

A Novel Operational Paradigm for Thermodynamically Reversible Logic

Adiabatic Transformation of Chaotic Nonlinear Dynamical Circuits

Michael P. Frank*, Erik P. DeBenedictis,

Center for Computing Research
Sandia National Laboratories
Albuquerque, NM 87185-1319
*mpfrank@sandia.gov

Abstract—Continuing to improve computational energy efficiency will soon require developing and deploying new operational paradigms for computation that circumvent the fundamental thermodynamic limits that apply to conventionally implemented Boolean logic circuits. In particular, Landauer’s principle tells us that irreversible information erasure requires a minimum energy dissipation of $kT \ln 2$ per bit erased, where k is Boltzmann’s constant and T is the temperature of the available heat sink. However, correctly applying this principle requires carefully characterizing what actually constitutes “information erasure” within a given physical computing mechanism. In this paper, we show that abstract combinational logic networks can validly be considered to contain no information beyond that specified in their input, and that, because of this, appropriately-designed physical implementations of even multi-layer networks can in fact be updated in a single step while incurring no greater theoretical minimum energy dissipation than is required to update their inputs. Furthermore, this energy can approach zero if the network state is updated adiabatically via a reversible transition process. Our novel operational paradigm for updating logic networks suggests an entirely new class of hardware devices and circuits that can be used to reversibly implement Boolean logic with energy dissipation far below the Landauer limit.

Keywords—thermodynamics of computation; reversible computing; adiabatic computing; nonlinear dynamics; chaotic computing

I. INTRODUCTION

The limits to the energy efficiency of conventional information processing technology are fast approaching, with thermal noise expected to become an important limiting factor within the next decade or two [1,2]. However, communication theorists have known since Shannon [3] that even a noisy channel can still be used to reliably communicate information at any rate up to the channel capacity

$$C = B \log \left(1 + \frac{P_S}{P_N} \right), \quad (1)$$

where B is the channel bandwidth and P_S , P_N denote signal and noise power, respectively. Even when the signal-to-noise ratio P_S/P_N is relatively small, the channel capacity may still be ample in a high-bandwidth communication medium; this observation forms the basis of our ubiquitous modern wireless telecommunication infrastructure. Analogously, reliable computation can theoretically still be performed even using noisy circuits, since a computational dataflow can be viewed as being a special case of a communication channel that just happens to transform the data in transit. This observation suggests that appropriately designed circuits may perform reliable computations at extremely low signal energies, operating in a regime where thermal noise dominates. Our goal, in this line of thought, is to identify a specific new computational mechanism and operating conditions under which the energy dissipation per useful computational operation is negligible.

A canonical example of a system having no (or negligible) energy dissipation is a system at (or near) thermal equilibrium, that is, a system occupying some thermally distributed ensemble of states at some uniform temperature T . Thermal states cannot dissipate energy because they already are at their maximum entropy. Of course, an equilibrium thermal state is statically distributed, and is not itself actively carrying out a useful computation. However, in a complex, spatially-extended system, the static *structure* of an equilibrium state of the system can reflect the logical structure of a computation. Furthermore, if we perturb the system in such a way as to carry out gradual, adiabatic transitions between near-equilibrium states, so that the system’s state is no longer static but *quasistatic*, the sequence of equilibrium states that is visited may represent a series of useful computations. This is the essential idea behind *reversible computing* [4,5].

As the size of our computational devices pushes further into the deep nanoscale, approaching molecular and atomic dimensions, and as the energy of their information-bearing signals approaches the characteristic thermal energy kT , the tendency is for the thermal and logical degrees of freedom of the system to become more closely intertwined with each other, and so un-

Approved for unclassified unlimited release SAND2016-XXXX.
Research supported by Sandia National Laboratories, a multiprogram laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the US Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000

derstanding the nature of thermal states and how they arise becomes increasingly important to the process of designing a new mechanism to carry out an intended logical function.

It is a very general feature of complex, nonlinearly-coupled dynamical systems that their degrees of freedom may often evolve chaotically [6,7,8], in which case the dynamical state converges towards a so-called strange attractor, in which the dynamical orbits exhibit a fractal dimension and the long-term state distribution resembles that of a thermal system in an equilibrium state at some nonzero effective temperature T , exhibiting characteristic thermal fluctuations (a.k.a. “noise) of RMS magnitude $kT/2$ in the energy (potential or kinetic) associated with each dynamical degree of freedom, which may be any generalized position or momentum coordinate variable within the system. In a conservative system (such as any closed physical system), there will be in general many different strange attractors, corresponding to different values of the total internal energy U and associated temperature T of the system, as well as to distinct large-scale system configurations. Different initial states will in general converge towards different strange attractors, allowing us to use such attractors to represent computational states corresponding to distinct logical input cases.

II. GENERAL APPROACH

To evaluate candidate new computing mechanisms for suitability within this new conceptual framework, we can use detailed circuit simulations (initially, say, using compact device models such as those provided by the NEEDS project¹), together with a detailed statistical analysis of simulation results, to identify methods for extracting useful computational results from the behavior of networks of dynamically-interacting, nonlinear nanodevices. In principle, such networks could implement reliable computation even at very low signal energies where thermal noise is dominant. Although individual degrees of freedom within the circuit may evolve chaotically on short timescales, over longer timescales, average properties of the trajectory can still reliably communicate results of the desired computation.

We anticipate that this new approach to achieving ultra-low energy dissipation may turn out to be simpler than traditional approaches to reversible computing using adiabatic retractile cascades, which apply larger signal energies and recover them adiabatically over the course of a long, complex, controlled staging sequence. In contrast, in the new approach, we can carry out multiple levels of logic reversibly within a single adiabatic transition of the overall state of the circuit.

Our approach relates somewhat to previous approaches to computing using the chaotic dynamics of nonlinear systems, such as, for example, “reservoir computing” [9,10,11]. However, as it has been framed in past work, reservoir computing typically invokes the application of artificial neural networks to learn how to interpret the dynamical trajectory of an arbitrarily-selected underlying chaotic system. These networks cannot necessarily be assumed to operate with very low energy cost. In contrast, our present proposal is to *design* the underlying system from the beginning so that its dynamical behavior, albeit

chaotic and noisy on short timescales, is nevertheless predictable and easy to interpret on longer timescales.

In other words, although the degrees of freedom within a noise-dominated nonlinear system will, in general, evolve chaotically, if we are careful, we can, in principle, design our computational mechanisms so that the particular attractor to which the dynamical trajectory converges still reliably depends on the input to the circuit and on the logical structure of the computation. In this vision, a desired result can be reliably obtained by appropriately measuring (e.g., by effectively averaging) the network outputs over longer time periods; this longer timescale effectively confines our attention to a narrowband slice of the noise spectrum, which therefore contains less noise power. This method is analogous to how, in communications, even a low-power (but narrow-band) signal can be reliably distinguished from higher-power background noise, if the noise is spread out over a relatively broad spectrum.

Our design strategy is that the nature of the interactions between devices in our model will be explicitly crafted in such a way that the statistical properties of the dynamical trajectory of the local degrees of freedom at specific points in the network will reflect the desired digital information that would be obtained at the corresponding points in a conventional logic circuit. Through detailed simulations and analysis over several design iterations, we will converge onto network-design techniques and measurement technologies that best implement such a mapping.

We expect this approach will prove to allow computations to be carried out with extremely low energy dissipation. Conventional wisdom says that digital information requires stored signal energies to be of a magnitude that is well above the thermal energy kT (which is ~ 26 meV at room temperature) to be reliably distinguished from thermal noise, and we know from fundamental entropy considerations that at least $kT \ln 2$ energy per bit must be dissipated whenever digital information is irreversibly discarded [12]. Although these observations limit the energy efficiency of conventionally-designed digital circuits, the novel insight that enables our new approach is that downstream logic nodes technically *contain no new (independent) information* since they depend deterministically on the input to the computation.

It may be the case that the energy scale of each input signal still has to be above the thermal noise energy, but because the low-power signal that is conveyed through the network from the input nodes as they are transitioned over a long time period is confined to a very narrow frequency band (perhaps at baseband or near DC), the downstream influences of the input, even after many stages of intermediate transformation by noisily fluctuating devices, can still remain reliably detectable, and as long as the system always remains close to the particular equilibrium state (a.k.a. strange attractor) selected by the instantaneous input configuration, the energy dissipated per useful device operation over the course of the transition and measurement process can in principle be *much less than $kT \ln 2$* without incurring the usual overhead, in terms of staged control sequencing, that is required by traditional approaches to multi-stage reversible logic [5].

¹See <https://nanohub.org/groups/needs>.

At this point, the above discussion still remains rather preliminary, in that it just outlines the essential seed of this new idea; these arguments still need to be validated by a more detailed investigation. However, this line of thought is already clarifying, for us, how to correctly interpret and reapply the known thermodynamic limits of computation. Although this line of work may be considered high-risk by industry engineers, it is highly innovative in that it shows how to use a deeper, more thorough understanding of how to correctly apply fundamental principles to (potentially) vastly improve the energy efficiency that is achievable in real computing systems.

In later sections of this paper, we will outline one potential initial approach for translating the above general insights into a particular (very simple) circuit model suitable for simulation and possible eventual implementation. However, the general points made above are not confined to the following material; rather, there is a very wide range of possible circuit models that could be explored as a basis for realizing the above concepts.

III. CLARIFYING LANDAUER'S PRINCIPLE

Before we describe our particular model in detail, we wish to clarify the nature of the fundamental thermodynamic limits of computation, so that it is more clear how our new paradigm will avoid running afoul of some of the limits on energy dissipation that apply to more conventional approaches.

In his landmark paper [12], Rolf Landauer of IBM analyzed the minimum energy dissipation during the erasure of a bit of information (considered in the abstract, as well as in the context of a simple, general model system, namely, a bistable potential energy well) and found that this energy loss could never be less than $kT \ln 2$, where k is Boltzmann's constant and T is the temperature of the system. This principle can be understood to apply generally to any possible mechanism for storing digital information, and in fact, it follows rigorously as a direct logical consequence of the microscopic invertibility of fundamental physical dynamics (including quantum time evolution), together with a pragmatic definition of physical entropy as comprising any and all information that cannot be *de-computed* (evolved back into a fixed standard microstate) by any invertible process that is practical for us to arrange.²

Given this perspective, “erasing” some digital information that is encoded in a physical state simply means undergoing a process that transforms the physically-encoded information into another form (*e.g.*, heat) from which a standard state is not reversibly recoverable, which puts this information in the category of non-invertibly-decomputable information or entropy. Meanwhile, thermodynamic temperature T is defined by

$$\frac{1}{T} = \frac{\partial S}{\partial E}, \quad (1)$$

²We can nevertheless imagine totally *impractical* processes that could, in principle, restore a standard state; for example, if we could somehow cause the entire microscopic time-evolution of the system (and its environment) to exactly reverse its direction, the system would eventually return to a (potentially standardized) initial state.

where ∂S is the infinitesimal entropy increase that results from adding infinitesimal energy ∂E to an equilibrium system (heat bath) at temperature T . It follows immediately from this definition that if an amount of entropy $\Delta S = k \ln 2 = 1$ bit will eventually be ejected from a given system into a relatively large external environment or heat sink at temperature T , then an amount of energy $\Delta E = T\Delta S = kT \ln 2$ must be accordingly invested into that environment (in the form of heat) in order to increase its entropy by the required amount. Thus, any bit of physical information that gets (by definition irrevocably) transformed into entropy must eventually result in this much energy dissipation, if entropy is not to build up within the computer indefinitely. This constitutes a simple proof of Landauer's principle that is completely general, in the sense that it is independent of the details of any particular physical implementation of an information-storing device.

However, when applying Landauer's principle to the analysis of the energy efficiency limits of physical systems that, themselves, implement rather complex computing structures, one must be careful to apply it appropriately. In particular, for systems that are more complex in structure than (say) a single two-state memory cell, one must be careful to quantify their information content appropriately.

How is the quantity of information characterized? As was first elucidated by Boltzmann [13] in the development of his pioneering H-theorem,³ and later elaborated upon and applied by Shannon [3] in the context of communications theory, the most natural measure H of entropy (and indeed, the only one exhibiting the desired properties) for a system X that has W possible ways of arranging its microscopic constituents, arrangements which we may label x_1, \dots, x_W , is:

$$H = -\sum_{i=1}^W p_i \log p_i, \quad (2)$$

where p_i denotes the probability that the system may be found in microstate x_i . This entropy measure H is maximized when every $p_i = 1/W$, in which case (assigning physical entropy $S = H$ and letting Boltzmann's constant $k = \log e$ denote the natural logarithmic unit), we have

$$S = k \ln W, \quad (3)$$

³Boltzmann actually defined his expression for H to be the negative of that in (2), and showed that it always decreases (rather than increases).

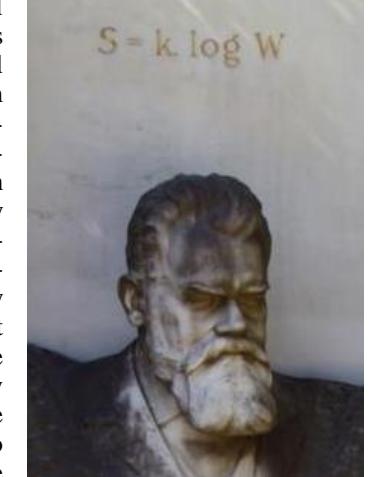


Fig. 1. Boltzmann's Legacy. Entropy S is a measure of information; it characterizes the number of ways W of arranging a system.

the equation which is famously carved on Boltzmann's tombstone (Fig. 1).

Now, let's drill down a little further into this definition. What is the physical meaning of *a way of arranging the system*? In modern physics, this refers to a possible pure quantum state of the system that is distinguishable from other such states in a given set (*i.e.*, their state vectors are pairwise mutually orthogonal). A given mixed state ρ of a quantum system is simply a statistical ensemble of pure states comprising an orthogonal set of basis states of the system, in the specific basis that diagonalizes the mixed state's density matrix.

For a system whose configuration may vary continuously (along various degrees of freedom, *e.g.*, the positions of constituent particles), in general the possible quantum states of that system include ones where the probability amplitude may be "smeared out" over a substantial volume in configuration space. In such states, the positions of the constituent particles are not individually well-defined, but may be entangled (quantum-mechanically correlated) with each other in complex ways.

Let us consider again what we just said above: In general, a distinguishable physical microstate may involve *a complex, correlated configuration of many degrees of freedom (many subsystems)*. In particular, it does not necessarily always consist of the state of just a single, small isolated system, such as (for example) a one-bit memory storage cell.

Why is this observation useful? Because it tells us that, if we can arrange for the distinguishable states of a physical system that we can measure in a given context to themselves be complex, spatially-extended objects, whose structure may reflect (for example) multiple consecutive stages of a desired computation, then the fact that this complex structure is made up of many smaller parts *does not imply that it actually contains many bits of information*—only that the individual states in question each have a complex form. Information, as always, is quantified by simply *counting* the distinguishable states, or more generally, by measuring the entropy of the set of states as in (2), in cases where the set of likely states is more constrained than the maximal set, and/or is nonuniform in terms of the states' statistical probabilities.

In subsequent sections of this paper, we elaborate further on how this insight allows us, in principle, to update, in a single step, an entire combinational logic network that may be multiple layers deep without necessarily requiring the usual Landauer-erasure minimum energy dissipation of $kT \ln 2$ for each internal logic node that is updated.

In the next section, we define a simple network model of combinational logic circuits for later reference. Then in secs. V-VI, we briefly discuss how logical states of nodes correspond to physical microstates, while logic gates correspond to dynamical interactions between nodes. Then in section VII, we restrict our attention to a simple conservative Hamiltonian dynamical model of logic circuits, which will allow us to illustrate the central idea of our approach.

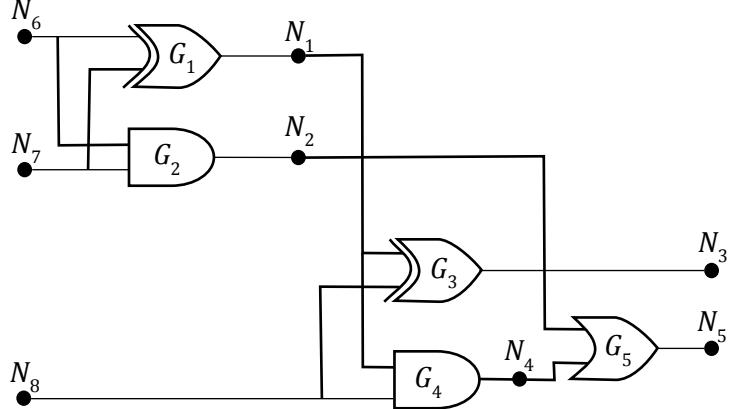


Fig. 2. A Simple Example Combinational Logic Network. (Full adder.)

IV. STATES OF A COMBINATIONAL LOGIC NETWORK

In the abstract, we can (quite conventionally) describe a *combinational logic network* in terms of a tuple (N, G) , where $N = \{N_1, \dots, N_n\}$ is a set of n logic *nodes* and $G = \{G_1, \dots, G_m\}$ is a set of m logic *gates* (meaning, gate instances), and where we stipulate that $n \geq m$, that is, there are at least as many nodes as there are gates, because each gate has (at least) one unique output node, and there may be other nodes that are only inputs to the network. Fig. 2 illustrates a simple example network, using the conventional graphic notation for a few basic Boolean gates. An individual deterministic logic gate G_j , where $1 \leq j \leq m$, can be described, in the most general sense, in terms of a tuple $G_j = (I, O, f)$, where $I \subseteq N$ is the gate's set of *input nodes*, $O \subseteq N$ is the gate's set of *output nodes*, and

$$f : \{0,1\}^{|I|} \rightarrow \{0,1\}^{|O|} \quad (4)$$

is a function that maps possible assignments of bit values $\{0,1\}$ to the gate's $|I|$ input nodes to the resultant assignment of bit values to the gate's $|O|$ output nodes. The output sets O of the various gates must have a null intersection; that is to say, each node can be the output of at most one gate. In addition, to say that the given network of logic gates is *combinational* means that the induced directed graph describing the possible paths through the network that pass only forwards through gates (from an input node to an output node) does not include any directed cycles; that is, there is no way to get from a given node back to itself while passing only forwards through gates.

With these definitions, it follows that the assignment of bit values to all nodes is uniquely determined by the assignment of values to those nodes that are inputs to the network (*i.e.*, those that are not outputs of any gates). Letting there be $i < n$ such nodes, the number of possible states of the network (consistent with the specified action of all gates) is actually only 2^i , not 2^n , and thus, the actual amount of information in the state of the network is only i bits, not n bits. This is our first clue that it ought to be theoretically possible to change the state of such a network while dissipating only $i \times (kT \ln 2)$ amount of energy, not $n \times (kT \ln 2)$, even when the inputs are destructively overwritten (prior input values are erased). In other words, *internal*

nodes of a combinational logic network are not necessarily subject to Landauer's principle because (when constrained by the gate behavior) they do not actually contain any additional information (independent of the network inputs). Note that this discussion assumes that the states of the network's internal and output nodes are in fact determined by the input nodes; we will see later that this can remain essentially true despite the effects of thermal noise and propagation delays through the network.

V. MICROSCALE NODE CONFIGURATION

What do we mean when we say that a bit value (0 or 1) is assigned to a node in a logic network? Simply that the node (in isolation) has at least two microstates (measurably distinguishable quantum states), and that (at least) one of these is labeled "0" and (at least) one of them is labeled "1"; the currently "assigned" bit value is then just the label corresponding to the state that would be found if the node were measured (using a measurement that can reliably distinguish the state) at a given time. Most generally, the digital states of interest might actually each correspond to a large set of (in principle) distinguishable quantum states, but where not all of the measurable distinctions are considered relevant in determining the bit values. So for example, in typical voltage-coded logic, a range of measurable node voltages between $[V_{0L}, V_{0H}]$ might all be considered to represent logic "0", while another range $[V_{1L}, V_{1H}]$ (where typically $V_{1L} > V_{0H}$) might all be considered to represent logic "1." Although in principle, the voltage might be able to be measured more accurately than is implied by just distinguishing these two sets, in practice, only the distinction between those two sets is used.

Now, the actual microstate of the node's physical hardware may in general be traversing a complex trajectory through a high-dimensional configuration space, comprising many degrees of freedom. For example, large numbers of conduction electrons may be flitting around chaotically through a solid conductor, scattering off of lattice defects and surface irregularities and atoms of trapped impurities. Even in a small deep-nanoscale device, such as a single-electron transistor or a rod in some molecular mechanical logic, the various degrees of freedom may in general be oscillating or fluctuating in complex patterns. However, regardless of the details, the system remains in principle characterizable by a density matrix which can be diagonalized into a set of measurably distinguishable quantum states. Each of these states may involve wavefunctions of electrons (or other moving particles) that are smeared out over a molecular orbital, or even over the entire volume of a large conductor, but that is fine. We can imagine that the individual electron positions are evolving along stochastic trajectories which are guided by the wavefunction amplitudes, in a Bohmian perspective [14]. Viewed on short timescales (if we could "snapshot" the electron motions) the evolution of the system would be noisy, but on longer timescales, where the short-timescale fluctuations are "averaged out," it is smooth. This is merely to say that the isolated node exhibits a reliably-measurable, persistent distinction between its "0" and "1" macrostates, which are (by assumption) distinguishable.

VI. CONNECTING NODES TOGETHER

Now, what happens when multiple logic nodes are connected together and made to interact with each other by logic gate hardware? The detailed physical behavior of course depends on the precise gate structure and mechanism, but a general picture is as follows. The gate introduces an interaction between the subsystems comprising the various nodes that impinge upon it. For functional or bidirectional-constraint type gates not exhibiting internal state, simple kinds of interactions could be represented by a Hamiltonian potential energy function.

More generally, certain kinds of gate device hardware may exhibit hysteretic effects, or some type of nonvolatile/persistent internal memory, such as memristors or Flash memory cells; such gates may introduce their own internal state variables which evolve on longer timescales; the new state variables may also contribute corresponding kinetic-type energy terms to the system's overall Hamiltonian, such as in, for example, the spin energy of a NEMS flywheel structure, or the magnetic field energy of a nanomagnet, a spintronic device, or a superconductive current loop. However, for the moment, we will set such possibilities aside.

The most general types of gates may also permit not just conservative (energy-conserving/elastic) interactions between node structures, but also dynamic flows of matter currents between them (e.g. short-circuit or leakage currents in electronic devices, spin currents in spintronic devices, and more generally flows down chemical gradients). Such flows typically are dissipative, and will cause a given structure to exceed the theoretical minimum energy dissipation that is of interest to us in this paper. For purposes of the present paper, we prescribe avoiding such behaviors; in future work, we will investigate to what extent reintroducing them may be helpful or necessary for improving the stability of a system. But for the moment, let our focus be on nodes that are interacting with each other through a simple, conservative Hamiltonian-type gate interaction.

VII. EXAMPLE HAMILTONIANS FOR STANDARD BOOLEAN GATES

In this section, we present a very basic dynamical model of networks of Boolean AND and NOT gates to illustrate our general picture, and serve as a starting point for later studies.

As a simple example, imagine that each node N_i in the logic network includes a single microscopic generalized-position degree of freedom (not necessarily a spatial position) represented by a continuous coordinate variable x_i . This could be a voltage, a superconducting current, or any other continuous variable. In general, this "position" variable can fluctuate dynamically; to enable such dynamics, we will include a kinetic energy term in our Hamiltonian of the simple nonrelativistic form

$$K_i = \frac{1}{2} m_i v_i^2, \quad (4)$$

where $v_i = dx_i/dt$ denotes the velocity of the generalized position coordinate, and m_i denotes a generalized effective mass associated with the given degree of freedom. For this abstract

example, units are arbitrary, so for simplicity, we can let each $m_i = 1$ unit, for the time being, without loss of generality.

Now, in the most general circumstances, the position x_i may be fluctuating around dynamically, but we can imagine that, in the system's lowest-energy state where all coordinates are at rest, the value of the position would tell us something about the node's logic value, in the sense that certain positions, if measured consistently, would imply a certain logic value.

Again, units are meaningless here, so for simplicity let $x_i = 1$ unit be interpreted as logic 1, and let $x_i = 0$ be interpreted as logic 0.

For our present purposes, we will assign the input nodes to the network as outputs of special 0-input, 1-output "gates" called *memory cells*, so that we may include these in our model. Therefore, the number of nodes equals the number of gates, $n = m$, and so without loss of generality, we can identify node and gate indices with each other; that is, let the output node of each gate G_i be node N_i .

Let a memory cell have a potential energy function of the form $E_i = \frac{1}{2} rkT(x_i - s_i)^2$, where s_i is the constant logic value (0 or 1) nominally being stored at a given time, and r is a factor determining the reliability with which the cell's value is maintained. Note that therefore the cell comprises a harmonic oscillator, whose ground state corresponds to the nominal bit value.

A more refined model of a memory cell for future consideration might be a bistable potential well with an energy barrier between 0 and 1 states of magnitude rkT ; such a cell could preserve its state over some timescale without the need for external control to maintain a constant bias.

Now, let a Boolean NOT gate or logical inverter have the potential energy function $E_i = \frac{1}{2} bkT(x_i + x_j - 1)^2$, where j is the index of the inverter's input node. Again this describes a quadratic potential, and the factor b determines the magnitude of the energy associated with the NOT-gate constraint that nominally, we should have $x_i = 1 - x_j$ (if the output were exact). Fig. 3 illustrates the shape of this potential energy surface, with the x and y axes corresponding to coordinates x_i and x_j .

Finally, let a Boolean AND gate have a potential energy function $E_i = \frac{1}{2} akT(x_i - x_j x_k)^2$, where j, k are the indices of the AND gate's input nodes. The factor a determines the magnitude of the energy associated with the AND-gate constraint that nominally, we should have $x_i = x_j x_k$ (if the output were exactly correct).

This gate set (AND and NOT) comprises a universal set, that is, it is sufficient for constructing arbitrary combinational Boolean logic functions.

The act of measuring the output of a logic cell and storing a corresponding new value in a memory cell is not yet modeled in terms of a detailed mechanism; first we want to understand better what the behavior of the above model looks like, in terms of its chaotic nonlinear dynamics, and in the presence of thermal noise. For the time being, updating the state of the memory cell can be simply modeled at an abstract level by

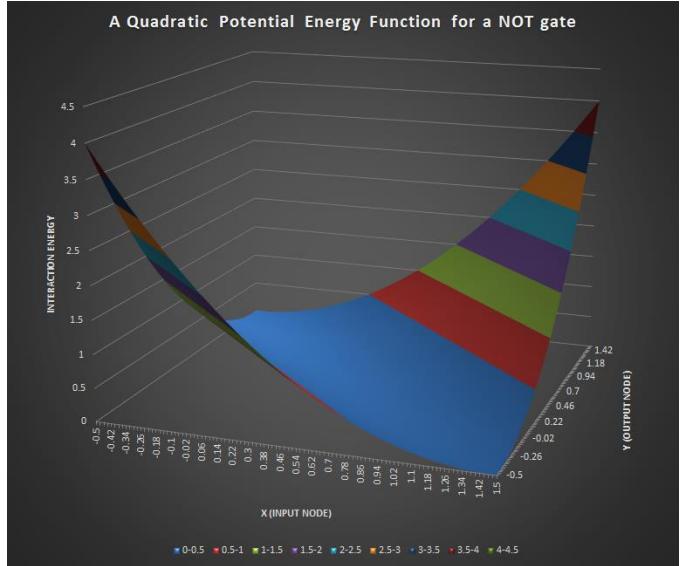


Fig. 3. Potential energy surface for an example NOT gate interaction term in our example Hamiltonian model of Boolean logic.

gradually transitioning the value of the input parameter s from old value to new value over a time t .

VIII. ANALYSIS AND SIMULATION PLAN

In principle, running our model simply consists of allowing all of the generalized coordinates in the system to simultaneously evolve, with all nodes in a complex combinational net fluctuating dynamically under the influence of thermal noise and neighboring nodes' configurations. We may consider how the distribution of final output configurations evolves over time as the model is run from a given initial system configuration. It would be interesting to study how the dynamical behavior changes as we scale down the magnitudes of the gate interaction energies a and b , relative to the energy scale r of the input memory cell.

Next, we may consider also what happens to the system on longer timescales, when the bias states s of the input memory nodes are transitioned between 0 and 1 (or vice-versa) over time t . Given our conservative Hamiltonian model, the adiabatic theorem implies that asymptotically, as $t \rightarrow \infty$, the energy dissipation (heating) of the system as a result of the transition should approach zero. But we may wish to examine how this behavior is affected by the relative energy scales of the logic gates versus the input memory, and how this affects the trade-offs between speed, energy dissipation, and reliability for this network model. For finite transition timescales, it would be interesting to study the precise trajectory along which the distribution of downstream nodes' values changes, to see how quickly reliable results can be inferred from output node values. It would also be interesting to study how, exactly, the asymptotically small thermal excitation of the system flows through the network from the sole energy input to the system, namely the transitioning of the memory node.

It would be straightforward to run simulations of this system to obtain results that cannot be easily inferred analytically. Even without explicit modeling of an external heat bath, the

system, with its many nonlinearly-interacting degrees of freedom, can be expected to evolve chaotically and to exhibit an effective temperature.

An easy way to initialize the system would be to set all of the generalized-position coordinates to their ideal values given the initial logic input, and assign their initial velocities randomly using the Maxwell-Boltzmann distribution (that is, letting the average energy per generalized-position coordinate be kT). Although initially the entire system energy will be found in the kinetic energy terms of the Hamiltonian (all of the potential energy terms being 0 in the ground state configuration), within a short time, the energy will become randomly distributed over all terms of the Hamiltonian as per the equipartition theorem, with an average energy of $\frac{1}{2}kT$ per microscopic degree of freedom (including both the position coordinates and their corresponding velocities).

It would be straightforward to generate phase diagrams of the various coordinates and to collect time-averaged statistics regarding the coordinates of the output nodes, and verify on the basis of these that, as expected, the correct digital value can be reliably determined via measurements over relatively long timescales.

As of this writing, a simulator along the lines described above, called DYNAMIC, is currently being prototyped in the Python language. Our core framework for simulating arbitrary networks of Hamiltonian interactions is functioning, but still needs more testing on complex networks, and the results need to be visualized and analyzed. We intend to release this software as open source once it has been completed.

IX. CONCLUSION

In this paper, we have argued that there is a potentially viable approach for implementing extremely energy-efficient Boolean logic (with potentially much less than kT energy dissipation per useful logic operation) that has not previously been explored in detail. This new approach can be understood as emerging from the convergence of several key insights:

1. As per Shannon [3], we can still communicate reliably at some rate even at signal power levels below the noise floor, and furthermore, computation can be viewed as merely a special case of communication;
2. Thermal equilibrium states are inherently nondissipative, and the static structure of a particular such state can reflect the logical structure of a particular computation;
3. Networks of conservatively-interacting nonlinear devices exhibiting chaotic dynamical behavior are effectively thermal systems, and thus are also nondissipative once their dynamical orbits have converged onto a particular strange attractor, which corresponds to a thermal equilibrium state. Yet, interactions between such devices can be tailored so as to embody desired computational structures, such that their time-averaged behavior reflects specific computational results.
4. A state reflecting the computational structure of a deep combinational Boolean logic network can be adiabatically updated in a single step, rather than through a

staged sequence of transitions. This does not violate Landauer's principle, because the values of downstream nodes do not represent independent bits of information being erased. Rather, the states of the network are spatially extended entities, and so the network can be updated all at once with no merging of states, and therefore no loss of information or entropy generation.

These insights open the door to the exploration of a fundamental new class of devices and circuits for carrying out computations reversibly while dissipating $\ll kT$ energy per useful operation. We can imagine working with assemblages of nonlinear devices coupled via conservative interactions. Even though the state of such networks may evolve chaotically and be unpredictable on short timescales, over longer timescales their statistical behavior can be arranged to embody desired computational results. This approach is inherently insensitive to noise, as we are dealing with thermal equilibrium states of the degrees of freedom of interest. Because of this, the information-bearing degrees of freedom do not even need to be perfectly isolated from their thermal environment, since at equilibrium (when the temperatures of the dynamical system and its environment are the same) any dissipative losses will be exactly counterbalanced by fluctuations induced by the thermal environment, as per the fluctuation-dissipation theorem.

Although this vision of a new paradigm for energy-efficient computing is still very preliminary, we have made initial steps towards its realization by describing and implementing a simple, classical Hamiltonian dynamical model which illustrates how the interactions between degrees of freedom could in principle be tailored so that their long-term statistical behavior will reflect computationally meaningful information. Further research is needed to characterize the limits and tradeoffs of this approach, explore variations on the underlying theme, translate this conceptual picture into concrete physical implementations leveraging manufacturable device and circuit structures, and determine whether these types of designs can yield practical benefits.

We also hope that the above line of work will contribute to a broader, deeper understanding of the more general notion of energy-efficient computation using nonlinear dynamics at the edge of chaos, which has been alluded to in recent works [15].

ACKNOWLEDGMENT

We thank Sapan Agarwal and John Aidun for helpful discussions and feedback on early write-ups of this concept.

REFERENCES

- [1] J. Izydorczyk, "Three steps to the thermal noise death of Moore's Law," *IEEE Trans. VLSI Sys.*, vol. 18, pp. 161–165, January 2010.
- [2] C. A. Mack, "Fifty years of Moore's Law," *IEEE Trans. Semicond. Mfg.*, vol. 24, pp. 202–207, May 2011.
- [3] C. E. Shannon, "Communication in the presence of noise," *Proc. IRE*, vol. 37, pp. 10–21, January 1949.
- [4] C. Bennett, "Logical reversibility of computation," *IBM J. Res. Dev.*, vol. 17, no. 6, pp. 525–532, 1973.
- [5] M. P. Frank, "Reversibility for efficient computing," Ph. D. thesis, Massachusetts Institute of Technology, 1999.

- [6] J. M. T. Thompson and H. B. Stewart, *Nonlinear Dynamics and Chaos*, 2nd ed. Wiley, 2002.
- [7] K. Mainzer and L. Chua, *Local Activity Principle*. London: Imperial College Press, 2013.
- [8] S. H. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering*, 2nd ed. Westview Press, 2015.
- [9] L. Appeltant, M.C. Soriano, *et al.*, “Information processing using a single dynamical node as complex system,” *Nat. Comm.*, vol. 2, no. 468, September 2011.
- [10] Y. Paquot, F. Duport, *et al.*, “Optoelectronic reservoir computing,” *Sci. Reports*, vol. 2, no. 287, February 2012.
- [11] D. Brunner, M. C. Soriano, *et al.*, “Parallel photonic information processing at gigabyte per second data rates using transient states,” *Nat. Comm.*, vol. 4, no. 1364, January 2013.
- [12] R. Landauer, “Irreversibility and heat generation in the computing process,” *IBM J. Res. Dev.*, vol. 5, no. 3, pp. 183–191, 1961.
- [13] L. Boltzmann, “Weitere studien über das wärmegleichgewicht unter gasmolekülen,” *Sitzungsberichte Akademie der Wissenschaften*, vol. 66, pp. 275-370, 1872.
- [14] D. Bohm, “A suggested interpretation of the quantum theory in terms of ‘hidden’ variables. I.” *Phys. Rev.*, vol. 85, no. 2, pp. 166–179, Jan. 1952.
- [15] R. S. Williams and E. P. DeBenedictis, “OSTP Nanotechnology-Inspired Grand Challenge: Sensible Machines,” IEEE whitepaper, <http://rebootingcomputing.ieee.org/archived-articles-and-videos/general/sensible-machine>, October 2015.