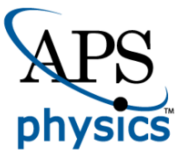

Maxwell Demon Dynamics: Deterministic Chaos, the Szilard Map, and the Intelligence of Thermodynamic Systems

Alexander B. Boyd and James P. Crutchfield

Phys. Rev. Lett. **116**, 190601 — Published 13 May 2016

DOI: 10.1103/PhysRevLett.116.190601

Santa Fe Institute Working Paper 15-06-019
arxiv.org:1506.04327 [cond-mat.stat-mech]

Demon Dynamics: Deterministic Chaos, the Szilard Map, and the Intelligence of Thermodynamic Systems

Alexander B. Boyd* and James P. Crutchfield†
Complexity Sciences Center and Department of Physics, University of California at Davis, One Shields Avenue, Davis, CA 95616
(Dated: April 11, 2016)

We introduce a deterministic chaotic system—the Szilard Map—that encapsulates the measurement, control, and erasure protocol by which Maxwellian Demons extract work from a heat reservoir. Implementing the Demon's control function in a dynamical embodiment, our construction symmetrizes Demon and thermodynamic system, allowing one to explore their functionality and recover the fundamental trade-off between the thermodynamic costs of dissipation due to measurement and due to erasure. The map's degree of chaos—captured by the Kolmogorov-Sinai entropy—is the rate of energy extraction from the heat bath. Moreover, an engine's statistical complexity quantifies the minimum necessary system memory for it to function. In this way, dynamical instability in the control protocol plays an essential and constructive role in intelligent thermodynamic systems.

Keywords: deterministic chaos, Lyapunov characteristic exponents, Markov partition, Landauer's Principle, measurement, erasure, memory, entropy rate, mutual information, dissipated work, control theory

PACS numbers: 05.45.-a 05.70.Ln 89.70.-a 02.30.Yy

Synthetic nanoscale machines [1–4], like their macromolecular biological counterparts [5–7], perform tasks that involve the simultaneous manipulation of energy, information, and matter. In this they are information engines—systems with two inextricably intertwined characters. The first aspect, call it "physical", is the one in which the system—seen embedded in a material substrate—is driven by, manipulates, stores, and dissipates energy. The second aspect, call it "informational", is the one in which the system—seen in terms of its spatial and temporal organization—generates, stores, loses, and transforms information. Information engines operate by synergistically balancing both aspects to support a given functionality, such as extracting work from a heat reservoir.

This is remarkable behavior. Though we can sometimes identify it—in a motor protein hauling nutrients across a cell's microtubule highways [5], in how a quantum dot transistor modulates current under the influence of an evanescent wave function [8, 9]—it is not well understood. Understanding calls on a thermodynamics of nanoscale systems that operate far out of equilibrium and on a physics of information that quantitatively identifies organization and function, neither of which has been fully articulated. However, recent theoretical and experimental breakthroughs [6, 7, 10–12] hint that we may be close to a synthesis which not only provides understanding but predicts quantitative, measurable functionalities.

We define an *information engine* as a system that performs information processing as it undergoes controlled thermodynamic transformations. We show that information engines are chaotic dynamical systems in the particular sense that energy extraction from the heat bath requires a spreading of ensemble trajectories and this leads to a positive Lyapunov characteristic exponent. (In a rather different setting, that demon-like behavior requires an overall chaotic dynamics was broached previously by Ref. [13].) Building this bridge to dynamical systems theory allows us to employ its powerful tools to analyze an engine's complex, nonlinear behavior. This includes not only monitoring instability via the Lyapunov exponents, but a thorough informational and structural analysis that leads to a measure of thermodynamic system intelligence.

By way of concretely illustrating the theory, we introduce an explicit implementation of Szilard's Engine [14] as an iterated composite map of the unit square that is a deterministic, but chaotic dynamical system. The result is a particularly simple and constructive view of the energetics and computation embedded in controlled nonlinear thermodynamical systems. That simplicity, however, gives a solid base for designing and analyzing real information engines. We end giving a concise statement of the general theory and applications.

Background The Szilard Engine is an ideal Maxwellian Demon for examining the role of information processing in the Second Law [14]. The engine consists of three components: a controller (the Demon),

* abboyd@ucdavis.edu
† chaos@ucdavis.edu

a thermodynamic system (a particle in a box), and a heat bath that keeps both thermalized to a temperature T. It operates by a simple mechanism of a repeating three-step cycle of measurement, control, and erasure. During measurement, a barrier is inserted midway in the box, constraining the particle either to box's left or right half, and the Demon memory changes to reflect on which side the particle is. In the thermodynamic control step, the Demon uses that knowledge to allow the particle to push the barrier, extracting $\int P \; dV = k_B T \ln 2$ work from the thermal bath. (Supplementary Materials review this and related thermodynamic calculations.) In the erasure step, the Demon resets its finite memory to a default state, so that it can perform measurement again. The periodic *protocol* cycle of measurement, control, and erasure repeats endlessly and deterministically.

The net result of the cyclic protocol is the extraction of work from the bath during control *balanced* by heat dissipated due to changes in the Demon's memory during measurement and erasure. Note that extracting $k_B T \ln 2$ work from a thermal reservoir was a paradox until the last century, when it was realized that the information processing steps of measurement and erasure have a compensating energy cost [15–17].

Rather than seeing the Demon and box as separate, though, we view it—an information engine—as the direct product of thermodynamic system and Demon memory [18]. Though we follow Szilard closely, he did not specify the Demon's physical embodiment. Critically, we choose the Demon's memory to be another spatial dimension of a particle in a box. Thus, we see the joint system as a single particle in a two-dimensional box, where one axis represents the thermodynamic System Under Study (SUS)—the original particle in a box—and the other axis represents the Demon memory. We now describe a deterministic protocol that implements the Szilard Engine, evolving a particle ensemble over the joint state space.

A Dynamical Engine The Szilard Engine's measurement-control-erasure barrier-sliding protocol is equivalent to a discrete-time two-dimensional map from unit square $\mathbb{I}^2 = [0,1] \times [0,1]$ to itself. The engine has two kinds of *mesoscopic states*—states of the particle's location $\{L \sim x \in (0,\delta], R \sim x \in (\delta,1)\}$ and states of the Demon's knowledge $\{A \sim y \in (0,\gamma], B \sim y \in (\gamma,1)\}$ of the location—that partition the joint states $(x,y) \in \mathbb{I}^2$.

The protocol cycle translates into a composite map $\mathcal{T}_{\text{Szilard}} = \mathcal{T}_{\text{E}} \circ \mathcal{T}_{\text{C}} \circ \mathcal{T}_{\text{M}}$ of $\mathbb{I}^2$; one map for each engine step; see Fig. 1(a). As they operate, they take the joint state space from one stage to another around the cycle:

Measurement: To correlate Demon memory with particle location $\mathcal{T}_{\text{M}}$ takes the $A \otimes L$ and the $B \otimes L$ mesostates to themselves, the $A \otimes R$ mesostate to $B \otimes R$, and $B \otimes R$ to $A \otimes R$:

$$\mathcal{T}_{\text{M}}(x,y) = \begin{cases} (x,y) & x < \delta, y < \gamma \text{ or } x < \delta, y \geq \gamma \ , \\ \left(x, \gamma + y\frac{1-\gamma}{\gamma}\right) & x \geq \delta, y \leq \gamma \ , \\ \left(x, \gamma\frac{y-\gamma}{1-\gamma}\right) & x \geq \delta, y > \gamma \ . \end{cases}$$

Thermodynamic control: To extract energy from the bath $\mathcal{T}_{\text{C}}$ expands both the A and B Demon memory mesostates over the SUS's whole interval:

$$\mathcal{T}_{\text{C}}(x,y) = \begin{cases} (\frac{x}{\delta}, y) & x < \delta \ , \\ \left(\frac{x-\delta}{1-\delta}, y\right) & x \geq \delta \ . \end{cases}$$

Erasure: $\mathcal{T}_{\text{E}}$ maps both the A and B Demon memory mesostates back to a selected Demon memory reset mesostate. If that reset state is A, then the mapping is:

$$\mathcal{T}_{\text{E}}^{A}(x,y) = \begin{cases} (x, y\delta) & y < \gamma \ , \\ \left(x, \delta\gamma + \frac{y-\gamma}{1-\gamma}\gamma(1-\delta)\right) & y \geq \gamma \ . \end{cases}$$

The resulting form of the whole measure, control, and erase cycle on the unit square is

$$\widehat{\mathcal{T}}_{\text{Szilard}}(x,y) = \begin{cases} (\frac{x}{\delta}, \delta y) & x < \delta, y < \gamma \ , \\ \left(\frac{x-\delta}{1-\delta}, \delta\gamma + y(1-\delta)\right) & x \geq \delta, y < \gamma \ , \\ \left(\frac{x}{\delta}, \delta\gamma + \frac{y-\gamma}{1-\gamma}\gamma(1-\delta)\right) & x < \delta, y \geq \gamma \ , \\ \left(\frac{x-\delta}{1-\delta}, \frac{y-\gamma}{1-\gamma}\gamma\delta\right) & x \geq \delta, y \geq \gamma \ . \end{cases}$$

This explicit construction establishes that the temporal behavior of Szilard's Engine can be modeled by a deterministic dynamical system whose component maps are thermodynamic transformations—a *piecewise thermodynamical system*. The mapping $\mathcal{T}_{\text{Szilard}}$ means we can avail ourselves of the analytical tools of dynamical systems theory [19, 20] to analyze the Szilard Engine mechanisms. This perforce suggests a number of more refined and quantitative questions about the engine dynamics ranging from the structural role of the stable and unstable submanifolds in supporting information and thermodynamic processing to the existence of asymptotic invariant distributions and measures of information generation, storage, and intelligence.

As shown in Fig. 1(leftmost), only the lower region $y \leq \gamma$ is occupied in the iteration of the Szilard Map. This is the Demon's default state, from which it starts every cycle. Over this region, the Szilard Map is a version of an asymmetric Baker's Map. As such, it is immediately clear that the Szilard Engine dynamics are chaotic [19, 20].

While the overall composite map $\mathcal{T}_{\text{Szilard}}$ is important, considering its complete-cycle behavior alone misses much. What is key are the component maps that nominally control a thermodynamic system, with each step

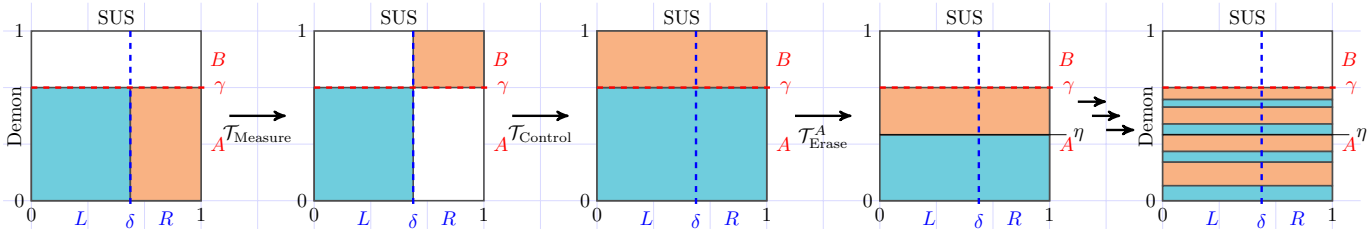

FIG. 1. Szilard Engine as a deterministic dynamical system: the Szilard Map $\mathcal{T}_{\text{Szilard}} = \mathcal{T}^{A}_{\text{Erase}} \circ \mathcal{T}_{\text{Control}} \circ \mathcal{T}_{\text{Measure}}$. Regions left and right of δ colored to aid visually tracking particle ensemble history. Rightmost: Action of $\mathcal{T}^{3}_{\text{Szilard}}$ resulting in self-similar (fractal) structure in density $\widehat{\rho}$; uniform $\widehat{\rho}$ requires $\eta = \gamma\delta$. These assume the Demon's reset memory state is A. (Supplementary Materials include an animation.)

corresponding to a different thermodynamic transformation. We now analyze the dynamics to see how the component maps contribute to information processing and thermodynamics. (Supplementary Materials give calculational details.)

Dynamical Systems Analysis What does chaos in the Szilard Engine mean? The joint system generates information—the information about particle position that the Demon must keep repeatedly measuring to stay synchronized to the SUS and so extract energy from the bath. On the one hand, it is generated by the heat bath through state-space expansion during $\mathcal{T}_{\text{C}}$. And, on the other, it is stored by the Demon (temporarily) and must be erased during $\mathcal{T}_{\text{E}}$. The latter's construction makes clear that it, dynamically, contracts state-space and so is locally dissipative.

With explicit equations of motion in hand, one can directly check, by calculating the Jacobian $\partial_{xy}\mathcal{T}_{\text{Szilard}}$, that the map is globally area preserving. Moreover, the invariant distribution $\widehat{\rho}$ can be determined from the Frobenious-Perron operator [19, 20]:

$$\widehat{\rho}(x', y') = \int_{\mathbb{I}^2} dxdy\ \delta\big((x', y') - \mathcal{T}_{\text{Szilard}}(x, y)\big)\widehat{\rho}(x, y)\ .$$

($\delta(\cdot)$ here, and only here, is the Dirac delta-function.) Calculation shows that $\widehat{\rho}$ has full support on the lower portion of the unit square $\mathbb{I} \otimes [0, \gamma]$ for all $\delta, \gamma \in (0, 1)$. That said, the action of $\mathcal{T}_{\text{Szilard}}$ builds up a self-similar interleaving within $\widehat{\rho}$; as shown on the far right of Fig. 1 via the third iterate of $\mathcal{T}_{\text{Szilard}}$ and in the online animation. In fact, the particle density is uniform when, during $\mathcal{T}_{\text{E}}$, the Demon's memory mesostate partition falls at $\eta = \gamma\delta$, where η is the iterate of $y = \gamma$ under $\mathcal{T}^{A}_{\text{Erase}}$. However, if we change γ or δ so that $\eta \neq \gamma\delta$, $\widehat{\rho}$ is no longer uniform, corresponding physically to a loss in efficiency of the Demon's information extraction during measurement.

The Szilard Map Jacobian also determines its local linearization and so one can easily calculate the spectrum of Lyapunov characteristic exponents (LCEs) for the overall cycle and so realize the contribution of each protocol step. This gives insight into the directions (submanifolds) of stability (convergence) and instability (divergence). There are two LCEs: one positive $\lambda^{+} = \text{H}(\delta)$ and one negative $\lambda^{-} = -\text{H}(\delta)$, where $\text{H}(\delta)$ is the (base 2) binary entropy function [21]. λ^{+} quantifies the exponential spreading of the distribution along the SUS axis, while λ^{-} quantifies its exponential contraction along the Demon axis. (See Supplementary Material for details.) Note that energy conservation ($\mathcal{T}_{\text{Szilard}}$'s area preservation) is reflected in the exact balance of instability and stability: $\lambda^{+} + \lambda^{-} = 0$. The unstable manifolds (parallel the x-axis) support the mechanism that amplifies small fluctuations from the heat bath to macroscopic scale during energy extraction ($\mathcal{T}_{\text{C}}$). The stable manifolds (parallel the y-axis) are the mechanism that dissipates energy into the ambient heat bath, during erasure ($\mathcal{T}_{\text{E}}$).

The overall information production rate is given by $\mathcal{T}_{\text{Szilard}}$'s Kolmogorov-Sinai entropy h_{μ}, which also quantifies the degree of chaos of the map [22]. This chaotic information production is necessary for an information engine's processing cycle. For the Szilard Engine, given the well behaved nature of $\widehat{\rho}$, $h_{\mu} = \sum_{\lambda > 0} \lambda = \lambda^{+}$ by Pesin's Theorem [20]. (That is, $h_{\mu} = \text{H}(\delta)$, directly verified shortly.) This entropy monitors the information generated in the SUS during the control step, as well as the information erased in the Demon in the measurement and erasure steps combined. In this way, it quantifies an effective flow of information from the SUS to the Demon. The physical consequence, simply stated, is perhaps striking: the degree of chaos determines the rate of energy extraction from the bath.

Computational Mechanics Analysis The Demon memory and particle location mesostates form Markov partitions for the Szilard Map dynamics [20, Chs. 7,9]: tracking sequences of symbols in $\{A, B\}$ or in $\{R, L\}$ (or all four pairs $\{AR, AL, BR, BL\}$) leads to a symbolic dynamics that captures all of the joint system's infor-

mation processing behavior. We now use this fact to analyze the various kinds of information processing and introduce a way to measure the Demon's "intelligence" or, more appropriately, that of the entire engine. We do this by calculating computational mechanics' ϵ-machines and ϵ-transducers from the engine's symbolic dynamics. The ϵ-machine for the Szilard Engine is a special kind of hidden Markov model—the minimal unifilar generator—of the observed symbol sequence. It's unique properties allow for exact calculation of many essential information-theoretic properties [23]. The ϵ-transducer is an extension that accepts control inputs, as well as generates outputs [24, 25]. The overall engine transducer is shown in Fig. 2(a), that for the SUS particle dynamics in Fig. S2(a), and for the Demon memory dynamics in Fig. S2(b).

In addition to explicitly expressing the effective mechanisms that support information processing, ϵ-machines allow us to quantify the effects of measurement, control, and erasure. The engine's Kolmogorov-Sinai entropy h_μ can be calculated directly from the ϵ-machines's causal-state averaged transition uncertainty. To quantitatively measure the minimal required memory—a key component of "intelligence"—for the information engine functioning, we employ the ϵ-machine's statistical complexity $C_\mu = \mathrm{H}[\Pr(\sigma)]$, where $\sigma \in \boldsymbol{\mathcal{S}}$ are the system's causal states [24] and $\mathrm{H}[\cdot]$ is the Shannon information [21].

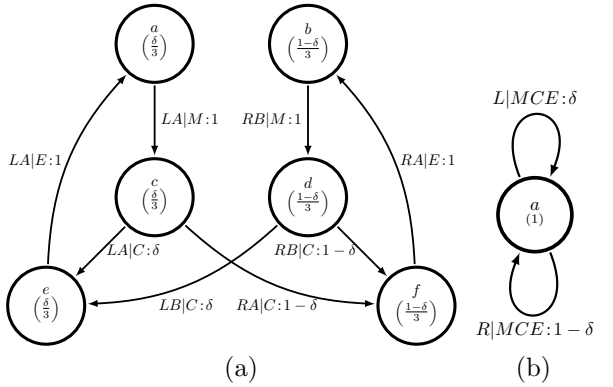

FIG. 2. ϵ-Transducers for the symbolic dynamics of the Szilard information engine from the Markov partition of its joint state space: (a) The ϵ-transducer for $\mathcal{T}_{\text{Szilard}}$ that reads in the periodic control signal for measure (M), control (C), and erase (E). (b) $\mathcal{T}_{\text{Szilard}}$ single-state ϵ-transducer: Memoryless over the full measure-control-erase protocol. Transitions $\beta|\alpha : p$ denote reading protocol symbol α, taking the transition with probability p, emitting symbol β. Asymptotic state probabilities are given in parentheses underneath state names.

It is important to emphasize an aspect of the information engine ϵ-machine construction: It is stage-dependent in that, to fully capture the component operations and their thermodynamic effects, the individual maps must be taken into account. This observation should be contrasted with the symbolic dynamics and particle position ϵ-machine for the *overall* Szilard map. The resulting process arises from stroboscopically observing the behavior after driving the engine with the three-symbol word MCE. As an example, the particle position process's ϵ-machine is shown in Fig. 2(b); it is a biased coin—a single-state ϵ-machine with no memory: $C_\mu(\mathcal{T}_{\text{Szilard}}) = 0$. This is as it should be: The overall cycle must return to the same state storing no memory of previous cycles.

Computational mechanics analysis shows that, over the three-step cycle, the Engine has an entropy rate of $h_\mu = \mathrm{H}(\delta)$ as seen above (or $\mathrm{H}(\delta)/3$ per map step) and a statistical complexity of $C_\mu = \log_2 3 + \mathrm{H}(\delta)$. (See Supplementary Materials for details, including analysis of SUS and Demon subsystems.) How predictable is the Engine's operation? The information in its future predictable from its past is given by the *excess entropy*: $\mathbf{E} = \mathrm{I}[\overleftarrow{Z}; \overrightarrow{Z}] = H[\overleftarrow{Z}] + H[\overrightarrow{Z}] - H[\overleftarrow{Z}, \overrightarrow{Z}]$, where $\overleftarrow{Z}$ is the past and $\overrightarrow{Z}$ is the future of the joint process over random variable $Z_t \in \{A, B\} \otimes \{R, L\}$. We see that the machines in Fig. 2(a) and 2(b), driven by the protocol, are counifilar and so $\mathbf{E} = C_\mu$ [26]. Thus, we see that while the Szilard Engine does not carry any information through one measure-control-erase cycle to the next, within the three steps of a single cycle, the engine stores $\log_2 3 + H(\delta)$ bits to operate.

Thermodynamics During each protocol step the Engine interacts thermodynamically with the heat bath. The Supplementary Materials calculate the average heat $\langle Q \rangle$ and work $\langle W \rangle$ flows between the Demon and the bath and between the SUS and the bath during each step. For this implementation of the Szilard Engine heat and work are equivalent, since there is no change in the average internal energy of the particle contained by the box during the isothermal measurement, control, and erasure protocol steps. Thus, we discuss only the heat, as energies $\langle Q_{\text{diss}} \rangle$ dissipated to the bath for each interaction. As we will see, although γ—the Demon memory partition—did not play a direct role in the informational properties, it does in the thermodynamics.

The expected heat flow during measurement is $\langle Q_{\text{measure}} \rangle = -k_B T(1-\delta) \ln((1-\gamma)/\gamma)$. Since $\gamma \in [0, 1]$, the dissipated heat can be negative or positive. It vanishes at $\gamma = 1/2$. Negative dissipated heat means that the engine absorbs energy from the heat bath and, in that case, turns it into work. The work $\int P\, dV$ done by the particle on the barrier is $k_B T\, \mathrm{H}(\delta) \ln 2$. Thus, the average heat absorbed by the engine from the heat bath during thermodynamic control is $\langle Q_{\text{control}} \rangle = -k_B T\, \mathrm{H}(\delta) \ln 2$, which is maximized when $\delta = 1/2$. During memory erasure the Demon shifts back to its default state, without affecting the SUS state. The barrier partitioning the

Demon's mesostates slides, compressing the contained particle ensemble into the default state A, say. The heat dissipated in this process is $\langle Q_{\text{erase}} \rangle = k_B T(1-\delta)\ln((1-\gamma)/\gamma) + k_B T\,\mathrm{H}(\delta)\ln 2$.

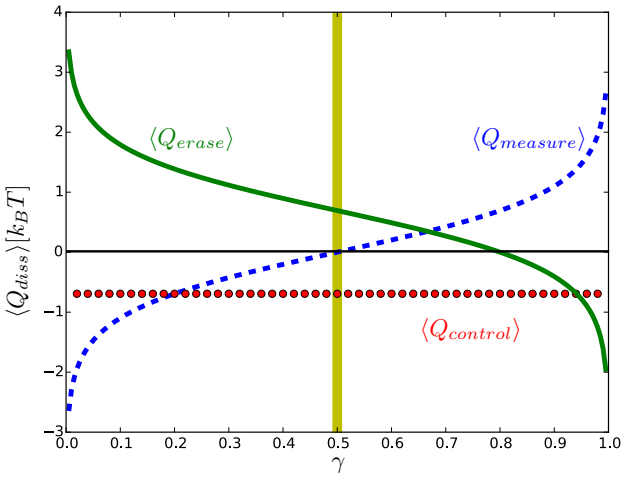

FIG. 3. Beyond Landauer's Principle: Thermodynamic costs (energy dissipation Q_{diss}) for measurement, control, and erasure in Szilard's information engine as a function of γ (Demon partition) with SUS barrier at $\delta = 1/2$. Landauer's Principle applies only at $\gamma = 1/2$ (vertical yellow band): measurements are thermodynamically free, erasure costs since heat is dissipated as a result of Demon resetting. Costs *exactly flip* at $\gamma = 4/5$, though.

While the heat dissipated during control is independent of γ, both measurement and erasure can dissipate any positive or negative amount of heat, depending on γ. Notably, for $\gamma > 1/2$, the Szilard Engine demonstrates an extension of Landauer's Principle [15, 16] in that $\langle Q_{\text{erase}} \rangle \leq k_B T \ln 2$, but this is balanced by an increase in $\langle Q_{\text{measure}} \rangle$. Indeed, for $\gamma = 4/5$, erasure is thermodynamically free and measurement takes on the usual cost of erasure.

Figure 3 illustrates the trade-offs in thermodynamic costs for each step. They sum to zero and so the Engine respects the Second Law over the whole range of δ and γ. The erasure and measurement steps together obey the relation: $\langle Q_{\text{erase}} \rangle + \langle Q_{\text{measure}} \rangle = k_B T\,\mathrm{H}(\delta)\ln 2$, recovering trade-offs noted previously [17, 27–29]. That is, the Szilard Engine achieves the lower bounds on energy dissipation during measurement and erasure. And so, it plays an analogous optimal role in the conversion of information into work as the Carnot Engine does for optimal efficiency when converting thermal energy to work.

Final Remarks We leveraged a straightforward observation to give a thorough dynamical systems, computational mechanics, and thermodynamic analysis of Szilard's Engine: an information engine's intrinsic computation is supported by the evolution of its joint state-space distribution and its thermodynamic costs monitor how those distributional changes couple energetically to its environment.

The Szilard Map construction is straightforward and easy to interpret. For these reasons, we selected it to illustrate the bridge between thermodynamics, information theory, and dynamical systems necessary to fully analyze information engines. The approach generalizes. We can now state *our central proposal*: (i) an information engine is the dynamic over a joint state space of a thermodynamic system and a physically embodied controller, (ii) the causal states of the joint dynamics, formed from the predictive equivalence classes of system histories, capture its information processing and emergent organization, (iii) a necessary component of the engine's effective "intelligence", its memory, is given by its statistical complexity C_μ, (iv) its dissipation is given by the dynamical system negative LCEs, and (v) the rate of energy extracted from the heat bath is governed by the Kolmogorov-Sinai entropy h_μ.

Sequels use this approach to analyze the information thermodynamics of more sophisticated engines, including the Mandal-Jarzynski ratchet [30], experimental nanoscale information processing devices, and intelligent macromolecules.

Supplementary Materials [36]: Calculational details, further discussion and interpretation, and animations illustrating a continuous-time embedding of $\mathcal{T}_{\text{Szilard}}$.

Acknowledgments We thank Gavin Crooks, David Feldman, Ryan James, John Mahoney, Sarah Marzen, Paul Riechers, and Michael Roukes for helpful discussions. JPC is an External Faculty member of the Santa Fe Institute. Work supported by the U.S. Army Research Laboratory and the U. S. Army Research Office under contracts W911NF-13-1-0390 and W911NF-12-1-0234.

[1] T. Rueckes, K. Kim, E. Joselevich, Greg Y. Tseng, C-L. Cheung, and C. M. Lieber. *Science*, 289:94–97, 2000.

[2] A. M. Fennimore, T. D. Yuzvinsky, W.-Q. Han, M. S. Fuhrer, J. Cumings, and A. Zettl. *Nature*, 424:408–410, 2003.

[3] Z. Zhong, D. Wang, Y. Cui, M. W. Bockrath, and C. M. Lieber. *Science*, 302:1377–1379, 2003.

[4] J. Chen, N. Jonoska, and G. Rozenberg, editors. *Nanotechnology: Science and Computation*, Natural Computing, New York, 2006. Springer-Verlag.

[5] F. Julicher, A. Ajdari, and J. Prost. *Rev. Mod. Phys.*, 69(4):1269–1281, 1997.

[6] C. Bustamante, J. Liphardt, and F. Ritort. *Physics Today*, 58(7):43–48, 2005.

[7] A. R. Dunn and A. Price. *Physics Today*, 68(2):27–32, 2015.

[8] L. Zhuang, L. Guo, and S. Y. Chou. *Appl. Phys. Lett.*, 72(10):1205–1207, 1998.

[9] F. Hetsch, N. Zhao, S. V. Kershaw, and A. L. Rogach. *Materials Today*, 16(9):312–325, 2013.

[10] S. Toyabe, T. Sagawa, M. Ueda, E. Muneyuki, and M. Sano. *Nature Physics*, 6:988–992, 2010.

[11] A. Berut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz. *Nature*, 483:187–190, 2012.

[12] R. Klages, W. Just, and C. Jarzynski, editors. *Nonequilibrium Statistical Physics of Small Systems: Fluctuation Relations and Beyond*. Wiley, New York, 2013.

[13] G. M. Zaslavsky. *Chaos*, 5:653–661, 1995.

[14] L. Szilard. *Z. Phys.*, 53:840–856, 1929.

[15] R. Landauer. *IBM J. Res. Develop.*, 5(3):183–191, 1961.

[16] C. H. Bennett. *Intl. J. Theo. Phys.*, 21:905, 1982.

[17] T. Sagawa. *Prog. Theo. Phys.*, 127, 2012.

[18] Cf. the schematic diagram in Ref. [16, Fig. 12], a more complicated product system, used to argue qualitatively that only erasure during logically irreversible operations dissipates energy. Figure 3 shows clearly how this is a restricted case.

[19] A. Lasota and M. C. Mackey. *Probabilistic Properties of Deterministic Systems*. Cambridge University press, Cambridge, United Kingdom, 1985.

[20] J. R. Dorfman. *An Introduction to Chaos in Nonequilibrium Statistical Mechanics*. Cambridge University Press, Cambridge, United Kingdom, 1999.

[21] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, second edition, 2006.

[22] Ja. G. Sinai. *Dokl. Akad. Nauk. SSSR*, 124:768, 1959.

[23] J. P. Crutchfield, P. Riechers, and C. J. Ellison. *Phys. Lett. A*, 380(9-10):998–1002, 2015.

[24] J. P. Crutchfield. *Nature Physics*, 8(January):17–24, 2012.

[25] N. Barnett and J. P. Crutchfield. *J. Stat. Phys.*, 161(2):404–451, 2015.

[26] C. J. Ellison, J. R. Mahoney, and J. P. Crutchfield. *J. Stat. Phys.*, 136(6):1005–1034, 2009.

[27] K. Shizume. *Phys. Rev. E*, 52(4):3495–3499, 1995.

[28] F. N. Fahn. *Found. Physics*, 26(1):71–93, 1996.

[29] M. M. Barkeshli. *arXiv:cond-mat*, 0504323.

[30] A. B. Boyd, D. Mandal, and J. P. Crutchfield. *New J. Physics*, 18:023049, 2016.

[31] M. v. Smoluchowski. *Physik. Zeit.*, 13:1069–1080, 1912.

[32] M. v. Smoluchowski. *Vortage uber die Kinetische Theorie der Materie und der Elektrizitat*. M. Planck, ed., 89–121, Teuber und Leipzig. Berlin, Germany, 1914.

[33] R. Feynman, R. B. Leighton, and M. Sands. *The Feynman Lectures on Physics–Volume 1*. Addison-Wesley, Reading, Massachusetts, 1963.

[34] G. Mayer-Kress and H. Haken. *Comm. Math. Phys.*, 111:63–74, 1987.

[35] J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney. *Phys. Rev. Lett.*, 103(9):094101, 2009.

[36] See Supplemental Material [url], which includes Refs. [31–35].