Article
A Cost/ Speed/ Reliability Tradeoff to Erasing[†]

Manoj Gopalkrishnan

1   Tata Institute of Fundamental Research; manoj.gopalkrishnan@gmail.com
†   This paper is an extended version of the paper entitled "Cost / speed / reliability tradeoff in erasing a bit", presented at the 14th International Conference on Unconventional and Natural Computation, Auckland, New Zealand

Abstract: We present a KL-control treatment of the fundamental problem of erasing a bit. We introduce notions of **reliability** of information storage via a reliability timescale τ_r, and **speed** of erasing via an erasing timescale τ_e. Our problem formulation captures the tradeoff between speed, reliability, and the Kullback-Leibler (KL) cost required to erase a bit. We show that rapid erasing of a reliable bit costs at least $\log 2 - \log\left(1 - e^{-\frac{\tau_e}{\tau_r}}\right) > \log 2$, which goes to $\frac{1}{2}\log\frac{2\tau_r}{\tau_e}$ when $\tau_r >> \tau_e$.

1. Motivation

Biological systems are remarkably ordered at multiple scales and dimensions, from the spatial order witnessed in the packing of DNA inside the nucleus, the arrangement of cells to form tissues, and organs, and whole organisms, to the temporal order witnessed in the execution of various cellular processes. Superficially such order might appear to violate the second law of thermodynamics which requires an increase in disorder with overwhelmingly high probability. In fact, there is no violation since biological systems expend energy to bring about and maintain this increase in order.

We would like to understand this "energy to order" conversion quantitatively. What are the fundamental limits to this conversion? Order can be measured in terms of information, by counting the number of bits required to describe that order. From this point of view, understanding how much energy is required to create order becomes an instance of the investigation of the connection between information processing and thermodynamics. The basic information processing operation that increases order is the operation of "erasing" or resetting a bit to state 0. To fix ideas, imagine erasing random chalk marks from a blackboard, to leave it in a neat and ordered state.

Szilard [1] and later Landauer [2] have argued from the second law of thermodynamics that erasing at temperature T requires at least $k_B T \log 2$ units of energy, where k_B is Boltzmann's constant. The **Szilard engine** is a simple illustration of this result. Imagine a single molecule of ideal gas in a cylindrical vessel. If this molecule is in the left half of the vessel, think of that as encoding the bit "0," and the bit "1" otherwise. Erasing this Brownian bit corresponds to ensuring that the molecule lies on the left half, for example by compressing the ideal gas to half its volume. For a heuristic analysis we may use the ideal gas law $PV = k_B T$, integrating the expression $dW = -PdV$ for work from limits V to $V/2$ to obtain $W = k_B T \log 2$. More rigorous and general versions of this calculation are known, which also clarify why this is a lower bound [3–5].

In practice, one finds that both man-made and biological instrumentation often require energy substantially more than $k_B T \log 2$ to perform erasing [6,7]. John von Neumann remarked on this large gap in his 1949 lectures at the University of Illinois [8]. (Bennett [9] has remarked that DNA polymerases come close to the $k_B T \log 2$ bound. To copy a single base, a DNA polymerase hydrolyzes

a triphosphate molecule to a monophosphate, which provides close to $18k_BT$ at temperature $T = 300K$. Note that this is still almost two orders of magnitude away from $k_BT \log 2$. Further, it is not clear whether the comparison is valid at all since copying and erasing are different operations.)

How does one explain this large gap? Note that the result of $k_BT \log 2$ holds only in the isothermal limit, which takes infinite time. In practice, we want erasing to be performed rapidly, say in time τ_e, which requires extra entropy production. For intuition, suppose one wants to compress a gas in finite time τ_e. The gas heats up, and pushes back, increasing the work required.

Several groups [10–12] have recognized that rapid erasing requires entropy production which pushes up the cost of erasing beyond $k_BT \log 2$, and have obtained bounds for this problem. A grossly oversimplified, yet qualitatively accurate, sketch of these various results is obtained by considering the energy cost of compressing the Szilard engine rapidly. Specializing a result from finite-time thermodynamics [13] to the case of the Szilard engine, one obtains an energy cost $\left(1 + \frac{k_B \log 2}{\sigma \tau_e - k_B \log 2}\right) k_BT \log 2$ where σ is the coefficient of heat conductivity of the vessel.

The bounds obtained by such considerations depend on technological parameters like the heat conductivity σ, and not just on fundamental constants of physics and the requirement specifications of the problem. If one varies over the technological parameters as well, e.g. allowing $\sigma \to \infty$, the energy cost tends to $k_BT \log 2$. Does there exist a more fundamental analysis for the cost of erasing that is independent of technological parameters, and improves on $k_BT \log 2$? This is the open question we address in this paper.

Our contribution: We follow up on von Neumann's suggestion [8] that the gap was "due to something like a desire for reliability of operation." Swanson [14] and Alicki [15] have also looked into issues of reliability. We introduce the notion of "reliability timescale" τ_r, and explicitly consider the three-way trade-off between speed, reliability, and cost.

The other novelty of our approach is in bringing the tools of Kullback-Leibler (KL) control [16,17] to bear on the problem of erasing a bit. The intuitive idea is that the control can reshape the dynamics as it pleases, but pays for the deviation from the uncontrolled dynamics. The cost of reshaping the dynamics is a relative entropy or KL divergence between the controlled and the uncontrolled dynamics, expressed as measures on path space.

We find the optimal control for rapid erasing of a reliable bit, and argue that it requires cost of at least $\log 2 - \log \left(1 - e^{-\frac{\tau_e}{\tau_r}}\right) > \log 2$, which goes to $\frac{1}{2} \log \frac{2\tau_r}{\tau_e}$ when $\tau_r >> \tau_e$. Importantly, our answer does not depend on any technological parameters, but only on the requirement specifications τ_r and τ_e of the problem.

2. The Erasing Problem

As a model of a bit, consider a two-state continuous-time Markov chain with states 0 and 1 and the **passive** or uncontrolled dynamics given by transition rates k_{01} from state 0 to state 1 and k_{10} from state 1 to state 0.
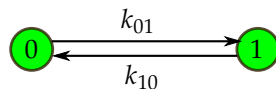
The transition rates k_{01} and k_{10} model spontaneous transitions between the states when no one is looking at the bit or trying to erase it. The time independence of these rates represents the physical fact that the system is not being driven.

Such finite Markov chain models often arise in physics by "coarse-graining." For example, for the case of the Szilard engine, the transition rate k_{10} models the rate at which the molecule enters the left side, conditioned on it currently being on the right side.

Apart from their importance in approximating the behavior of real physical systems, finite Markov chains are also important to thermodynamics from a logical point of view. They may be viewed as finite models of a mathematical theory of thermodynamics. The terms "theory" and

"model" are to be understood in their technical sense as used in mathematical logic. We develop this remark no further here since doing so would take us far afield.

Suppose the distribution at time t is $(p_0(t), p_1(t))$ with $p_1(t) = 1 - p_0(t)$. Then the time evolution of the bit is described by the ODE

$$\dot{p}_0(t) = -k_{01}p_0(t) + k_{10}(1 - p_0(t)). \tag{1}$$

Setting $\pi_0 = k_{10}/(k_{01} + k_{10})$ and the **reliability timescale** $\tau_r := 1/(k_{01} + k_{10})$, this admits the solution

$$p_0(t) = \pi_0 + e^{-t/\tau_r}(p_0(0) - \pi_0) \tag{2}$$

Here τ_r represents the time scale on which memory is preserved. The smaller the rates k_{01} and k_{10}, the larger is the value of τ_r, and the slower the decay to equilibrium, so that the system remembers information for longer.

Fix a **required erasing time** τ_e. Fix $p(0) = \pi_0$. We want to control the dynamics with transition rates $u_{01}(t)$ and $u_{10}(t)$ to achieve $p(\tau_e) = (1, 0)$, where

$$\dot{p}_0(t) = -u_{01}(t)p_0(t) + u_{10}(t)(1 - p_0(t)) \tag{3}$$

We want to find the cost of the optimal protocol $u_{01}^*(t)$ and $u_{10}^*(t)$ to achieve this objective, according to a cost function which we introduce next. In particular, when $k_{01} = k_{10}$, the equilibrium distribution $\pi = (\pi_0, 1 - \pi_0)$ takes the value $(1/2, 1/2)$ and we can interpret this task as erasing a bit of reliability $\tau_r = 1/(k_{01} + k_{10})$ in time τ_e.

2.1. Kullback Leibler Cost

Define the **path space** $\mathcal{P} := \{0, 1\}^{[0, \tau_e]}$ of the two-state Markov chain. This is the set of all paths in the time interval $[0, \tau_e]$ that jump between states 0 and 1 of the Markov chain. Each path can also be succinctly described by its initial state, and the times at which jumps occur. We can also effectively think of the path space as the limit as $h \to 0$ of the space $\mathcal{P}_h := \{0, 1\}^{\{0, h, 2h, \dots, Nh = \tau_e\}}$ corresponding to the discrete-time Markov chain that can only jump at clock ticks of h units.

Once the rates $u_{01}(t), u_{10}(t)$ and the initial distribution $p(0) = p$ for the Markov chain are fixed, there is a unique probability measure $\mu_{u,p}$ on path space which intuitively assigns to every path the probability of occurrence of that path according to the Markov chain evolution (Equation 3) with initial conditions p.

For pedagogic reasons, we first describe the discrete-time measure $\mu_{u,p}^h$ for a single path $i = (i_0, i_1, \dots, i_N) \in \mathcal{P}_h$. First we describe the transition probabilities of the discrete-time Markov chain. For $a, b \in \{0, 1\}$ with $a \neq b$, for all times t, define $u_{aa}^h(t) := 1 - hu_{ab}(t)$ and $u_{ab}^h(t) := hu_{ab}(t)$ as the probability of jumping to a and to b respectively in the time step t, conditioned on being in state a. Then the probability of the path i under control u is given by:

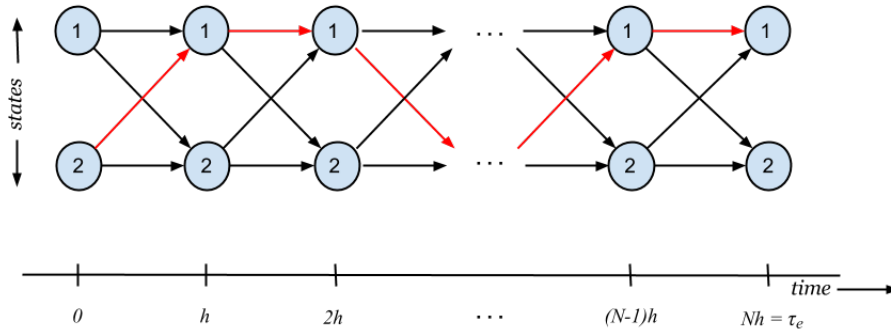$$\mu_{u,p}(i) := p_{i_0} \prod_{j=0}^{N-1} u_{i_j, i_{j+1}}^h(jh)$$

Figure 1. The discrete-time path space $\mathcal{P}_h$. A specific path is labeled in red.

We describe the continous-time case now. We could obtain the measure $\mu_{u,p}$ from $\mu_{u,p}^h$ by sending $h \to 0$, but it can also be described more directly. Fix $i_0 \in \{0,1\}$, and consider the set of paths $\mathcal{S} = \mathcal{S}_{i_0,t_1,t_2,\ldots,t_n}$ starting at i_0 with jumps occurring around times $t_1 < t_2 < \cdots < t_n$ within infinitesimal intervals $dt_1, dt_2, \ldots, dt_n$ and leading to the trajectory $(i_0, i_1, \ldots, i_n) \in \{0,1\}^{n+1}$. Setting $t_0 = 0$:

$$\mu_{u,p}(\mathcal{S}) := p_{i_0} \prod_{j=0}^{n-1} e^{-\int_{t_j}^{t_{j+1}} u_{i_j i_{j+1}}(s)ds} u_{i_j i_{j+1}}(t_{j+1})dt_{j+1}$$

where p_{i_0} is the probability of starting at i_0, $e^{-\int_0^{t_1} u_{i_0 i_1}(s)ds}$ is the probability of not jumping in the time interval $(0, t_1)$, $u_{i_0 i_1}(t_1)dt_1$ is the probability of jumping from i_0 to i_1 in the interval $(t_1, t_1 + dt_1)$ and so on. This is the well-known Feynam-Kac formula for this Markov chain.

Specializing to $u_{01}(t) = k_{01}$ and $u_{10}(t) = k_{10}$, we obtain the probability measure $\mu_{k,p}$ induced on $\mathcal{P}$ by the passive dynamics (Equation 1) with initial conditions p.

We declare the Kullback Leibler (KL) cost $D(\mu_{u,p} \| \mu_{k,p})$ as the cost for implementing the control u. More generally, for a physical system with path space $\mathcal{P}$, passive dynamics corresponding to a measure ν on $\mathcal{P}$, and a controlled dynamics with a control corresponding to a measure μ on $\mathcal{P}$, we declare $D(\mu \| \nu)$ as the cost for implementing the control. This cost function has been widely used in control theory [16–27]. In Section 4 we will explore some other interpretations of this cost function.

3. Solution to the Erasing Problem

Out of all controls $u(t)$, we want to find a control $u^*(t)$ that starts from

$$p(0) = \pi = \left(\frac{k_{10}}{(k_{01} + k_{10})}, \frac{k_{01}}{k_{01} + k_{10}} \right)$$

and achieves $p(\tau_e) = (1, 0)$ while minimizing the relative entropy $D(\mu_{u^*,\pi} \| \mu_{k,\pi})$.

$$\text{To find: } u^* = \arg \inf_u D(\mu_{u,\pi} \| \mu_{k,\pi})$$

$$\text{Subject to: } \mu_{u,\pi}(\tau_e) = (1, 0)$$

Question **??** can be described within the framework of a well-studied problem in optimal control theory that has a closed-form solution [16,17,28]. Following Todorov [16], we introduce the **optimal cost-to-go** function $v(t) = (v_0(t), v_1(t))$. We intend $v_i(t)$ to denote the expected cumulative cost for starting at state i at time $t < \tau_e$, and reaching a distribution close to $(1, 0)$ at time τ_e.

To discourage the system from being in state 1 at time τ_e, define $v_1(\tau_e) = +\infty$ and $v_0(\tau_e) = 0$. Suppose the control performs actions $u_{01}(t)$ and $u_{10}(t)$ at time t. Fix a small time $h > 0$. Define the transition probability $u_{ij}^h(t)$ as the probability that a trajectory starting in state i at time t will be found

in state j at time $t + h$. When $i \neq j$, $u_{ij}^h(t) \approx h u_{ij}(t)$, whereas $u_{ii}^h(t) \approx 1 - u_{ij}^h(t)$ ignoring terms of size $O(h^2)$. We define k_{ij}^h similarly.

Let "log" denote the natural logarithm. To derive the law satisfied by the optimal cost-to-go $v(t)$, we approximate $v(t)$ by the backward recursion relations:

$$
\begin{aligned}
v_0(t) &= \min_{u_{01}(t)} \mathbb{E} \left[v_i(t+h) + \log \frac{u_{0i}^h(t)}{k_{0i}^h} \right] \\
v_1(t) &= \min_{u_{10}(t)} \mathbb{E} \left[v_i(t+h) + \log \frac{u_{1i}^h(t)}{k_{1i}^h} \right]
\end{aligned}
\tag{4}
$$

where the first expectation is over $i \sim_{\text{law}} (u_{00}^h(t), u_{01}^h(t))$, and the second is over $i \sim_{\text{law}} (u_{10}^h(t), u_{11}^h(t))$, and the approximation ignores terms of size $O(h^2)$. As $h \to 0$ the second terms $\mathbb{E} \log \frac{u_{ji}^h(t)}{k_{ji}^h}$ approach the relative entropy cost in path space over the time interval $(t, t+h)$.

In words, Equation 4 says that the cost-to-go from state 0 at time t equals the cost of the control $u(t)$ plus the expected cost-to-go in the new state i reached at time $t + h$. The cost of the control is measured by relative entropy of the control dynamics relative to the passive dynamics, over the time interval $(t, t+h)$.

Define the **desirability** $z_0(t) = e^{-v_0(t)}$ and $z_1(t) = e^{-v_1(t)}$. Define

$$
\begin{aligned}
G_0[z](t) &= k_{00}^h z_0(t) + k_{01}^h z_1(t), \\
G_1[z](t) &= k_{10}^h z_0(t) + k_{11}^h z_1(t).
\end{aligned}
$$

We can rewrite Equation 4 as

$$
\begin{aligned}
\log z_0(t) &= \log G_0[z](t+h) - \min_{u_{01}(t)} \mathbb{E} \left[\log \frac{u_{0i}^h(t) G_0[z](t+h)}{k_{0i}^h z_i(t+h)} \right] \\
\log z_1(t) &= \log G_1[z](t+h) - \min_{u_{10}(t)} \mathbb{E} \left[\log \frac{u_{1i}^h(t) G_1[z](t+h)}{k_{1i}^h z_i(t+h)} \right]
\end{aligned}
\tag{5}
$$

Since the last term is the relative entropy of $(u_{j0}^h(t), u_{j1}^h(t))$ relative to the probability distribution $(k_{j0}^h z_0(t+h) / G_j[z](t+h), k_{j1}^h z_1(t+h) / G_j[z](t+h))$, its minimum value is 0, and is achieved by the protocol u^* given by:

$$
\frac{u_{ji}^*(t)}{k_{ji}} = \lim_{h \to 0} \frac{e^{-v_i(t+h)}}{G_j[z](t+h)} = \frac{e^{-v_i(t)}}{e^{-v_j(t)}}
\tag{6}
$$

when $i \neq j$.

It remains to solve for $z(t)$ and the optimal cost. From Equation 5, at the optimal control u^* the desirability $z(t)$ must satisfy the equation $-\log z(t) = -\log G[z](t+h) + 0$, so that:

$$
\begin{pmatrix} z_0(t) \\ z_1(t) \end{pmatrix} = \begin{pmatrix} 1 - k_{01}h & k_{01}h \\ k_{10}h & 1 - k_{10}h \end{pmatrix} \begin{pmatrix} z_0(t+h) \\ z_1(t+h) \end{pmatrix}
$$

which simplifies to $\frac{dz}{dt} = -Kz$ in the limit $h \to 0$, where $K = \begin{pmatrix} -k_{01} & k_{10} \\ k_{10} & -k_{10} \end{pmatrix}$ is the infinitesimal generator of the Markov chain. This equation has the formal solution $z(\tau_e - t) = e^{Kt} z(\tau_e)$ where $z(\tau_e) = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$. In the symmetric case $k_{01} = k_{10}$,

$$z(t) = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix} + e^{-\frac{\tau_e - t}{\tau_r}} \begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}$$

where $\tau_r = 1/(k_{01} + k_{10})$. Substituting $t = 0$ and taking logarithms, we find the cost-to-go function at time 0:

$$v(0) = \begin{pmatrix} \log 2 \\ \log 2 \end{pmatrix} - \begin{pmatrix} \log\left(1 + e^{-\tau_e/\tau_r}\right) \\ \log\left(1 - e^{-\tau_e/\tau_r}\right) \end{pmatrix}$$

When $v(0) = (1/2, 1/2)$ with $k_{01} = k_{10}$, the cost $C_{\text{erase}}(\tau_r, \tau_e, T)$ required for erasing a bit of reliability $\tau_r = 1/(k_{01} + k_{10})$ in time τ_e is at least:

$$\log 2 - \frac{1}{2} \log\left(1 - e^{-2\tau_e/\tau_r}\right) \tag{7}$$

Note that $C_{\text{erase}} \geq \log 2$ with equality when $\tau_e/\tau_r \to \infty$, since $1 - e^{-2\tau_e/\tau_r} \leq 1$.

From Equation 7, $C_{\text{erase}} \geq \frac{1}{2} \log \frac{2\tau_r}{\tau_e}$ when $\tau_r >> \tau_e$.

4. Interpreting the KL cost

One motivation for our cost function comes from the field of KL control theory. We now compare other possible meanings to this cost function.

4.1. Path space Szilard-Landauer correspondence

The correspondence between information and thermodynamics was revealed in the work of Szilard, and clarified by Landauer. More rigorous and general treatments of this correspondence have been worked out recently [3–5]. We first recall this result, and then show how our cost function is a formal extension of this result.

Consider a physical system with finite state space S and **energy** $E : S \to \mathbb{R}$. (More general state spaces S can be handled by replacing the sum by an appropriate integral. For our present purposes, it suffices to assume S is finite.) Define the **Gibbs distribution** π at temperature T by

$$\pi(i) = \frac{e^{-E_i/k_B T}}{\sum_{j \in S} e^{-E_j/k_B T}}$$

for all $i \in S$. Define the **free energy**

$$F(p) := \sum_{i \in S} p_i E_i - k_B T \sum_{i \in S} p_i \log \frac{1}{p_i}$$

where p is a probability distribution.

Define the **relative entropy** $D(p\|q) = \sum_{i \in S} p_i \log \frac{p_i}{q_i}$ with Euler's constant for the base of the logarithm. Following Jaynes [29], assume that equilibrium π corresponds to a maximally uninformative state of the system, so that we have zero information about the system when it is at equilibrium. Recall that a nat is the unit of information when logarithms are taken to the base of Euler's constant. 1 bit = log 2 nats. Then the relative entropy $D(p\|\pi)$ has an axiomatic identification with the amount of information in nats that we know about the system when it is in a nonequilibrium state p [4].

The following identity is easily verified:

$$F(p) - F(\pi) = k_B T D(p \| \pi).\tag{8}$$

The conceptual significance of this simple identity is that it supplies a dictionary between thermodynamics and information theory [4]. In particular, erasing a bit corresponds to increasing relative entropy which in turn corresponds — via the identity — to increasing available free energy $F(p) - F(\pi)$ by $k_B T \log 2$, recovering the classical result of Szilard as an alternative statement of the second law of thermodynamics. In the other direction, charging a battery corresponds to increasing available free energy which in turn corresponds — via Identity 8 — to erasing of information. This relates the energy efficiency of charging a battery to the energy required to erase a bit.

Now consider our cost function $D(\mu \| \nu)$. The relative entropy $D(\mu \| \nu)$ counts the number of nats erased by the control in path space, relative to the passive dynamics. Since the Szilard-Landauer principle asserts that erasing one bit requires at least $k_B T \log 2$ units of energy, our cost function may be viewed as a **Path Space Szilard-Landauer Principle**, formally extending Identity 8 to path space.

4.2. Thermodynamic interpretation

We wish to compare the cost $D(\mu \| \nu)$ with the usual thermodynamic expected work ΔW. We will quickly outline how thermodynamic quantities can be defined for a two-state Markov chain.

4.2.1. Thermodynamics on a two-state Markov chain

The ideas we present here are well-known in the nonequilibrium thermodynamics community, for example see Propp's thesis [30]. The construction can be carried out more generally, but the generalization is not necessary for our present purposes.

1. Consider again the two-state continuous-time Markov chain with passive dynamics given by transition rates k_{01} and k_{10}.
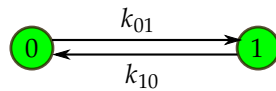
   Let E_0 and E_1 denote the **internal energy** of states "0" and "1" respectively. Then the equilibrium distribution is given by $\pi_0 \propto e^{-E_0/k_B T}$ and $\pi_1 \propto e^{-E_1/k_B T}$. We also have $k_{01}\pi_0 = k_{10}\pi_1$ from detailed balance. Together this yields

$$E_0 - E_1 = k_B T \log \frac{k_{01}}{k_{10}}.\tag{9}$$

2. Now consider the same two-state system with a control applied to it by means of a field of potential $\phi(t) = (\phi_0(t), \phi_1(t))$ so that the potential energy in state i becomes $E_i + \phi_i(t)$. The transition rates due to the control become $u_{01}(t)$ and $u_{10}(t)$. By a reasoning similar to how we derived Equation 9, we get

$$E_0 + \phi_0 - E_1 - \phi_1 = k_B T \log \frac{u_{01}}{u_{10}}.$$

   Combining with Equation 9 this yields

$$\phi_0 - \phi_1 = k_B T \log \frac{u_{01}k_{10}}{u_{10}k_{01}}\tag{10}$$

3. Given a distribution $p = (p_0, p_1)$ on the states, we can define the following thermodynamic quantities:

  - **Expected internal energy** $E(p) = p_0 E_0 + p_1 E_1$.
  - **Entropy** $S(p) = -p_0 \log p_0 - p_1 \log p_1$.
  - **Nonequilibrium free energy** $F(p) = E(p) - k_B T S(p)$.

4. Given a transition from state i to state j in the presence of the control field, we can define the following thermodynamic quantities:

  - **Heat** dissipated $Q_{ij}(t) = E_i + \phi_i(t) - E_j - \phi_j(t)$.
  - **Work** done by the control $W_{ij}(t) = \phi_i(t) - \phi_j(t)$. This expression for work can be traced back to Sekimoto [31], and is commonly employed in the field of Stochastic Thermodynamics to describe the work done by switching on a control field [32].
  - **Entropy increase** of the system $S_{ij}(t) = \log \frac{p_i(t)}{p_j(t)}$.

The first law of thermodynamics manifests as $W_{ij} = Q_{ij} + E_i - E_j$.

5. Suppose the system is described at time t by a distribution $p(t) = (p_0(t), p_1(t))$. Define the **Current** $J_{ij}(t) = p_i(t)u_{ij}(t) - p_j(t)u_{ji}(t)$ so that $\dot{p}_0(t) = -J_{01}(t)$.

6. We can further compute

$$\frac{dE}{dt} = J_{01}(t)(E_1(t) - E_0(t)) = k_B T J_{01}(t) \log \frac{k_{10}}{k_{01}}$$

$$\frac{dW}{dt} = J_{01}(t)W_{01}(t) = J_{01}(t)(\phi_0(t) - \phi_1(t)) = k_B T J_{01}(t) \log \frac{u_{01}(t)k_{10}}{u_{10}(t)k_{01}}$$

$$\frac{dQ}{dt} = J_{01}(t)Q_{01}(t) = J_{01}(t)(E_0 + \phi_0(t) - E_1 - \phi_1(t)) = k_B T J_{01}(t) \log \frac{u_{01}(t)}{u_{10}(t)}$$

$$\frac{dS}{dt} = J_{01}(t)S_{01}(t) = J_{01}(t) \log \frac{p_0(t)}{p_1(t)}$$

$$\frac{dF}{dt} = k_B T J_{01}(t) \log \frac{k_{10}p_1(t)}{k_{01}p_0(t)}$$

7. Define **Total Entropy Production** $S_{\text{tot}}(t)$ to be the total entropy produced from time 0 to time t. In other words, $S_{\text{tot}}(0) = 0$ and

$$\frac{dS_{\text{tot}}(t)}{dt} = \frac{1}{k_B T} \frac{dQ}{dt} + \frac{dS}{dt}.$$

After simplification,

$$\frac{dS_{\text{tot}}(t)}{dt} = (p_0(t)u_{01}(t) - p_1(t)u_{10}(t)) \log \frac{p_0(t)u_{01}(t)}{p_1(t)u_{10}(t)} \geq 0 \tag{11}$$

which is a statement of the second law of thermodynamics.

8. The following identity is immediate

$$\frac{dW}{dt} = \frac{dF}{dt} + \frac{dS_{\text{tot}}}{dt}$$

and is another form of the first law.

4.2.2. Thermodynamic cost for rapid erasing of a reliable bit

How much does it cost for rapid erasing of a reliable bit, with the cost function equal to ΔW? We claim that it costs $k_B \log 2$. In particular, neither the reliability timescale τ_r nor the erasing timescale τ_e appear in this answer.

Suppose we can erase a (τ_r, τ_e) bit for work W. First note that $\frac{dW}{dt}$ is a function of k_{01}/k_{10}, $u_{01}(t)/u_{10}(t)$ and $J_{01}(t)$ as in (6). In particular, simultaneously sending the rates k_{01} and k_{10} as low as possible while keeping their ratio the same has no effect on the work. So if we can erase a (τ_r, τ_e) bit for work W, then we can erase a $(A\tau_r, \tau_e)$ bit for work W, for an arbitrarily large constant A. In particular it is enough for us to demonstrate a protocol when $\tau_r = 1$.

Now note that $\frac{dW}{dt}$ depends only on the ratio $u_{01}(t)/u_{10}(t)$ and not on the actual values of the rates. We can also erase a $(\tau_r, \tau_e/2)$ bit for work W by taking the (τ_r, τ_e)-protocol $(u_{ij}(t))$ and defining new rates $v_{ij}(t) = 2u_{ij}(2t)$. Since $v_{01}(t)/v_{10}(t) = u_{01}(2t)/u_{10}(2t)$, it follows from a simple calculation that the work required does not change.

By taking a limit of this time scaling argument, we only need to erase a $(1, +\infty)$ bit. Here the infinite-time isothermal protocol, which proceeds by raising the '1' well infinitesimally, waiting for the system to equilibrate, and repeating, erases for a total work of $k_B T \log 2$ since that is the free energy difference between the initial and final state, and there is no extra dissipation. This establishes our claim.

A more detailed version of this calculation can be found in [33]. This work assumes that there is a maximum energy limit $E_{\max}$ to which a state can be raised, so that there will be some small error to erasing. It also makes another assumption about thermalization timescale which translates in our setting to assuming that there is a maximum value to the rates $u_{01}(t)$ and $u_{10}(t)$. With these assumptions, they show that the cost of rapid erasing is slightly more than $k_B \log 2$ and goes to $k_B T \log 2$ very quickly as the timescale of thermal relaxation becomes smaller and $E_{\max}$ goes to infinity.

4.2.3. Link between KL-cost and thermodynamic work

We will now characterize entropy production in terms of time reversal. This will allow us to make a link between KL-cost and the thermodynamic work W.

We first recall the notion of **time reversal** of a Markov chain. Usually time-reversal is defined for time-homogeneous Markov chains. However, for the purposes of characterizing entropy production in terms of time reversal, we will work with a definition that applies to time-inhomogeneous Markov chains also. Instead of giving this definition in full generality, we work with a Markov chain with a finite state space. This is sufficient for our purposes, and allows us to avoid dealing with certain technical issues.

Note that given a matrix U with positive non-diagonal entries and $U \cdot 1 = 0$, there is a **nonnegative** vector v such that $v \cdot U = 0$. This can be shown by applying the Perron-Frobenius theorem to the exponential matrix e^U.

Definition 1 (Time-reversal). Consider a continuous-time time-inhomogeneous Markov chain with state space $[n] = \{1, 2, \ldots, n\}$ described by a time-dependent transition matrix $U(t) = (u_{ij}(t))_{i,j \in [n]}$ (so that at time t, $u_{ij}(t)$ denotes the rate of jumping to state j given that the system is in state i). Let $\pi(t)$ be a sequence of stationary probability distributions on $[n]$, i.e., $\pi(t)U(t) = 0$ and $\pi_i \geq 0$ for all $i \in [n]$ and $\sum_i \pi_i = 1$. Then the **time-reversal** Markov chain is described by the time-dependent transition matrix $\hat{U}(t) = (\hat{u}_{ij}(t))_{i,j \in [n]}$ where

$$\hat{u}_{ij}(t) = \frac{\pi_i(t)u_{ij}(t)}{\pi_i(t)}$$

A Markov chain is **reversible** iff $\hat{U} = U$.

The justification for considering time-reversal comes from Bayes' rule. Reversible Markov chains are well-known to be characterized by the conditions of existence of a detailed balanced equilibrium, as well as by the Kolmogorov chain conditions. In particular, two-state Markov chains are always reversible.

For the special case of Equation 3 in particular, given a distribution q at time τ_e, the time reversal Markov chain evolves in time according to the ODE:

$$\dot{q}_0(t) = -u_{01}(t)q_0(t) + u_{10}(t)(1 - q_0(t)),$$
$$q(\tau_e) = q.$$

(12)

We see that the difference is that in Equation 3, the boundary condition was specified at time 0, whereas here the boundary condition is specified at time τ_e.

We define the **time-reversed measure** $\mu_{u,q}^{\text{rev}}$ as the measure on path space corresponding to the process described by Equation 12. Strictly speaking, we should write $\mu_{u,q,\tau_e}^{\text{rev}}$ to denote the time at which the boundary condition is provided to the differential equation, but we will avoid this by using the convention that we're always going to set the boundary condition at time τ_e when considering the time-reversal.

The following result is key to our comparison.

Theorem 2. *Run the control dynamics Equation 3 forward from initial condition $p(0)$ upto time τ_e to obtain the distribution $p(\tau_e)$. Consider the measure $\mu_{u,p(\tau_e)}^{rev}$. Then the total entropy production $S_{tot}(\tau_e)$ from time 0 to time τ_e equals*

$$S_{tot}(\tau_e) = D(\mu_{u,p(0)} \| \mu_{u,p(\tau_e)}^{rev}).$$

Proof. We will show that the time derivative of the RHS equals the rate of entropy production. This will prove the theorem.

Fix a time $t \in [0, \tau_e]$. Let the probability distribution at time t be represented by $p(t) = (p_0(t), p_1(t))$. Let $f_{ij} := p_i(t)u_{ij}(t)$ denote the flow rate from state i to state $j \neq i$ at time t. Then $p_i(t + h) = p_i(t) + h(f_{ji} - f_{ij}) + o(h)$ where $i \neq j$ and $o(h)$ denotes terms $g(h)$ such that $\lim_{h \to 0} g(h)/h \to 0$.

We will consider the probabilities of the four Markov chain transitions $0 \to 0, 0 \to 1, 1 \to 0$ and $1 \to 1$ in the interval $(t, t + h)$ in the limit $h \to 0$ according to $\mu_{u,p(0)}$ and according to $\mu_{u,p(\tau_e)}^{\text{rev}}$. Up to terms of size $o(h^2)$, we have for $i \neq j$:

$$\mu_{u,p(0)}(i \to j) = hp_i(t)u_{ij}(t) = hf_{ij}$$
$$\mu_{u,p(\tau_e)}^{\text{rev}}(i \to j) = hp_j(t + h)u_{ji}(t) = h(f_{ji} + h(f_{ij} - f_{ji}))$$

The increment in the relative entropy in the time interval $(t, t + h)$ equals, upto $o(h^2)$ terms:

$$(1 - hf_{01})\log\frac{1 - hf_{01}}{1 - hf_{10} - h^2(f_{01} - f_{10})} + hf_{01}\log\frac{hf_{01}}{hf_{10} + h^2(f_{01} - f_{10})}$$
$$+ hf_{10}\log\frac{hf_{10}}{hf_{01} + h^2(f_{10} - f_{01})} + (1 - hf_{10})\log\frac{1 - hf_{10}}{1 - hf_{01} - h^2(f_{10} - f_{01})}$$

The off-diagonal terms contribute:

$$(1 - hf_{01})\log\frac{1 - hf_{01}}{1 - hf_{10} - h^2(f_{01} - f_{10})} + (1 - hf_{10})\log\frac{1 - hf_{10}}{1 - hf_{01} - h^2(f_{10} - f_{01})}$$
$$\approx (1 + o(h))(hf_{10} - hf_{01} + o(h^2)) + (1 + o(h))(hf_{01} - hf_{10} + o(h^2))$$
$$\approx o(h^2)$$

Divide by h, and take the limit $h \to 0$. We can ignore the off-diagonal terms. The diagonal terms sum to the rate of entropy production $\frac{dS_{\text{tot}}}{dt}$ as in Equation 11, and we are done. $\square$

By the First Law of Thermodynamics and Theorem 2,

$$\Delta W = \Delta F + k_B T D(\mu_{u,p} \| \mu^{\text{rev}}_{u,p(\tau_e)}) \tag{13}$$

where the increase in free energy of the system

$$\Delta F = k_B T \left(D(p(\tau_e) \| \pi) - D(p(0) \| \pi) \right)$$

by Equation 8. Now to compare our cost function with ΔW.

Theorem 3. *The KL-cost equals change in free energy by $k_B T$ plus a path-space relative entropy term that resembles entropy production:*

$$k_B T D(\mu_{u,p} \| \mu_{k,p}) = \Delta F + k_B T D(\mu_{u,p} \| \mu^{rev}_{k,p(\tau_e)}) \tag{14}$$

where $p(\tau_e)$ is — as in Equation 13 — the solution to the control dynamics Equation 3 at time τ_e.

Proof. Using Equation 8, we can rewrite the claim as

$$D(\mu_{u,p} \| \mu_{k,p}) + D(p(0) \| \pi) = D(\mu_{u,p} \| \mu^{\text{rev}}_{k,p(\tau_e)}) + D(p(\tau_e) \| \pi)$$

Both LHS and RHS equal $D(\mu_{u,p} \| \mu_{k,\pi})$. The assertion for the LHS is straightforward. The assertion for the RHS is true because time-reversal dynamics was defined to keep the stationary distribution π remain stationary under time reversal. $\square$

Comparing 13 and 14, a KL control treatment replaces the total entropy production $D(\mu_{u,p} \| \mu^{\text{rev}}_{u,p(\tau_e)})$ in 13 by the new term $D(\mu_{u,p} \| \mu^{\text{rev}}_{k,p(\tau_e)})$ which compares the control dynamics with the time reversal of the passive dynamics. This suggests an interpretation as follows. If we applied the control during some time interval $[0, \tau_e]$, and remembered what control we applied, then the entropy production is correctly given by $D(\mu_{u,p} \| \mu^{\text{rev}}_{u,p(\tau_e)})$. However, the information that a control was applied also needs to be stored somewhere. If we forget that a control was applied, and if application of the control is very rare, then our default model for the dynamics should be much closer to the passive dynamics. In this case, entropy production may be closer to the value $D(\mu_{u,p} \| \mu^{\text{rev}}_{k,p(\tau_e)})$.

4.3. Large deviations interpretation

Our cost function $D(\mu \| \nu)$ also admits a large deviation interpretation which was, remarkably, already noted by Schrödinger in 1931 [34–37]. Motivated by quantum mechanics, Schrödinger asked: conditioned on a more or less astonishing observation of a system at two extremes of a time interval, what is the least astonishing way in which the dynamics in the interval could have proceeded? Specializing to our problem of erasing, suppose an ensemble of two-state Markov chain with passive dynamics given by Equation 1 was observed at time 0 and at time τ_e. Suppose the empirical state distribution over the ensemble was found to be the equilibrium distribution π at time 0, and $(1, 0)$ at time τ_e respectively. This would be astonishing because no control has been applied, yet the ensemble has arrived at a state of higher free energy. Conditioned on this rare event having taken place, what is the least unlikely measure μ^* on path space via which the process took place?

By a statistical treatment of multiple single particle trajectories, Schrödinger found that the likelihood of an empirical measure μ on path space falls exponentially fast with the relative entropy

$D(\mu\|\nu)$ where ν is the measure induced by the passive dynamics. In particular, the least unlikely measure μ^* is that measure which — among all μ whose marginals at time 0 and time τ_e respect the observations — minimizes $D(\mu\|\nu)$. So for the problem of erasing where $k_{01} = k_{10}$, the measure μ varies over all measures that have marginal $(1/2, 1/2)$ at time 0 and marginal $(1, 0)$ at time τ_e, and μ^* is that measure among all such μ that minimizes $D(\mu\|\mu_{k,(1/2,1/2)})$. Thus our optimal control produces in expectation the least surprising trajectory among all controls that perform rapid erasing.

4.4. Gibbs measure

Equation 6 is not accidental for this example, but is in fact a general feature when the cost function is relative entropy [28]. More abstractly, the Radon-Nikodym derivative (i.e., "probability density") $\frac{d\mu^*}{d\nu}$ of the measure μ^* induced on path space by the optimal control u^* with respect to the measure ν induced by the passive dynamics is a **Gibbs measure**, with the cost-to-go function $v(t)$ playing the role of an energy function. In other words, mathematically our problem is precisely the free energy minimization problem so familiar from statistical mechanics. There is also a possible physical interpretation: we are choosing paths in $\mathcal{P}$ as microstates, instead of points in phase space. The idea of paths as microstates has occurred before [38].

5. Concluding remarks

Since charging a battery can also be thought of as erasing a bit [4], our result may also hold insights into the limits of efficiencies of rapidly charging batteries that must simultaneously hold their energy for a long time.

So long as the noise is Markovian, we conjecture that the KL cost for erasing the two-state Markov chain is a lower bound for more general cases – for example for bits with Langevin dynamics [39], which is a stochastic differential equation expressing Newton's laws of motion with Brownian noise perturbations.

References

1. Leo Szilard. On the reduction of entropy in a thermodynamic system by the interference of intelligent beings. *Z Phys*, 53:840–856, 1929.
2. Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5:183–191, 1961.
3. Massimiliano Esposito and Christian Van den Broeck. Second law and Landauer principle far from equilibrium. *EPL (Europhysics Letters)*, 95(4):40004, 2011.
4. Manoj Gopalkrishnan. The Hot Bit I: the Szilard-Landauer correspondence. *CoRR*, abs/1311.3533, 2013.
5. David Reeb and Michael M Wolf. An improved Landauer principle with finite-size corrections. *New Journal of Physics*, 16(10):103011, 2014.
6. Simon B Laughlin, Rob R de Ruyter van Steveninck, and John C Anderson. The metabolic cost of neural information. *Nature neuroscience*, 1(1):36–41, 1998.
7. Trevor Mudge. Power: A first-class architectural design constraint. *Computer*, 34(4):52–58, 2001.
8. John von Neumann. *Theory of self-reproducing automata*, page 66. University of Illinois Press, Urbana, 1966. lecture delivered at University of Illinois in December, 1949.
9. Charles H. Bennett. The thermodynamics of computation – a review. *Int. J. Theo. Phys.*, 21(12):905–940, 1982.
10. Erik Aurell, Krzysztof Gawędzki, Carlos Mejía-Monasterio, Roya Mohayaee, and Paolo Muratore-Ginanneschi. Refined second law of thermodynamics for fast random processes. *Journal of statistical physics*, 147(3):487–505, 2012.
11. Giovanni Diana, G Baris Bagci, and Massimiliano Esposito. Finite-time erasing of information stored in fermionic bits. *Physical Review E*, 87(1):012111, 2013.

12. Patrick R Zulkowski and Michael R DeWeese. Optimal finite-time erasure of a classical bit. *Physical Review E*, 89(5):052140, 2014.

13. Peter Salamon and Abrahan Nitzan. Finite time optimizations of a Newton's law Carnot cycle. *The Journal of Chemical Physics*, 74:3546, 1981.

14. JA Swanson. Physical versus logical coupling in memory systems. *IBM Journal of Research and Development*, 4(3):305–310, 1960.

15. Robert Alicki. Information is not physical. *arXiv preprint arXiv:1402.2414*, 2014.

16. Emanuel Todorov. Efficient computation of optimal actions. *Proceedings of the national academy of sciences*, 106(28):11478–11483, 2009.

17. Wendell H Fleming and Sanjoy K Mitter. Optimal control and nonlinear filtering for nondegenerate diffusion processes. *Stochastics: An International Journal of Probability and Stochastic Processes*, 8(1):63–77, 1982.

18. Hilbert J Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of statistical mechanics: theory and experiment*, 2005(11):P11011, 2005.

19. Hilbert J Kappen. Linear theory for control of nonlinear stochastic systems. *Physical review letters*, 95(20):200201, 2005.

20. Evangelos A Theodorou. *Iterative path integral stochastic optimal control: Theory and applications to motor control*. PhD thesis, University of Southern California, 2011.

21. Evangelos Theodorou and Emanuel Todorov. Relative entropy and free energy dualities: Connections to path integral and KL control. In *CDC*, pages 1466–1473, 2012.

22. Freek Stulp, Evangelos A Theodorou, and Stefan Schaal. Reinforcement learning with sequences of motion primitives for robust manipulation. *Robotics, IEEE Transactions on*, 28(6):1360–1370, 2012.

23. Krishnamurthy Dvijotham and Emanuel Todorov. A unified theory of linearly solvable optimal control. *Artificial Intelligence (UAI)*, page 1, 2011.

24. Hilbert J Kappen, Vicenç Gómez, and Manfred Opper. Optimal control as a graphical model inference problem. *Machine learning*, 87(2):159–182, 2012.

25. Bart van den Broek, Wim Wiegerinck, and Bert Kappen. Graphical model inference in optimal control of stochastic multi-agent systems. *J. Artif. Intell. Res.(JAIR)*, 32:95–122, 2008.

26. Wim Wiegerinck, Bart van den Broek, and Hilbert Kappen. Stochastic optimal control in continuous space-time multi-agent systems. *arXiv preprint arXiv:1206.6866*, 2012.

27. Matanya Benasher Horowitz. *Efficient methods for stochastic optimal control*. PhD thesis, California Institute of Technology, 2014.

28. Paul Dupuis and Richard S Ellis. *A weak convergence approach to the theory of large deviations*, volume 902. John Wiley & Sons, 2011.

29. Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.

30. Michael Benjamin Propp. *The thermodynamic properties of Markov processes*. PhD thesis, Massachusetts Institute of Technology, 1985.

31. Ken Sekimoto. Kinetic characterization of heat bath and the energetics of thermal ratchet models. *Journal of the physical society of Japan*, 66(5):1234–1237, 1997.

32. Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.*, 75(12):126001, 2012.

33. Browne, Cormac, Andrew JP Garner, Oscar CO Dahlsten, and Vlatko Vedral. Guaranteed energy-efficient bit reset in finite time. *Physical review letters*, 113(10):100603, 2014.

34. E Schrödinger. Uber die umkehrung der naturgesetze, sitzung ber preuss. *Akad. Wissen., Berlin Phys. Math*, 144, 1931.

35. Arne Beurling. An automorphism of product measures. *Annals of Mathematics*, pages 189–200, 1960.

36. Hans Föllmer. Random fields and diffusion processes. In *École d'Été de Probabilités de Saint-Flour XV–XVII, 1985–87*, pages 101–203. Springer, 1988.

37. Robert Aebi. *Schrödinger diffusion processes*. Springer, 1996.

38. AD Wissner-Gross and CE Freer. Causal entropic forces. *Physical review letters*, 110(16):168702, 2013.

39. R. Zwanzig. *Nonequilibrium Statistical Mechanics*. Oxford University Press, USA, 2001.