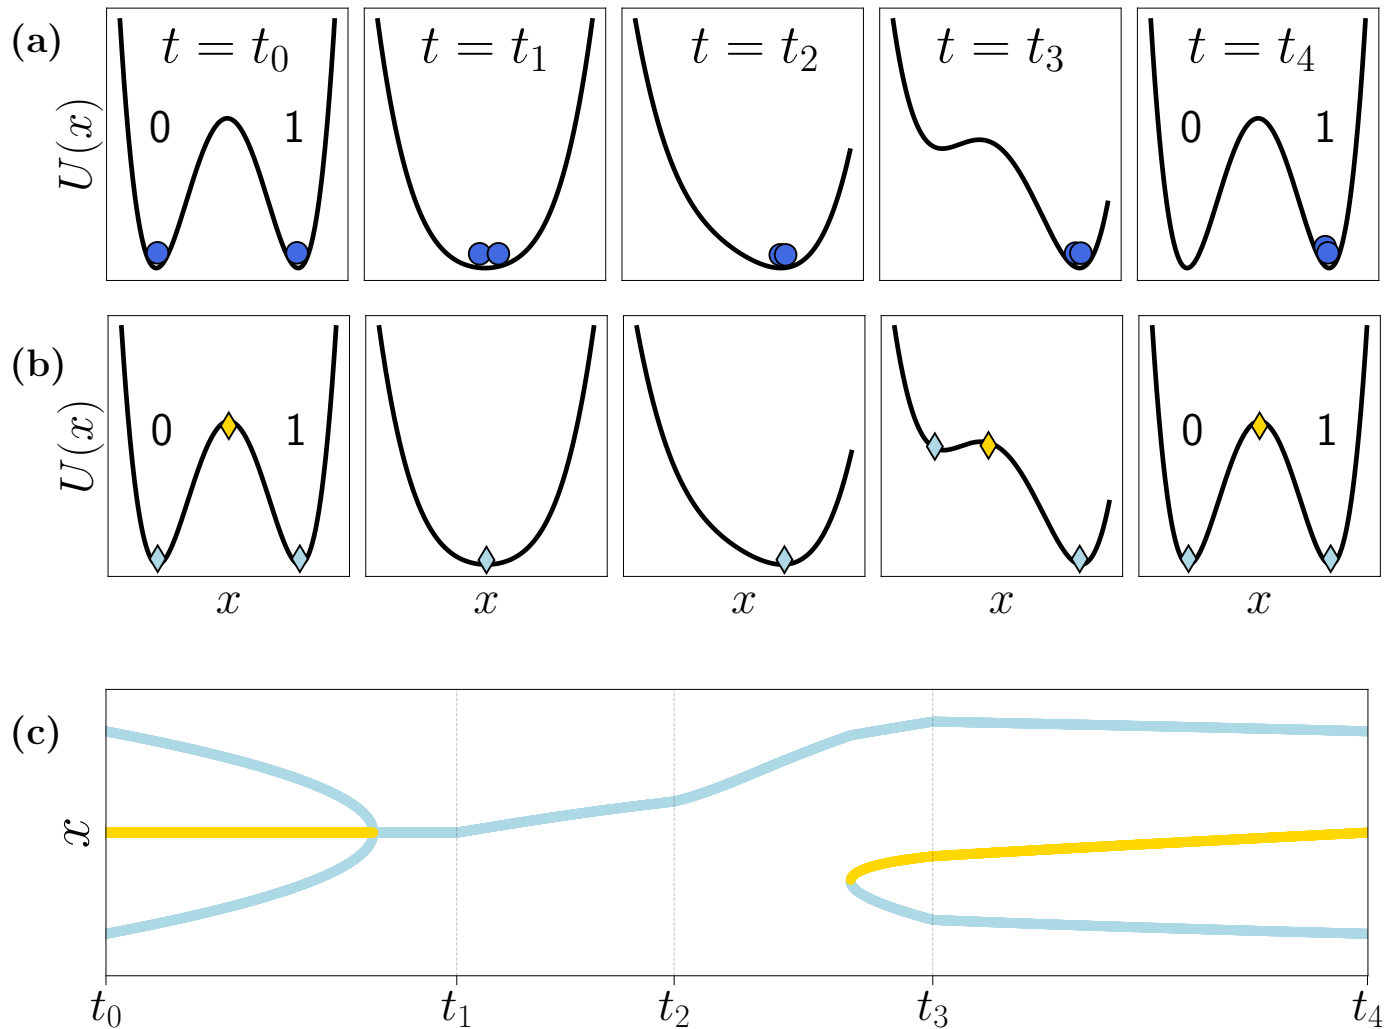


FIG. 2. Pitchfork-bifurcation information erasure protocol implementing Landauer's Restore-to-One (RT1) operation from different dynamical perspectives, taken at different points in time. (a) RT1 protocol showing both possible initial particle locations that end in the 1 state. (b) RT1 protocol only illustrating the potential's fixed points. (c) A dynamical view of information erasure, demonstrating that erasure protocols are easily understood using bifurcation theory. The blue (yellow) lines represent the stable (unstable) fixed point trajectories throughout the protocol. In dynamical-systems language, overall the protocol is quite simple—a reverse pitch-fork bifurcation followed by a reverse saddle-node bifurcation.

Fig. 2(a) at $t = t_1$. In dynamical-systems language, this corresponds to a pitchfork bifurcation that forms one single stable fixed point in Fig. 2(b) at $t = t_1$. We choose the pitchfork bifurcation since we are able to perform this step adiabatically in the quasi-static limit ($\tau_{\text{protocol}} \rightarrow \infty$ relative to τ_{well} , while still holding $\tau_{\text{protocol}} \ll \tau_{\text{escape}}$), where each particle remains close to its respective potential energy minima. In this limit, a particle stays arbitrarily close to local equilibrium if the protocol is done arbitrarily slowly. The next subsection showcases a protocol that fundamentally differs in this step and discusses the relative energetic efficiencies.

Next, to ensure that the particle ends in the 1 state regardless of its initial start state, the potential tilts towards the positive x -axis. Fig. 2(a) at $t = t_2$ demon-

strates this, where the particle moves into the right half-plane. Correspondingly, the landscape's single stable fixed point is translated towards the positive x -axis, as illustrated in Fig. 2(b) at $t = t_2$.

Now, to maintain the particle's location while completing the computation, we raise the energy barrier to its original height while keeping the potential tilted, as seen in Fig. 2(a) at $t = t_3$. This recreates the well in the left half-plane. Due to the large barrier height, the particle is likely to stay in right half-plane. This protocol step can be viewed as the creation of one stable and one unstable fixed point, bringing the total number of fixed points to match the start of the protocol as illustrated in Fig. 2(b) at $t = t_3$. This is a reverse saddle-node bifurcation.

The final step returns the landscape back to its original configuration, completing the protocol cycle. The potential first untilts to bring the landscape back to its initial configuration. In this way and independent of initial location, particles now reside in the 1 state as desired; see Fig. 2(a) at $t = t_4$. In Fig. 2(b) at $t = t_4$, this corresponds to the fixed points converging to their original locations. This final step completes the RT1 protocol. Note that the protocol's cyclicity ensures accounting for all relevant information and thermodynamic costs related to the operation—analogue to forming a closed engine cycle in classical thermodynamics.

Introducing the bifurcation theory perspective of the RT1 protocol emphasizes that we need to track the entire protocol only from the fixed point paths in the state space. Using this vantage point, we are able to further simplify our understanding of what a computation is—we no longer need to track the full trajectory of the microstate distribution. Fig. 2(c) illustrates the entire protocol in terms of the trajectories of the potential's fixed points.

B. Saddle-node protocol

As an alternative to relying on a pitchfork bifurcation, now perform a RT1 protocol that uses a saddle-node bifurcation. Here, the potential's parameters are tuned differently than with pitchfork erasure, but all possible input states are still erased to the 1 state.

In Fig. 3(a) at $t = t_0$, the potential and particles are initialized in the same manner as the previous protocol. Likewise, Fig. 3(b) at $t = t_0$ shows the same starting locations of the fixed points as Fig. 2(b) at $t = t_0$.

At $t = t_1$, we drop the barrier slightly, as well as tilt the potential towards the positive x -axis, as shown in Fig. 3(a) at $t = t_1$. Then Fig. 3(a) at $t = t_2$ (Fig. 3(b) at $t = t_2$) shows the completion of this step from a particle (fixed point) perspective. Manipulating the potential in this way annihilates the unstable fixed point with the stable fixed point on the negative x -axis, while maintaining the fixed point on the positive x -axis. This is the essence of the saddle-node bifurcation: We are able to store information on one side of an axis throughout the protocol, while annihilating other fixed points in the landscape.

The final step re-raises the energy barrier and untilts the potential, as shown in Fig. 3(a) at $t = t_3$ and Fig. 3(b) at $t = t_3$. This completes a protocol cycle, returning the potential back to its original configuration and bringing both microstate distributions to the 1 state.

Now, having completed a saddle-node erasure, let's detail the differences between this protocol and the pitchfork erasure. Due to how the saddle-node protocol controls the landscape, the particle that was originally stored in the 0 state will gain kinetic energy as it travels from the bifurcation point where its local minima vanished, to the well associated with the 1 state. In order to settle

into local equilibrium, the particle must dissipate this energy. This manifests as more irreversible dissipation into the heat bath, along with the Landauer bound cost—the inevitable work cost needed to perfectly erase an unbiased input bit, that being $k_B T \ln 2$ [5]. As a consequence, the saddle-node bifurcation cannot be carried out adiabatically, even in the quasi-static limit. Due to this, to accurately calculate thermodynamic quantities, the time it takes for the particle to settle into local equilibrium must be accounted for.

This fundamentally differs from the pitchfork protocol. There, since the particles remain in the phase space region surrounding their respective local minima throughout the entire protocol, the process can be made adiabatic when performed infinitely slowly, allowing it to reach the Landauer limit. From a dynamical-systems theory perspective, because we form one stable fixed point from the bifurcation of two stable and one unstable fixed point, the pitchfork protocol can be more energy efficient.

By viewing these computations in terms of the landscape's fixed point dynamics, thermodynamic costs can be more readily understood. Additionally, considering fixed-point creations and annihilations allows for analyzing dynamical computations in higher dimensions, which is the subject of the next section.

IV. 2-BIT DYNAMICAL COMPUTATIONS

Let's extend the landscape computing paradigm to implement processing two bits of information. Previously, to represent one bit of information required a double-well landscape. This suggests that a landscape that contains four wells provides four different memory states.

Fig. 4(a) shows an example potential that contains four wells, each separated by sufficiently large energy barriers, and connected to a thermal environment, as required. Similar to Eq. (3), its form can be created from a general quartic function, but in two dimensions:

$$\begin{aligned}
 U(x, y) = & \frac{1}{4}ax^4 - \frac{1}{2}bx^2 + cx \\
 & + \frac{1}{4}dy^4 - \frac{1}{2}ey^2 + fy \\
 & + gxy .
 \end{aligned} \tag{6}$$

The landscape's control parameters are now a , b , c , d , e , f , and g . Allowing $a = d = 1$, and $c = f = g = 0$ produces the example landscape displayed in Fig. 4(a); changing the values of b and e determines the barrier height along the y and x axes, respectively. Note that if $b \neq e$, then the potential has asymmetrical barrier heights with respect to either axis.

Performing a bifurcation analysis in a similar way to Sec. II yields four stable fixed points, corresponding to the four wells in the potential; four saddle points, associated with the transitions between adjacent memory states; and an unstable point at the origin. We coarse-

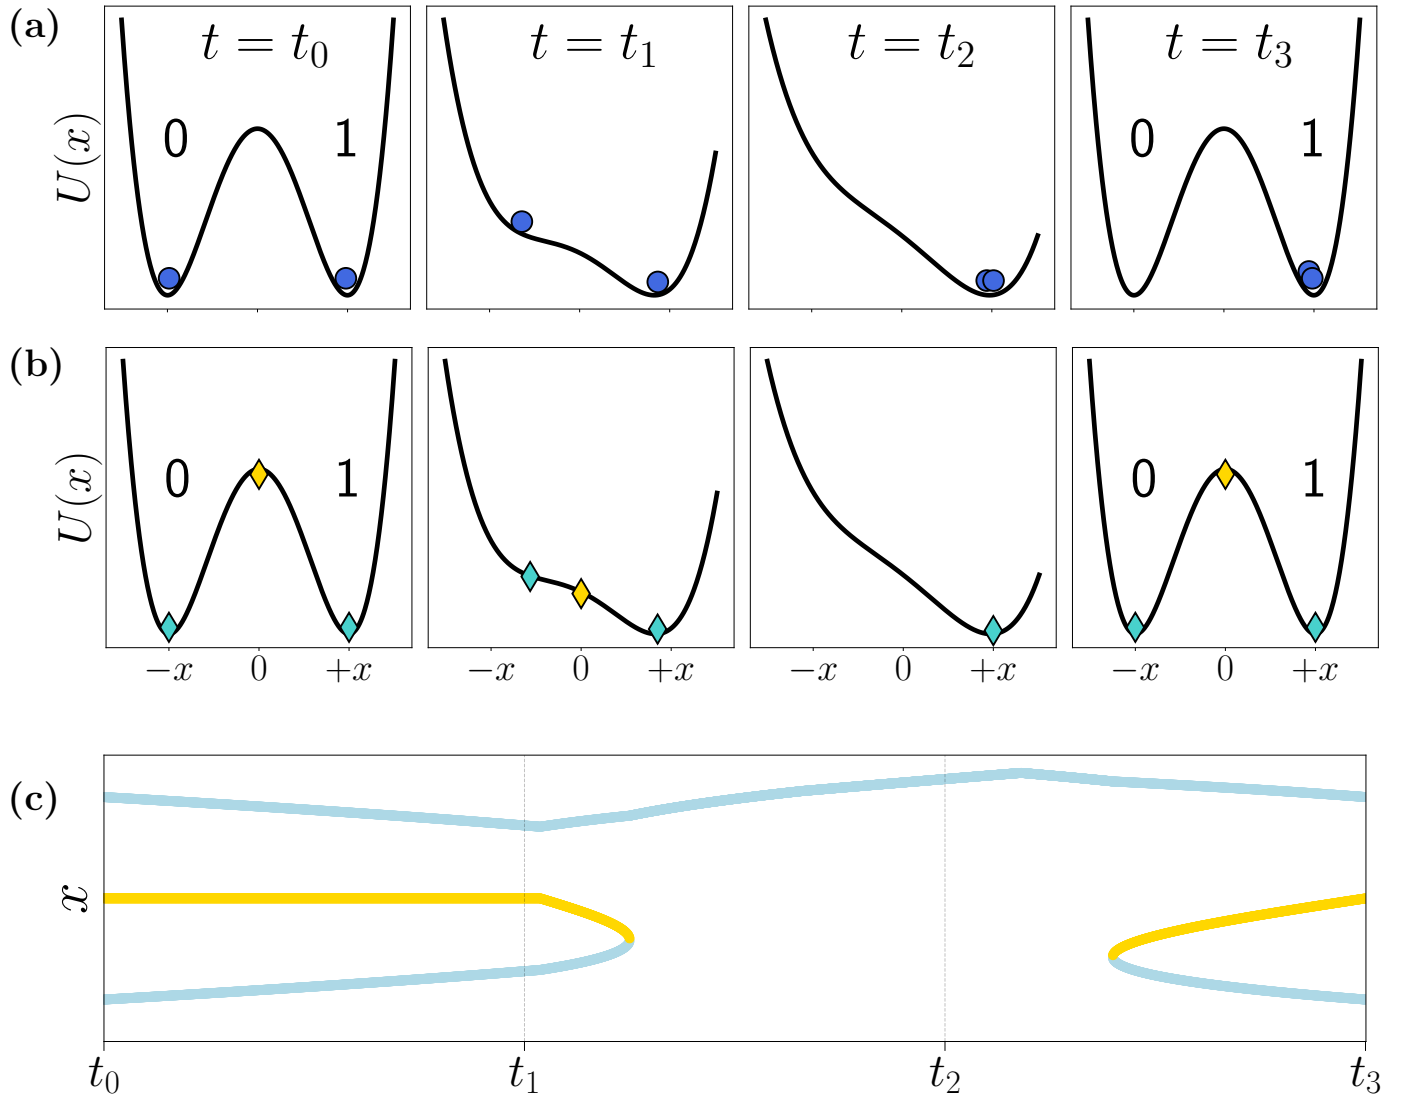


FIG. 3. A saddle-node bifurcation-inspired information-erasure protocol. Differing from the pitchfork protocol, it cannot be carried out infinitely slowly to make it adiabatic and energy efficient. Throughout the entire protocol, there are two saddle-node bifurcations. (a) The evolution of the potential during the protocol when keeping track of the microstate distributions. (b) The same protocol, but following the landscape's fixed points. (c) The saddle-node erasure protocol in terms of the landscape's fixed points only.

grain the potential so that its memory states are distinguished according to the minima's location with respect to the x and y axes:

$$x[y] = \begin{cases} 0 & \text{if } x[y] < 0, \\ 1 & \text{if } x[y] > 0. \end{cases} \quad (7)$$

This choice of memory-state assignment is shown in Fig. 4(b).

To illustrate an example of a computation with this landscape, we first detail the saddle-node erasure protocol that was introduced in Ref. [16]. There, a particular physical device was assumed, which provides a potential energy surface that is qualitatively similar to Eq. (6),

but it has a different mathematical form and control parameters.

We start by considering the landscape's fixed points, which are depicted as colored diamonds in Fig. 5(a)(i-iii). Further coarse graining, the dynamical point of view directly reveals the potential's fixed points and flow fields, shown in Fig. 5(b)(i-iii). Here, blue (red) arrows represent stable (unstable) fixed points along with a gray maxima. This "dynamical skeleton" illustrates where microstate distributions are favored to travel without being concerned about the potential's exact shape.

Analogous to Sec. III B, we need to lower the barrier separating a minimum, but because we have a four-well potential, we must subsequently tilt the potential in a




FIG. 4. Potential energy landscape that supports carrying out 2-bit computations: (a) Illustrating the two-dimensional landscape $U(x, y)$ via a three-dimensional landscape plot above its projected contour plot. (b) Contour plot of $U(x, y)$ showing example memory-state instantiations.

way that erases some information on one side of an axis, but stores the rest on the other axis. The simplest, and most energy-efficient, option is to implement a pitchfork bifurcation erasure to merge the two wells in the negative x semi-plane. However, such a pitchfork protocol is not possible with the form of the potential in Eq. (6) without destroying stability between the other pair of wells; see Appendix A for the mathematical details.

Instead, we implement the erasure using a saddle-node

bifurcation, as in Fig. 5(a), by lowering the barrier in the left half-plane and then tilting the potential in such a way to annihilate the red and green fixed points while maintaining all other fixed points. The situations just before and after the annihilation are shown in Fig. 5(a)(ii) and Fig. 4(a)(iii), as well as in Fig. 5(b)(ii) and Fig. 5(b)(iii), respectively. Here, the information that was in the third quadrant is erased into the second quadrant, while all other information is preserved. Completing the protocol involves re-raising and un-tilting the potential in an analogous way to the RT1 computation. This erasure protocol is called a *control erasure* (CE) [16]: Its functionality involves controlling what input information is erased or preserved in order to perform the 2-bit computation. Interestingly, there are eight possible CEs that can be performed with the quadruple well potential [16].

V. CONCLUSION

With the attenuation of Moore's law colliding with the accelerating demand for computational resources, now is the time to investigate alternative computing paradigms. This is critical for future progress in computing. Here, to address this and show a promising way forward, we described a design strategy for performing computations with dynamical potential-energy landscapes. This perspective of computing brings dynamical-systems theory, statistical mechanics, and nonequilibrium thermodynamics together in a coarse-grained approach to understand thermally-activated classical computations. This computing framework differs from other related dynamical-systems paradigms that assume a physical substrate, such as chaos computing [24, 25] and nonlinear memristive computing [26, 27]. Detailed investigations of thermodynamic quantities and higher-dimensional computations are readily accessible with this framework, opening up new options for optimizing computing systems both in terms of raw performance and energy costs.

ACKNOWLEDGMENTS

We thank Greg Wimsatt and Kuen Wai Tang for insightful feedback discussions. The authors also thank the Telluride Science Research Center for its hospitality during visits and the participants of the Information Engines workshop there for their valuable feedback. J.P.C. acknowledges the kind hospitality of the Santa Fe Institute. This material is based on work supported by, or in part by, the U.S. Army Research Laboratory and U.S. Army Research Office under Grant No. W911NF-21-1-0048, and by the Art and Science Laboratory via a gift to UC Davis.

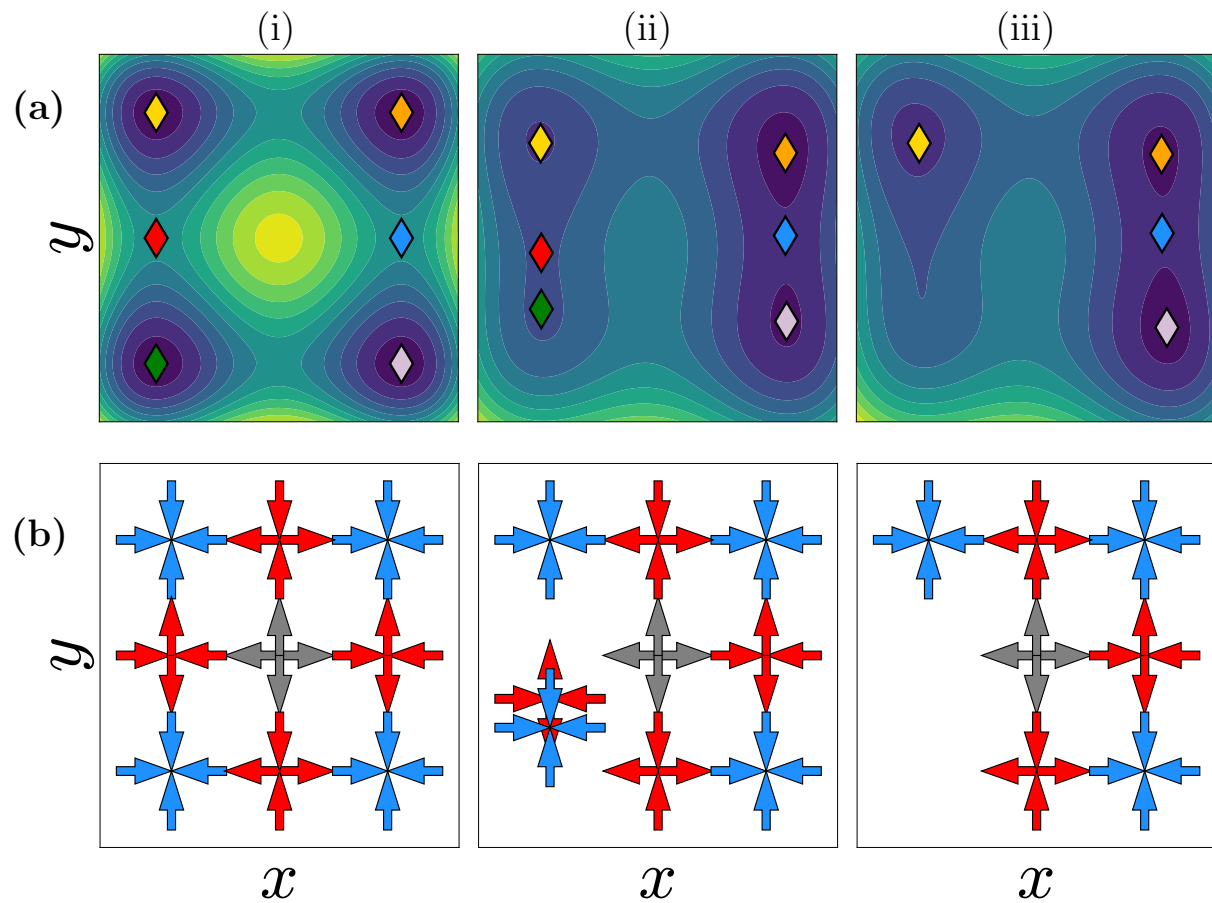


FIG. 5. Dynamical view of the control erasure (CE) protocol, as motivated by Ref. [16]: (a) The CE characterized by the annihilation of the red and green fixed points in the third quadrant, while maintaining all other fixed points. (b) A dynamical skeleton of the CE that only shows the landscape's fixed points and flow fields. Blue (red) corresponds to stable (unstable) fixed points, as well as a grey local maxima. The steps of the protocol are as follows: (i) Initialization. (ii) Drop barrier separating the green (red) stable (unstable) fixed points, while maintaining all other fixed points. (iii) Just after the red-green fixed point annihilation.

AUTHOR DECLARATIONS

The authors have no conflicts to disclose.

AUTHOR CONTRIBUTIONS

All authors developed thermodynamic landscape computing. CZP was responsible for generating the simulation data, formal analysis, and visualizations. KJR and JPC verified the theoretical framework and supervised the development. JPC provided funding. All authors contributed to the manuscript writing and editing, and approved the final version for submission.

COMPETING INTERESTS

The authors declare no competing interests.

DATA AVAILABILITY

The data that support this review are available from the corresponding author upon reasonable request.

Appendix A: Impossibility of a Pitchfork Bifurcation in a Specific Landscape

Following Refs. [28, 29] and as discussed in Section III A, a pitchfork bifurcation occurs when two fixed points on either side of an axis converge to a point of annihilation at the coordinate axis. We will now show the impossibility of performing a pitchfork bifurcation on a pair of two adjacent wells in the potential generated by Eq. (6) and displayed in Figure 4 without destroying stability of the other pair of wells as well. As a consequence, the control erasure (CE) described in Section IV can only be performed using a saddle-node bifurcation.

First, consider a pitchfork bifurcation in a one-dimensional potential $U(x)$, which is useful for analyzing the two-dimensional case. A common ordinary differential equation that produces a one-dimensional pitchfork bifurcation in the landscape $U(x)$ takes on the polynomial form [29]:

$$\dot{x} = -\frac{dU(x)}{dx} = rx - x^3. \quad (\text{A1})$$

Upon observation, the bifurcation parameter appears in the linear term, and Eq. (A1) is odd. Said another way, a pitchfork bifurcation is found when the local expansion around a bifurcation point is an odd function with non-vanishing linear and cubic terms [28, 29].

Now, the mathematical form of Eq. (6) does not permit carrying out a CE using a pitchfork bifurcation. To see this, assume that a pitchfork bifurcation can occur at the point $(x, y) = (x_0, y_0)$ in Eq. (6). Then, expand the first derivative of the potential $U(x, y)$ around x_0 :

$$\begin{aligned} \partial_x U(x_0 + h, y_0) &= \partial_x^2 U(x_0, y_0)h + \partial_x^3 U(x_0, y_0)h^2/2 + \partial_x^4 U(x_0, y_0)h^3/6 \\ &= (3ax_0^2 - b)h + 3ax_0h^2 + ah^3. \end{aligned}$$

Enforcing that this function takes the form of Eq. (A1)

means that the linear and cubic terms remain while the quadratic one vanishes. If $a = 0$ then the cubic term vanishes, so we find that $x_0 = 0$ is the only way for the quadratic term to vanish. Now, because $x_0 = 0$, this means that either $b \neq 0$ or the linear term vanishes. In order to satisfy the criteria for a pitchfork bifurcation, then $x_0 = 0$ and $a, b \neq 0$. Accordingly, we restrict our analysis only to those fixed points at $(x = 0, y = y_0)$.

As detailed in Sec. II, the stability of a candidate fixed point is controlled by the curvature of the potential at $x = x_0$. This gives $\partial_x^2 U(0, y) = -b$. For this to be a stable fixed point, then $b < 0$. That said, suppose this pitchfork bifurcation happens with the 00 and 01 wells in the negative y semi-plane from Fig. 4. At the same time, to perform a CE, we must maintain the 10 and 11 states, as well as the respective energy barrier separating them. This requires $b > 0$, which contradicts our original restriction of $b < 0$. As a result, the bifurcation cannot happen at $x_0 = 0$, and so it is not a pitchfork bifurcation [29].

In other words, we are unable to form a stable fixed point from an annihilation of the two lower wells, while also maintaining the energy barrier in the upper half-plane. As a result, to perform a CE with the potential in Eq. (6), we are limited to the less efficient saddle-node bifurcation as discussed in Sec. IV.

-
- [1] 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(3):031064, 2021.
 - [2] C. Y Gao and D. T. Limmer. Principles of low dissipation computing from a stochastic circuit model. *Phys. Rev. Res.*, 3(3):033169, 2021.
 - [3] N. Takeuchi, T. Yamae, C. L. Ayala, H. Suzuki, and N. Yoshikawa. Adiabatic quantum-flux-parametron: A tutorial review. *IEICE Trans. Elec.*, E105.C(6):251–263, 2022.
 - [4] Intel Corporation. Intel Xeon Platinum 8280 Processor.
 - [5] R. Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.*, 5(3):183–191, 1961.
 - [6] S. Kamatala, A. K. Jonnalagadda, and P. Naayini. Transformers beyond NLP: Expanding horizons in machine learning. *IRE J.*, 8(7), 2025.
 - [7] N. Jones. How to stop data centres from gobbling up the world's electricity. *Nature*, 561(7722):163–166, 2018.
 - [8] C. H. Bennett. The thermodynamics of computation—a review. *Intl. J. Theo. Physics*, 21(12):905–940, 1982.
 - [9] K. Sekimoto. Kinetic characterization of heat bath and the energetics of thermal ratchet models. *J Phys. Soc. Japan*, 66(5):1234–1237, 1997.
 - [10] M. B. Plenio and V. Vitelli. The physics of forgetting: Landauer's erasure principle and information theory. *Contemporary Physics*, 42(1):25–60, 2001.
 - [11] V. Blickle, T. Speck, L. Helden, U. Seifert, and C. Bechinger. Thermodynamics of a colloidal particle in a time-dependent nonharmonic potential. *Phys. Rev. Lett.*, 96(7):070603, 2006.
 - [12] Y. Jun, M. Gavrilov, and J. Bechhoefer. High-precision test of Landauer's principle in a feedback trap. *Phys. Rev. Lett.*, 113(19):190601, 2014.
 - [13] P. M. Riechers. *Transforming Metastable Memories: The Nonequilibrium Thermodynamics of Computation*, page 353–381. SFI Press, 2019.
 - [14] P. M. Riechers, A. B. Boyd, G. W. Wimsatt, and J. P. Crutchfield. Balancing error and dissipation in computing. *Phys. Rev. Res.*, 2(3):033524, 2020.
 - [15] 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. Physics*, 187(2):17, 2022.
 - [16] C. Z. Pratt, K. J. Ray, and J. P. Crutchfield. Controlled erasure as a building block for universal thermodynamically robust superconducting computing. *Chaos: Interdisc. J. Nonlin. Sci.*, 35(4), 2025.
 - [17] K. J. Ray and J. P. Crutchfield. Gigahertz sub-Landauer momentum computing. *Phys. Rev. Applied*, 19(1):014049, 2023.
 - [18] S. Han, J. Lapointe, and J. E. Lukens. Effect of a two-dimensional potential on the rate of thermally induced escape over the potential barrier. *Phys. Rev. B*, 46(10):6338–6345, 1992.
 - [19] K. Sekimoto. Langevin equation and thermodynamics. *Prog. Theo. Physics Supp.*, 130:17–27, 1998.
 - [20] N. Shiraishi. An introduction to stochastic thermodynamics. *Fundamental Theories of Physics. Springer, Singapore*, 2023.
 - [21] P. Hanggi. Escape from a metastable state. *J. Stat.*

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0327461

- Phys.*, 42(1–2):105–148, 1986.
- [22] S. Han, J. Lapointe, and J. E. Lukens. Thermal activation in a two-dimensional potential. *Phys. Rev. Lett.*, 63(16):1712–1715, 1989.
- [23] U. Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75(12):126001, 2012.
- [24] T. Munakata, S. Sinha, and W. L. Ditto. Chaos computing: Implementation of fundamental logical gates by chaotic elements. *IEEE Trans. Circ. Sys. I: Fund. Th. App.*, 49(11):1629–1633, 2002.
- [25] B. Majumder, S. Hasan, M. Uddin, and G. S. Rose. Chaos computing for mitigating side channel attack. In *2018 IEEE Intl. Symp. Hardware-oriented Security Trust (HOST)*, pages 143–146. IEEE, 2018.
- [26] J. J. Yang, D. B. Strukov, and D. R. Stewart. Memristive devices for computing. *Nature Nanotechn.*, 8(1):13–24, 2013.
- [27] K. Yang, J. J. Yang, R. Huang, and Y. Yang. Nonlinearity in memristors for neuromorphic dynamic systems. *Small Science*, 2(1):2100049, 2022.
- [28] S. Wiggins. *Introduction to Applied Nonlinear Dynamical Systems and Chaos*, volume 2 of *Texts in Applied Mathematics*. Springer-Verlag, New York, 2003.
- [29] Steven H. Strogatz. *Nonlinear Dynamics and Chaos*. CRC Press, second edition, May 2018.