

Quantum Foundations of Classical Reversible Computing

Michael P. Frank ^{1,*}  and Karpur Shukla ^{2,*} ¹ Center for Computing Research, Sandia National Laboratories, Albuquerque, NM 87185, USA² Department of Electrical and Computer Engineering, Brown University, Providence, RI 02906, USA

* Correspondence: mpfrank@sandia.gov (M.P.F.); karpur_shukla@brown.edu (K.S.);

Tel.: +1-505-284-4103 (M.P.F.); +1-646-580-5277 (K.S.)

† These authors contributed equally to this work.

Abstract: The reversible computation paradigm aims to provide a new foundation for general classical digital computing that is capable of circumventing the thermodynamic limits to the energy efficiency of the conventional, non-reversible digital paradigm. However, to date, the essential rationale for, and analysis of, classical reversible computing (RC) has not yet been expressed in terms that leverage the modern formal methods of non-equilibrium quantum thermodynamics (NEQT). In this paper, we begin developing an NEQT-based foundation for the physics of reversible computing. We use the framework of Gorini-Kossakowski-Sudarshan-Lindblad dynamics (a.k.a. *Lindbladians*) with multiple asymptotic states, incorporating recent results from resource theory, full counting statistics and stochastic thermodynamics. Important conclusions include that, as expected: (1) Landauer's Principle indeed sets a strict lower bound on entropy generation in traditional non-reversible architectures for deterministic computing machines when we account for the loss of correlations; and (2) implementations of the alternative *reversible* computation paradigm can potentially avoid such losses, and thereby circumvent the Landauer limit, potentially allowing the efficiency of future digital computing technologies to continue improving indefinitely. We also outline a research plan for identifying the fundamental minimum energy dissipation of reversible computing machines as a function of speed.

Keywords: non-equilibrium quantum thermodynamics; thermodynamics of computing; Landauer's principle; Landauer limit; reversible computing; resource theory of quantum thermodynamics; Gorini-Kossakowski-Sudarshan-Lindblad dynamics; Lindbladians; von Neumann entropy; Rényi entropy; open quantum systems



Citation: Frank, M.P.; Shukla, K. Quantum Foundations of Classical Reversible Computing. *Entropy* **2021**, *23*, 701. <https://doi.org/10.3390/e23060701>

Academic Editor: Neal G. Anderson

Received: 24 April 2021

Accepted: 27 May 2021

Published: 1 June 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The concept of *reversible computation*, or computation without information loss (even locally), played a centrally important role in the historical development of the thermodynamics of computation [1–7]. It remains critically important today in the field of quantum computing, where it is necessary for maintaining coherence in quantum algorithms [8]. However, the original motivation for reversible computation, which was to circumvent the $k_B T \ln 2$ Landauer limit¹ on energy dissipation in *classical* digital computing, is less often remembered today. Some authors have critiqued the original arguments for Landauer's limit and reversible computing as relying on equilibrium assumptions (e.g., [10]), but in fact, no such assumption beyond the existence of an *external* heat sink at some temperature T is required. When properly stated and interpreted, Landauer's limit holds regardless of whether the computing system is at (or even *close* to) equilibrium *internally*. This statement follows directly from elementary statistical physics and information theory [11,12].

Indeed, Landauer's limit has also been derived directly for systems *out* of equilibrium [13,14]. This nonequilibrium limit is expressed purely in terms of the non-unitality of

¹ In this expression, $k = k_B \approx 1.38 \times 10^{-23}$ J/K is Boltzmann's constant, which is the natural logarithmic unit of entropy [9], and T is temperature.

the quantum channel evolving the system and heat bath. In other words, Landauer's limit has been derived *solely* as a consequence of thermal operations (as defined in NEQT) acting on the joint quantum mechanical evolution of a system and a bath. This directly reinforces the motivation for reversible computing, which is to avoid the Landauer cost of ejecting correlated information into the environment. The free energy² cost of operations that do *not* eject correlated information can be made arbitrarily small, a fact rigorously proven using resource theoretic techniques in NEQT [15]. We discuss these connections in some detail in later sections. Further, the enterprise of recasting the classic understanding of the thermodynamics of computing in more modern terms offers other benefits. In particular, it allows the theoretical apparatus of the modern NEQT formalism to be brought to bear on the problem of analyzing the potential capabilities of, and fundamental limits on, classical reversible computational processes.

This problem is of far more than just academic interest. Today, an increasingly serious concern is that the conventional non-reversible paradigm for general digital computation is approaching firm limits to its energy efficiency and cost efficiency. These limits ultimately trace back to the $k_B T$ thermal energy scale. Since reversible computing is, broadly speaking, the only non-conventional computing paradigm that can potentially offer a sustainable path forward capable of circumventing the efficiency limits associated with that energy scale in *general digital* computing, it is therefore critically important to the prospects for medium- and long-term improvement in the efficiency and economic utility of general digital computing to determine exactly what the potentialities and limitations of reversible computational mechanisms may be, according to fundamental theory.

In this paper, we aim to carry out the essential groundwork for this enterprise, laying down low-level theoretical foundations upon which a comprehensive NEQT-based treatment of physical mechanisms for reversible computation may be based. It is essential for any such effort to identify the most appropriate definitions for key concepts, and we do this throughout, taking special care with the definitions of the appropriate physical concepts corresponding to classical digital computational states and operations. On this ground, we advocate for our position that the most appropriate understanding of Landauer's Principle is to view it as comprising, most essentially, a statement about the strict entropy *increase* that is required due to the loss of mutual information that necessarily occurs whenever (nontrivially) deterministically computed (ergo correlated) bits are thermalized in isolation. There are other, oft-cited forms of the Principle that deal only with a *transfer* of entropy between computational and non-computational forms; but we instead refer to these as *The Fundamental Theorem of the Thermodynamics of Computation* to avoid confusion, since it has long been known that simple transfers of entropy between different forms can occur in a thermodynamically reversible way [7]. The inappropriate conflation of Landauer's Principle proper, as we identify it, with the Fundamental Theorem is what we believe has been the root cause of a great deal of confusion in the thermodynamics of computing field. As we suggest, simply appropriately distinguishing these concepts permits the straightforward resolution of many long-standing controversies.

Another central aim of this work is to go beyond discussions of Landauer's limit, to develop a first-principles model of classical RC operations using information-theoretic techniques in nonequilibrium quantum thermodynamics. These techniques allow us to understand the fundamental quantum mechanical expressions of, and restrictions on, classical RC operations in several ways. From resource theory and fluctuation theorems [15,16], these techniques provide us with a way of understanding the overall limitations of state transitions, including those that correspond to classical RC operations. In addition to constraints, the framework of Gorini-Kossakowski-Sudarshan-Lindblad operators (GKSL operators, also known as Lindbladians) with multiple asymptotic states [17–19] offer a framework by which *explicit* nonequilibrium quantum thermodynamic expressions of

² The free energy referred to here is the $\alpha = 1$ free energy in particular; that is, the (nonequilibrium) Helmholtz free energy. In the words of [15] directly, that work shows that the most general possible type of catalytic thermal operation (and thus the most general type of transition possible in quantum thermodynamics) “restores the distinguished role of the Helmholtz free energy.”

classical RC operations can be realized. As such, these techniques offer a natural language for expressing the dynamics of RC operations, and provide us with an understanding of the fundamental quantum mechanical restrictions on the way these operations can manifest in physical systems. Fundamental bounds on quantities of interest, such as the dissipation of an operation as a function of its speed, will necessarily have to arise from NEQT.

Here, we provide a description of RC operations via the theory of open quantum systems. In this formulation, we can examine the joint evolution of the computing system with a thermal environment (a.k.a. heat bath), using the machinery of completely positive trace preserving maps (CPTP maps, a.k.a. quantum channels). In particular, we rely on the framework of GKSL operators with multiple asymptotic states [17–19] to develop representations of classical reversible information processing operations. Quite powerfully, this framework can directly give us bounds on dissipation quantities of interest not only for RC operations, but for quantum computation (QC) operations as well—since we express RC operations in terms of quantum channels, the results we derive for RC operations can be directly extended to QC operations in future work.

Nonequilibrium quantum thermodynamics is, we believe, a natural and proper language for understanding the fundamental principles of reversible computing, and for expressing RC operations. As such, a broader aim of this work is to bridge outstanding gaps between the NEQT and RC communities. By providing RC practitioners a feel for some of the modern tools used in NEQT and expressing familiar RC concepts in this language, and by providing NEQT practitioners a sense of how RC principles arise naturally from familiar NEQT frameworks, we hope to achieve this synthesis. As such, this presentation is intended to be a brief and self-contained exposure to some of the basic concepts of NEQT. For further reading, [20] provides a highly pedagogical introduction, while [21] gives a clear and comprehensive overview of current major topics.

Our framework of expressing RC in NEQT relies on the quantum information perspective of thermodynamics (comprehensively reviewed in [22]), the resource theory of quantum thermodynamics (comprehensively reviewed in [23–25]), and the theory of open quantum systems (reviewed in [26], comprehensively discussed in [27,28], extended to multiple asymptotic states in [17–19]). The concepts of quantum speed limits and shortcuts to adiabaticity are not discussed here, but will appear in future work; these are comprehensively reviewed in [29,30], respectively. Readers interested in greater detail on this framework are highly encouraged to read these references. Readers unfamiliar with quantum information theory are also encouraged to refer to [8,31–35].

We would also like to emphasize that, in this paper, we are manifestly taking a stance towards the thermodynamics of information that treats systems as evolving *autonomously*, that is, without invoking the concept of an “observer” outside the system that is performing measurements on and/or controlling the system. This is necessary in order to construct a complete, coherent treatment of self-contained physical computing systems. Other examples of work that takes a self-contained/autonomous perspective regarding the physics of information include [36–38].

The structure of the remainder of the paper is as follows: Section 2 describes materials and methods, including outlining a broad theoretical framework in Section 2.1, relating that broad framework to the more detailed tools and methods of NEQT in Section 2.2, and reviewing a variety of existing and proposed physical implementation technologies for reversible computing in Section 2.3. Section 3 presents our early results, specifically reviewing how a few classic theorems can be easily proven in our framework. These include (Section 3.1) The Fundamental Theorem of the Thermodynamics of Computing, which we distinguish from (Section 3.2) Landauer’s Principle (properly stated); (Section 3.3) fundamental theorems of traditional and generalized reversible computing; and (Sections 3.4 and 3.5) the representation of classical reversible computational operations via the frameworks of catalytic thermal operations and GKSL dynamics. Section 4 gives some general discussion of results and outlines our research plan looking forwards, and Section 5 concludes.

2. Materials and Methods

As this article presents theoretical, rather than experimental, work, there is no laboratory apparatus to speak of; however, we provide a brief review of some of the existing and proposed physical implementation technologies for reversible computing in Section 2.3. But first, we present the key foundational definitions of our broad theoretical picture in Section 2.1. Please note that this presentation roughly follows, but expands upon, that given in [11,12,39,40]. Then in Section 2.2, we tie this broad picture to the much more detailed theoretical apparatus of NEQT.

2.1. Broad Theoretical Foundations

In this subsection, we present and review a number of important low-level definitions that form the broad foundation upon which our overall approach to the physics of reversible computing is built. This includes (Section 2.1.1) our overall picture based on a framework of open quantum systems; (Section 2.1.2) our definition of classical digital computational states and their physical representation, which invokes what we call a *proto-computational basis*, which may in general be time-dependent; (Section 2.1.3) our definitions for classical computational operations, and different types of operations, which are expressible in terms of (Section 2.1.4) primitive computational state transitions; and (Section 2.1.5) the appropriate definition of what it means for a given unitary quantum time evolution to *implement* a classical computational operation.

2.1.1. Open Quantum Systems Framework

In this subsection, we briefly review the broad outlines of the open quantum systems based picture that we are using in this paper. Further details are developed in Section 2.2.

2.1.1.1. System and Environment

We begin with a fairly conventional picture of a physical computer system based on an open quantum systems perspective. At the highest level, we assume that the model universe \mathcal{U} under study can be described as a composition of two subsystems \mathcal{S} , \mathcal{E} , where \mathcal{S} is the physical computer system in question, and \mathcal{E} is its external environment (i.e., the rest of \mathcal{U} , outside of \mathcal{S}). As an example, one could define the “computer system” \mathcal{S} as consisting of everything (i.e., all quantum fields) encompassed within some region of $(3 + 1)D$ spacetime circumscribed by some closed $(2 + 1)$ -dimensional bounding surface. For simplicity, one could think of a spatial boundary that is unchanging in time over some interval. Typically, our analyses will treat the environment \mathcal{E} as an (effectively infinitely large) uniform thermal reservoir (heat bath) that is internally at thermal equilibrium, at some (effectively constant over time) temperature T . The temperature may be treated as effectively constant when the environment is large enough that its temperature is negligibly affected by heat transferred from \mathcal{S} .³

Meanwhile, we will treat the computer system \mathcal{S} as an (in general) *non-equilibrium* system which includes its own internal supply of free energy (e.g., this could be a battery or a fuel reservoir). This is just a simplification of the overall picture, for our present purposes, to avoid the need to explicitly represent a flow of work or free energy in from a separate “power supply” system; that is, the power supply is treated as internal to the computer. However, we will allow that the system \mathcal{S} is able to exchange thermal energy (and entropy) with its environment \mathcal{E} through (all or some portion of) its boundary. Typically, the computer would be assumed to expel waste heat to its external environment \mathcal{E} during operation in order to maintain its (\mathcal{S} ’s) own internal operating temperature (which will generally be non-uniform) within some reasonable bounds. The mechanisms for managing the needed thermal flows are generally assumed to be contained within \mathcal{S} . See Figure 1.

³ Of course, this model is already somewhat of an idealization, since a typical real environment would attain a nonuniform temperature profile under a steady-state thermal flow with constant power output from \mathcal{S} , but it can nevertheless be considered an adequate model for an initial study.

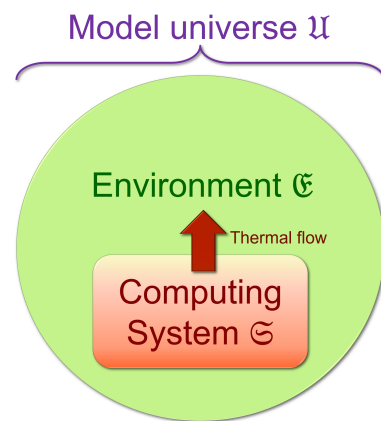


Figure 1. Simplified picture of our model universe \mathcal{U} in an open quantum systems framework. Power supplies and waste heat removal mechanisms are assumed to be included within the physical computer system \mathcal{S} . In general, we may assume there is a flow of waste heat from the system out to an (assumed very large) external heat bath \mathcal{E} .

2.1.1.2. Decoherence Model

A further important simplifying assumption is that we, as modelers, cannot effectively track any (classical or quantum) correlations between the detailed states of \mathcal{S} and \mathcal{E} , or internal correlations between different parts of \mathcal{E} , on any practical timescales. Note, this is not to say that such correlations do not exist physically, as they do under unitary time evolution, but rather just that they are not reflected in our state of knowledge as modelers. A typical assumption, which we adopt, is that it is also safe, or reasonable, to ignore any such correlations that may exist.⁴

Stated slightly more formally, we first assume, for simplicity, that the Hilbert space $\mathcal{H}_{\mathcal{U}}$ of the model universe \mathcal{U} factorizes neatly into separate Hilbert spaces for the system \mathcal{S} and environment \mathcal{E} , that is,

$$\mathcal{H}_{\mathcal{U}} = \mathcal{H}_{\mathcal{E}} \otimes \mathcal{H}_{\mathcal{S}}. \quad (1)$$

Given this, we can imagine that, within a negligible thermalization timescale subsequent to the emission of any small increment ΔQ of waste heat out of \mathcal{S} , the mixed state $\rho_{\mathcal{U}}$ of the model universe would quickly degrade, for all practical purposes, into a (correlation-free) product state $\rho_{\mathcal{U}} = \rho_{\mathcal{E}} \otimes \rho_{\mathcal{S}}$, where $\rho_{\mathcal{E}}$ is the maximum-entropy (equilibrium) mixed state of energy $Q_{\mathcal{E}}$ which includes the heat increment ΔQ after it has diffused into the environment, and $\rho_{\mathcal{S}}$ is a reduced density matrix for the mixed state of the computer system after one has traced out any lingering correlations it may have had with the environment initially upon emission of the waste heat. Note that in the absence of totally separable dynamics (i.e., a Hamiltonian over $\mathcal{H}_{\mathcal{U}}$ given at all times by $\hat{H}_{\mathcal{U}} = \hat{H}_{\mathcal{E}} \otimes \hat{H}_{\mathcal{S}}$), a strict entropy increase is implied by taking the trace over \mathcal{E} (compared to the entropy of an immediately-prior joint state $\rho'_{\mathcal{U}}$ briefly entangling the environment with the system) in the instant just after the emission of the heat ΔQ . Thus, simply performing this state reduction results in global entropy increases in the model universe even when all other dynamics (including the internal dynamics within \mathcal{S}) is taken to be unitary. This state reduction process models the effective decoherence of the system \mathcal{S} as a result of its interaction with the (modeled as thermal) environment \mathcal{E} [41].

A slightly more general model (with weaker assumptions) can be provided by stipulating that we take the trace over \mathcal{E} only once, at the very *end* of an evolution of interest, rather than continuously after each incremental emission ΔQ of heat into this environment. Postponing the state reduction allows for the possibility that correlations/entanglements between the system \mathcal{S} and environment \mathcal{E} , and within \mathcal{E} may persist for some period of

⁴ Whether this is, in fact, a valid assumption is a broad question about the applicability of this open-systems perspective which we do not address in this paper.

time, and affect the evolution to some extent. However, it is not expected that this change will make very much difference in practice.⁵

Modeling the allowed thermal transformations of open quantum systems in detail is the topic of the *resource theory of quantum thermodynamics* (RTQT), which we review briefly in Section 2.2.1 below. First, we continue outlining our broad framework for studying the physics of RC.

2.1.2. Computational States and the Proto-Computational Basis

We now discuss how to formally model, in both abstract and physical terms, the digital computational states of a computer system \mathfrak{S} . Note that our emphasis, in this paper, is on *classical*, not quantum, reversible computation. Furthermore, we wish the scope of our model to include the usual case in real engineered digital computing systems, which is that digital computational states may be encoded by extended physical objects, whose detailed microstate is, in general, not fully determined by the computational state being represented. As an example of this, consider a logic node (connected conductor) within a digital CMOS circuit, where typically the digital symbols ‘0’ and ‘1’ may nominally be represented by node voltages within certain pre-specified, non-overlapping low-to-high ranges $[V_{0L}, V_{0H}]$ and $[V_{1L}, V_{1H}]$, respectively.

We will make the above, informally stated notion (regarding the correspondence between the abstract digital state, and the more detailed physical microstates that are interpreted as encoding it) more formal and precise in Section 2.1.2.2 below, and then in Section 2.1.2.3 we will show how this formal structure allows us to systematically subdivide any given computing system into what we call *computational* versus *non-computational* subsystems. These formalize what Bennett [7] refers to as the *information-bearing* versus *non-information-bearing* degrees of freedom within the system. (However, we will avoid using the latter terminology in this paper, since we consider the non-computational subsystem to still carry *physical* information.)

2.1.2.1. Designated Times

Given that we wish to model active computing machines in which the abstract computational state of the machine *changes* over the course of some time interval, we can expect to encounter the difficulty that the classical digital state of the machine, which, as a discrete entity, takes on values that range over some merely countable set, may not be well-defined (in traditional terms, at least) at all moments during the (physically continuous) transition from one state to the next. To avoid this difficulty, while maintaining simplicity in our model, we will declare, for purposes of the present paper, that there exists some countable set $\{\tau_\ell\}$ of time points $\tau_\ell \in \mathbb{R}$, labeled with integers $\ell \in \mathbb{Z}$, which we will call the *designated times* at which the classical digital computational state of the machine is well-defined.

Note that this model is somewhat oversimplified, since a real engineered computing system is typically not monolithic, but is broken down into subsystems, and it may be the case that the larger system is globally *asynchronous*, in the sense that some subsystems may be in the middle of state transitions while others are in well-defined states. Indeed, depending on the system architecture, there may be *no* moments at which the *entire* machine is simultaneously in a nominally well-defined digital state. However, we will postpone elaboration upon methods to handle this more general case to a later time, as it does not affect anything essential in the present paper.⁶

⁵ As a thought experiment, consider a computer system \mathfrak{S} in deep space, such that any thermal photons emitted from the system into the environment would be expected to mostly just propagate to infinity, with only an astronomically tiny probability of reflecting off of interplanetary gas or dust in such a way as to convey correlated quantum information back into the system. The analysis of such cases, at least, would clearly be only insignificantly affected by postponing the state reduction.

⁶ However, to just briefly preview one way in which a resolution of this problem can work, we can augment the concept of a *well-defined state* of a classical computation with that of a *well-defined state transition*, as we do in Section 2.1.4; this can be meaningful even for non-reversible and/or stochastic operations. Then, at any moment across an extended, asynchronous machine, we can say that each local subsystem is either in a well-defined computational state, or is partway through a well-defined state transition.

2.1.2.2. Computational States Correspond to Sets of Orthogonal Microstates

Regardless of what precise physical encoding of digital computational states is used, we take it as fundamental to the concept of a classical digital computational state that at any designated time $t = \tau_\ell \in \mathbb{R}$, there exists some set $\mathcal{C}(t) = \{c_i(t)\}$ of abstract entities comprising all of the possible alternative *well-defined computational states* $c_i(t)$ of the computer system \mathfrak{S} that the machine could occupy at the time t , and further, that, for any given such state $c = c_i(t)$, there exists some corresponding set $\mathcal{B}_c \subset \mathcal{H}_{\mathfrak{S}}$ of mutually orthogonal, normalized basis vectors $\vec{b} \in \mathcal{B}_c$, each of which represents a pure quantum state of \mathfrak{S} that is unambiguously interpretable as representing the state c . In other words, there is some orthonormal basis $\mathcal{B} \supset \mathcal{B}_c$ for $\mathcal{H}_{\mathfrak{S}}$ such that, if one were to hypothetically perform a complete projective measurement⁷ of the state of the entire computer system \mathfrak{S} down onto the basis \mathcal{B} at time t , and the measured state $|\psi\rangle$ in that basis were found to be one of the basis states $\vec{b} \in \mathcal{B}_c$, then the computational state is unambiguously interpreted to be c . It follows from this that any superposition of the $\vec{b} \in \mathcal{B}_c$ must also be unambiguously interpreted as c , since such a superposition is not distinguishable from the members of \mathcal{B}_c ; if any such superposition state were measured in the basis \mathcal{B} , the projected state would necessarily be contained in the set \mathcal{B}_c .

We note that, when the physical state of a machine is dynamically evolving over (continuous) time, the basis states making up the set \mathcal{B}_c could also generally be evolving; for example, consider an information-bearing signal pulse propagating down a transmission line, which may be convenient to represent in terms of a basis that propagates down the line along with the pulse. When discussing such cases, we can write $\mathcal{B}_c(t)$ to explicitly denote the possible time-dependence of the physical basis-set representation of a given computational state.

However, as mentioned above, we will often assume, for simplicity, that there exists some discrete set of designated time points $\tau_\ell \in \mathbb{R}$ (where $\ell \in \mathbb{Z}$) at which the computational states are well-defined, and focus our attention on those. This will then allow us to characterize non-reversible and stochastic computational evolutions in between those designated time points, in which there is merging or splitting of computational states, at a more abstract level in our model, without having to specify all details of the transition process, such as when, exactly, the computational states split or merge (and indeed, physically, these transitions will in general not be sharp).

Further, it follows from the assumption that, at designated time points $t = \tau_\ell$, each $\mathcal{B}_c(\tau_\ell)$ identifies $c(t)$ unambiguously, that, at least at these times, all of the \mathcal{B}_c are mutually orthogonal to each other, and thus can be taken to be disjoint subsets of a single “master” orthonormal basis $\mathcal{B}(t)$; that is, $\forall c \in \mathcal{C}(t) : \mathcal{B}_c(t) \subset \mathcal{B}(t)$. We call such a master basis a *proto-computational basis* for the system \mathfrak{S} ; “proto” because the basis states unambiguously determine the computational state, but are not in general uniquely determined by the computational state. They are lower-level, physical entities, defined prior to the computational state itself.

Note that any particular proto-computational basis $\mathcal{B}(t)$ at a designated time point $t = \tau_\ell$, since it is defined to be a *complete* basis for the Hilbert space $\mathcal{H}_{\mathfrak{S}}$ of the physical system comprising our computer, may in general include some basis states \vec{b} that do not fall into *any* of the sets $\mathcal{B}_c(t)$. These are microstates of the system that do not correspond to well-defined computational states. Such microstates could arise in practice for any number of reasons; for example, such as if the machine has not yet been powered on and initialized, or it has broken down, or simply gone out of spec. Regardless of the cause, we will group these “invalid” basis states together into a special set $\mathcal{B}_\perp = \mathcal{B} - \bigcup_{c \in \mathcal{C}} \mathcal{B}_c$ meaning that the computational state is undefined; for convenience, we can also define an extra “dummy” computational state c_\perp representing this undefined condition, and an

⁷ Please note that this definition of computational states does not require us to actually be able to do these complete projective measurements in practice; it is sufficient, for purposes of the definition, that they could be done *in principle*, by (we can imagine) applying a suitable abstract operator that measures some complete set of commuting observables of the system.

augmented computational state set $C_{\perp} = C \cup \{c_{\perp}\}$, so that then we can say that the system \mathfrak{S} always has *some* computational state $c \in C_{\perp}$, although it may be the undefined state c_{\perp} . With this change, note that the set $\{B_c\}$ of basis sets corresponding to all of the computational states $c \in C_{\perp}$ (in the augmented set) now corresponds to a proper set-theoretic *partition* of the full proto-computational basis \mathcal{B} . See Figure 2.

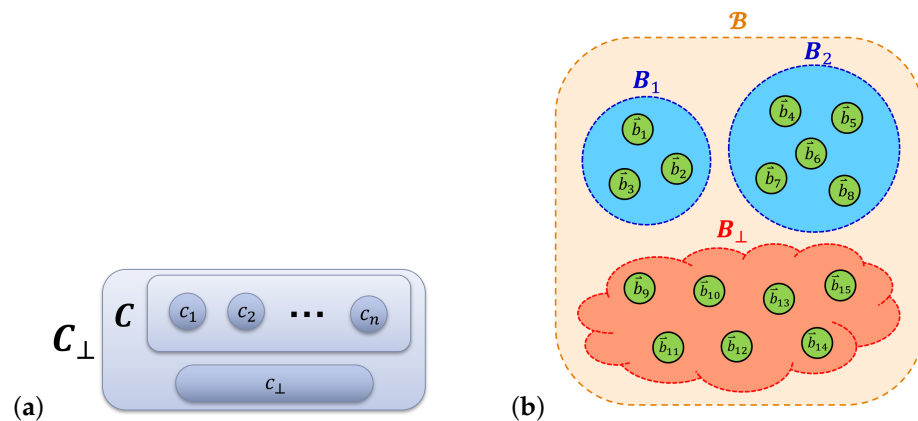


Figure 2. Model of classical digital computational states (at some particular time $t = \tau_{\ell} \in \mathbb{R}$). (a) Abstract computational states of a physical computer system \mathfrak{S} with n distinct states. The “catch-all” state c_{\perp} represents the condition that the physical state of \mathfrak{S} is such that the computational state is not otherwise well-defined. (b) Basis sets B_i corresponding to the computational states c_i , where $i \in \{\perp, 1, \dots, n\}$. Here, $n = 2$. These basis sets partition the complete proto-computational basis \mathcal{B} .

Note that the foregoing treatment of computational states is really no different, fundamentally, from the case of identifying any other (potentially macroscale) classical discrete state variable. In other words, a classical computational state, in our formulation, can be viewed as simply corresponding to a discrete physical macrostate that we happen to consider as carrying some informational significance within a computational system.

2.1.2.3. Computational and Non-Computational Subsystems

As an additional, but inessential assumption that will be useful in some derivations, we can suppose that the Hilbert space $\mathcal{H}_{\mathfrak{S}}$ of the system can be factored as a product of subspaces corresponding to what we call *computational* and *non-computational* subsystems $\mathcal{C}, \mathfrak{N}$ of the computer system \mathfrak{S} . That is, we write $\mathcal{H}_{\mathfrak{S}} = \mathcal{H}_{\mathcal{C}} \otimes \mathcal{H}_{\mathfrak{N}}$, with the idea being that the computational states c correspond to basis vectors of $\mathcal{H}_{\mathcal{C}}$, which are tensored with the basis vectors of $\mathcal{H}_{\mathfrak{N}}$ to obtain the protocomputational basis \mathcal{B} for the entire system \mathfrak{S} .

However, this factorizability assumption is really just a special case, which only holds when the basis sets B_c of the whole system are identically sized. More generally, we can express $\mathcal{H}_{\mathfrak{S}}$ as a *subspace sum*:

$$\mathcal{H}_{\mathfrak{S}} = \bigoplus_{c \in C_{\perp}} \mathcal{H}_{\mathfrak{N}}^c, \tag{2}$$

where $\mathcal{H}_{\mathfrak{N}}^c$ denotes the Hilbert space of the non-computational subsystem \mathfrak{N} , when restricted to the case that the computational state is c ; that is, it is the subspace spanned by the basis vectors in B_c . See Figure 3.

2.1.2.4. Rapid Collapse of Superpositions

In the course of the real physical evolution of the system \mathfrak{S} , it is of course possible, in general, that quantum states could arise that are superpositions of basis states \vec{b} from *different* basis sets B_c and exist in the system briefly, yielding an indeterminate computational state at such moments. However, since our primary focus, in the present paper, is on the analysis of machines that are not even designed to carry out quantum computing, it is reasonable to suppose that such superposition states will spontaneously decohere on very short timescales, as they would naturally tend to do anyway in most large-scale systems.

In other words, we expect that, most of the time, our computational subsystem \mathcal{C} will be living in a *decoherence-free subspace* (DFS), such that the computational states are naturally stable, as part of the system’s “pointer states” [41] towards which the system is continually being decohered by its interactions with its environment. (In this context, the environment of the computational subsystem \mathcal{C} can include portions of the non-computational subsystem \mathcal{N} of the computer system, as well as the machine’s external environment \mathcal{E}).

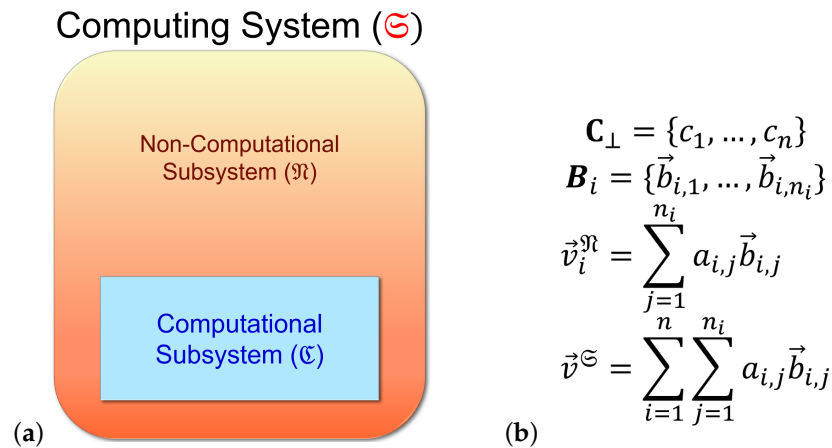


Figure 3. (a) Breakdown of the computing system \mathcal{S} into computational (\mathcal{C}) and non-computational (\mathcal{N}) subsystems. The Hilbert space of \mathcal{S} can be expressed as either a product $\mathcal{H}^{\mathcal{S}} = \mathcal{H}^{\mathcal{C}} \otimes \mathcal{H}^{\mathcal{N}}$ of subsystem subspaces, or more generally as a subspace sum $\mathcal{H}^{\mathcal{S}} = \bigoplus_{c \in \mathbf{C}_\perp} \mathcal{H}_c^{\mathcal{N}}$. (b) A general vector $\vec{v}^{\mathcal{S}}$ in a system Hilbert space $\mathcal{H}^{\mathcal{S}}$ that is defined as a sum of the non-computational Hilbert spaces $\mathcal{H}_i^{\mathcal{N}}$ corresponding to the n individual computational states $c_i \in \mathbf{C}_\perp$ can be obtained by simply summing general vectors $\vec{v}_i^{\mathcal{N}}$ within the individual subspaces $\mathcal{H}_c^{\mathcal{N}}$, treating the basis vectors $\vec{b}_{i,j}$ across all n subspaces as mutually orthogonal. Note that the dimensionalities n_i of the component subspaces do not all have to be the same in general.

Thus, at least in the case of a logically deterministic (non-stochastic) computational process starting from a well-defined initial computational state $c(\tau_0)$, we assume that each of the system’s pointer states will, at any designated time τ_ℓ (for $\ell > 0$), have all (or nearly all) of its probability mass concentrated within a single well-defined computational state $c(\tau_\ell)$. Furthermore, even in a stochastic computation, we will obtain a classical statistical mixture of computational states, not a quantum superposition over them. The challenge, in reversible computing, is then to arrange for the system’s already naturally stable pointer states to (notwithstanding their stability) still remain subject to undergoing a physically natural (if engineered) dynamics in which they will evolve, relatively quickly over time, translating themselves (eventually) one-to-one into new computational states $c(\tau_{\ell+1})$, which may bear new semantic interpretations. Such a system could thereby carry out useful computations at useful speeds.

We will discuss computational operations, and their physical correlates, in more detail in Sections 2.1.3–2.1.5. These can be very directly embedded in open quantum systems exhibiting GKSL dynamics, which we discuss in detail in Section 3.5. However, even just the above definitions already suffice to prove what we call *The Fundamental Theorem of the Thermodynamics of Computing*; this is discussed in Section 3.1.

2.1.2.5. Timing Variables

One nearly ubiquitous feature of engineered physical computing systems is the concept of a *timing variable*, that is, some non-computational, non-equilibrium degree of freedom that influences *when* transitions between computational states will occur, and possibly how long they will take. As an example, an ordinary synchronous digital computer normally includes at least one *clock oscillator*, which outputs a periodic *clock* waveform at a prespecified or controllable frequency which is used to control the timing of digital

operations. In such a situation, we can take the phase θ of the oscillator as a timing variable. In adiabatic circuits (see Section 2.3.1), not only the frequency but also the *speed* (quickness) of digital state transitions is controlled by the clock speed $\omega = d\theta/dt$. Furthermore, even non-synchronous computing systems typically still have physical degrees of freedom that influence the timing of transitions. For example, in the novel BARC (Ballistic Asynchronous RC) computing paradigm being developed at Sandia, discussed further in Section 2.3.5 below, individual bits propagate ballistically as flux solitons (*fluxons*) traveling along interconnect lines between devices; the position x of a given fluxon (of a given velocity) along the length of its interconnect can be considered a timing variable.

It is important to note that, while the values of timing variables are not digitally discretized, they are also generally *not* entirely random, or uncorrelated to other parts of the machine, unlike thermal state variables. Thus, timing variables will be the one common exception, in digital computing systems, to our general rule that non-computational degrees of freedom will be assumed to rapidly thermalize.

Next, we define some key concepts of classical computational operations.

2.1.3. Computational Operations

In order to discuss in detail the thermodynamic implications and limits of performing classical digital computational operations (including reversible operations), we first present some basic terminology and definitions relating to such operations in this subsection.

As mentioned, in general the set $\mathcal{C}(\tau_\ell)$ of well-defined computational states could be different at different designated time points τ_ℓ , but, to simplify our presentation, we will temporarily focus on the case where it is unchanging, that is, $\forall \ell \in \mathbb{Z} : \mathcal{C}(\tau_\ell) = \mathbf{C}$.

To permit treatment of stochastic (randomizing) computational operations, we define some related notation. Let $\mathcal{P}(\mathbf{C})$ denote the set of all (normalized) probability distributions over \mathbf{C} . For simplicity, to avoid having to deal with normalizability issues, we can assume that \mathbf{C} is finite.⁸

Then, a (possibly stochastic) *computational operation* O on \mathbf{C} simply refers to some arbitrary function $O : \mathbf{C} \rightarrow \mathcal{P}(\mathbf{C})$ mapping each *initial state* $c_i \in \mathbf{C}$ to a probability distribution over the possible *final states* $c_f \in \mathbf{C}$. For a given initial computational state c_i , we can write $O(c_i) = P_i \in \mathcal{P}(\mathbf{C})$ where $P_i : \mathbf{C} \rightarrow [0, 1]$ denotes the resulting probability distribution over final states. We can also allow O to be a partial function, for example, when discussing operations that are not defined over all states $c \in \mathbf{C}$, which can be useful if the operation will only ever be applied to states $c \in \text{dom}[O] \subset \mathbf{C}$.

Note that it is sufficient, for our present purposes, to use probabilities in the above definition instead of complex amplitudes, since, for classical reversible computing systems, we are going to assume that the system is highly decoherent in any case; any superposition over different computational states would soon decohere to a classical statistical mixture.⁹

2.1.3.1. Deterministic Operations

A particular computational operation O is called (*fully*) *deterministic* (meaning, non-stochastic) if and only if all of its final-state distributions P_i have zero entropy, that is, $\forall c \in \mathbf{C} : H(O(c)) = 0$, where here we reference the standard (Shannon) entropy functional $H(\cdot) : \mathcal{P}(\mathbf{C}) \rightarrow \mathbb{R}^{0+}$, that is,

$$H(p) = - \sum_{c \in \mathbf{C}} p(c) \log p(c), \quad (3)$$

in generic logarithmic units [9]. (Note this is 0 only in the limit of a point distribution.)

⁸ Indeed, due to the holographic bound [42], if the minimal bounding surface of \mathfrak{S} has finite area, then the Hilbert space $\mathcal{H}_{\mathfrak{S}}$, and therefore also $\mathcal{H}_{\mathcal{C}}$, must be finite-dimensional in any event.

⁹ However, in the future, we anticipate that the present line of work may usefully be extended to explore dissipation limits for quantum computation as well; at that point, it would be appropriate to replace the probability distribution $P_i(c)$ with a more general density operator ρ_i .

If an operation is not fully deterministic, we say it is *stochastic*. We could also have that O is *deterministic over a subset* $A \subset C$ of initial states, whilst not being deterministic over the entire set C . Such an O can also be called *conditionally deterministic under the precondition that the initial state $c \in A$* .

2.1.3.2. Reversible Operations

We say that an operation O is (*unconditionally, logically, fully*) *reversible* if and only if there is no state $c_k \in C$ such that for two different i, j ($i \neq j$), both $P_i(c_k) > 0$ and $P_j(c_k) > 0$. Otherwise, we say that O is *logically irreversible*. We say that O is *conditionally (logically) reversible under the precondition that $c \in A$* , for some $A \subset C$, if and only if there is no state $c_k \in C$ such that, for two different i, j ($i \neq j$) with $c_i, c_j \in A$, it is the case that $P_i(c_k) > 0$ and $P_j(c_k) > 0$. In such a case, we could also say that O is *reversible over A* .

2.1.3.3. Time-Dependent Case

Note that it is easy to generalize the above definitions to situations in which the set C of computational states may be different at different designated times. Let $s = \tau_\ell$, $t = \tau_m$ be two different designated times, with $s < t$; then we can write O_s^t to denote a computational operation being performed over the time interval from start time s to end time t . Then we have that $O_s^t : C(s) \rightarrow \mathcal{P}(C(t))$, and the remaining definitions (for determinism, reversibility, etc.) also change accordingly, in the natural way. For an operation taking place between times s and t ($s < t$), we can define $d = t - s$ as the *delay* or *latency* of the operation, and $q = d^{-1}$ as its *quickness* or *speed*.

The above definitions are illustrated in Figure 4 below.

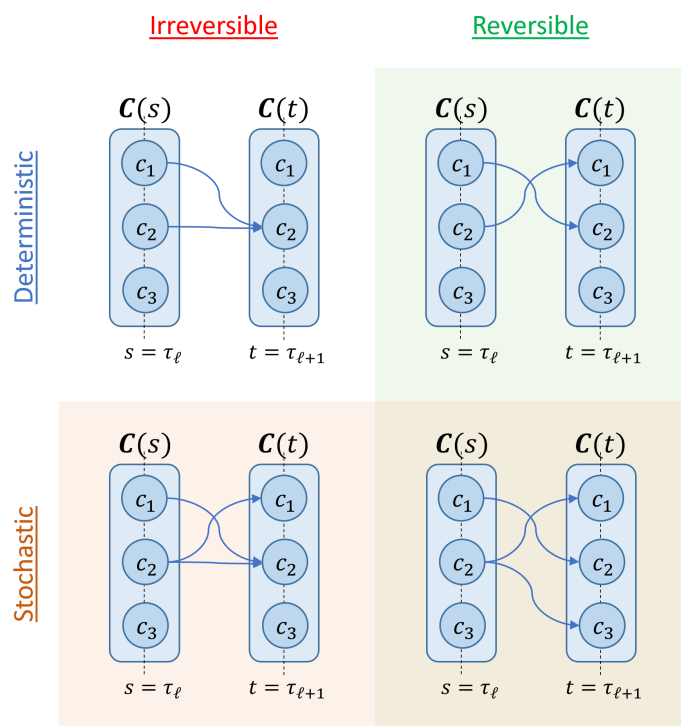


Figure 4. Illustration of different types of computational operations O_s^t on a set C of 3 computational states. Examples shown here are partial functions— $O(c_3)$ is not defined. At upper-left is a conventional (deterministic but non-reversible) computational operation which merges two initially distinct computational states. At upper-right is a deterministic, reversible operation which is injective (one-to-one) over the subset $A = \{c_1, c_2\}$ of initial states for which it is defined. At lower-right is a stochastic but reversible operation which does not merge any states, but splits the state c_2 (with some nonzero probability to transition to either c_1 or c_3). Finally, at lower-left is a stochastic, irreversible operation which includes both splits and merges.

2.1.4. Computational State Transitions

We can describe the computational operations from the previous subsection as a combination of various primitive *computational state transitions*, such as those illustrated in Figure 5 below. Other types of transitions may be described as combinations of these. For example, the lower-left operation in Figure 4 includes a transition of $\{c_1, c_2\}$ from time s to t that exhibits both splits and merges. However, as suggested by the arrows in the diagram, it could be decomposed into a sequence of a split of c_2 into two (unlabeled) states, followed by a merge of c_1 into one of those states.

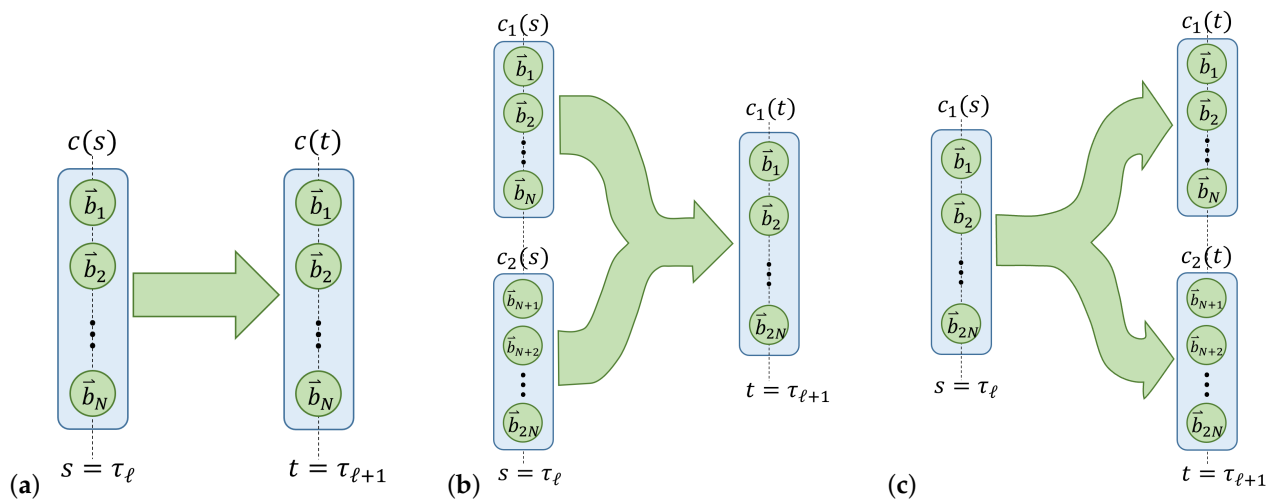


Figure 5. Illustration of different types of computational state transitions between two designated time points s and t , showing the protocomputational basis state sets B_c corresponding to computational states c . All of the individual basis states $\vec{b} = \vec{b}(\tau)$ are also implicitly time-dependent in general. Reversible computational operations include only one-to-one transitions such as (a). In such transitions, the initial and final basis state sets may be the same size N , even in a closed system. Irreversible computational operations include at least some examples of many-to-one transitions (merges), such as (b). In a closed system, the basis state set resulting from a merge must be (at least) as large as the sum of the merged sets, due to unitarity. Stochastic operations include one-to-many transitions (splits), such as (c). Each of the basis state sets resulting from a split may be smaller than the original set, even in a closed system, although their aggregate size must be at least as large.

2.1.5. Correspondence between Classical Operations and Quantum Evolution

In this subsection, we give a general theoretical picture regarding how a real (ergo, quantum-mechanical) physical process may effectively implement classical computational state transitions, and computational operations, such as described above.

2.1.5.1. Unitary Dynamics

As before, we focus our attention on a computational process taking place between two designated time points $s = \tau_\ell$ and $t = \tau_{\ell+1}$ (where $t > s$). Consider, now, the joint Hilbert space \mathcal{H}_Ω of the model universe (the environment together with the computer). Whatever is happening physically in the universe over the time interval $[s, t]$ (including the performance of the computational operation) will be encompassed, in a theoretical perspective assuming perfect knowledge of the universe’s dynamics, by the overall time evolution operator, which we will denote $\hat{U} = \hat{U}_s^t(\Omega)$, that applies between those times in \mathcal{H}_Ω . Formally, if we describe the initial quantum state of the model universe as a mixed state using an initial density matrix ρ_s , then the final density matrix ρ_t is given by

$$\rho_t = \hat{U}\rho_s\hat{U}^\dagger. \tag{4}$$

This overall time evolution process includes activities such as the dynamical details of the computation process itself, together with the incremental delivery of some needed free

energy from the power supply (e.g., battery) into that process, and the transport of some incremental amount of dissipated energy (waste heat) away from that process, or more precisely, the incremental progression of a continuous *flow* of waste heat that is propagating away from the computational mechanism, and out towards the environment \mathfrak{E} —since in general, the waste heat that resulted from prior operations will still be traveling outwards when subsequent operations occur. We call this picture the *open system* case.

Now, let us restrict attention temporarily to the subspace of $\mathcal{H}_{\mathfrak{U}}$ that is the Hilbert space $\mathcal{H}_{\mathfrak{S}}$ of the closed spacelike hypersurface (slice of spacetime volume) enclosing the computer system. Ignoring, for the moment, the flow of waste heat through the system’s bounding surface, let us pretend for a moment that the dynamics within the surface itself can also be described by a unitary time-evolution operator $\hat{U}_s^t(\mathfrak{S})$ over $\mathcal{H}_{\mathfrak{S}}$, the quantum subsystem contained within the boundary. We call this the *closed system* case.

Of course, the closed-system picture is a simplification, since in reality, no thermal isolation is perfect, and so there will also be interactions across the surface, to transport heat out. However, we expect that theoretical developments for the closed-system case can generally be preserved when re-expanding the model to include the outward thermal flow, since the net effect of that flow will just be to maintain a reasonable temperature inside the boundary, by exporting excess thermal entropy to the environment.

An easy way to see that this translation from the closed-system to the open-system case ought to work, in general, is simply to note that if the bounding surface of \mathfrak{S} is taken to be extremely remote to begin with, then there will be negligible *practical* difference between the open-system and closed-system cases. that is, a real computer with an internal power supply would operate just fine, at least for a while, even if enclosed in a very large, but finite, perfectly thermally insulated box.

Thus, henceforth in this subsection, we will take the time-evolution \hat{U}_s^t to be the one for the computer system \mathfrak{S} , in the closed-system picture, while remembering that we can revert to the open-system view when necessary.

Earlier, we noted that the protocomputational basis \mathcal{B} may, in general, be time-dependent, so that the two bases $\mathcal{B}(s)$ and $\mathcal{B}(t)$ may not correspond to exactly the same set of physical quantum states. However, the effect of any change in the protocomputational basis \mathcal{B} between times s and t can also just be represented as a unitary operator, which we denote $\hat{U}_{\mathcal{B}(s)}^{\mathcal{B}(t)}$. Then, we can define a suitably “basis-corrected” version of $\hat{U}_s^t(\mathfrak{S})$ as:

$$\hat{U}_s^t(\mathfrak{S}, \mathcal{B}) = \hat{U}_{\mathcal{B}(s)}^{\mathcal{B}(t)} \cdot \hat{U}_s^t(\mathfrak{S}). \tag{5}$$

2.1.5.2. Quantum Statistical Operating Contexts

Next, we need to define a *computational process* in a statistically-contextualized form. Earlier, we abstractly defined computational state transitions and computational operations, but this definition said nothing whatsoever about the statistics of the initial state (either computational or physical) before the operation was performed. We require a formalism for describing such information in order to speak meaningfully about the informational or thermodynamic effect of performing a computational operation within a particular, statistically-defined scenario. Note that the following presentation just generalizes the discussion of (classical) statistical operating contexts that can be found in, for example, [11], to a quantum context.

Since we want to produce a quantum-mechanical model of classical computation (including reversible computation), we require a quantum statistical picture. Thus, let us define ρ_s to be a mixed quantum state (i.e., a statistical mixture of orthogonal pure states, in some diagonal basis) that encompasses all of our uncertainty, as modelers, regarding what the initial quantum state of the physical computational system \mathfrak{S} is at time s , prior to performing the desired computational operation O_s^t .

We further require that ρ_s must have a *block-diagonal* structure in the initial proto-computational basis $\mathcal{B}(s)$, such that the blocks correspond to the partition $\{ \mathcal{B}_c \}$ of basis vectors corresponding to the (augmented) initial computational state set $\mathcal{C}_{\perp}(s)$. Stated

more formally, the density matrix representation of ρ_s in the $\mathcal{B}(s)$ basis must not include any nonzero, off-diagonal terms between basis states $\vec{b}_p, \vec{b}_q \in \mathcal{B}(s)$ such that $\vec{b}_p \in \mathbf{B}_i$ and $\vec{b}_q \in \mathbf{B}_j$ where $\mathbf{B}_i, \mathbf{B}_j$ are the subsets of $\mathcal{B}(s)$ corresponding to two distinct computational states $c_i, c_j \in \mathcal{C}_\perp(s)$, that is, with $i \neq j$. See Figure 6. This block-diagonal structure models our assumption, mentioned earlier, that a classical computer is highly *decoherent*; thus, there are no quantum coherences between the blocks corresponding to different computational states. (In the terminology coined by Zurek [43], the digital computational states would be considered natural *pointer states* of the computing apparatus.) However, note that it is permissible for coherences to exist *within* blocks. This is just another way of saying that the choice of protocomputational basis vectors is completely arbitrary within the subspace corresponding to each block; the sub-basis for any block can be freely rotated within its subspace, and we still will have a valid protocomputational basis for time s .

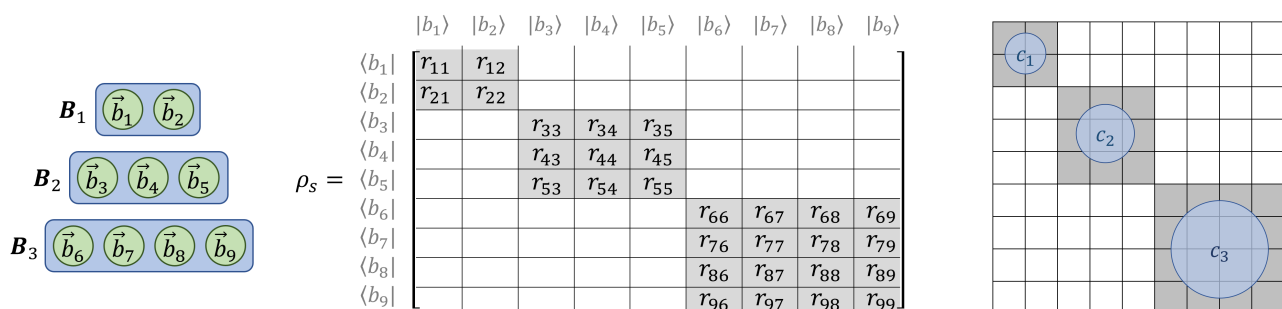


Figure 6. Block-diagonal density matrix for an initial quantum statistical operating context ρ_s . In this example, we imagine there are 3 computational states c_1, c_2, c_3 (and let $c_\perp = c_3$, say) with corresponding basis state sets $\mathbf{B}_1, \mathbf{B}_2, \mathbf{B}_3$ (left). At the center, we illustrate a corresponding block-diagonal *quantum statistical operating context* or initial density matrix ρ_s . Rows and columns are labeled in gray with the corresponding basis vectors; note that $\langle b_i | \rho_s | b_j \rangle = r_{ij}$. This is an Hermitian matrix, so $r_{ij} = r_{ji}^*$, also $\text{Tr}[\rho_s] = 1$. Matrix entries left blank are 0. On the right is a simplified depiction of ρ_s .

2.1.5.3. Quantum Contextualized Computations

Now that we have defined the quantum version of a statistical operating context, we can define what a “quantum-contextualized computation” means. This generalizes the discussion of (statistically contextualized) computations that can be found in [11].

A *quantum-contextualized computational process* or just *quantum-contextualized computation*, denoted $\mathcal{C}_s^t(O_s^t, \rho_s)$, refers to the act of carrying out a specified computational operation O_s^t from time s to $t > s$ within the computer system (\mathfrak{S}) in a quantum statistical operating context wherein the initial mixed state of \mathfrak{S} at time s is given by ρ_s ; where ρ_s meets the conditions (i.e., block-diagonal structure) described above, given the protocomputational basis $\mathcal{B}(s)$ and computational state set $\mathcal{C}(s)$. (Note that $\mathcal{B}(s)$ and $\mathcal{C}(s)$ are left implicit in the $\mathcal{C}_s^t(\cdot)$ notation for brevity.)

2.1.5.4. Implementation of Classical Computation by Unitary Dynamics

Given the above definitions, we can now formally define what it means for a system’s unitary dynamics to implement a given (classical) computation.

We say that the basis-adjusted time-evolution operator $\hat{U}_s^t(\mathfrak{S}, \mathcal{B})$ implements the quantum contextualized computational process $\mathcal{C}_s^t(O_s^t, \rho_s)$, written

$$\hat{U}_s^t(\mathfrak{S}, \mathcal{B}) \Vdash \mathcal{C}_s^t(O_s^t, \rho_s), \tag{6}$$

if and only if the final density matrix

$$\rho_t = \hat{U}_s^t(\mathfrak{S}, \mathcal{B}) \rho_s \hat{U}_s^t(\mathfrak{S}, \mathcal{B})^\dagger, \tag{7}$$

generated by applying $\hat{U}_s^t(\mathfrak{S}, \mathcal{B})$ to ρ_s has the property that, for any initial computational state $c_i(s) \in \mathcal{C}(s)$ that has nonzero probability under ρ_s , if we were to zero out all elements

of ρ_s outside of the rows/columns corresponding to $c_i(s)$'s basis set $\mathbf{B}_i(s)$ and renormalize, and then apply $\hat{U}_s^t(\mathfrak{S}, \mathbf{B})$ to this restricted ρ'_s , the resulting final mixed state ρ'_t would imply the same probability distribution $P_i(t)$ over final computational states in $\mathbf{C}(t)$ as is specified by applying the stochastic map O_s^t to the initial computational state, that is, $P_i = O_s^t(c_i(s))$.

One can see by inspection that this is a very straightforward and natural definition. Since, by assumption, the initial quantum statistical operating context ρ_s has no coherences between different initial computational states, it is impossible for the transition amplitudes from initial to final basis states to interfere with each other in ways that would disrupt the overall probability distribution over final computational states from what one would obtain by simply combining the results from treating the initial computational states separately.¹⁰

Note that the above definition *does not* by itself immediately require that the unitary evolution $\hat{U}_s^t(\mathfrak{S}, \mathbf{B})$ cannot introduce any *immediate* coherences between different computational states $c_i(t), c_j(t)$, where $i \neq j$, but, *this is not a problem*, since one of our background assumptions throughout this treatment is that the system will naturally decohere very quickly to a definite computational state, so, any off-diagonal matrix elements between different computational states that may arise will naturally decay by themselves very quickly. This can happen via the usual Zurek process [41], wherein the decoherent state variables entangle with nearby non-computational degrees of freedom, which then—at least, in the open-system version of this treatment—carry the associated quantum information out to the external thermal environment \mathfrak{E} . Once it is in that environment, taking the trace over the environment state to reflect our ignorance about the environment's detailed evolution then effectively erases the entanglement between the system \mathfrak{S} and the environment, and decays the coherences between the different naturally-stable “pointer states” of the computer. In Zurek's terms, the natural interaction between the computer system and its environment effectively “observes” the state of the system, and this effective measurement of the system by the environment collapses the system down to (what is then effectively just a classical statistical mixture of) the observably distinct classical computational states.

At this point, having described what it means, in quantum physical terms, to perform classical digital computational operations, our problem in building quantum physical models of reversible computing has now been reduced to:

1. Finding specific closed-system time-evolution unitaries $\hat{U}_s^t(\mathfrak{S}, \mathbf{B})$ that meet the above definition of the “implements” operator \Vdash for the case of desired *reversible* (and/or conditionally reversible) operations O_s^t in specific physical setups—and, it is easy to see that there's no essential loss of generality in starting with the closed-system case, since, for large enough systems, closed-system evolution should work just fine for a while, until the system runs out of effective free energy¹¹ or overheats.
2. Showing that the closed-system definition of $\hat{U}_s^t(\mathfrak{S})$ can then be extended appropriately to the open-system case where there may be a heat flow out from the system's bounding surface, for consistency with the existence of a global unitary evolution $\hat{U}_s^t(\mathfrak{U})$ for the model universe that includes the process of heat outflow to the environment—but this part is expected to be a relatively easy formal technicality. And finally:
3. Showing that some such unitaries can indeed be implemented via realistic, buildable physical computing mechanisms. Of these three steps, this one is expected to be the most difficult one to accomplish in practice.

However, the supposition that the above physical picture of classical reversible computing *can*, in fact, be realistically implemented is supported by the illustration of a number of existing and proposed examples of concrete physical implementation technologies which appear to accomplish this, which are briefly reviewed in Section 2.3 below.

¹⁰ As discussed in Section 2.1.5.2, each computational state corresponds to an orthogonal subspace; indeed, for *classical* computing, each computational state *must* be an orthogonal subspace, to ensure distinguishability of different computational states. Then, this statement is a direct consequence of the fact that the full space can be decomposed into the sum of projectors over all of its orthogonal subspaces.

¹¹ Technically speaking, free energy in closed systems is conserved, but if part of \mathfrak{S} treated thermally, there can be an increase in effective entropy.

First, we now review some relevant tools and methods from NEQT which can be used to flesh out the general theoretical framework presented above in more detail.

2.2. Tools and Methods from Non-Equilibrium Quantum Thermodynamics

In this section, we review some key theoretical tools and methods from non-equilibrium quantum thermodynamics (NEQT) that we believe will prove to be invaluable in the effort to arrive at a more complete understanding of the physics of reversible computing, and relate them to the more general picture presented above in Section 2.1.

2.2.1. Resource Theory of Quantum Thermodynamics

First, we review several theoretical tools relating to what is known as *the resource theory of quantum thermodynamics* (RTQT), in order to relate them to the broad framework presented above.

2.2.1.1. Stinespring Dilation Theorem and Thermomajorization

We briefly summarized the overall open quantum systems perspective in Section 2.1.1 earlier. The rules of quantum thermodynamics let us turn the broad intuitions summarized there into specific statements about the types of transformations allowable on the system \mathfrak{S} . The evolution of a general density matrix ρ is given by a *completely positive trace-preserving* (CPTP) map $\rho \mapsto \Lambda_t[\rho]$, also known as a *quantum channel* or (*quantum*) *dynamical map*. The map Λ_t maps the initial density matrix to a final density matrix. (Here, t represents the time interval from the initial time t_0 to the final time t_f .¹²

$\Lambda_t[\rho]$ represents the most generic type of transformation that we can apply to ρ . In general, the density matrices ρ and $\Lambda_t[\rho]$ are not required to be taken over the same Hilbert space. In our setup, however, we stipulate that the initial and final Hilbert spaces are the same (namely, $\mathcal{H}_{\mathfrak{S}}$). As the name “CPTP map” suggests, in order for the map $\rho \mapsto \Lambda_t[\rho]$ to satisfy the laws of quantum mechanics, we need Λ_t to preserve the trace of ρ , to preserve the positivity of ρ , and to preserve the positivity of ρ even when ρ is part of a larger system (which Λ_t acts on as a whole). Furthermore, we need Λ_t to be Hermitian, to be linear, and finite under the trace norm.

The *Stinespring dilation theorem* [44] provides a very natural representation of this channel. From this theorem, the action of $\Lambda_t[\rho]$ can always be represented by embedding ρ in a larger Hilbert space, where the dynamics corresponding to $\Lambda_t[\rho]$ is now unitary, and then tracing out the auxiliary part of the larger space. Then, we can express the evolution of an open quantum system \mathfrak{S} in terms of the unitary joint evolution of the system and the environment together, which together comprise the entire universe (i.e., $\mathfrak{U} = \mathfrak{S}\mathfrak{E}$). If \mathfrak{S} starts in the state $\rho_{\text{in},\mathfrak{S}}$ and \mathfrak{E} starts in the state $\rho_{\mathfrak{E}}$, the evolution appears as:

$$\begin{aligned} \rho_{\text{in},\mathfrak{S}} \mapsto \Lambda_t[\rho_{\text{in},\mathfrak{S}}] &:= \text{Tr}_{\mathfrak{F}} \left[\widehat{U}_{t,\mathfrak{S}\mathfrak{E}} (\rho_{\text{in},\mathfrak{S}} \otimes \rho_{\mathfrak{E}}) \widehat{U}_{t,\mathfrak{S}\mathfrak{E}}^\dagger \right] \\ &= \text{Tr}_{\mathfrak{F}} \left[e^{-i\widehat{H}(t_f - t_0)} (\rho_{\text{in},\mathfrak{S}} \otimes \rho_{\mathfrak{E}}) e^{i\widehat{H}(t_f - t_0)} \right]. \end{aligned} \quad (8)$$

As we noted earlier, the final state may not be in the same Hilbert space as the initial state; that is, Λ_t may not necessarily map \mathfrak{S} to itself. This is reflected in the fact that we take the final trace over $\mathfrak{F} \subseteq \mathfrak{U}$, where \mathfrak{F} may not necessarily be the same space as \mathfrak{E} . However, for our purposes, we will always have $\mathfrak{F} = \mathfrak{E}$.

In (8), we defined $\widehat{U}_{t,\mathfrak{S}\mathfrak{E}} := e^{-i\widehat{H}(t_f - t_0)}$ as the global unitary evolution operator over all of $\mathfrak{U} = \mathfrak{S}\mathfrak{E}$. Here, $\widehat{H} = \widehat{H}_{\mathfrak{S}} + \widehat{H}_{\mathfrak{E}} + \widehat{H}_{\mathfrak{S}\mathfrak{E}}$ is the global Hamiltonian over all of $\mathfrak{S}\mathfrak{E}$, divided into the Hamiltonian $\widehat{H}_{\mathfrak{S}}$ over \mathfrak{S} alone, the Hamiltonian $\widehat{H}_{\mathfrak{E}}$ over \mathfrak{E} alone, and the interaction Hamiltonian $\widehat{H}_{\mathfrak{S}\mathfrak{E}}$ between \mathfrak{S} and \mathfrak{E} . This representation forms the basis for

¹² Most generally, a t subscript for a time interval can be taken to specify an ordered pair (t_0, t_f) of start and end times: that is, the notation does not need to assume time translation invariance *a priori*; but, in the usual case when it is present, for t to specify just the time difference $d = t_f - t_0$ is sufficient.

both the existing NEQT results on Landauer’s principle, as well as the GKSL framework for examining open quantum systems. Beyond the rules of quantum mechanics, the only additional assumptions here are that \mathfrak{S} is coupled to some environment \mathfrak{E} which it jointly evolves with unitarily, and that the initial state of $\mathfrak{S}\mathfrak{E}$ can be factorized [45].

In terms of the dilation theorem, the set of transformations on $\rho_{\text{in},\mathfrak{S}}$ allowed by thermodynamics is simply the set of unitary transformations $\hat{U}_{t,\mathfrak{S}\mathfrak{E}}$ that preserves the total energy over all of $\mathfrak{S}\mathfrak{E}$. These transformations are explicitly described by the *resource theory of quantum thermodynamics* (RTQT) [23,24]. In general, quantum resource theories (QRT) provide a information-theoretic framework for describing all possible operations on a given state ρ , by describing the information cost of operations and states (in terms of new information we require about the system) [25]. In particular, QRTs describe the conditions on operations to act at no additional information cost and provide the conditions on the types of states of new systems that can be prepared at no additional information cost and appended to the overall system. These are respectively known as the *free operations* and *free states*. In addition to these, the quantum resource theory provides the conditions on transformations on ρ , known as the *state conversion conditions*. The nature of free operations, free states, and the conversion conditions depends on the specific resource theory.¹³

In RTQT, we start with the system Hamiltonian $\hat{H}_{\mathfrak{S}}$ and the (inverse) environment temperature $\beta = 1/T$. The thermal (Gibbs) states $\tau := e^{-\beta\hat{H}} / \text{Tr}[e^{-\beta\hat{H}}]$ are the maximum-entropy states, which must necessarily be preserved by energy-preserving unitary operations [46]. Thus, these are the free states of the environment. As such, if we examine a system using the dilation theorem, it takes no additional information to set the initial state of the environment \mathfrak{E} to be the thermal state $\tau_{\mathfrak{E}}$. (Conversely, selecting any other state *does* involve extra information not specified in the resource theory; namely, information about the distribution of states over \mathfrak{E} .) Setting $\mathfrak{F} = \mathfrak{E}$, this gives us a direct expression for the free operations, which are known as the *thermal operations*:

$$\rho_{\text{in},\mathfrak{S}} \mapsto \Xi_t[\rho_{\text{in},\mathfrak{S}}] := \text{Tr}_{\mathfrak{E}} \left[\hat{U}_{t,\mathfrak{S}\mathfrak{E}} (\rho_{\text{in},\mathfrak{S}} \otimes \tau_{\mathfrak{E}}) \hat{U}_{t,\mathfrak{S}\mathfrak{E}}^\dagger \right]. \tag{9}$$

The necessary conditions for these transformations to occur are called the *thermomajorization conditions*. When the commutator $[\Xi_t[\rho_{\text{in},\mathfrak{S}}], \hat{H}_{\mathfrak{S}}] = 0$ (i.e., when the final state of \mathfrak{S} has a definite energy value, as is the case for all of the systems we will be considering), these conditions are both necessary and sufficient.¹⁴

These conditions are defined in terms of the β -ordering of a state $\rho_{\mathfrak{S}}$, which has eigenvalues $\{p_i\}_{i=1}^n$ and corresponds to a Hamiltonian $\hat{H}_{\mathfrak{S}}$ (with $[\rho_{\mathfrak{S}}, \hat{H}_{\mathfrak{S}}] = 0$). The β -ordering $p^\downarrow = (p_1^\downarrow, \dots, p_n^\downarrow)$ is defined [47] as an ordering of the p_i that satisfies $p_i^\downarrow e^{-\beta E_i} \geq p_j^\downarrow e^{-\beta E_j}$ for all $i < j$, where E_i is the energy corresponding to p_i . (Thus, the β -ordering of the $\{p_i\}$ s is defined by decreasing values of $p_i^\downarrow e^{-\beta E_i}$.) From this ordering, we can define the *thermomajorization curve* as the curve defined by the points:¹⁵

$$\left\{ (0, 0), \left(e^{-\beta E_1}, p_1^\downarrow \right), \left(e^{-\beta E_1} + e^{-\beta E_2}, p_1^\downarrow + p_2^\downarrow \right), \dots, \left(\sum_{i=1}^n e^{-\beta E_i}, \sum_{i=1}^n p_i^\downarrow \right) \right\}. \tag{10}$$

¹³ As an example, the resource theory of asymmetry tells us the free operations and states of a system with an overall symmetry described by a compact Lie group G , with the free operations and conversion conditions given in terms of the unitary representations of G . Operations that are covariant with G and states that are invariant under G require no additional information beyond the group already specified; thus, these are respectively the free operations and free states of this resource theory. This example is expanded upon in [25], which also provides the illustrative example of the resource theory of bipartite entanglement.

¹⁴ The case of $[\Xi_t[\rho_{\text{in},\mathfrak{S}}], \hat{H}_{\mathfrak{S}}] \neq 0$ remains an open question [25]; fortunately, that is also beyond the scope of our model.

¹⁵ It is worth noting that the ordering of the $\{p_i^\downarrow\}$ s may *not* be unique when some of the $p_i^\downarrow e^{-\beta E_i}$ values are equal to each other, but even in this case the thermomajorization curve is *always* unique.

Then, finally, the thermal operation $\rho_{in,\mathfrak{S}} \mapsto \Xi_t[\rho_{in,\mathfrak{S}}]$ can occur if the thermomajorization curve of $\Xi_t[\rho_{in,\mathfrak{S}}]$ is below or equal to the thermomajorization curve of $\rho_{in,\mathfrak{S}}$ everywhere. Collectively, the thermal states, thermal operations, and thermomajorization conditions determine the complete set of states we can generate and transformations we can perform in quantum thermodynamics.

2.2.1.2. Catalytic Thermal Operations and Correlated Systems

The concept of a thermal operation can be extended to the case of a *catalytic thermal operation* (CTO), in which a component of the system is a so-called *catalyst* subsystem which cycles back to the initial state. This can be an appropriate model for certain types of subsystems in a computer—for example, a periodic clock signal, such as a resonant clock-power oscillator for an adiabatic circuit (see Section 2.3). Further, every digital data signal in a typical reversible logic technology (e.g., [48]) cycles from a standard “neutral” or no-information state to an information-bearing state, and then back to neutral; thus, every node in a typical reversible circuit effectively acts like a catalyst. (This will be discussed in more detail in Sections 3.4 and 4.2.)

We now present the most general type of CTOs explicitly, following the presentation in [15]. (Note that this examines a *more* general class of transformations than the ones traditionally examined in the “second laws of thermodynamics” framework; the relationship between this presentation and the “second laws” is discussed in Section 2.2.1.3). If we divide the overall system \mathfrak{S} into the subsystems \mathfrak{T} and \mathfrak{R} , the catalyst \mathfrak{R} is defined as a subsystem which is required within the overall dynamics of the system for the state transition $\rho_{in,\mathfrak{T}} \mapsto \Xi_t[\rho_{in,\mathfrak{T}}]$.¹⁶ If the state of the catalyst is given as $\sigma_{\mathfrak{R}}$, then the transition of the state ($\rho_{in,\mathfrak{T}} \otimes \sigma_{\mathfrak{R}}$) is given [15] by:

$$(\rho_{in,\mathfrak{T}} \otimes \sigma_{\mathfrak{R}}) \mapsto \zeta_{\mathfrak{T}\mathfrak{R}} := \text{Tr}_{\mathfrak{E}} \left[\hat{U}_{t,\mathfrak{T}\mathfrak{R}\mathfrak{E}} (\rho_{in,\mathfrak{T}} \otimes \sigma_{\mathfrak{R}} \otimes \tau_{\mathfrak{E}}) \hat{U}_{t,\mathfrak{T}\mathfrak{R}\mathfrak{E}}^\dagger \right]. \tag{11}$$

Here, we have $[\hat{U}_{t,\mathfrak{T}\mathfrak{R}\mathfrak{E}}, \hat{H}_{\mathfrak{T}\mathfrak{R}\mathfrak{E}}] = 0$, and $\text{Tr}_{\mathfrak{R}\mathfrak{E}}[\zeta_{\mathfrak{T}\mathfrak{R}}]$ is arbitrarily close to $\Xi_t[\rho_{in,\mathfrak{T}}]$ under the trace norm: for all $\epsilon \in \mathbb{R}^+$, there are allowed transformations with:

$$\|\text{Tr}_{\mathfrak{R}\mathfrak{E}}[\zeta_{\mathfrak{T}\mathfrak{R}}] - \Xi_t[\rho_{in,\mathfrak{T}}]\|_1 < \epsilon. \tag{12}$$

Meanwhile, the catalytic condition requires $\text{Tr}_{\mathfrak{T}\mathfrak{E}}[\zeta_{\mathfrak{T}\mathfrak{R}}] = \sigma_{\mathfrak{R}}$. This is the most general type of CTO, which can be realized if and only if the final Helmholtz free energy \mathcal{F} is less than or equal to the initial Helmholtz free energy. Out of equilibrium, the Helmholtz free energy of a state ρ in a system governed by a Hamiltonian \hat{H} is given by [49,50]:

$$\mathcal{F}(\rho) := \text{Tr} \left[\hat{H}\rho \right] - k_B T S(\rho). \tag{13}$$

Here, $S(\rho)$ is the Rényi-1 entropy (i.e., the von Neumann entropy) of ρ . In terms of \mathcal{F} , the condition we require for the CTO in (11) to be realizable is:¹⁷

$$\mathcal{F}(\text{Tr}_{\mathfrak{R}\mathfrak{E}}[\zeta_{\mathfrak{T}\mathfrak{R}}]) \leq \mathcal{F}(\rho_{in,\mathfrak{T}}). \tag{14}$$

Notably, these CTOs do not impose any additional major constraints on the shape of the correlations between \mathfrak{T} and \mathfrak{R} : for *any* $\delta \in \mathbb{R}^+$, there exists some \mathfrak{R} and $\zeta_{\mathfrak{T}\mathfrak{R}}$ such that $\hat{H}_{\mathfrak{R}} = 0$ and the *quantum mutual information* $I(\mathfrak{T} : \mathfrak{R})$ between \mathfrak{T} and \mathfrak{R} is bounded by δ :

$$I(\mathfrak{T} : \mathfrak{R}) := S(\zeta_{\mathfrak{T}\mathfrak{R}} \| \text{Tr}_{\mathfrak{T}\mathfrak{E}}[\zeta_{\mathfrak{T}\mathfrak{R}}] \otimes \text{Tr}_{\mathfrak{R}\mathfrak{E}}[\zeta_{\mathfrak{T}\mathfrak{R}}]) < \delta. \tag{15}$$

¹⁶ The reasons the catalyst is required are *not* explored in the CTO framework; rather, the CTO framework merely assumes that a catalyst is required.

¹⁷ In the case that $\mathcal{F}(\rho_{in,\mathfrak{T}}) - \mathcal{F}(\text{Tr}_{\mathfrak{R}\mathfrak{E}}[\zeta_{\mathfrak{T}\mathfrak{R}}]) = E$, the condition (14) can be extended by including an uncorrelated source \mathfrak{W} of free energy. In this case, if \mathfrak{W} has the states $\lambda_{\mathfrak{W}}$ and $\kappa_{\mathfrak{W}}$ with $\mathcal{F}(\kappa_{\mathfrak{W}}) - \mathcal{F}(\lambda_{\mathfrak{W}}) \geq E$, then we can realize the CTO $\rho_{in,\mathfrak{T}} \otimes \sigma_{\mathfrak{R}} \otimes \kappa_{\mathfrak{W}} \mapsto \zeta_{\mathfrak{T}\mathfrak{R}} \otimes \lambda_{\mathfrak{W}}$ when $\mathcal{F}(\rho_{in,\mathfrak{T}} \otimes \sigma_{\mathfrak{R}} \otimes \kappa_{\mathfrak{W}}) \geq \mathcal{F}(\zeta_{\mathfrak{T}\mathfrak{R}} \otimes \lambda_{\mathfrak{W}})$. Here, the condition (15) remains. Naturally, if we generalize to the case where correlations between $\mathfrak{T}\mathfrak{R}$ and \mathfrak{W} are permitted, we return to the CTO given in (11).

In practical terms, this means that we can achieve state transitions from $\rho_{\text{in},\mathfrak{T}}$ to $\text{Tr}_{\mathfrak{R}}[\zeta_{\mathfrak{T},\mathfrak{R}}]$ by engineering the catalyst and the CPTP map Ξ_t to minimize the correlation $I(\mathfrak{T} : \mathfrak{R})$. This process of *correlation engineering* [11,15] lies at the heart of reversible computing: By engineering interacting subsystems bearing computational degrees of freedom and the transformations Ξ_t applied on them, we can achieve the CTOs given in (11), with the net energy dissipation given by the free energy difference in (14).

2.2.1.3. Uncorrelated Catalytic Thermal Operations

The expressions in Section 2.2.1.2 may come as a surprise to those familiar with thermal operations and catalytic thermal operations. Conventionally, CTOs are defined [23] by the transformation:

$$(\rho_{\text{in},\mathfrak{T}} \otimes \sigma_{\mathfrak{R}}) \mapsto \text{Tr}_{\mathfrak{E}} \left[\hat{U}_{t,\mathfrak{T},\mathfrak{R},\mathfrak{E}} (\rho_{\text{in},\mathfrak{T}} \otimes \sigma_{\mathfrak{R}} \otimes \tau_{\mathfrak{E}}) \hat{U}_{t,\mathfrak{T},\mathfrak{R},\mathfrak{E}}^\dagger \right] = \Pi_t[\rho_{\text{in},\mathfrak{T}}] \otimes \sigma_{\mathfrak{R}}. \tag{16}$$

The data processing inequality (DPI) [51] can give us necessary conditions for these CTOs to be realized, which become necessary and sufficient when $[\Pi_t[\rho_{\text{in},\mathfrak{T}}] \otimes \sigma_{\mathfrak{R}}, \hat{H}_{\mathfrak{T},\mathfrak{R}}] = 0$ [23,52]. For any information distance function $f(\rho\|\sigma)$ of the density matrices ρ and σ , and for any CPTP map Λ_t , the DPI gives:

$$f(\rho\|\sigma) \geq f(\Lambda_t(\rho)\|\Lambda_t(\sigma)). \tag{17}$$

In other words, the DPI is a requirement that must be satisfied for all functions f for Λ_t to be a valid CPTP map. One family of such functions are the α -relative Rényi entropies (α -RRE), which are defined [46,53,54] as:

$$S_\alpha(\rho\|\sigma) := \begin{cases} \frac{\text{sgn } \alpha}{\alpha - 1} \ln \frac{\text{Tr}[\rho^\alpha \sigma^{1-\alpha}]}{\text{Tr } \rho} & \alpha \in (-1, 0) \cup (0, 1); \\ \frac{\text{sgn } \alpha}{\alpha - 1} \ln \left\{ \frac{\text{Tr} \left[\left(\sigma^{(1-\alpha)/2\alpha} \rho \sigma^{(1-\alpha)/2\alpha} \right)^\alpha \right]}{\text{Tr } \rho} \right\} & \alpha \in (-\infty, -1) \cup (1, \infty); \\ \text{Tr} [\rho (\ln \rho - \ln \sigma)] & \lim \alpha \rightarrow 1. \end{cases} \tag{18}$$

The $\alpha \rightarrow 1$ limit¹⁸ provides us with the familiar expression for the *quantum relative divergence* (QRD). As such, the DPI imposes the requirement that the CTO (16) must satisfy

$$S_\alpha(\rho_{\text{in},\mathfrak{T}} \otimes \sigma_{\mathfrak{R}}\|\tau_{\mathfrak{T},\mathfrak{R}}) \geq S_\alpha(\Pi_t[\rho_{\text{in},\mathfrak{T}}] \otimes \sigma_{\mathfrak{R}}\|\tau_{\mathfrak{T},\mathfrak{R}}), \tag{19}$$

for all α as a necessary condition.¹⁹ In the case that we have transition from the product state $\rho_{\text{in},\mathfrak{T}} \otimes \sigma_{\mathfrak{R}}$ to the product state $\Pi_t[\rho_{\text{in},\mathfrak{T}}] \otimes \sigma_{\mathfrak{R}}$, these are in fact *sufficient* conditions, beyond being simply necessary ones [57–59]. Thus, (19) tells us the constraints we need to satisfy the CTOs (16). The $S_\alpha(\rho\|\sigma)$ in turn define [46] the α -Helmholtz “free energies:”

$$\mathcal{F}_\alpha(\rho_{\text{in},\mathfrak{S}}) := -k_B T \ln Z + S_\alpha(\rho_{\text{in},\mathfrak{S}}\|\tau_{\mathfrak{S}}). \tag{20}$$

¹⁸ The $\alpha \rightarrow 1$ limit can be taken either with a monotonically increasing sequence from the $\alpha \in (0, 1)$ case, or with a monotonically decreasing sequence from the $\alpha \in (1, \infty)$ case. As with the QRDs, other familiar entropies are recovered as limiting cases of the α -RREs: the 0 -RRE $S_0(\rho\|\sigma) = \ln \text{Tr}_{\text{supp } \rho} [\sigma]$ is given by the $\alpha \searrow 0$ limit, the *max-RRE* $S_\infty(\rho\|\sigma) = \inf\{\lambda \in \mathbb{R} \mid \rho \leq e^\lambda \sigma\}$ is given by the $\alpha \nearrow \infty$ limit, and the $S_{-1}(\rho\|\sigma)$ and $S_{-\infty}(\rho\|\sigma)$ cases are given by interchanging ρ and σ in the arguments. It is also worth noting that the expression for $\alpha \in (1, \infty)$ in (18) is the conventional form for the α -RRE at these values of α ; this expression is called the *sandwiched RRE*. However, because in general ρ and σ do not commute with each other, there are an infinite number of ways to arrange powers of ρ and σ that satisfy the Rényi entropy axioms [55] and retrieve the appropriate limiting cases. These can all be expressed as a single two-parameter family of entropies [56], known as the α -z-RREs.

¹⁹ Indeed, the DPI tells us that a necessary condition for the CTO (16) to be valid is that $f(\rho_{\text{in},\mathfrak{T}} \otimes \sigma_{\mathfrak{R}}\|\tau_{\mathfrak{T},\mathfrak{R}}) \geq f(\Pi_t[\rho_{\text{in},\mathfrak{T}}] \otimes \sigma_{\mathfrak{R}}\|\tau_{\mathfrak{T},\mathfrak{R}})$ for all functions f of $\rho_{\text{in},\mathfrak{T}} \otimes \sigma_{\mathfrak{R}}$ and $\tau_{\mathfrak{T},\mathfrak{R}}$.

We can immediately recognize the $\alpha = 1$ case as equivalent to the expression (13). The CTOs defined in (16) are realized when we have

$$\mathcal{F}_\alpha(\Pi_t[\rho_{\text{in},\mathfrak{T}}]) \leq \mathcal{F}_\alpha(\rho_{\text{in},\mathfrak{T}}) \quad (21)$$

for all α . These conditions are known as the “second laws of thermodynamics” [46].

2.2.1.4. Correlated vs. Uncorrelated CTOs

The expressions for the α -RREs in Section 2.2.1.3 might initially be cause for some concern, since a CTO must satisfy (21) for *all* alpha to be a viable transition. This concern may escalate to alarm when we consider that the α -RREs (and thus the α -free energies) are monotone in α ; that is, $F_\beta(\rho) \leq F_\gamma(\rho)$ for all $\beta \leq \gamma$. Beyond the standard Helmholtz free energy (F_1), two notable cases are the *extractable work* F_0 and the *work of formation* F_∞ . As their names imply, these are respectively the amount of work we can extract from a given state and the amount of work it takes to form that same state. Since in most cases we have $F_0 < F_\infty$, it would appear that the energy difference $F_\infty - F_0$ is simply dissipated in the process of creating a state and then extracting work from that state. As a corollary, this would imply that our only hope for a viable reversible computing framework in this formulation is to find sets of states $\{\rho_i\}$ where the equality between $F_\alpha(\rho_i)$ is satisfied for all α , which may be a highly restrictive condition.

As discussed in [15], however, the “second laws of thermodynamics” (and these attendant issues) arise from an additional assumption about the shape of CTOs. Specifically, the CTOs (16) that give rise to the “second laws of thermodynamics” assume that the final state of \mathfrak{S} after the thermal operation is a product state of \mathfrak{T} and \mathfrak{R} , that is, that catalytic thermal operations transform the state $\rho_{\text{in},\mathfrak{T}} \otimes \sigma_{\mathfrak{R}}$ to the state $\Pi_t[\rho_{\text{in},\mathfrak{T}}] \otimes \sigma_{\mathfrak{R}}$. However, by *definition* of the catalytic thermal operation, we needed the presence of $\sigma_{\mathfrak{R}}$ to induce the transformation to begin with. Thus, the CTO on $\rho_{\text{in},\mathfrak{T}} \otimes \sigma_{\mathfrak{R}}$ *necessitates* an increase in the QMI between \mathfrak{T} and \mathfrak{R} , specifically given by (15). Indeed, as proven in [15], this mutual information can be made to be infinitesimally small, but *cannot* be zero. Thus, the CTO in (16), in which we demand that the final state of \mathfrak{S} be in the product state $\Pi_t[\rho_{\text{in},\mathfrak{T}}] \otimes \sigma_{\mathfrak{R}}$, can be thought of as performing the general CTO (11) and then *ejecting* the QMI (15). A direct consequence of this is that, as proven in [15], in the general CTO (11) where we permit correlations to develop between the system and catalyst, the ($\alpha = 1$) Helmholtz free energy *uniquely* specifies the condition required for the transition to take place.²⁰

Consequently, if we seek to develop a framework for computing which reduces energy dissipation by avoiding the energy cost of expelling the built up QMI, our computing operations must follow the CTO expression given in (11). Since reversible computing is precisely this framework, (11) provides an explicit expression for the shape of reversible computing operations in terms of CTOs. As a trade-off, we achieve these operations via a buildup of QMI (15), which can be made arbitrarily small but cannot be precisely zero. In the framework of reversible computing, this is an acceptable (indeed, *preferred*) trade-off to make.²¹

2.2.2. Quantum Mechanical Models of the Landauer Bound

The general CTOs given in [15] and discussed in Section 2.2.1.2 further give us a conceptual framework for understanding the nonequilibrium Landauer bound [13] and the difference between conditional and unconditional Landauer state reset [61]. First, we briefly review the Landauer principle, following the excellent presentation found in [61],

²⁰ Thus, [15] verifies the conjecture first provided in [60]: in the words of [60] directly, “the [$\alpha = 1$] Helmholtz] free energy is singled out as a measure of athermality.”

²¹ As pointed out in [15], the more general CTOs (11) are not necessarily an *improved* form of the CTOs (16), but rather simply offer a different setup. In lieu of the unavoidable “free energy” differences $F_\infty - F_0$, we have accepted the unavoidable buildup of QMI as a trade-off. For our purposes here, building up QMI and engineering the system and CTO to minimize the QMI and the difference $\|\text{Tr}_{\mathfrak{R}}[\zeta_{\mathfrak{T}\mathfrak{R}}] - \Xi_t[\rho_{\text{in},\mathfrak{T}}]\|$ is preferred, but the optimal type of CTO will in general be a function of the type of process we are interested in.

before connecting these to the nonequilibrium Landauer bound and the general CTOs found in [13,15].

2.2.2.1. Conditional vs. Unconditional Landauer Erasure

Here, and in much of what follows, we restrict our attention to a subsystem $\mathfrak{S}' = \mathfrak{M}$ of the entire computer system \mathfrak{S} that plays the role of passively registering data; we can generically call such a subsystem a “memory,” without implying any particular architectural structure (i.e., it could be any information-bearing set of signals in the machine). This is to be distinguished from other components that actively manipulate the state of the machine and control the timing of operations, which we will assume are separated into another subsystem \mathfrak{Z} which will usually be left implicit. See Figure 7a. However, note that even \mathfrak{M} , as a physical system, can still be separated into computational and non-computational subsystems as in Figure 3, and thus (assuming, as usual, no coherences between digital states) still has states in block-diagonal form as in Figure 6. Each computational state $c_i \in \mathcal{C}$ of \mathfrak{M} thus has a unique corresponding representation $\rho_{i,\mathfrak{M}}$ as a density matrix if we assume minimal information about the non-computational part of the state, that is, taking it to have maximum entropy, given the specifications as to what constitutes the set B_i of microstates of \mathfrak{M} validly representing the given computational state c_i in a given technological scenario. See Figure 7b. Note that the following discussion blurs the distinction between these density matrices and the abstract states c_i that they represent, and calls them “computational states” even though, in the density matrix form, they are also manifestly *physical* entities.

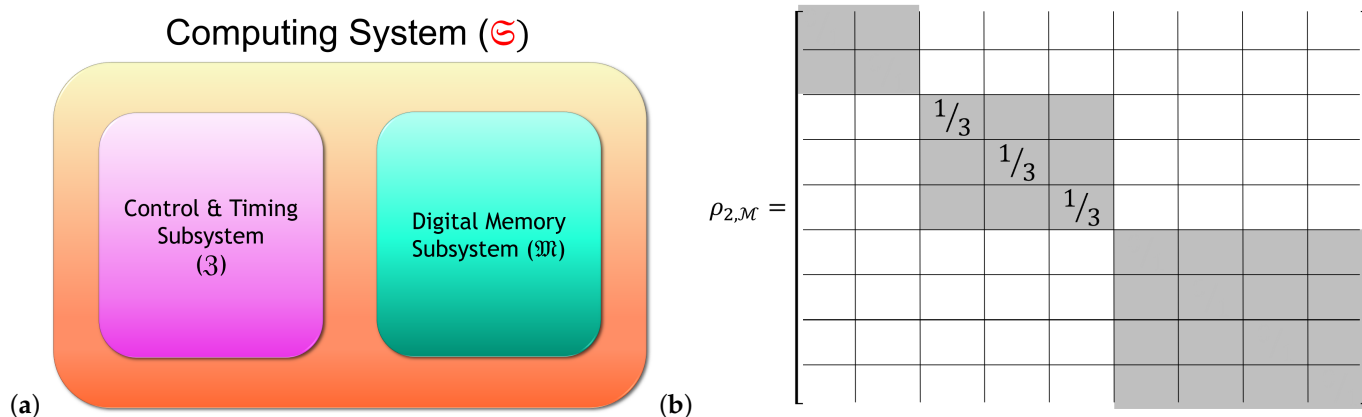


Figure 7. Breakdown of subsystems for purposes of Section 2.2.2.1, and so forth. (a) For our purposes in this section, and in much of what follows, we focus our attention on a subsystem $\mathfrak{S}' = \mathfrak{M}$ of the entire physical computer system (a “memory”) that exists for the purpose of passively registering some computational data of interest, but does not include any active mechanisms for controlling the timing and performance of state transitions. (b) An example of a density matrix representation of a computational state of \mathfrak{M} in the block-diagonal picture from Figure 6. In this example, the non-computational subsystem of \mathfrak{M} is assumed to be in a maximum-entropy mixed state conditioned on the computational state being $c = c_2$.

Now, for a subsystem $\mathfrak{S}' = \mathfrak{M}$ carrying some computational degrees of freedom, the Landauer state reset process, following [61], is the process by which the state $\rho_{\ell,\mathfrak{S}'}$ is set to some standard reset state $\rho_{r,\mathfrak{S}'}$. In a practical implementation of general computational operations on the system \mathfrak{S}' which bears computational degrees of freedom, the reset state is important for describing the operation of typical reversible computing systems such as described in Section 2.3. The reset state is a standard, known reference state; when we perform operations on the system which correspond to operations, we typically transform the system from $\rho_{r,\mathfrak{S}'}$ in a known way such that the operations we perform on $\rho_{r,\mathfrak{S}'}$ correspond to sensible computational operations. The end result of this series of operations will be the final computational state, which we then typically need to reset to the standard state in order to perform a new set of operations.

As discussed in Section 2.1, in general, we expect the computer to have a possibly very large, but finite number N of total possible computational states $\{\rho_{\ell, \mathcal{S}'}\}_{\ell=1}^N$ that it could be in at any time. Then, the reset process we are interested in is the set of the CPTP maps $\{\rho_{\ell, \mathcal{S}'}\} \mapsto \rho_{r, \mathcal{S}'}$ for all $\ell \in \{1, \dots, N\}$. We can start by taking these operations to be thermal operations of the form (9):

$$\begin{aligned} \rho_{\ell, \mathcal{S}'} \mapsto \Xi_{\ell, t}[\rho_{\ell, \mathcal{S}'}] &:= \text{Tr}_{\mathcal{E}} \left[\widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}} (\rho_{\ell, \mathcal{S}'} \otimes \tau_{\mathcal{E}}) \widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}}^\dagger \right] \\ &= \text{Tr}_{\mathcal{E}} [\rho_{\text{fc}, \mathcal{U}}] = \rho_{r, \mathcal{S}'} \end{aligned} \tag{22}$$

(As before, $\tau_{\mathcal{X}}$ denotes the thermal/Gibbs state in the system/subsystem \mathcal{X} .) Here, we have defined $\rho_{\text{fc}, \mathcal{U}}$ as the final *global* state of the entire universe $\mathcal{U} = \mathcal{S}' \mathcal{E}$ (ignoring \mathcal{J} for now) following application of the $\widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}}$ evolution to $(\rho_{\ell, \mathcal{S}'} \otimes \tau_{\mathcal{E}})$:

$$\rho_{\text{fc}, \mathcal{U}} := \widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}} (\rho_{\ell, \mathcal{S}'} \otimes \tau_{\mathcal{E}}) \widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}}^\dagger \tag{23}$$

Crucially, the effect of the unitary evolution operators $\widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}}$ over all of \mathcal{U} is to transform the set of N initial states into the *same* final state over all of \mathcal{U} , given by $\rho_{\text{fc}, \mathcal{U}}$. The overall unitary evolution operators $\widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}}$ are given by:²²

$$\begin{aligned} \widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}} &= \mathcal{T} \exp \left\{ -\frac{i}{\hbar} \int_{t_0}^{t_f} dt' \widehat{H} \right\} \\ &= \mathcal{T} \exp \left\{ -\frac{i}{\hbar} \int_{t_0}^{t_f} dt' \left[\widehat{H}_{\mathcal{S}'} + \widehat{H}_{\mathcal{E}} + \widehat{H}_{\mathcal{I}, \mathcal{S}' \mathcal{E}} + \widehat{V}_{\ell r, \mathcal{S}'} \right] \right\} \end{aligned} \tag{24}$$

Here, in addition to the terms contributing to \widehat{H} we outlined in (8), we also explicitly pulled out the reset Hamiltonian $\widehat{V}_{\ell r, \mathcal{S}'}$. Note that the reset Hamiltonian is applied *solely* to \mathcal{S}' .

An extremely important feature of the set $\{\widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}}\}$ of unitary operators is that we have a *unique* operator for each state $\rho_{\ell, \mathcal{S}'}$, and that the distinctiveness of each of these operators comes *solely* from the fact that the reset Hamiltonian $\widehat{V}_{\ell r, \mathcal{S}'}$ is individualized for each $\rho_{\ell, \mathcal{S}'}$. Each of these operators gives us a distinct CPTP map $\Xi_{\ell, t}(\rho_{\ell, \mathcal{S}'})$. As a result, the expressions (22) and (24) correspond to *conditional Landauer reset*; that is, the process of resetting the computational state $\rho_{\ell, \mathcal{S}'}$ to the standard state $\rho_{r, \mathcal{S}'}$ where the reset protocol $\widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}}$ is conditioned on the specific state $\rho_{\ell, \mathcal{S}'}$. In other words, the process of conditional Landauer erasure involves selecting $\widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}}$ (and, even more specifically, selecting $\widehat{V}_{\ell r, \mathcal{S}'}$) for each $\rho_{\ell, \mathcal{S}'}$ such that the final state $\rho_{\text{fc}, \mathcal{U}}$ is the same for *any* initial state $\rho_{\ell, \mathcal{S}'} \otimes \tau_{\mathcal{E}}$ we choose.

A central quantity of interest in the Landauer reset process is the lower bound on the amount of energy transfer (a.k.a. the *dissipation*) from the system to the environment. This can be calculated by examining the change in the environment energy $\Delta \langle E_{\ell, \mathcal{E}} \rangle$ during the evolution (22). In terms of this evolution, the final state of the environment is given by:

$$\rho_{\text{fc}, \mathcal{E}} = \text{Tr}_{\mathcal{S}'} [\rho_{\text{fc}, \mathcal{U}}] = \text{Tr}_{\mathcal{S}'} \left[\widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}} (\rho_{\ell, \mathcal{S}'} \otimes \tau_{\mathcal{E}}) \widehat{U}_{\ell, t, \mathcal{S}' \mathcal{E}}^\dagger \right] \tag{25}$$

²² Note that the exponentials in (24) are implicitly required to be time-ordered if the potential $\widehat{V}_{\ell r, \mathcal{S}'}$ is time-dependent, as is the general case.

Because $\rho_{f,c,\mathfrak{E}}$ is the same for all ℓ , we can directly examine the energy increase of the environment as a result of conditional Landauer reset protocol applied to any of the initial states:

$$\begin{aligned} \Delta \langle E_{\ell,\mathfrak{E}} \rangle_c &= \text{Tr} [\rho_{f,c,\mathfrak{E}} \hat{H}_{\mathfrak{E}}] - \text{Tr} [\tau_{\mathfrak{E}} \hat{H}_{\mathfrak{E}}] \\ &= \text{Tr}_{\mathfrak{E}} \left[\text{Tr}_{\mathfrak{S}'} \left[\hat{U}_{\ell,t,\mathfrak{S}'\mathfrak{E}} (\rho_{\ell,\mathfrak{S}'} \otimes \tau_{\mathfrak{E}}) \hat{U}_{\ell,t,\mathfrak{S}'\mathfrak{E}}^\dagger \right] \hat{H}_{\mathfrak{E}} \right] - \text{Tr} [\tau_{\mathfrak{E}} \hat{H}_{\mathfrak{E}}]. \end{aligned} \tag{26}$$

(Here, the subscript c indicates that this is specifically for the conditional Landauer reset.) For a pair of interacting systems a and b in which b is initially in a thermal state, we can straightforwardly derive the inequality $\Delta(S_b - \beta U_b) \leq 0$ (where $U_b := \text{Tr} [\rho_b \hat{H}_b]$) from the basic definition of entropy and its convexity property [62]. (This is sometimes referred to in the literature as *Partovi's inequality*.) For this system, this gives $\Delta \langle E_{\mathfrak{E}} \rangle \geq -k_B T \ln 2 \Delta S_{\mathfrak{E}}$. Combining this inequality with the strong subadditivity of the von Neumann entropy, $\Delta \langle E_{\ell,\mathfrak{E}} \rangle_c$ has a lower bound given by [61]:

$$\Delta \langle E_{\ell,\mathfrak{E}} \rangle_c \geq -k_B T \ln 2 \Delta S_{\ell,\mathfrak{S}'} := -k_B T \ln 2 [S(\Xi_{\ell,t}(\rho_{\ell,\mathfrak{S}'})) - S(\rho_{\ell,\mathfrak{S}'})]. \tag{27}$$

Here, $\Delta S_{\ell,\mathfrak{S}'}$ is the change in the von Neumann entropy between the initial and final states of \mathfrak{S}' . Thus, when the reset protocol is given as in (22), the sole contribution to the bound on dissipation into the environment is given by the change in entropy in \mathfrak{S}' induced by the overall unitary evolution $\hat{U}_{\ell,t,\mathfrak{S}'\mathfrak{E}}$.²³ Notably, when the initial states and final state have the same von Neumann entropies, the expression $\Delta S_{\ell,\mathfrak{S}'}$ is zero, and thus the lower bound on dissipation is zero in this case.

By contrast to the conditional Landauer reset, we can also define the *unconditional* Landauer reset protocol, in which transitions from each of the states $\{\rho_{\ell,\mathfrak{S}'}\}$ to the reset state $\rho_{r,\mathfrak{S}'}$ are achieved by applying a single, standard potential $\hat{V}_{r,\mathfrak{S}'}$ to *any* of the states $\rho_{\text{in},\mathfrak{S}'}$ that \mathfrak{S}' may be in. Thus, in lieu of the set of N unitary operators given in (22), we have a single unitary operator for all of the states defined by:

$$\begin{aligned} \hat{U}_{u,\mathfrak{S}'\mathfrak{E}} &= \mathcal{T} \exp \left\{ -\frac{i}{\hbar} \int_{t_0}^{t_f} dt' \hat{H} \right\} \\ &= \mathcal{T} \exp \left\{ -\frac{i}{\hbar} \int_{t_0}^{t_f} dt' \left[\hat{H}_{\mathfrak{S}'} + \hat{H}_{\mathfrak{E}} + \hat{H}_{\text{I},\mathfrak{S}'\mathfrak{E}} + \hat{V}_{r,\mathfrak{S}'} \right] \right\}. \end{aligned} \tag{28}$$

The corresponding set of thermal operations in this case are given by:

$$\rho_{\ell,\mathfrak{S}'} \mapsto \Xi[\rho_{\ell,\mathfrak{S}'}] := \text{Tr}_{\mathfrak{E}} \left[\hat{U}_{u,\mathfrak{S}'\mathfrak{E}} (\rho_{\ell,\mathfrak{S}'} \otimes \tau_{\mathfrak{E}}) \hat{U}_{u,\mathfrak{S}'\mathfrak{E}}^\dagger \right] = \rho_{r,\mathfrak{S}'}. \tag{29}$$

The set of evolutions given in (29) provide a sharp contrast with those given in (22). In (22), we chose $\hat{V}_{r,\mathfrak{S}'}$ such that $\hat{U}_{\ell,t,\mathfrak{S}'\mathfrak{E}}$ mapped every $\rho_{\ell,\mathfrak{S}'} \otimes \tau_{\mathfrak{E}}$ to the same final state $\rho_{f,\mathfrak{M}}$. By contrast, in (29) we have only a single unitary operator for every possible state under consideration. As a consequence, $\hat{U}_{u,\mathfrak{S}'\mathfrak{E}}$ maps each $\rho_{\ell,\mathfrak{S}'} \otimes \tau_{\mathfrak{E}}$ to a *different* global final state:

$$\left\{ \rho_{\ell,f,\mathfrak{M}} \right\}_{\ell=1}^N := \left\{ \hat{U}_{u,\mathfrak{S}'\mathfrak{E}} (\rho_{\ell,\mathfrak{S}'} \otimes \tau_{\mathfrak{E}}) \hat{U}_{u,\mathfrak{S}'\mathfrak{E}}^\dagger \right\}_{\ell=1}^N. \tag{30}$$

In other words, for each $\rho_{\ell,\mathfrak{S}'}$, the result of the evolution given by $\hat{U}_{u,\mathfrak{S}'\mathfrak{E}}$ is to produce a *distinct* final state over all of \mathfrak{M} . The only constraints on the evolutions in (29) (and thus, on the states $\rho_{\ell,f,\mathfrak{M}}$) beyond the laws of quantum mechanics and quantum thermodynamics

²³ As a reminder, the unitary dynamics over $\mathfrak{S}'\mathfrak{E}$ by construction cannot increase the entropy over $\mathfrak{S}'\mathfrak{E}$ as a whole [41,63]. Since we do not *a priori* assume that $\hat{U}_{\ell,t,\mathfrak{S}'\mathfrak{E}}$ maps $\mathfrak{S}'\mathfrak{E}$ product states to $\mathfrak{S}'\mathfrak{E}$ product states, however, this can still change the subsystem entropy of \mathfrak{S}' .

is that the final *subsystem* state of \mathfrak{S}' must be the reset state: we require $\text{Tr}_{\mathfrak{E}}[\rho_{\ell, f, \mathfrak{M}}] = \rho_{r, \mathfrak{S}'}$ for all ℓ .

Because each of the final global states is different, the energy increase of the environment (calculated as in (26) and (27)) will be a distinct expression for each initial state. However, we can collect these expressions together by examining the *average* energy increase of the conditional and unconditional Landauer resets, over a collection of reset operations performed over a set of individual states. If the states $\rho_{\ell, \mathfrak{S}'}$ appear in our collection a fraction of p_{ℓ} times, then the average energy increase of the environment will be given by:²⁴

$$\langle \Delta \langle E_{\ell, \mathfrak{E}} \rangle_c \rangle = \text{Tr}_{\mathfrak{E}} \left[\text{Tr}_{\mathfrak{S}'} \left[\sum_{\ell=1}^N p_{\ell} \hat{U}_{\ell, t, \mathfrak{S}' \mathfrak{E}} (\rho_{\ell, \mathfrak{S}'} \otimes \tau_{\mathfrak{E}}) \hat{U}_{\ell, t, \mathfrak{S}' \mathfrak{E}}^{\dagger} \hat{H}_{\mathfrak{E}} \right] - \text{Tr} [\tau_{\mathfrak{E}} \hat{H}_{\mathfrak{E}}] \right]; \quad (31)$$

$$\langle \Delta \langle E_{\ell, \mathfrak{E}} \rangle_c \rangle \geq -k_B T \ln 2 \sum_{\ell=1}^N p_{\ell} \Delta S_{\ell, \mathfrak{S}'}. \quad (32)$$

We can compare these expressions to the average energy increase and average energy bound of the unconditional Landauer reset protocol across all of the $\{\rho_{\ell, \mathfrak{S}'}\}$ s. Since convex linear combinations of density matrices form another density matrix, we can express the weighted sum of the $\{\rho_{\ell, \mathfrak{S}'}\}$ s as a new fiducial density matrix $\rho_{\text{in}, \mathfrak{S}'}$:

$$\rho_{\text{in}, \mathfrak{S}'} := \sum_{\ell=1}^N p_{\ell} \rho_{\ell, \mathfrak{S}'}. \quad (33)$$

Then, the average energy increase of the unconditional Landauer reset protocol corresponds [61] to the average energy increase of $\rho_{\text{in}, \mathfrak{S}'}$:

$$\begin{aligned} \langle \Delta \langle E_{\ell, \mathfrak{E}} \rangle_u \rangle &= \text{Tr}_{\mathfrak{E}} \left[\text{Tr}_{\mathfrak{S}'} \left[\hat{U}_{t, \mathfrak{S}' \mathfrak{E}} \left(\sum_{\ell=1}^N p_{\ell} \rho_{\ell, \mathfrak{S}'} \otimes \tau_{\mathfrak{E}} \right) \hat{U}_{t, \mathfrak{S}' \mathfrak{E}}^{\dagger} \hat{H}_{\mathfrak{E}} \right] - \text{Tr} [\tau_{\mathfrak{E}} \hat{H}_{\mathfrak{E}}] \right] \\ &= \text{Tr}_{\mathfrak{E}} \left[\text{Tr}_{\mathfrak{S}'} \left[\hat{U}_{t, \mathfrak{S}' \mathfrak{E}} (\rho_{\text{in}, \mathfrak{S}'} \otimes \tau_{\mathfrak{E}}) \hat{U}_{t, \mathfrak{S}' \mathfrak{E}}^{\dagger} \hat{H}_{\mathfrak{E}} \right] - \text{Tr} [\tau_{\mathfrak{E}} \hat{H}_{\mathfrak{E}}] \right]. \end{aligned} \quad (34)$$

(In this expression, since $\hat{U}_{t, \mathfrak{S}' \mathfrak{E}}$ is independent of ℓ , we were able to move the sum inside the expression).

As with (27), the strong subadditivity of the von Neumann entropy and Partovi’s inequality gives [61] a lower bound on $\langle \Delta \langle E_{\ell, \mathfrak{E}} \rangle_u \rangle$ in terms of the entropy:

$$\begin{aligned} \langle \Delta \langle E_{\ell, \mathfrak{E}} \rangle_u \rangle &\geq -k_B T \ln 2 \sum_{\ell=1}^N p_{\ell} (\Delta S_{\ell, \mathfrak{S}'} + \log_2 p_{\ell}); \\ \langle \Delta \langle E_{\ell, \mathfrak{E}} \rangle_u \rangle &\geq -k_B T \ln 2 \left(\sum_{\ell=1}^N p_{\ell} \Delta S_{\ell, \mathfrak{S}'} - \Delta I_{\text{er}, \mathfrak{S}'} \right). \end{aligned} \quad (35)$$

Here, we recognize $\sum_{\ell=1}^N p_{\ell} \log_2 p_{\ell} =: -H(\{p_{\ell}\}) = \Delta I_{\text{er}, \mathfrak{S}'}$ as the Shannon entropy of the distribution $\{p_{\ell}\}$, which is equivalent to the information quantity transferred from \mathfrak{S}' to \mathfrak{E} . As before, when the initial states and final state have the same von Neumann entropies, the expression $\Delta S_{\ell, \mathfrak{S}'}$ is zero. In this case, the lower bound on the unconditional Landauer reset protocol is given *entirely* by the amount of information $\Delta I_{\text{er}, \mathfrak{S}'}$ transferred from \mathfrak{S}' to \mathfrak{E} .

The only difference between the evolutions (24) and (28) is whether or not the reset potential \hat{V} is conditioned on the initial state of \mathfrak{S}' . This seems to indicate that $\Delta I_{\text{er}, \mathfrak{S}'}$

²⁴ We can alternately think [61] of this as a series of conditional or unconditional Landauer resets, respectively, over a single joint system $\mathfrak{S}' \mathfrak{E}$, with the population fractions representing the number of times $\mathfrak{S}' \mathfrak{E}$ is set to that state.

contains within it some correlated information (i.e., QMI) between \mathfrak{S}' and whatever implemented the potential \widehat{V} . In fact, as rigorously proven in [64], the *entirely* of $\Delta I_{er,\mathfrak{S}'}$ is the QMI that arises specifically from the process of conditioning (or not conditioning) \widehat{V} on the initial state of \mathfrak{S}' . The correlated nature of $\Delta I_{er,\mathfrak{S}'}$ plays a central role in understanding the conditional and unconditional Landauer bounds. Likewise, the details of where the reset potentials $\widehat{V}_{lr,\mathfrak{S}'}$ and $\widehat{V}_{r,\mathfrak{S}'}$ come from will play a key role in understanding the distinction between these two. These issues will be discussed in detail in Section 3.4, as we tie in this model to the CTO framework.

2.2.2.2. Nonequilibrium Landauer Bound

We can understand the expressions in Section 2.2.2.1 very straightforwardly from the NEQT point of view, both from the point of view of quantum thermodynamic fluctuation relations [13,14] and from the point of view of the general CTOs discussed in the previous section. We start with a general CPTP map given in terms of the dilation theorem (8) with $\mathfrak{F} = \mathfrak{E}$ and a general environment state $\rho_{\mathfrak{E}}$; that is, a CPTP map given by (henceforth just writing \mathfrak{S} for \mathfrak{S}'):

$$\rho_{in,\mathfrak{S}} \mapsto \Lambda_t[\rho_{in,\mathfrak{S}}] = \text{Tr}_{\mathfrak{E}} \left[\widehat{U}_{t,\mathfrak{S}\mathfrak{E}} (\rho_{in,\mathfrak{S}} \otimes \rho_{\mathfrak{E}}) \widehat{U}_{t,\mathfrak{S}\mathfrak{E}}^\dagger \right]. \tag{36}$$

If we label the eigenstates of $\rho_{\mathfrak{E}}$ as $|e_a\rangle$, the eigenvalues of $\rho_{\mathfrak{E}}$ as e_a , and perform the partial trace over the basis $|v_b\rangle$ of \mathfrak{E} , then the expression of $\Lambda_t[\rho_{in,\mathfrak{S}}]$ expands to give:

$$\Lambda_t[\rho_{in,\mathfrak{S}}] = \left(\mathbb{1}_{\mathfrak{S}} \otimes \sum_b \langle v_b | \right) \widehat{U}_{t,\mathfrak{S}\mathfrak{E}} \left(\rho_{in,\mathfrak{S}} \otimes \sum_a e_a |e_a\rangle \langle e_a| \right) \widehat{U}_{t,\mathfrak{S}\mathfrak{E}}^\dagger \left(\mathbb{1}_{\mathfrak{S}} \otimes \sum_b |v_b\rangle \right). \tag{37}$$

The distributivity of the tensor product allows us to write this expression solely in terms of operators on \mathfrak{S} . This defines the *system Kraus operators* (usually simply the *Kraus operators* [65]) as:

$$\widehat{M}_{ab} := \sum_{a,b} \sqrt{e_a} \langle v_b | \widehat{U}_{t,\mathfrak{S}\mathfrak{E}} | e_a \rangle. \tag{38}$$

It is worth noting that the Kraus operators are *dependent* on the global operator $\widehat{U}_{t,\mathfrak{S}\mathfrak{E}}$ and the environment expressions $|e_a\rangle$, e_a , and $|v_b\rangle$, but as operators themselves solely map density matrices over $\mathcal{H}_{\mathfrak{S}}$ to $\mathcal{H}_{\mathfrak{S}}$. In other words, even though we have \widehat{M}_{ab} dependent on quantities outside of \mathfrak{S} , nevertheless we have $\widehat{M}_{ab} \in \text{Aut}(\mathcal{D}(\mathcal{H}_{\mathfrak{S}}))$ when considering it as an operator. Note that a given set $\{\widehat{M}_{ab}\}$ of Kraus operators is emphatically not unique: any unitary rotation of the basis $|v_b\rangle$ defines a new set of Kraus operators.

Any given set of Kraus operators satisfies the completeness relation:

$$\sum_{a,b} \widehat{M}_{ab}^\dagger \widehat{M}_{ab} = \mathbb{1}. \tag{39}$$

The Kraus operators in turn give the *operator-sum representation* of the CPTP map $\Lambda_t[\rho_{in,\mathfrak{S}}]$:

$$\Lambda_t[\rho_{in,\mathfrak{S}}] = \sum_{a,b} \widehat{M}_{ab}^\dagger \rho_{in,\mathfrak{S}} \widehat{M}_{ab} = \mathbb{1}. \tag{40}$$

From the Kraus operator completeness relation (39), the (Hölder) dual of any CPTP map is always *unital*; that is, for any CPTP map Λ_t , we always have $\Lambda_t^\dagger(\mathbb{1}_{\mathfrak{S}}) = \mathbb{1}_{\mathfrak{S}}$. The same may not necessarily be true for Λ_t itself; instead, the unitality condition for Λ_t is given in terms of the Kraus operators by the condition:

$$\sum_{a,b} \widehat{M}_{ab} \widehat{M}_{ab}^\dagger = \mathbb{1}. \tag{41}$$

Unital channels are notable since they map the identity $\mathbb{1}_{\mathfrak{S}}$ to itself (and thus the maximally mixed state $\mathbb{1}_{\mathfrak{S}}/\text{Tr}[\mathbb{1}]$ to itself): we have $\Lambda_t[\mathbb{1}_{\mathfrak{S}}] = \mathbb{1}_{\mathfrak{S}}$ only when Λ_t is unital (by definition). It is worth noting that even though the Kraus operators themselves can be arbitrarily changed by a unitary transform, this sum is invariant under such a transform, so the unitality condition is independent of the specific basis we evaluate the Kraus operators in.

We can very straightforwardly understand the difference between the conditional and unconditional Landauer reset, and in particular the terms in the unconditional Landauer bound (35), in terms of the Kraus operators and the unitality condition. In the same way as we defined the system Kraus operators, we can define the *environment Kraus operators* $\hat{N}_{cd} \in \text{Aut}(\mathcal{D}(\mathcal{H}_{\mathfrak{E}}))$ as Kraus operators on \mathfrak{E} . Labelling the eigenstates of $\rho_{\text{in},\mathfrak{S}}$ as $|s_c\rangle$, the eigenvalues of $\rho_{\text{in},\mathfrak{S}}$ as s_c , and the basis of \mathfrak{S} as $|w_d\rangle$, we can define \hat{N}_{cd} as:

$$\hat{N}_{cd} := \sum_{c,d} \sqrt{s_c} \langle w_d | \hat{U}_{t,\mathfrak{S}\mathfrak{E}} | s_c \rangle. \tag{42}$$

Then, as with (9), (22), and (29), we examine the evolution of $\mathfrak{S}\mathfrak{E}$ when we couple \mathfrak{S} initially in the state $\rho_{\text{in},\mathfrak{S}}$ to the environment \mathfrak{E} , initially in the thermal state:

$$\rho_{\text{in},\mathfrak{S}} \otimes \tau_{\mathfrak{E}} \mapsto \hat{U}_{t,\mathfrak{S}\mathfrak{E}}(\rho_{\text{in},\mathfrak{S}} \otimes \tau_{\mathfrak{E}}) \hat{U}_{t,\mathfrak{S}\mathfrak{E}}^\dagger. \tag{43}$$

As with (25), we are interested in the final environment state of \mathfrak{E} , which can tell us the bound on the energy increase of \mathfrak{E} . The final state of the environment as a result of the transformation (43) is given by:

$$\rho_{t,\mathfrak{E}} = \text{Tr}_{\mathfrak{S}} \left[\hat{U}_{t,\mathfrak{S}\mathfrak{E}}(\rho_{\text{in},\mathfrak{S}} \otimes \tau_{\mathfrak{E}}) \hat{U}_{t,\mathfrak{S}\mathfrak{E}}^\dagger \right] = \sum_{c,d} \hat{N}_{cd} \tau_{\mathfrak{E}} \hat{N}_{cd}^\dagger. \tag{44}$$

From this, and using the two-time measurement formalism [66], the probability distribution $P(Q)$ of the environment heat Q in the eigenbasis $|e_a\rangle$ of $\tau_{\mathfrak{E}}$ is given [13] by:

$$\begin{aligned} P(Q) &= \sum_{c,d;g,h} \langle e_g | \hat{N}_{cd} | e_h \rangle \langle e_h | \tau_{\mathfrak{E}} | e_h \rangle \langle e_h | \hat{N}_{cd}^\dagger | e_g \rangle \delta(Q - (E_h - E_g)) \\ &= \sum_{c,d;g,h} \langle e_g | \hat{N}_{cd} | e_h \rangle \left\langle e_h \left| \frac{e^{-\beta \hat{H}_{\mathfrak{E}}}}{\text{Tr}_{\mathfrak{E}}[e^{-\beta \hat{H}_{\mathfrak{E}}}]} \right| e_h \right\rangle \langle e_h | \hat{N}_{cd}^\dagger | e_g \rangle \delta(Q - (E_h - E_g)). \end{aligned} \tag{45}$$

This gives the moment-generating function of the dissipated heat given by:

$$\langle e^{-\beta Q} \rangle = \sum_{c,d} \text{Tr} \left[\hat{N}_{cd}^\dagger \tau_{\mathfrak{E}} \hat{N}_{cd} \right] = \sum_{c,d} \text{Tr} \left[\hat{N}_{cd} \hat{N}_{cd}^\dagger \tau_{\mathfrak{E}} \right]. \tag{46}$$

Then, a direct consequence of Jensen’s inequality is that the energy increase in this process is given in terms of the Kraus operators:

$$\Delta \langle E_{\mathfrak{E}} \rangle \geq -k_B T \ln \text{Tr} \left[\sum_{c,d} \hat{N}_{cd} \hat{N}_{cd}^\dagger \tau_{\mathfrak{E}} \right]. \tag{47}$$

The expression (47) immediately helps us understand the conditional Landauer bound (27) and the unconditional Landauer bound (35): the overall evolution (43) corresponds to a CPTP map (a.k.a. quantum channel) over \mathfrak{E} .²⁵ This channel may or may not be unital over \mathfrak{E} , and the degree to which this channel fails to be unital is exactly the degree to which the channel increases the overall entropy of \mathfrak{S} and expels the information

²⁵ It is worth noting that although (47) is slightly unclear compared to (27) and (35), in the sense that Kraus operator \hat{N}_{cd} expression mixes both the entropy increase contribution and the correlated information ejection contribution, it serves as the most clear expression from a quantum information theory point of view, and is also the tightest bound available [13].

quantity $\Delta I_{er,\mathfrak{S}}$ to \mathfrak{E} . The fact that unital quantum channels map maximally mixed states to maximally mixed states is essential: the degree to which this channel fails to be unital tells us the extent to which the channel perturbs the maximally mixed state $\tau_{\mathfrak{E}}$ of the environment. Indeed, we see that for a perfectly unital channel, the sum of the Kraus operators retrieves $\mathbb{1}_{\mathfrak{E}}$, and the energy bound is zero.

The degree of unitality stands out as a key quantity of interest in examining the nonequilibrium Landauer bound in a given system. Using the technique of full counting statistics [67], the expressions (46) and (47) can be extended [14] to a one-parameter family of expressions (replacing β with a more general parameter). This technique gives an explicit way to quantify the non-unitality of $\hat{N}_{cd}\hat{N}_{cd}^\dagger$ in the above expressions:

$$\mathcal{N}_{\mathfrak{E}} := \left\| \sum_{c,d} \hat{N}_{cd} \hat{N}_{cd}^\dagger - \mathbb{1}_{\mathfrak{E}} \right\|_2. \tag{48}$$

Here, $\|\hat{A}\|_2$ represents the Hilbert-Schmidt norm. Finally, from (46), we have the average energy dissipated into the environment given [13,14,67–69] by:

$$\begin{aligned} \beta\langle Q \rangle(t) &= -\Delta S_{\mathfrak{S}} - I(\mathfrak{S} : \mathfrak{E}) - S(\rho_{\mathfrak{E}}(t) \| \tau_{\mathfrak{E}}) \\ &= -S(\Lambda_t[\rho_{in,\mathfrak{S}}]) + S(\rho_{in,\mathfrak{S}}) - \Delta I_{er,\mathfrak{S}} - I(\mathfrak{S} : \mathfrak{E}) - S(\rho_{t,\mathfrak{E}}(t) \| \tau_{\mathfrak{E}}). \end{aligned} \tag{49}$$

We can immediately recognize this expression as simply the extension of the expressions (27) and (35) to include the possibility of initial correlations between \mathfrak{S} and \mathfrak{E} and the possibility that the environment may not start out in the thermal state. For our setup, neither of these conditions are applicable, and thus the last two terms vanish.

We would expect that the environment is not a “special” subsystem in terms of these derivations, and that an equivalent expression can be derived by considering the system. From each subsystem’s point of view, the other serves as the ancillary system in the dilation theorem sense. Indeed, expanding the Kraus operators in terms of $|s_c\rangle$ and $|w_d\rangle$, and rearranging terms in the overall trace, provides us with an equivalent expression to (46):

$$\langle e^{-\beta Q} \rangle = \text{Tr}_{\mathfrak{S}} \left[\text{Tr}_{\mathfrak{E}} \left[\hat{U}_{t,\mathfrak{S}\mathfrak{E}}^\dagger (\mathbb{1}_{\mathfrak{S}} \otimes \tau_{\mathfrak{E}}) \hat{U}_{t,\mathfrak{S}\mathfrak{E}} \right] \rho_{in,\mathfrak{S}} \right]. \tag{50}$$

As with the Kraus operators, the expression $\text{Tr}_{\mathfrak{E}} \left[\hat{U}_{t,\mathfrak{S}\mathfrak{E}}^\dagger (\mathbb{1}_{\mathfrak{S}} \otimes \tau_{\mathfrak{E}}) \hat{U}_{t,\mathfrak{S}\mathfrak{E}} \right]$ is an operator which depends on properties outside of \mathfrak{S} , but as an operator lives in $\text{Aut}(\mathcal{D}(\mathcal{H}_{\mathfrak{S}}))$; that is, it maps density matrices in \mathfrak{S} to density matrices in \mathfrak{S} . The connection between these expressions to the conditional and unconditional Landauer reset protocols is apparent, but the connection between both of these to the CTO framework is slightly more subtle. The connection between all three is discussed in Section 3.4.

As mentioned, the expectation value (46) derived in [13], and its extension derived in [14], rely on the two-time measurement formalism [66]. This might cause some trepidation—when considering the final energy, we generally must also consider the impact on the system of performing the measurement itself [8]. Conventionally, measuring the system in the state $|k\rangle$ corresponds [70,71] to a projection $\hat{\Pi}_k$ upon the pre-measurement state ρ onto $|k\rangle$. This is given by Born’s rule, which in terms of density matrices we can express as:

$$\rho \mapsto \frac{\hat{\Pi}_k \rho \hat{\Pi}_k}{\text{Tr}[\hat{\Pi}_k \rho]}. \tag{51}$$

This corresponds to a change in the von Neumann entropy given by:

$$\Delta S_m = - \sum_k \text{Tr} \left[\hat{\Pi}_k \rho \hat{\Pi}_k \ln \hat{\Pi}_k \rho \hat{\Pi}_k \right] - \text{Tr}[\rho \ln \rho]. \tag{52}$$

As we have just seen in (49), this corresponds to a change in energy. In fact, “ideal” projective measurements (i.e., those in which the measurements reproduce the measured statistics of the system, those which exhibit a one-to-one correspondence between measurement states and measured states, and which do not change the measurement statistics after measurement) cost an *infinite* amount of energy [72].

Quite fortuitously, an alternate formulation of quantum work can be developed [73] which completely avoids this issue, by focusing on the change in the expectation values of the energy eigenstates. For a time-evolving Hamiltonian starting at \hat{H}_0 with an initial eigenstate $|k_0\rangle$, this formulation is defined by:

$$\tilde{W}_k := \langle k_0 | \hat{U}_t^\dagger \hat{H}(t) \hat{U}_t | k_0 \rangle \tag{53}$$

Notably, this formulation retrieves the *same* average work as the two-time measurement formalism:

$$\begin{aligned} \langle \tilde{W} \rangle &= \sum_{k_0} \langle k_0 | \hat{U}_t^\dagger \hat{H}(t) \hat{U}_t | k_0 \rangle \frac{e^{-\beta E_{k_0}}}{\text{Tr}[e^{-\beta \hat{H}_0}]} - \text{Tr}[\tau_0 \hat{H}_0] \\ &= \text{Tr}[\rho(t) \hat{H}(t)] - \text{Tr}[\tau_0 \hat{H}_0] = \langle W \rangle. \end{aligned} \tag{54}$$

(Here, E_{k_0} is the energy eigenvalue of $|k_0\rangle$, τ_0 is the thermal (Gibbs) state corresponding to \hat{H}_0 , and $\rho(t)$ is the system state at time t). As a direct consequence of the $\langle \tilde{W} \rangle = \langle W \rangle$ equality, we see that although (46)–(50) are derived using the two-time measurement formalism, they are compatible with any alternate definitions of quantum work which provide the same expectation values.

2.2.3. Gorini-Kossakowski-Sudarshan-Lindblad (GKSL) Dynamics

Beyond providing NEQT justifications for Landauer’s principle and providing a new explanation for the difference between the conditional and unconditional Landauer recent protocols, another central aim of this work is to lay out the foundations for representing classical RC operations explicitly in terms of open quantum systems. Here, we discuss the framework of *GKSL equations with multiple asymptotic states* [17–19], which we apply in Section 3.5 to model reversible computing operations.

2.2.3.1. Markov Assumption

In a closed system governed by a Hamiltonian \hat{H} and whose state at time t is given by $\rho(t)$, the dynamics are given by the *Liouville-von Neumann (LvN) equation*:

$$\frac{d\rho(t)}{dt} = -i[\hat{H}, \rho]. \tag{55}$$

By analogy with the Liouville theorem of classical statistical mechanics and symplectic geometry, we can define [27,28] the *Liouville superoperator* as $\hat{\mathcal{L}}[\rho(t)] := -i[\hat{H}, \rho]$. This gives the superoperator version of the LvN equation as:

$$\frac{d\rho(t)}{dt} = \hat{\mathcal{L}}[\rho(t)] := -i[\hat{H}, \rho]. \tag{56}$$

(This is also known in the literature as the *quantum master equation* or the *Liouvillian*). As with the unitary evolution of states, the formal solution to this is given by a Volterra integral equation:

$$\rho(t) = \mathcal{T} \exp \left\{ \int_{t_0}^t dt' \hat{\mathcal{L}}(t') \right\} \rho(t_0). \tag{57}$$

In the specific case that $\hat{\mathcal{L}}$ is independent of time, this simplifies to $\rho(t) = e^{t\hat{\mathcal{L}}}\rho(t_0)$. In general, this may not be guaranteed to converge, let alone have a closed-form solution [74–76].²⁶ Nevertheless, this is the formal solution to the LvN Equation (56). Using the dilation theorem, we expect the time evolution of $\rho_{\mathfrak{S}}(t)$ to follow the same principle; that is, that we can determine the dynamics of $\rho_{\mathfrak{S}}(t)$ by examining the time evolution of the closed system $\mathfrak{U} = \mathfrak{S}\mathfrak{E}$ and taking the partial trace over \mathfrak{E} . Thus, we have the LvN equation for $\rho_{\mathfrak{S}}(t)$ given by:

$$\frac{d\rho_{\mathfrak{S}}(t)}{dt} = \text{Tr}_{\mathfrak{E}} \left[-i \left[\hat{H}_{\mathfrak{S}\mathfrak{E}}, \rho_{\mathfrak{S}\mathfrak{E}}(t) \right] \right] = -i \text{Tr}_{\mathfrak{E}} \left[\hat{H}_{\mathfrak{S}\mathfrak{E}}, \hat{U}_{t,\mathfrak{S}\mathfrak{E}}(\rho_{\mathfrak{S}}(t_0) \otimes \rho_{\mathfrak{E}}) \hat{U}_{t,\mathfrak{S}\mathfrak{E}}^\dagger \right]. \quad (58)$$

(Just to be clear about the notation, in the last expression we are taking the partial trace over \mathfrak{E} of the commutator of $\hat{H}_{\mathfrak{S}\mathfrak{E}}$ with the state given by the unitary time evolution of $\rho_{\mathfrak{S}}(t_i) \otimes \rho_{\mathfrak{E}}$.)

To find a solution to this equation for $\rho_{\mathfrak{S}}$, we would need to evaluate the Volterra integral Equation (57) for the global evolution over $\mathfrak{S}\mathfrak{E}$ and then trace over \mathfrak{E} . This has the exact same problems of convergence and closed form as before, since we have not changed the problem itself. Instead, as a first step to determining the dynamics of $\rho_{\mathfrak{S}}$, we can make the simplifying assumption of *Markovian dynamics*; that is, that over a differential time evolution $t \mapsto t + dt$, the properties of $\rho(t + dt)$ are determined *entirely* by the properties of $\rho(t)$. Since this assumption explicitly states that $\rho_{\mathfrak{S}}(t + dt)$ depends only on $\rho_{\mathfrak{S}}(t)$, we *must* make the Markov approximation in order to write down a differential evolution equation $\rho_{\mathfrak{S}}$ that’s first-order in time. We might be concerned that this is an overly restrictive assumption for a sensible model of reversible computing; fortunately, this assumption is in fact entirely in line with some of the key assumptions we make in our generalized models of reversible computing. The relation between these assumptions, and their suitability, is discussed in Section 4.3.

The map $\rho_{\mathfrak{S}}(t_0) \mapsto \rho_{\mathfrak{S}}(t)$ is a quantum channel $\rho_{\text{in},\mathfrak{S}} \mapsto \Lambda_t[\rho_{\text{in},\mathfrak{S}}]$; thus, we can express $\rho(t + dt)$ in terms of the operator-sum representation (40) of the CPTP map. In this representation, the Markov approximation appears as:

$$\rho_{\mathfrak{S}}(t + dt) = \sum_{c,d} \hat{N}_{cd}(dt) \rho_{\mathfrak{S}}(t) \hat{N}_{cd}^\dagger(dt). \quad (59)$$

We can retrieve a Liouville-type superoperator in the Markov approximation by examining the differential evolution of a quantum channel:

$$\Lambda_{dt}[\rho_{\mathfrak{S}}] = \hat{\mathcal{L}} + dt \lim_{dt \rightarrow 0} \frac{\Lambda_{dt} - \hat{\mathcal{L}}}{dt} + \dots := \hat{\mathcal{L}} + dt \hat{\mathcal{L}} + \dots \quad (60)$$

By expanding the Kraus operators and keeping the terms up to order $\mathcal{O}(dt)$, we get the *Gorini-Kossakowski-Sudarshan-Lindblad (GKSL) superoperator/equation* [77,78]:

$$\frac{d\rho_{\mathfrak{S}}}{dt} = \hat{\mathcal{L}}[\rho_{\mathfrak{S}}(t)] := -i \left[\hat{H}_{\mathfrak{S}\mathfrak{E}}, \rho_{\mathfrak{S}} \right] + \frac{1}{2} \sum_{a,b>0} \kappa_{ab} \left(2\hat{F}_{ab} \rho_{\mathfrak{S}} \hat{F}_{ab}^\dagger + \hat{F}_{ab}^\dagger \hat{F}_{ab} \rho_{\mathfrak{S}} + \rho_{\mathfrak{S}} \hat{F}_{ab}^\dagger \hat{F}_{ab} \right). \quad (61)$$

(These are also referred to in the literature as *Lindbladians* or *quantum Markov equation*). Here, \hat{F}_{cd} are the so-called *jump operators*. These induce “quantum jumps”; that is, the quantum state transitions that are distinct from the (closed-system) evolution of $\rho_{\mathfrak{S}}$ under $\hat{H}_{\mathfrak{S}\mathfrak{E}}$.²⁷ κ_{ab} are the rates corresponding to the *ab*th jump.

The Markov approximation made in (59) has important consequences for the time scales we consider. By definition, the Markov approximation assumes that the state

²⁶ This is of course true for the unitary evolution of states as well; convergence is only guaranteed when the algebra of the argument of the integral has a commuting structure (i.e., when the Volterra integral equation is over c-numbers).

²⁷ More precisely, $\hat{F}_{ab} \rho_{\mathfrak{S}} \hat{F}_{ab}^\dagger$ induces the jumps, and $\hat{F}_{ab}^\dagger \hat{F}_{ab} \rho_{\mathfrak{S}} + \rho_{\mathfrak{S}} \hat{F}_{ab}^\dagger \hat{F}_{ab}$ normalizes the evolution in the case that there are no jumps.

$\rho_{\mathfrak{E}}(t + dt)$ at time $t + dt$ depends *only* on the state $\rho_{\mathfrak{E}}(t)$ at time t . In doing so, we explicitly preclude [27,28,35] the possibility of *fluctuations* that can take information from \mathfrak{S} to \mathfrak{E} during an intermediate time period and then have that information return to \mathfrak{S} at time $t + dt$. Instead, this assumption is equivalent to saying that any information that is ejected from the system to the environment cannot be returned to the system. As a result, this approximation requires a separation between the time scales of these fluctuations, the time scales available at our resolution (and the dynamics of interest), and the relaxation time of the system. If we denote τ_F as the time scale of these fluctuations, τ_S as the time scale available to us at the resolution we are capable of (and thus, the time scale of the dynamics we are interested in), and τ_R as the relaxation time of the system, the Markov approximation requires a clean separation between all three time scales, corresponding to:

$$\tau_F \ll \tau_S \ll \tau_R. \tag{62}$$

Before continuing, it is worth mentioning that a change in the relationship between τ_S and τ_F leads to a substantial change in the dynamics of the system. By contrast with the $\tau_F \ll \tau_S$ condition, the condition $\tau_S \ll \tau_F \ll \tau_R$ leads to *quantum Brownian motion (QBM)* [27,28,79]. The most famous example of QBM is the *Caldeira-Leggett model* [80], which examines the quantum Brownian dynamics of a particle coupled to a bath described by a set of harmonic oscillators. At high temperatures, this gives rise to a different Markovian master equation than the GKSL Equation (61); at low temperatures, this gives rise to substantial non-Markovianities.

The GKSL evolution Equation (61) provide the evolution of the system in the Schrödinger picture; that is, when the operators $\hat{A} \in \mathcal{B}(\mathcal{H}_{\mathfrak{E}})$ on $\mathcal{H}_{\mathfrak{E}}$ are stationary and the states (and thus the density matrices) evolve with time.²⁸ We can examine the Heisenberg picture (where the states are stationary and the operators evolve with time) using the adjoint differential evolution expression:

$$\hat{A}(t + dt) = \sum_{c,d} \hat{N}_{cd}^\dagger(dt) \hat{A}(t) \hat{N}_{cd}(dt). \tag{63}$$

As with the GKSL Equation (61), we can expand the Kraus operators (or alternately take the adjoint of the GKSL equation directly) to get the *adjoint GKSL equation* governing the time evolution of operators in the Heisenberg picture:

$$\frac{d\hat{A}}{dt} = \hat{\mathcal{L}}^\dagger[\hat{A}(t)] := i[\hat{H}_{\mathfrak{E}}, \hat{A}] + \frac{1}{2} \sum_{a,b>0} \kappa_{ab} \left(2\hat{F}_{ab}^\dagger \hat{A} \hat{F}_{ab} + \hat{F}_{ab}^\dagger \hat{F}_{ab} \hat{A} + \hat{A} \hat{F}_{ab}^\dagger \hat{F}_{ab} \right). \tag{64}$$

The adjoint GKSL superoperator provides us with the time evolution of operators, including the conserved quantities of the system; these are discussed in more detail in Section 2.2.3.2. The formal solution to this evolution equation is, as we'd expect:

$$\hat{A}(t) = \mathcal{T} \exp \left\{ \int_{t_0}^t dt' \hat{\mathcal{L}}^\dagger(t') \right\} \hat{A}(t_0). \tag{65}$$

As always, this simplifies to $\hat{A}(t) = e^{t\hat{\mathcal{L}}^\dagger} \hat{A}(t_0)$ when $\hat{\mathcal{L}}^\dagger$ is independent of time. Unfortunately, the expressions of $\hat{\mathcal{L}}$ and $\hat{\mathcal{L}}^\dagger$ defined in (61) and (64) are not unique: we can have unitary transformations which redefine the jump operators or mix the jump operators and Hamiltonian (or that do both) while leaving the overall forms of $\hat{\mathcal{L}}$ and $\hat{\mathcal{L}}^\dagger$ invariant.²⁹

The GKSL Equation (61) also serves as the generator of the *quantum dynamical semigroups* [26,81,82]. Starting with the dilation theorem (8), we as always select $\mathfrak{F} = \mathfrak{E}$. As dis-

²⁸ The notation $\mathcal{B}(\mathcal{H})$ represents the set of bounded operators on \mathcal{H} with finite trace norm.

²⁹ This is unfortunate but not surprising; we saw the same kind of ambiguity in the definition of the Kraus operators. Indeed, the first kind of unitary transform is due to precisely that ambiguity from earlier.

cussed before, seeking a first-order differential equation of the form (61) automatically imposes the Markov property. If we consider the set of all of the CPTP maps which start at the same starting environment state $\rho_{\mathfrak{E}}$ and are evolved to various times t , the Markov property is equivalent to the semigroup property $\Lambda_{t_1}\Lambda_{t_2} = \Lambda_{t_1+t_2}$. Thus, we define the quantum dynamical semigroup as the family $\{\Lambda_t \mid t \geq 0\}$ that satisfies the Markov property.³⁰

As always, the differential Equation (61) is solved by (57), which simplifies to $\Lambda_t[\rho_{\mathfrak{E}}(t_0)] = \rho_{\mathfrak{E}}(t) = e^{t\hat{\mathcal{L}}}$ when $\hat{\mathcal{L}}$ is independent of t . Since the family $\{\Lambda_t \mid t \geq 0\}$ is a one-parameter semigroup, we can recognize $\hat{\mathcal{L}}$ as the generator [26,85] of this semigroup. As usual, we can then define $e^{t\hat{\mathcal{L}}}$ in terms of its Taylor-Madhava series expansion:

$$e^{t\hat{\mathcal{L}}} := \lim_{n \rightarrow \infty} \left(1 - \frac{t\hat{\mathcal{L}}}{n} \right)^n. \tag{66}$$

We see that the Markovian approximation gives a vital tool for explicitly calculating the dynamics of the evolution of $\rho_{\mathfrak{E}}$: it gives us a first-order differential equation in time for the evolution of $\rho_{\mathfrak{E}}$, entirely in terms of the known quantities $\hat{H}_{\mathfrak{E}}$ and \hat{N}_{ab} . Furthermore, by cleverly engineering \mathfrak{E} , these may even be controllable expressions in experiments.

2.2.3.2. GKSL Dynamics with Multiple Asymptotic States

Central quantities of interest in GKSL evolution are the *asymptotic states*, which are the set of states that $\rho_{\mathfrak{E}}(t)$ evolves into in the infinite time limit under the $e^{t\hat{\mathcal{L}}}$ evolution:

$$|\rho_{\infty}\rangle\rangle := \lim_{t \rightarrow \infty} e^{t\hat{\mathcal{L}}} |\rho_{\text{in},\mathfrak{E}}\rangle\rangle. \tag{67}$$

(Here, we have employed the “double-ket” notation that appears in the vectorization of $\mathcal{B}(\mathcal{H}_{\mathfrak{E}})$; this is discussed in more detail in Appendix B. This notation will be used for the remainder of the text).

We can determine the asymptotic states by examining the spectral decomposition of $\hat{\mathcal{L}}$. Here, we follow the excellent presentation in [19], which is the thesis corresponding to the original papers [17,18] which develop this framework. (This is simply an abbreviated version; the reader interested in more details of the framework is highly encouraged to read these references.) In general, $\hat{\mathcal{L}}$ is not necessarily Hermitian. If it is still unitarily diagonalizable, we denote the eigenvalues by λ_a , the (right) eigenvectors by $|p_a\rangle\rangle$, and the left eigenvectors by $\langle\langle q_a|$. Denoting $\hat{D} = \text{diag}(\lambda_a)$, \hat{P} as the matrix formed by setting $\{|p_a\rangle\rangle\}$ as rows, and \hat{P}^{-1} as the matrix formed by setting $\{\langle\langle q_a|\}$ as columns, we have $|\rho_{\mathfrak{E}}(t)\rangle\rangle$ in terms of the spectral decomposition of $\hat{\mathcal{L}}$ given by:

$$\begin{aligned} |\rho_{\mathfrak{E}}(t)\rangle\rangle &= e^{t\hat{\mathcal{L}}} |\rho_{\text{in},\mathfrak{E}}\rangle\rangle = \hat{P} e^{t\hat{D}} \hat{P}^{-1} |\rho_{\text{in},\mathfrak{E}}\rangle\rangle \\ &= \sum_a |p_a\rangle\rangle e^{t\lambda_a} \langle\langle q_a|\rho_{\text{in},\mathfrak{E}}\rangle\rangle = \sum_a c_a e^{t\lambda_a} |p_a\rangle\rangle. \end{aligned} \tag{68}$$

Here, $c_a = \langle\langle q_a|\rho_{\text{in},\mathfrak{E}}\rangle\rangle$ are simply the coefficients of the (Hilbert-Schmidt) inner product $\langle\langle q_a|\rho_{\text{in},\mathfrak{E}}\rangle\rangle = \text{Tr}_{\mathfrak{E}}[q_a^\dagger \rho_{\text{in},\mathfrak{E}}]$. Since $t \in \mathbb{R}$, the eigenvalues can only satisfy $\Re\{\lambda_a\} < 0$ or

³⁰ According to Lindblad’s theorem [77], any quantum operation that satisfies the semigroup property will satisfy the GKSL equation and vice versa. However, note that not every CPTP map satisfies the GKSL equation [83,84]; it just so happens that we are only concerned with the ones that do. Oftentimes, the requirement that $\text{Tr}[(\Lambda_t[\rho])\hat{A}]$ is a continuous function of t for all trace-norm-bounded operators $\hat{A} \in \mathcal{B}(\mathcal{H}_{\mathfrak{E}})$ is also specified. However, because we already specified that Λ_t is a CPTP map for all t , that $\hat{A} \in \mathcal{B}(\mathcal{H}_{\mathfrak{E}})$ is a bounded function, and that the semigroup property is specified where $t \in \mathbb{R}$ is a continuous parameter, this requirement is a direct consequence of what we already have. Finally, we note that the notation $\text{Op}(\mathbb{H})$ is sometimes used for $\mathcal{B}(\mathcal{H}_{\mathfrak{E}})$, for example, as in [17–19].

$\Re\{\lambda_a\} = 0$ to have finite $|\rho_\infty\rangle\rangle$. The eigenstates for $\Re\{\lambda_a\} < 0$ are the *damped* or *decaying states*, while the eigenstates for $\Re\{\lambda_a\} = 0$ are the asymptotic states.³¹

In the case that $\hat{\mathcal{L}}$ is not diagonalizable, it still has a decomposition in Jordan normal form. Each Jordan block has a specific eigenvalue, but eigenvalues could be spread over multiple Jordan blocks. If we remove the assumption of unitary diagonalizability, the decomposition given by (68) still holds for diagonal Jordan blocks, where λ_a is now specifically the eigenvalue of $|\mathbf{p}_a\rangle\rangle$. (Since the left eigenvectors are the same as the right eigenvectors of the adjoint operator, the eigenvalue of $\langle\langle \mathbf{q}_a|$ is just λ_a^* . The overall spectral decomposition of the operators now use the right eigenvalue specifically.) For non-diagonal Jordan blocks, the expression $e^{t\hat{\mathcal{L}}}$ pulls down a factor of $t^n/n!$ for each instance of a nonzero superdiagonal entry. Thus, for a given non-diagonal Jordan block with right eigenvalue λ_a and N generalized eigenvectors, if we index the generalized eigenvectors by $\mu, \nu \in \mathbb{N}_N^+$, we have the decomposition of $e^{t\hat{\mathcal{L}}}|\rho_{\text{in},\mathfrak{G}}\rangle\rangle$ for that specific block given by [17,19]:

$$\begin{aligned} |\rho_{\mathfrak{G}}(t)\rangle\rangle &= e^{t\hat{\mathcal{L}}}|\rho_{\text{in},\mathfrak{G}}\rangle\rangle = \sum_{\nu \leq \mu} \frac{t^{(\nu-\mu)}}{(\nu-\mu)!} |\mathbf{p}_a\rangle\rangle e^{t\lambda_a} \langle\langle \mathbf{q}_a|\rho_{\text{in},\mathfrak{G}}\rangle\rangle \\ &= \sum_{\nu \leq \mu} \frac{c_a e^{t\lambda_a} t^{(\nu-\mu)}}{(\nu-\mu)!} |\mathbf{p}_a\rangle\rangle. \end{aligned} \tag{69}$$

A direct consequence is that the Jordan normal form of the asymptotic state eigenvalue blocks (i.e., the eigenvalue blocks with pure imaginary eigenvalues) are all diagonal, since the factor of $t^{(\nu-\mu)}$ would blow up otherwise.

As mentioned earlier, the factor of $e^{t\lambda_a}$ in the spectral decompositions (68) or (69) of $|\mathbf{p}_a\rangle\rangle$ tells us that the set of pure imaginary eigenvalues correspond to the asymptotic states; we can denote these as $\Lambda_a = \Im\{\lambda_a\}$. The corresponding right eigenvectors form a subspace of $\mathcal{B}(\mathcal{H}_{\mathfrak{G}})$, called the *asymptotic subspace* and denoted $\text{As}(\mathcal{H}_{\mathfrak{G}})$. The asymptotic left and right eigenvectors are respectively the asymptotic states and the (asymptotic) conserved quantities of the GKSL system we are examining.³² Indexing these by their value of Λ and their corresponding degeneracy μ , we will denote the right and left eigenvectors as $|s_{\Lambda\mu}\rangle\rangle$ and $\langle\langle J_{\Lambda\mu}|$ respectively. (For clarity, this means that we have $\hat{\mathcal{L}}|s_{\Lambda\mu}\rangle\rangle = i\Lambda|s_{\Lambda\mu}\rangle\rangle$ and $\hat{\mathcal{L}}^\dagger|J_{\Lambda\mu}\rangle\rangle = -i\Lambda|J_{\Lambda\mu}\rangle\rangle$.)

Under the action of $e^{t\hat{\mathcal{L}}}$, all initial states $|\rho_{\text{in}}\rangle\rangle$ converge to $\text{As}(\mathcal{H}_{\mathfrak{G}})$ in the $t \rightarrow \infty$ limit. Naturally, the GKSL dynamics maps states outside of $\text{As}(\mathcal{H}_{\mathfrak{G}})$ —the states that do not survive in the infinite time limit—to some linear combination of states in $\text{As}(\mathcal{H}_{\mathfrak{G}})$. However, at a *finite* time t , there's no requirement that if $|\rho_{\mathfrak{G}}(t)\rangle\rangle = e^{t\hat{\mathcal{L}}}$ is in $\text{As}(\mathcal{H}_{\mathfrak{G}})$, then it must stay in that state for all future times. Indeed, the GKSL dynamics permits the time evolution of a state in $\text{As}(\mathcal{H}_{\mathfrak{G}})$ to leave that subspace altogether at some other time, as long as the state returns to $\text{As}(\mathcal{H}_{\mathfrak{G}})$ at $t \rightarrow \infty$. Two examples of GKSL state evolution are provided in Figure 8.

³¹ Most examinations of GKSL dynamics examine the case of a single asymptotic state, given by the right eigenvector corresponding to $\lambda = 0$. The corresponding left eigenvector is $\mathbb{1}$. GKSL systems with multiple asymptotic states form a set of measure zero [19] over the set of all possible GKSL systems. However, we are representing an actual engineered system: we are not interested in the set of *all* possible GKSL systems, we are interested in the one that actually represents our system, so this is not a problem.

³² These conserved quantities are slightly different from some of the more familiar conserved currents in classical and quantum mechanics; for instance, the conserved quantity given by the $\langle\langle \mathbb{1}|$ eigenvector is the trace of ρ_∞ .




Figure 8. Two examples of evolution of states under GKSL dynamics to $As(\mathcal{H}_{\mathfrak{E}})$ in the $t \rightarrow \infty$ limit. Here, the gray blocks denote further subspaces of $As(\mathcal{H}_{\mathfrak{E}})$. $|\rho_{in,2}\rangle\rangle$ starts outside of $As(\mathcal{H}_{\mathfrak{E}})$ and settles into one such subspace of $As(\mathcal{H}_{\mathfrak{E}})$ as $t \rightarrow \infty$, whereas $|\rho_{in,3}\rangle\rangle$ starts inside one subspace of $As(\mathcal{H}_{\mathfrak{E}})$, leaves it but stays within $As(\mathcal{H}_{\mathfrak{E}})$, and then returns to the same subspace as $t \rightarrow \infty$. Many further dynamics are possible: for instance, a state in one of these subspaces could leave $As(\mathcal{H}_{\mathfrak{E}})$ altogether, and then settle into the same or even a different subspace of $As(\mathcal{H}_{\mathfrak{E}})$ as $t \rightarrow \infty$; meanwhile, a state could start in a subspace of $As(\mathcal{H}_{\mathfrak{E}})$, leave it but stay within $As(\mathcal{H}_{\mathfrak{E}})$ overall, and settle into a different subspace as $t \rightarrow \infty$. There are, indeed, no restrictions on where the initial state can start from, as long as it is somewhere in the universe and settles to $As(\mathcal{H}_{\mathfrak{E}})$ as $t \rightarrow \infty$: as with $|\rho_{in,2}\rangle\rangle$, a state could start outside of $As(\mathcal{H}_{\mathfrak{E}})$ altogether, or a state could alternately start inside $As(\mathcal{H}_{\mathfrak{E}})$ but outside one of the subspaces. These subspaces are defined in the infinite time dynamics; thus, they only have relevance *after* the GKSL dynamics has finished.

From $|s_{\Lambda\mu}\rangle\rangle$ and $|J_{\Lambda\mu}\rangle\rangle$, we can construct a superoperator projector $\hat{\mathcal{P}}_{\infty}$, known as the *asymptotic projection*, which projects solely onto $As(\mathcal{H})$:³³

$$\hat{\mathcal{P}}_{\infty} := \sum_{\Lambda,\mu} |s_{\Lambda\mu}\rangle\rangle\langle\langle J_{\Lambda\mu}| = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{\Lambda} \int_0^T dt e^{t(\hat{\mathcal{L}} - i\Lambda\hat{\mathcal{L}})} = \frac{1}{2\pi i} \oint_C \frac{dz e^{t\hat{\mathcal{L}}}}{z\hat{\mathcal{L}} - \hat{\mathcal{L}}}. \quad (70)$$

(Here, C is the contour which encloses only the $\{\lambda\}$ s.) Using $\hat{\mathcal{P}}_{\infty}$, we can write $|\rho_{\infty}\rangle\rangle$ as:

$$|\rho_{\infty}\rangle\rangle = e^{-i\hat{H}_{\infty}s} \left(\hat{\mathcal{P}}_{\infty} |\rho_{in}\rangle\rangle \right) e^{i\hat{H}_{\infty}s}. \quad (71)$$

Here, \hat{H}_{∞} is the Hamiltonian that governs the asymptotic dynamics, parametrized by the time parameter $s \gg t$; that is, s *solely* describes the dynamics of the system after it has already equilibrated. From this expression, we can directly see that $As(\mathcal{H}_{\mathfrak{E}}) = \hat{\mathcal{P}}_{\infty}[\mathcal{B}(\mathcal{H}_{\mathfrak{E}})]$.

At this point, an important subtlety about time scales must be addressed. As mentioned in Section 2.2.3.1, the Markov assumption involves the separation of multiple time scales. Three time scales in particular are relevant: the time scale τ_F of the fluctuations between \mathfrak{S} and \mathfrak{E} , the time scale $\tau_{\mathfrak{S}}$ of the GKSL dynamics, and the relaxation time scale τ_R ; all of which are separated according to (62). The introduction of multiple asymptotic states gives rise to the possibility of dynamics within the asymptotic space *after* the relaxation. By contrast, the GKSL Equation (61) examines the process by which a system initially in a state $|\rho_{in}\rangle\rangle$ relaxes to one or more of the asymptotic state(s) $|\rho_{\infty}\rangle\rangle$. This assumes that the relaxation occurs on a timescale much faster than the timescale involved in the dynamics of the asymptotic state(s). Thus, implicit in this expression is the notion that t is a parameter that *only* sees timescales on the order of the relaxation timescale. In other words, the GKSL dynamics is entirely *before* the dynamics of the asymptotic states, and the limit as $t \rightarrow \infty$ is

³³ The last two equalities come from recognizing $\hat{\mathcal{P}}_{\infty}$ as the projector onto the peripheral spectrum of $e^{t\hat{\mathcal{L}}}$. The second-to-last equality expresses $\hat{\mathcal{P}}_{\infty}$ as the Cesàro mean of $e^{t\hat{\mathcal{L}}}$, and the last equality expresses $\hat{\mathcal{P}}_{\infty}$ as the Dunford-Taylor integral of $e^{t\hat{\mathcal{L}}}$ via the resolvent $(z\hat{\mathcal{L}} - \hat{\mathcal{L}})^{-1}$.

then still *before* the asymptotic state dynamics. Labelling the asymptotic dynamics timescale s_{As} , we now modify (62) to include s_{As} , giving:

$$\tau_F \ll \tau_S \ll \tau_R \ll s_{As}. \tag{72}$$

Intuitively, the GKSL dynamics is interested in the limit as $t \rightarrow \tau_R$, but due to the parametrization of the dynamics, we take the limit as $t \rightarrow \infty$. In order to describe the asymptotic dynamics, we use a separate parameter s .

The spectral properties of $\hat{\mathcal{L}}$ play a central role in understanding the GKSL dynamics of the system in question. As such, the projector decomposition of these dynamics which separates the asymptotic dynamics from the dissipative dynamics serves as an essential tool to understand the overall structure of GKSL evolution. This is given by the *four-corners decomposition*, first developed in [18]. (As before, we follow the presentation in [19].) We can define the operator $\hat{P}_A \in \mathcal{H}_G$ as the projector onto the asymptotic states. Explicitly, \hat{P}_A is defined in terms of ρ_∞ by:

$$\begin{aligned} \hat{P}_A \rho_\infty \hat{P}_A &= \rho_\infty; \\ \text{Tr}_G [\hat{P}_A] &= \max_{\rho_\infty} [\text{rank}(\rho_\infty)]. \end{aligned} \tag{73}$$

(The second expression helps us to ensure that \hat{P}_A is defined so that it *only* projects onto the asymptotic states). Meanwhile, the complement of \hat{P}_A is given by: $\hat{Q} := \mathbb{1}_G - \hat{P}_A$, with $\hat{Q} \rho(t) \hat{Q} \rightarrow 0$ as $t \rightarrow \infty$. Together, \hat{P}_A and \hat{Q} provide the four-corners projections of operators $\hat{A} \in \mathcal{B}(\mathcal{H}_G)$:

$$\begin{aligned} \hat{A}_{\blacksquare} &:= \hat{P}_A \hat{A} \hat{P}_A; & \hat{A}_{\blacklozenge} &:= \hat{Q} \hat{A} \hat{Q}; \\ \hat{A}_{\blacktriangleleft} &:= \hat{P}_A \hat{A} \hat{Q}; & \hat{A}_{\blacktriangleright} &:= \hat{Q} \hat{A} \hat{P}_A. \end{aligned} \tag{74}$$

$$\hat{A} = \begin{pmatrix} \hat{A}_{\blacksquare} & \hat{A}_{\blacktriangleleft} \\ \hat{A}_{\blacktriangleleft} & \hat{A}_{\blacklozenge} \end{pmatrix}$$

Thus, \hat{P}_A and \hat{Q} provide a decomposition of every $\hat{A} \in \mathcal{B}(\mathcal{H}_G)$. (74) also provides a definition of the *four corners projection superoperators* $\{\hat{\mathcal{P}}_{\blacksquare}, \hat{\mathcal{P}}_{\blacklozenge}, \hat{\mathcal{P}}_{\blacktriangleleft}, \hat{\mathcal{P}}_{\blacktriangleright}\}$, which act on \hat{A} as indicated.

(73) and (74) serve as the foundation for examining the properties of GKSL systems with multiple asymptotic states, as well as the geometric properties of their quantum state spaces. These in turn are fundamental for examining the properties of classical (including reversible) computing operations in GKSL systems. However, a detailed discussion of the properties of the four-corners decomposition is far beyond the scope of this paper; the interested reader is highly encouraged to examine [17–19]. For our purposes here, what’s relevant is that $As(\mathcal{H}_G)$ forms an identifiable subspace, which we can project onto using the four corners projection superoperators.

As an important note, $\hat{\mathcal{P}}_{\blacksquare}$ does *not* project onto $As(\mathcal{H}_G)$ directly. Rather, the \blacksquare subspace contains $As(\mathcal{H}_G)$ in its entirety: $As(\mathcal{H}_G) \subseteq \blacksquare$. The difference between these lies in the asymptotic dynamics governed by \hat{H}_∞ : $As(\mathcal{H}_G)$ describes the states that survive in the infinite-time limit (as given in (70) and (71)). If there are no further dephasing dynamics within \blacksquare , then $As(\mathcal{H}_G) = \blacksquare$; conversely, if there are, then $As(\mathcal{H}_G) \subsetneq \blacksquare$. This notation also serves as a visual indication for the framework: each operator in $\mathcal{B}(\mathcal{H}_G)$ can be subdivided into four regions, corresponding to these projections. We can freely gather $As(\mathcal{H}_G)$ into the top-left corner. Then, $\hat{P}_A \hat{A} \hat{P}_A$ projects into the top-left corner; that is, $\hat{P}_A \hat{A} \hat{P}_A$ projects into \blacksquare .

The subspace \blacksquare is a full Hilbert space in its own right, supporting quantum mechanical evolution under \hat{H}_∞ in the $t \rightarrow \infty$ limit. Thus, it supports *any* possible dynamics that can

be governed by a Hamiltonian. Indeed, this framework provides a way to describe an open system extension of *any* finite-dimensional system that can be governed by the laws of non-relativistic quantum mechanics, as long as the open system relaxation is governed by Markov dynamics.³⁴ Thus, we can directly model a system of computational states as discussed in Section 2.1.5.2: each computational state corresponds to a DFS within an overall Hilbert space. The overall Hilbert space \mathbb{H} is then the direct sum of the individual DFS spaces, known as the *von Neumann algebra* [86–90]:

$$\mathbb{H} = \bigoplus_i \mathcal{H}_{\text{DFS},i}. \tag{75}$$

The fact that this is identical to (2) (at the level of the vectorized space) shows us that the von Neumann algebra is a natural framework to represent reversible computing operations. This representation is discussed in detail in Section 3.5. The embedding of a von Neumann algebra within the four corners representation of an operator evolving under GKSL dynamics is given in Figure 9.




Figure 9. An example of the type of overall operator algebra $\mathcal{B}(\mathcal{H}_{\mathbb{C}})$ that can support classical (including reversible) computing operations. (Note this matrix operates on the space of vectorized density matrices). Here, the upper left corner subspace is a von Neumann algebra corresponding to the direct sum of decoherence-free subspaces, with the gray regions representing the individual decoherence-free subspaces.

A striking feature of GKSL dynamics with multiple asymptotic states is that the shape of \mathbb{H} corresponds to substantially different expressions for the quantum geometric tensor (QGT) over \mathbb{H} . This is a key aspect of GKSL dynamics with multiple asymptotic states, and will also serve as an essential feature of understanding the properties of RC operations in open quantum systems. The dependence of the dynamics on the QGT of \mathbb{H} is central to the framework developed in [18], and is discussed in detail there and in [19]. Unsurprisingly, because the framework of classical reversible computing operations in open quantum systems relies at its core on GKSL dynamics with multiple asymptotic states, the dependence of GKSL dynamics on the QGT over \mathbb{H} is an indispensable part of classical RC operations in open quantum systems as well. The quantum geometric properties of RC operations are briefly mentioned in Section 3.5. A more detailed analysis of these properties, and conclusions regarding RC operations, are the central theme of a forthcoming work which follows up on these discussions.

2.3. Existing and Proposed Implementation Technologies

In this section, we briefly survey a number of conceptual examples of concrete physical mechanisms of operation that may be suitable, to varying degrees, for performing reversible computations. The detailed performance characteristics for these example technologies

³⁴ In principle, infinite-dimensional Hilbert spaces, relativistic quantum mechanical systems, and quantum field theories should be describable as well; however, the details of these descriptions are still in progress.

(e.g., the exact energy dissipation per operation as a function of speed) depend on a great many design details, so we will not attempt to derive those characteristics here. Rather, this survey is just to give the reader an idea regarding the range of physical mechanisms for reversible computing that may be possible. It is likely that many other, much more efficient mechanisms can be invented with further research.

First, here is a concise list of the technologies we will survey, with abbreviations noted (some of which are coined here):

1. Reversible adiabatic CMOS (RA-CMOS).
2. Reversible quantum flux parametron (RQFP).
3. Reversible quantum-dot cellular automaton (R-QCA).
4. Reversible nanomechanical rod logics (RNRL).
5. Ballistic asynchronous reversible computing in superconductors (BARC or BARCS).

These particular examples will be described in a bit more detail in the following subsections. These are not the *only* physical mechanisms for reversible computing to have been proposed, but (except for BARCS) are some of the most well-developed implementation concepts so far.

2.3.1. Reversible Adiabatic CMOS

This class of implementation technologies for reversible computing refers to a logic design discipline based on ordinary CMOS (complementary metal-oxide-semiconductor) field-effect transistors [48,91–96]. To approach physical reversibility in these types of circuits requires several conditions to be met (see Appendix A for some key derivations):

1. The *on/off conductance ratio* $r_{\text{on/off}} = G_{\text{on}}/G_{\text{off}}$ of the device channel (at the specified operating points) should diverge, as the technology is improved. The quantity G_{on} refers to the typical effective peak source-drain conductance through the channel of a device (transistor) when it is in the “on” state (with gate voltage set accordingly, for example, $V_g = V_{\text{dd}} = \text{logic HIGH}$ for an n-type FET). Meanwhile, G_{off} refers to the maximum conductance through the device for off-state “leakage” current (including both gate leakage and subthreshold current) when the device is nominally turned off (e.g., $V_g = 0 = \text{logic LOW}$ for an nFET). Roughly speaking, $1/G_{\text{on}}$ ends up being proportional to the characteristic *relaxation timescale* $\tau_r = R_{\text{on}}C$ of the circuit, while $1/G_{\text{off}}$ ends up being proportional to the characteristic *equilibration timescale* $\tau_e = R_{\text{off}}C$ of the circuit when its non-equilibrium state is not being actively maintained. One of the classic results of physical reversible computing theory, the roots of which can be traced back to Feynman’s lectures on computation, delivered in the early 1980s [97], is that in general, at least for any classic “adiabatic” reversible computing technology, the maximum energy recovery efficiency for a reversible device is ultimately limited as a function of the ratio of these two timescales, for example, as $\eta_{\text{er}} \leq 1 - c \cdot \sqrt{\tau_r/\tau_e}$ for $\tau_e \gg \tau_r$. (See Appendix A) That is, the minimum fraction of signal energy dissipated per operation cycle scales like $\sqrt{\tau_r/\tau_e}$, quite generally. For CMOS, this means that, to attain high energy efficiency, we want to make the leakage conductance G_{off} as small as possible to extend the equilibration timescale τ_e , and doing this well in practice requires some combination of various engineering refinements (e.g., higher threshold voltages, thicker gate oxides, lower operating temperatures, higher materials purity). Identifying the most economical manufacturing process to minimize G_{off} in practice is not a simple optimization problem by any means. However, there appears to be no fundamental reason why the ratio $r_{\text{on/off}}$ cannot be made as large as desired with further refinement of the technology over time. Thus, it seems that this class of circuits can approach ideal reversibility with continued development.
2. Since in CMOS, the relaxation timescale τ_r is subject to lower bounds, the *transition time* t_{tr} for the adiabatic logic transitions should also diverge. For a given technology, the minimum dissipation per cycle will be found when the transition time t_{tr} is (within a small constant factor) roughly at the geometric mean $\tau_m = \sqrt{\tau_r\tau_e}$ between

the relaxation and equilibration timescales (Appendix A). However, as long as we can arrange to keep extending the equilibration timescale τ_e , the useful transition time $t_{tr} \approx \tau_m$ can continue increasing as well.

3. The *effective quality factor* Q_{eff} of any external resonant oscillatory element serving as the clock-power supply driving the adiabatic circuit should also diverge. For our purposes, Q_{eff} can be defined as the ratio between the peak electrostatic energy $E_{\text{load}} = \frac{1}{2}C_{\text{load}}V_{\text{dd}}^2$ stored transiently on the logic nodes, and the energy dissipated by the resonant oscillator per cycle, $Q_{\text{eff}} = E_{\text{load}}/E_{\text{odiss}}$.

In addition to the above, two design rules that must always be obeyed in (conditionally) reversible adiabatic CMOS circuits in order for them to be able to approach physically-reversible operation in the above limits are the following:

1. Never turn on a transistor when there is a nonzero source-drain voltage across it.
2. Never turn off a transistor when there is a nonzero source-drain current through it.

If either of these rules is ever broken in the design, this can lead to substantial non-adiabatic dissipation, and the physical computational process as a whole no longer qualifies as being asymptotically physically reversible. This is discussed further in [48,93,98].

A brief description of the overall normal mechanism of reversible operation for these kinds of circuits is as follows. Periodic voltage waveforms are supplied by a resonant oscillatory circuit that is customized to provide quasi-trapezoidal wave shapes (that is, with roughly flat waveform tops and bottoms). The flat regions are needed in order to avoid pushing current through devices while they are being switched on or off. The provided waveforms exist in several different (mutually-offset) phases. Each phase drives a corresponding section (subset) of adiabatic logic circuits. The choice of which circuit nodes to charge up in a given section is determined using series-parallel networks of devices, whose gate (control) electrodes are connected to the (quiescent) nodes controlled by a neighboring phase. After the supplied waveforms have finished causing the desired transitions between valid logic voltage levels for a given section of the circuit, now those circuit nodes can be used to control the adiabatic transitions for the neighboring sections in adjacent clock phases. The correct architectural design of these kinds of circuits can become somewhat involved, but is conceptually straightforward. See [48] for an example.

2.3.1.1. Description of RA-CMOS in Terms of Our General Framework

In reversible adiabatic CMOS technology, a given (time-dependent) computational state $c(\tau_\ell)$ has a very simple description in terms of physical microstates. Essentially, in a given computational state, at a given time, each circuit node exhibits a given well-defined, relatively uniform voltage level, within some tolerances. Of course, there will be local fluctuations about that average level. Physical states in which voltages depart substantially (over a broad region) from any of the computationally-meaningful levels can be relegated to the catch-all “invalid” computational state c_\perp , but during normal operation of a well-engineered circuit, such states should have an astronomically close to zero probability of arising in any case.

In terms of computational operations, a certain computational operation O_s^t is carried out each time one of the supply waveforms executes a voltage-level transition between two distinct valid logic levels, that is, from an initial level V_i to a final level $V_f \neq V_i$, over the time interval between the two time points s (start time) and $t = s + t_{tr}$ (end time). During this transition, the voltage levels on the set of circuit nodes that are connected to that particular supply line themselves undergo (with a slight delay, and modulo voltage offsets due to leakage and other non-idealities) the same transition between voltage levels. In this process, some transistors (e.g., ones whose gates are controlled by the transitioning nodes) may be turned on or turned off, causing the source and drain nodes of those transistors to become connected to or disconnected from each other. These connection and disconnection events result in the set of accessible computational states changing over time (since even the

number of independent connected components will be changing, thus, so will the number of available computational states).

In any case, as long as the two rules of adiabatic design are respected throughout a given transition, the operation O_s^t that is performed will be both (conditionally) logically reversible (under the condition that the rules are respected), as well as asymptotically thermodynamically reversible, in the limit described above where $G_{\text{off}} \rightarrow 0$ and $t_{\text{tr}} \rightarrow \infty$, while keeping $t_{\text{tr}} \ll \tau_e$. Thus, we can allow both $r_{\text{on/off}}$ and t_{tr} to continue increasing as the technology develops, and approach perfect reversibility over time given continued development of this family of technologies.

The above discussion glosses over the important issue of the effective quality factor Q_{eff} of the driving power-clock resonator, which will also limit the overall degree of reversibility, but, as far as we know at present, there is no reason why this quantity cannot diverge as well, with continued engineering refinements. (The development of high-quality trapezoidal resonators suitable for driving adiabatic circuits is in the scope of engineering R&D work being performed at Sandia.)

2.3.2. Reversible Quantum Flux Parametron

The Reversible Quantum Flux Parametron (RQFP) logic family [99–101] is a logically reversible variant of the well-developed superconducting logic family AQFP (Adiabatic Quantum Flux Parametron), which has been being developed primarily at Yokohama National University in Japan. RQFP (and its not-necessarily-reversible generalization AQFP) rely on adiabatic transformation of the abstract potential energy surface (PES) that obtains within Josephson-junction-based superconducting circuits. The independent variables for the PES describe the current distribution in the circuit, and the phase (order parameter) differences across the junctions. The PES is manipulated in such a way that the occupied potential energy valley of the system is transformed adiabatically to configurations representing different computational states.

RQFP circuits are controlled by externally supplied waveforms, similarly to the case in RA-CMOS, except that the supplies are providing current signals, not voltage signals (since voltages, except for inductive transients, are normally zero in superconducting circuits). Except for the fact that the state of the circuit and the driving signals at a given time is described in terms of currents instead of voltages, and that the physics of superconductivity dominates the charge transport, the representation of RQFP in terms of computational and physical states is, roughly speaking, qualitatively similar to the case in RA-CMOS. That is to say, the higher-level principles of pipelined reversible logic are roughly comparable between the two technologies.

One advantage, however, of RQFP compared to CMOS is that, due to Meissner-effect trapping of flux quanta, the natural equilibration timescale τ_e is extremely large (effectively infinite) in RQFP, and as a result, scaling to extreme ultra-low levels of dissipation may ultimately prove far easier to do in RQFP than in adiabatic CMOS. The primary disadvantages of RQFP, compared to CMOS, are its lower density and accordingly higher manufacturing cost per-device, together with its requirement for low-temperature operation.

2.3.3. Reversible Quantum-Dot Cellular Automaton

The Quantum-dot Cellular Automata (QCA or QDCA) [102–104] family of technologies operate using single electrons confined to quantum dots, dipole configurations of two such electrons confined to four such dots in a square layout separated by tunnel barriers (a.k.a. a “cell”), and linear/branching arrays of such cells interacting through dipole-dipole Coulombic interactions. As in RA-CMOS and RQFP, externally supplied signals are used to adiabatically raise and lower potential energy barriers that separate neighboring regions of the physical state space, in patterns that (in the technology’s reversible variant, here dubbed R-QCA) allow the overall computational state to evolve reversibly—which, as usual, means in a (conditionally) logically reversible and asymptotically physically reversible way.

An interesting note about QDCA is that it was recently shown [105] that exponential scaling of adiabaticity with speed (as in Landau-Zener transitions with a missed level crossing) exists in this system, apparently implying that there is no fundamental lower bound on dissipation-delay product. In [106], Pidaparthi and Lent investigate this phenomenon in more detail using a Lindbladian analysis, finding that when there is weak thermal coupling to the environment, these systems can exhibit substantially suppressed dissipation within a certain regime of speeds. This is a promising result, and we expect that this type of behavior likely generalizes to a wider variety of quantum systems.

2.3.4. Reversible Nanomechanical Rod Logics

This is a concept that goes back to K. Eric Drexler's work in the 1980s leading up to his dissertation on molecular nanotechnology at MIT [107–109]. The original idea was that logical bits are encoded in the linear displacements of atomically-precise nanorods that move within sleeves at the ends of nano-springs. The nods are pushed back and forth (by externally supplied mechanical signals, following the same kind of quasi-trapezoidal waveforms we have talked about previously) to adiabatically transform them between computational states, using nano-scale bumps on the rods to sterically hinder each other's motion in ways that allow them to perform (conditionally reversible) Boolean logic. The whole scheme is closely analogous to RA-CMOS, except that it uses mechanical rather than electrical state variables.

Drexler's rod logic concept was updated more recently [110,111] by a group led by Ralph Merkle (a pioneering cryptographer and early nanotechnologist). The new concept eliminated the sleeve bearings, whose friction had dominated the dissipation in Drexler's earlier concept. In the new scheme, the only bearings are *rotary* bearings implemented by single carbyne bonds, whose orbitals are circularly symmetrical. Frictional losses in this system were assessed in simulations [112] to be so low that individual joints (operated reversibly) would dissipate $\sim 70,000\times$ less than Landauer's $k_B T \ln 2$ limit even when operating at frequencies as high as 100 MHz. This example illustrates that in principle, dissipation-delay products for reversible operations can be far smaller than is the case in RA-CMOS. (This particular dissipation-delay value is roughly $10^6\times$ improved versus projected *end-of-roadmap* CMOS—and at room temperature)!

The main problem with the Drexler-Merkle family of nanomechanical rod logic concepts for reversible computing is simply that building them would seemingly require a very general, sophisticated, atomically-precise and fast technology for nano-fabrication and assembly, which does not yet exist and may continue to not exist for some decades.

2.3.5. Ballistic Asynchronous Reversible Computing in Superconductors

Ballistic Asynchronous Reversible Computing (BARC, previously called ABRC) [113] is a fundamentally new physical model of reversible computing in which the computational degrees of freedom evolve *ballistically* (i.e., under their own inertia) rather than being dragged along adiabatically as a side effect of the oscillatory evolution of an external resonator. This change may provide certain advantages in terms of, for example, allowing us to avoid having to worry about accidentally exciting undesired modes of the resonator and of the distribution network for the driving signal. The BARC model is required to be asynchronous as a means to prevent the nonlinear interactions between subsystems from chaotically amplifying uncertainties in the subsystem trajectories.

In a current project at Sandia, we are attempting to implement the BARC model in superconducting electronic circuits [114]. (We call such realizations BARCS.) The computational subsystems are individual polarized flux solitons (or *fluxons*) propagating near-ballistically along long Josephson junction (LJJ) transmission lines. In our circuits, fluxons are conserved and interact asynchronously with stored flux quanta at (stateful) interaction sites or *circuit elements*, transforming the local digital state reversibly, in a deterministic sort of elastic "scattering" interaction.

BARC is an extremely novel concept, and has only been developed to a very preliminary level to date. So far, we have a single (very simple) “working” BARCS circuit element (i.e., it simulates correctly in SPICE) [115]; a test chip for it has been fabricated, and experimental tests of it are in progress. However, a wider variety of useful elements, leading up to a complete logic family, still need to be developed and optimized.

We are also collaborating with a group at the University of Maryland which has been working on a similar ballistic approach which they call Reversible Fluxon Logic (RFL) [116–122]. The original RFL concept envisioned synchronous ballistic logic, but the Maryland group is also now also developing asynchronous elements which fall into the BARCS paradigm [123].

3. Results

Much work remains to be done, in terms of fleshing out a complete physical theory of reversible computing informed by NEQT, but in this section, we review some important preliminary results in this area that can be, or have already been, obtained.

First, we view it as important, for resolving some of the long-standing controversies in the thermodynamics of computation, to distinguish a couple of different results that have historically been associated with Landauer’s Principle:

1. First is a simple result regarding the interchangeability of entropy between computational and non-computational forms. This one follows directly from the association of computational states to sets of microstates discussed in Section 2.1.2. However, it is such an important result that we call it *The Fundamental Theorem of the Thermodynamics of Computation*. We review it in Section 3.1 below. This result implies that non-computational or “physical” entropy must be increased when computational (“information”) entropy is reduced, but *does not require that total entropy be increased*.
2. Second is a result (Section 3.2) showing that a strict entropy *increase* is required whenever there is a loss of known information (which by itself is not surprising, since entropy increase effectively *means* that known information is reduced), and furthermore, that an example of this necessarily occurs when one of two mutually-correlated subsystems is *obviously erased*, meaning that, in isolation, its reduced subsystem entropy is ejected to its local thermal environment without regards to its existing correlations. To the extent that the ejected information is then thermalized, with its correlations to the other subsystem being lost, this then corresponds to a strict increase in total entropy. This result follows directly from unitarity, information theory, and the definitions in Section 2.1.

We argue that it is the second result, and not the first one, that is most properly understood as being *Landauer’s Principle*, because Landauer’s Principle is most properly taken to concern the consequences of information loss *in a computer*—since that was the subject of Landauer’s original paper [1]. At least in an ordinary, deterministic computer, it is normally the case that *computed bits are correlated*—meaning, there is mutual information between them (and/or, between them and the inputs that they were computed from)—since in fact, one can say that it is the generation of specific desired patterns of correlation between different computational subsystems that is exactly the entire point of what computation, *per se*, is all about.

In particular, we show in Section 3.2 that, for any deterministically computed bit (or larger computational subsystem), the amount of *new entropy that is generated* when that bit is obviously erased is strictly lower-bounded by the prior reduced subsystem entropy of that bit (or subsystem).³⁵ Note that, in that statement, we are talking about an absolute *increase* in the *total* entropy of the model universe (including computational *and* non-computational entropy, defined below), and *not* just a *transfer* of entropy from computational to non-computational form. Thus, Landauer’s Principle, when it is properly

³⁵ The form of this argument that we present was previously made explicit in [12], but we reprise it here.

understood in this way, really does provide a lower bound on new entropy *generation*, and not simply on entropy transfer, as has sometimes been claimed.

In addition to the above clarification of Landauer's Principle, we also (in Section 3.3) review two fundamental theorems of physical reversible computing theory. (These were previously presented in [39,40], but we reprise them here).

1. *The Fundamental Theorem of Traditional Reversible Computing*, whose proof is summarized in Section 3.3.1 below, states that the only deterministic computational operations that *always* avoid ejecting computational entropy to non-computational form (and thus, can avoid Landauer's lower limit on entropy increase when operating in isolation on computed bits) are the *unconditionally logically reversible operations* (e.g., Toffoli gate operations) traditionally studied in reversible computing theory.
2. *The Fundamental Theorem of Generalized Reversible Computing*, whose proof is summarized in Section 3.3.2 below, states that, in a statistically *contextualized* computation, it can suffice (in a properly designed mechanism) to avoid entropy ejection (and the resultant entropy increase due to Landauer's limit) if a computational operation is simply reversible on the *subset* of initial states having nonzero probability in the given statistical operating context [39,40].

Taking the latter observation (the generalized theorem) into account is essential in order for the scope of the theory to adequately encompass the state of the art of the existing best practices (Section 2.3) in the engineering of reversible computing hardware. The generalized theorem significantly expands the class of computational mechanisms that can be seen to be capable of approaching thermodynamic reversibility when appropriate constraints are met. In particular, all of the actual implementation technologies for reversible computing described in Section 2.3 can *only* be understood properly in the light of the generalized form of the theorem. In other words, all of the real reversible computing technologies that have been implemented to date are only *conditionally* reversible, and so they rely, for their ability to achieve asymptotic reversibility in practice, on the fact that their preconditions for reversibility have been met by design, within the architectures of those machines, and (implicitly) on the *Generalized Theorem of RC*.

Next, the framework of NEQT provides us with several new perspectives from which to understand Landauer's principle, and to begin to characterize the properties of RC operations in open quantum systems. In Section 3.4, we relate Landauer's principle and the structure of reversible computing to the CTOs discussed in Section 2.2.1.2. By examining the difference between conditional and unconditional Landauer reset using the structure of these CTOs, we find a general motivation and rationale for the structure of RC directly from RTQT. Quite spectacularly, we see that by treating elements of the physical computer system as the "catalyst" in the CTO sense, we can directly represent repeated cycles of computation and Landauer reset. In particular, we see that using the structure of CTOs outlined in Section 2.2.1.2, we can cycle through the operation of computational systems as long as we wish, with minimal buildup of QMI. Finally, in Section 3.5, we begin the foundational work of representing RC operations in terms of open quantum systems from a first-principles level, using the properties of GKSL dynamics with multiple asymptotic states. In particular, we note specific properties of computational and noncomputational operations, and briefly discuss implications in terms of their quantum geometric signatures.

Let us now review these results in a bit more detail.

3.1. The Fundamental Theorem of the Thermodynamics of Computation

First, starting from the basic conception of (classical, digital) computational states presented in Section 2.1.2, we can easily derive what we call *The Fundamental Theorem of the Thermodynamics of Computation*. This theorem formalizes the necessary relationship between so-called "information entropy" (that is, entropy of the computational state) and physical entropy.

To start, let ϕ be a variable representing the (complete, micro-) physical state of the computer system \mathfrak{S} , specified by a choice of one of the protocomputational basis

vectors $\vec{b} \in \mathcal{B}$. Assume that, at a given point in time τ , the probability mass over the different possible physical states ϕ is distributed according to a probability distribution $p(\phi)$, as given in the usual way using the Born rule, or equivalently, by the diagonal elements of the system’s instantaneous density matrix $\rho(\tau)$ in the \mathcal{B} basis.

We can then derive an implied probability distribution $P(c)$ over the computational states $c_i \in \mathcal{C}$, by simply summing p over the various physical states $\phi = \phi_j \in \mathcal{B}_i$, where \mathcal{B}_i denotes the specific basis set $\mathcal{B}_c \subset \mathcal{B}$ corresponding to computational state $c = c_i$:

$$P(c_i) = \sum_{\phi_j \in \mathcal{B}_i} p(\phi_j). \tag{76}$$

It is then trivial to show that the system’s total (von Neumann/Shannon) entropy $S(\Phi)$ (where Φ is a random variable ranging over values ϕ) can always be partitioned as:

$$S(\Phi) = H(C) + S(\Phi|C), \tag{77}$$

where $H(C)$ (with C a random variable ranging over values c) refers to the computational entropy or “information entropy,” meaning the entropy of the computational state variable C according to the above-derived probability distribution $P(c)$, and meanwhile $S(\Phi|C)$ refers to the conditional entropy of the physical state variable Φ if the value of the computational state variable C is given. If we then define *non-computational entropy* as $S_{nc} = S(\Phi|C)$, then we can just say, “total entropy equals computational entropy plus non-computational entropy.” Note that this is true always—no matter how the protocomputational basis \mathcal{B} is defined and partitioned into basis sets corresponding to computational states. See Figure 10. (For more details, see [11,12]).

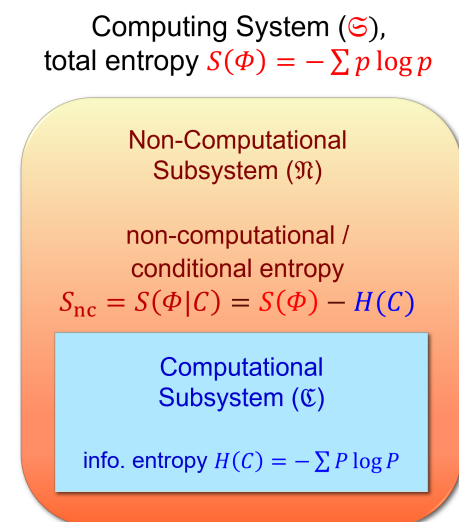


Figure 10. Fundamental Theorem of the Thermodynamics of Computing, illustrated using the picture of Figure 3. No matter how we choose the protocomputational basis \mathcal{B} of the computing system \mathfrak{S} and partition it into subsets \mathcal{B}_c for distinct computational states $c \in \mathcal{C}$, we can always express the total physical entropy $S(\Phi)$ of the system as the sum of the information entropy $H(C)$ of the computational state (state of the computational subsystem \mathfrak{C}), and the non-computational entropy $S_{nc}(\Phi)$ of \mathfrak{S} , which is equal to the conditional entropy $S(\Phi|C)$ of the physical state when the computational state is given.

This fact, together with the Second Law of Thermodynamics (i.e., $\partial S/\partial t \geq 0$ globally³⁶), implies that one cannot ever reduce computational entropy (for example, by merging two computational states, like in Figure 5b) without also increasing non-computational

³⁶ It is worth noting that the entropy production rate can be negative in local systems due to non-Markovian dynamics [124] or systems with delayed response [125].

entropy by (at least) a corresponding amount. Of course, this works both ways—meaning, if you *increase* computational entropy (e.g., by *splitting* a computational state, in a stochastic computational operation, like in Figure 5c), you can thereby also *reduce* non-computational entropy accordingly. This is done in practice, for example, in paramagnetic cooling [126–128], if we think of the (relatively stable) randomized magnetic domains that form during the cooling process as constituting “computational” bits.

A widespread perception³⁷ is that the theorem corresponding to the above observations constitutes a form of Landauer’s Principle, but we argue that, although it is indeed a very important basic theorem of the thermodynamics of computation, calling it “Landauer’s Principle” creates confusion, because it entirely misses what we argue is the central, most important point of Landauer’s Principle proper, which has to do more specifically with the loss of computed, *correlated* information, such as typically exists within a computer. This viewpoint is discussed at great length in [11,12], and more concisely in the following subsection.

3.2. Landauer’s Principle Proper

Since we wish to argue that Landauer’s Principle is, most centrally, a theorem about the consequences of information loss in *computation*, specifically, it behooves us to say a little bit more about what we mean by that.

A quite general picture of computation involves the concept of function evaluation, for example, computing $y = f(x)$, given x . In fact, historically, the very first formal model of universal computation, due to Alonzo Church, defined general computations in terms of recursive function evaluation [131]. It is of course well-known today that arbitrarily complex computations can be composed out of simple function evaluations (e.g., Boolean logic operations).

Let us then consider, as an example, two subsystems $\mathfrak{X}, \mathfrak{Y}$ of a computational system \mathfrak{C} , that exist for purposes of holding the input value x and output value y of some function $f(\cdot)$ to be evaluated. Let us assume that subsystems \mathfrak{X} and \mathfrak{Y} have separable corresponding computational state spaces $C_{\mathfrak{X}}, C_{\mathfrak{Y}}$, which is to say, the computational states of subsystems \mathfrak{X} and \mathfrak{Y} are independently measurable. There is then a joint computational state space $C_{\mathfrak{X}\mathfrak{Y}} = C_{\mathfrak{X}} \times C_{\mathfrak{Y}}$. Suppose initially we have some distribution $P(C_{\mathfrak{X}})$ over the initial state of \mathfrak{X} . Suppose, then, a deterministic computational operation $O_{\mathfrak{X}\mathfrak{Y}}$ is performed on the joint system $\mathfrak{X}\mathfrak{Y}$ which leaves $C_{\mathfrak{X}}$ unchanged, but results in $c_{\mathfrak{Y}} = f(c_{\mathfrak{X}})$, which is to say, the computational state $c_{\mathfrak{Y}}$ of \mathfrak{Y} becomes a state representing the value $y = f(x)$, where x is the value represented by the state $c_{\mathfrak{X}}$ of \mathfrak{X} . Note that this operation $O_{\mathfrak{X}\mathfrak{Y}}$ could also be reversible, for example, if $C_{\mathfrak{Y}}$ contained a known value initially (e.g., is “cleared memory”).

It then follows from the above setup that:

1. First, the *reduced* computational entropy of \mathfrak{Y} , written $H(C_{\mathfrak{Y}})$, after performing $O_{\mathfrak{X}\mathfrak{Y}}$, is *entirely* accounted for by the mutual information between \mathfrak{Y} and \mathfrak{X} ; that is, $H(C_{\mathfrak{Y}}) = I(C_{\mathfrak{X}}; C_{\mathfrak{Y}})$. In other words, \mathfrak{Y} contains an exactly *zero* amount of *independent* entropy, relative to \mathfrak{X} , since $H(C_{\mathfrak{Y}}|C_{\mathfrak{X}}) = 0$. (I.e., \mathfrak{Y} is completely determined by \mathfrak{X} .) This just follows from the fact that, as is typically the case in traditional digital computation, function evaluation is a deterministic operation.
2. Second, now suppose that, next, an *irreversible* computational operation O_{erase} is performed *locally* on \mathfrak{Y} in *complete isolation* from \mathfrak{X} , that is, without any *influence* from the state of \mathfrak{X} , or even any applied *knowledge* about the state of \mathfrak{X} (beyond

³⁷ As an example, [129] discusses Landauer’s Principle and logical reversibility in a context that *only* considers independent systems, rather than subsystems of correlated systems, and concludes that logical reversibility is not required for thermodynamic reversibility, an observation which was already noted explicitly for this case in earlier work such as [7]. Although formally correct, such analyses are misleading, in that they neglect to mention the key point that, when obliviously erasing parts of *correlated* systems, such as deterministically computed bits, there is necessarily a loss of mutual information that does indeed result in a required thermodynamic irreversibility. Thus, these analyses have been misinterpreted by some (e.g., [10,130]) as evidence by the fundamental rationale for reversible computing is incorrect, but this is not in fact the case, since, when the correlated case which is relevant to computing is considered, the connection between logical and thermodynamic reversibility is recovered. We discuss this in more detail in Sections 3.2 and 4.4.

our prior distribution $P(C_x)$, and suppose, further, that the overall output-state distribution $P(C_{\mathfrak{Y}})$ resulting from O_{erase} has zero entropy. This resultant distribution $P(C_{\mathfrak{Y}})$ is found by computing a weighted sum of $O_{\text{erase}}(c_{\mathfrak{Y}})$ over the set of all input computational states of \mathfrak{Y} with probability $P(c_{\mathfrak{Y}}) > 0$. For this distribution to have zero entropy implies that all such states of \mathfrak{Y} map to the same value, $c_{\mathfrak{Y}} = c_0$, which is why O_{erase} can be considered an “erasure” operation.

3. If we now simply assume that *the non-computational entropy in \mathfrak{S} will shortly be thermalized*—which is to say, the entropy ejected from \mathfrak{Y} is not being preserved in a stable or predictable form elsewhere in the physical state of the system—then it follows that the correlation previously embodied by the mutual information $I(C_x; C_{\mathfrak{Y}})$ has now been lost, and therefore, *the total entropy of the model universe \mathfrak{U} is immediately (i.e., after a thermalization timescale) increased by (at minimum) the prior value of the reduced (marginal) subsystem entropy $H(C_{\mathfrak{Y}})$, just before the erasure*. An example is illustrated in Figure 11.




Figure 11. Landauer’s Principle as entropy increase from thermalization of mutual information. Red shading denotes probability density. **(Left)** Two perfectly-correlated computational bit-systems \mathfrak{X} and \mathfrak{Y} ; their states could have been prepared by computing \mathfrak{Y} ’s value y deterministically from x , for example, using $y := x$. **(Middle)** When the variable Y (for the computational state of \mathfrak{Y}) is obliviously erased, this amounts to merging the two computational states in each column; we can say that now $Y = 0$ (say) in each merged state. Note that now, there briefly exists a correlation between X and the non-computational part of the physical state. **(Right)** Very quickly (over a thermalization timescale), we lose track of the probabilities of the different physical states making up each computational state, thus losing this correlation information. This is where the absolute increase of total entropy from Landauer’s Principle necessarily occurs. We cannot, of course, then undo this entropy increase by simply reversing the first step (un-merging the Y states), because the correlation information between X and Y has already been irrevocably lost by this point.

We argue that the resulting theorem constitutes what is the *most appropriate* statement of Landauer’s principle: Namely, that to erase *any* deterministically-computed information in an isolated computational subsystem *obliviously*, that is, without regards to its correlations with other information that may exist, and if this is followed by allowing the reduced subsystem entropy that was thereby ejected to subsequently thermalize, results in turning that subsystem’s previous mutual information (which was *not* independent entropy) into *true* entropy (real, new uncertainty), and thereby results in a permanent increase in the total entropy of the universe by the corresponding amount.

Finally, perhaps the easiest way to see that a loss of mutual information generally implies entropy increase is simply to note that, for any random variables X, Y ,

$$H(X, Y) = H(X) + H(Y) - I(X; Y), \tag{78}$$

and thus, if we hold the marginal subsystem entropies $H(X), H(Y)$ steady while their mutual information $I(X; Y)$ is reduced, the total entropy $H(X, Y)$ of the joint distribution of the XY system must increase accordingly. More broadly, one can say that mutual information is a part of the total *known information* $K = C - H$ in a system, where C is the *information capacity* or maximum entropy of the system, and H is its present entropy (which

we can consider *unknown information*). Thus, a loss of mutual information is really just a special case of the more general process of the transformation of physical information from a ‘known’ to an ‘unknown’ status, which is entropy increase.

Additional details of the argument in this subsection can be found in [12].

3.3. Fundamental Theorems of Reversible Computing

In this section, we review what we call *The Fundamental Theorems of Reversible Computing (RC)* [39,40], which show that, in order for a deterministic computing system to avoid entropy increases due to Landauer’s Principle (when understood properly, as in Section 3.2 above), logically reversible computational operations must be utilized. The Fundamental Theorem of RC comes in two versions: The traditional version, which shows that traditional *unconditionally* logically reversible operations are required in order to avoid entropy increase from Landauer’s Principle in *all* possible input circumstances, and the (less often recognized) *generalized* version, which shows that a broader class of *conditionally* reversible operations suffice, for use in systems that are properly designed to ensure that the preconditions for reversibility of those operations are met.

3.3.1. Fundamental Theorem of Traditional Reversible Computing

Before we review the traditional RC theorem, we first present a simple definition, based on the discussion of Sections 3.1 and 3.2 above, which will be helpful for stating it. (The following is presented in time-independent terms, but can easily be made time-dependent.)

3.3.1.1. Entropy-Ejecting Operations

A computational operation O on a computational state set \mathbf{C} is called (*potentially*) *entropy-ejecting* if and only if there exists some possible prior distribution $P(\mathbf{C}) \in \mathcal{P}(\mathbf{C})$ such that, when the operation O is applied within that context, the increase ΔS_{nc} in the non-computational entropy required by the Fundamental Theorem of the Thermodynamics of Computing (Section 3.1) is greater than zero. If an operation O is not (even potentially) entropy-ejecting, we call it *non-entropy-ejecting*.

Note that if an operation is entropy-ejecting, and it is performed *in isolation* (by which we implicitly also mean *obliviously*, without external knowledge of the state being applied) on a subsystem that contains mutual information with other subsystems (and if we assume that any non-computational entropy will not be preserved in a predictable form, but will be thermalized), then this entropy ejection will furthermore result in a *global entropy increase*, by a straightforward generalization of Landauer’s Principle (in its proper form, stated above in Section 3.2). This leads us to:

The Fundamental Theorem of Traditional Reversible Computing. If a deterministic computational operation O is non-entropy-ejecting (by the above definition), then it follows that O must be unconditionally logically reversible.

The proof of this statement is trivial, but can be found in [40].

Note, also, that an immediate corollary of this theorem is that, if we wish to perform computational operations *in isolation* (i.e., obliviously) on subsystems that contain any mutual information with other systems (such as subsystems whose state was computed deterministically from those other systems), then we can *only* avoid a global entropy increase from Landauer’s Principle in the *general* case (i.e., for *any* distribution $P(\mathbf{C})$ over initial computational states, and when the non-computational state does not preserve information in a predictable form) if those operations are *unconditionally logically reversible*.

3.3.2. Fundamental Theorem of Generalized Reversible Computing

The traditional theorem, above, is in essence about how we can avoid Landauer losses in the *worst case*—that is, when we assume that we have no control over what the initial computational state $c_1 \in \mathbf{C}$ may be, and thus, any statistical mixture of initial states is

possible. However, in a real computer, the initial state prior to a given computational operation may be (and usually is) a resultant state from a previous operation. Thus, it is frequently the case that we can, *by design* in a computer, *restrict* the set of possible initial states to a proper subset $A \subset \mathbf{C}$ of *allowed* states. This then makes it possible to design computing mechanisms that avoid Landauer losses by transforming just the *subset* A of allowed states reversibly. This is, in fact, how typical real engineered technologies for reversible computing (including those described in Section 2.3 above) work—since it turns out, in general, to be much easier, in practice, to design mechanisms that only transform restricted subsets of computational states reversibly, rather than the full set of all potentially describable states. To show why doing this is sufficient, we need a more general version of the fundamental theorem of RC, one that properly models the case where the set of initial states is restricted.

To do this, we also need to extend the concept of an entropy-ejecting operation from Section 3.3.1 as follows:

3.3.2.1. Entropy-Ejecting Computations

For purposes of the below theorem, let a (statistically contextualized) *computation* $\mathcal{C} = \mathcal{C}(O, P)$ refer to the concept of performing a computational operation O over its computational state space \mathbf{C} , *given* a particular *initial* probability distribution $P = P(C)$, where C is a random variable ranging over computational states $c \in \mathbf{C}$. (The *quantum contextualized computation* concept of Section 2.1.5 is just a straightforward generalization of this concept to a quantum context.) We say that a (deterministic) computation \mathcal{C} is (*specifically*) *entropy-ejecting* if and only if the increase ΔS_{nc} in the non-computational entropy required by the Fundamental Theorem of the Thermodynamics of Computation (i.e., due to a reduction in computational entropy $H(C)$) is greater than zero. If the computation \mathcal{C} is not specifically entropy-ejecting, we call it non-entropy-ejecting.

As before, this then allows us to immediately state the corresponding theorem:

Fundamental Theorem of Generalized Reversible Computing. A deterministic computation $\mathcal{C}(O, P)$ is non-entropy-ejecting if and only if *at least one* of its preconditions for reversibility is satisfied with probability 1 under the initial probability distribution P .

As with the traditional theorem, the proof of this is easy, but may be found in [40].

Like with the traditional theorem, the generalized theorem has an immediate corollary, which is that if we wish to perform the computation $\mathcal{C}(O, P)$ in isolation (obviously) on a subsystem bearing mutual information with other systems (such as a subsystem whose computational state was deterministically computed from those outside systems), then we can only avoid a global entropy increase from Landauer's Principle for that specific computation (assuming, as usual, that the non-computational state does not preserve information in a predictable form) if the operation O is conditionally reversible, under (at least) the precondition that $c \in A$, where $A = \{c_i \in \mathbf{C} \mid P_i > 0\}$.

The significance of the two RC theorems together is that, in order to avoid the otherwise-necessary entropy increase resulting from Landauer's Principle when performing isolated computational operations on subsystems in the context of larger deterministic computations, one must confine oneself to the above two cases (unconditionally reversible operations, and/or conditionally reversible operations that have a satisfied condition for reversibility).

The significance of the *generalized* reversible computing theorem, as opposed to the traditional one, is to observe that it is a sufficient logical-level requirement, to avoid requiring an entropy increase from Landauer's Principle, if simply those initial states having *nonzero probability* in the given statistical operating context $P(C)$ are mapped one-to-one onto final states.

Of course, in any event, even when these conditions for reversibility are satisfied, to avoid entropy increase in reality *also* requires that the physical mechanisms *implementing* the given computation must be designed to approach thermodynamic reversibility in

practice—but, the import of the RC theorems is to say that, when the conditions of either theorem are satisfied, Landauer’s Principle, at least, does not preclude doing this.

Additional discussion of these two theorems can be found in [39], with detailed proofs available in the associated postprint [40]. We should note that, although the particular proofs of these theorems presented in that earlier work did not yet explicitly utilize the quantum generalization of the concept of a statistical operating context that we presented in Section 2.1.5 above, all of the constructions in Section 2.1.5 were specifically designed to guarantee that the exact same proofs will go through essentially unmodified in the quantum version of the theory (given our assumptions about rapid decoherence of final states). Thus, the above theorems remain completely valid within the quantum framework of the present paper.

3.4. Representations of Reversible Computing by Catalytic Thermal Operations

The results above follow from a *static* analysis just based on the overall starting and ending states of a given computational process; however, obtaining more detailed results (e.g., about minimum energy dissipation as a function of speed) will require more detailed attention to the dynamics of computational transitions. This then requires engaging more detailed theoretical methods, such as the resource-theoretic tools we reviewed in Section 2.2.1. In this section, we discuss how to think about reversible computational operations in terms of those more detailed methods.

3.4.1. Reconsidering the Notion of a Catalyst

The nonequilibrium Landauer results in Sections 2.2.2.1 and 2.2.2.2 emphasize how essential it is to have a map as close to unital as possible in order to minimize the energy cost of the operation (i.e., the importance of a map that minimizes the entropy difference between $\rho_{\text{in},\mathfrak{S}}$ and $\Lambda_t[\rho_{\text{in},\mathfrak{S}}]$). Ideally, we’d like to connect these to the theory of thermal operations, so that we can begin to identify the thermal processes that are relevant for reversible computing. As discussed in Section 2.2.2.1, the bound is zero when the Landauer reset protocol applied to the subsystem bearing computational degrees of freedom is conditioned on the state of that subsystem *and* when the von Neumann entropies of the subsystem state before and after reset are equal. Although achieving this in practice is a nontrivial engineering challenge, there is no *fundamental* (quantum mechanical) barrier to this constraint, so we can consider this case specifically (i.e., that $S(\rho_i) = S(\rho_j) = S(\rho_r)$ where ρ_i and ρ_j can be any possible computational state).

In order to select the right TO, we again return to the Landauer reset protocols. These were described earlier as the process of resetting a state bearing some computational degrees of information, either with a conditioned or unconditioned potential. Without loss of generality, we can consider the system carrying these degrees of freedom as part of a larger system. Then, we can label the larger system \mathfrak{S} and the subsystem carrying these degrees of freedom Ω , as shown in Figure 12. In terms of Ω , the Landauer reset protocol is the process of transforming the state $\rho_{\ell,\Omega}$ to the reset state $\rho_{r,\Omega}$. Until now, we have been content to consider the Hamiltonians acting on all of the systems and subsystems as background Hamiltonians. Ordinarily, this would not pose any issues, and indeed helps us keep the properties of the system as general as possible.

However, as we saw in Section 2.2.2.1, whether or not $\widehat{V}_{r,\Omega}$ can be conditioned on $\rho_{\ell,\Omega}$ plays a vital role in calculating the free energy bound on the reset process, and in particular on the information contribution to the bound. Thus, by contrast to the other contributions to the overall Hamiltonian, by having \widehat{V}_{Ω} as an ambient background potential we in fact lose a vital piece of context of the overall process. This context is relevant when trying to identify the correct TO we want to use to examine the process. Thus, we need to think of \widehat{V}_{Ω} as not an ambient potential that acts on Ω , but rather a potential that acts on Ω from a *different* subsystem \mathfrak{B} —some relevant part of the previously unaddressed control/timing subsystem \mathfrak{J} from Figure 7—within the overall computer system \mathfrak{S} . This echoes the discussion at the end of Section 2.2.2.1, when we saw that the ΔI_{er} term in

the unconditional Landauer bound (35) contained within it mutual information between the state of the system implementing the Landauer reset potential \hat{V} and the state of the subsystem \mathfrak{S}' bearing computational degrees of freedom.

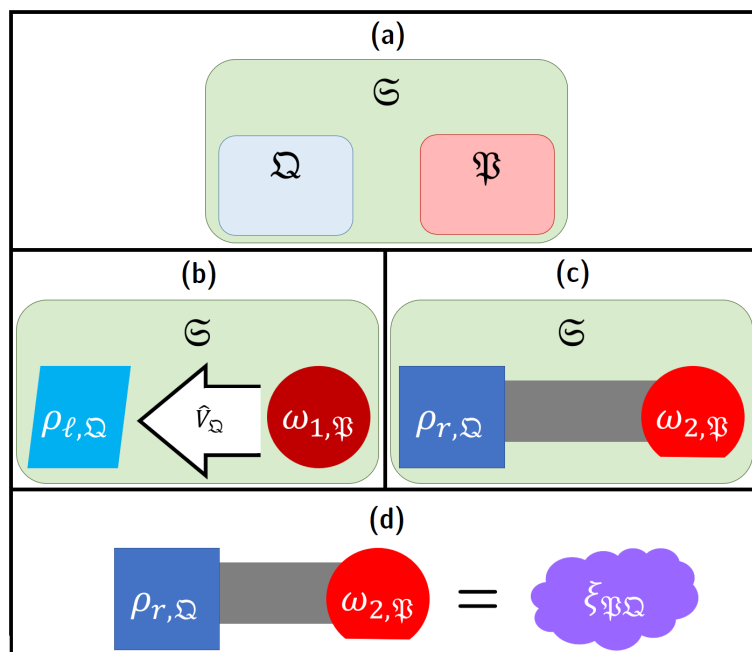


Figure 12. (a) The embedding of a system Ω carrying computational degrees of freedom inside a larger system \mathfrak{S} , which also contains a subsystem \mathfrak{P} used to induce state transformations on Ω . (b) An example of the process discussed in Section 2.2.2.1, in terms of this embedding. Ω is in one of the N states $\{\rho_{\ell, \Omega}\}$. A potential \hat{V}_{Ω} is applied to $\{\rho_{\ell, \Omega}\}$ to transform it into the reset state $\rho_{r, \Omega}$. Crucially, since \hat{V}_{Ω} is part of the system, it must be contained in \mathfrak{S} (i.e., it must *not* be part of the environment); meanwhile, it also must be *outside* of all of Ω . We can without loss of generality consider the subsystem that applies \hat{V}_{Ω} onto Ω to be part of (or all of) \mathfrak{P} . (c) The application of this potential transforms the local state of Ω as $\rho_{\ell, \Omega} \mapsto \rho_{r, \Omega}$. Assuming that the dynamics can be described by a CTO of the form (11), the transformation must also induce some state transformation $\omega_{1, \mathfrak{P}} \mapsto \omega_{2, \mathfrak{P}}$, to preserve unitarity of the overall dynamics on $\mathfrak{P}\Omega$. (The sole exception to this is the identity transformation). The application of \hat{V}_{Ω} by \mathfrak{P} upon Ω also gives rise to a correlation between Ω and \mathfrak{P} , given by the QMI (15). (d) The correlated states $\rho_{r, \Omega}$ and $\omega_{2, \mathfrak{P}}$ correspond to a single state $\xi_{\mathfrak{P}\Omega}$ over $\mathfrak{P}\Omega$ as a whole, with $\text{Tr}_{\mathfrak{P}}[\xi_{\mathfrak{P}\Omega}] = \rho_{r, \Omega}$ and $\text{Tr}_{\Omega}[\xi_{\mathfrak{P}\Omega}] = \omega_{2, \mathfrak{P}}$.

We can immediately recognize \mathfrak{S}' in Section 2.2.2.1 as identical to Ω here, and \mathfrak{P} as the system implementing \hat{V} . If we require that the local state of \mathfrak{P} (i.e., the state of \mathfrak{P} when we trace out every other subsystem) be the same to within some value $\epsilon \in \mathbb{R}^+$ under the trace distance, then we can treat \mathfrak{P} as a catalyst which is necessary to induce the local state transformation $\rho_{\ell, \Omega} \mapsto \rho_{r, \Omega}$. In this framework, the natural TO to examine this transformation is the general CTO given by (11). Indeed, we can recognize the ΔI_{er} term in the unconditional Landauer bound (35) as precisely the same as the QMI (15) built up between the transformed subsystem and the catalyst.

Thus, in terms of TOs, the correlation-preserving generalized CTOs discussed in Section 2.2.1.2 are precisely the conditioned Landauer reset protocols discussed in Section 2.2.2.1. By contrast, the requirement in (16) that the final state of the catalyst remain uncorrelated with the final state of the transformed subsystem is identical to applying a single unconditioned potential for the Landauer reset. As was the case when comparing the most general CTO (11) and the more traditional CTO (16), the correlated information $\Delta I_{\text{er}} = I(\Omega : \mathfrak{P})$ is ejected from the overall system $\mathfrak{S} = \Omega\mathfrak{P}$ after the unconditional protocol. Alternately, and equivalently, the unconditional Landauer reset protocol (resp., the more traditional CTO) can be realized by performing the condi-

tional Landauer reset protocol (resp., the general CTO) and then ejecting the correlated information afterwards.

Thus far, we have used the general CTO to describe the process of a subsystem being transformed with the help of a catalyst, and to obtain the quantum thermodynamic restrictions on this process. However, the word “catalyst” is simply a convention based on our modelling; mathematically, the catalyst is simply another subsystem, which (within an $\epsilon \in \mathbb{R}^+$) returns locally to the same state before and after the overall state transformation on \mathfrak{S} . Nothing *a priori* tells us that the subsystem \mathfrak{K} in (11) must be a catalyst; as far as the mathematics is concerned, this is simply a statement about a specific type of transformation on the system $\mathfrak{S} = \mathfrak{T}\mathfrak{K}$, where \mathfrak{T} and \mathfrak{K} are subsystems and \mathfrak{K} has the additional requirement (12) that the state before and after the transformation is locally the same up to an infinitesimal ϵ . This offers an interesting extension of the discussion above regarding the reset of Ω due to a subsystem \mathfrak{P} : what happens if we allow \mathfrak{P} to change its state before and after the transition, but we require that the subsystem Ω bearing computational degrees of freedom must return to the same state? *A priori* this seems like it would make no sense: how can we meaningfully talk about computation if the subsystem with computational degrees of freedom is left unchanged?

To answer this, we will consider the decomposition of the transition $\rho_\Omega \otimes \omega_{\mathfrak{P}} \mapsto \zeta_{\Omega\mathfrak{P}}$ into two distinct (non-catalytic) TOs over the subsystems Ω and \mathfrak{P} . Here, we enforce $\text{Tr}_{\mathfrak{P}}[\zeta_{\Omega\mathfrak{P}}] = \rho_\Omega$ and $\|\text{Tr}_\Omega[\zeta_{\Omega\mathfrak{P}}] - \Xi(\omega_{\mathfrak{P}})\|_1 < \epsilon$ for some infinitesimal $\epsilon \in \mathbb{R}^+$. In other words, the overall transition $\rho_\Omega \otimes \omega_{\mathfrak{P}} \mapsto \zeta_{\Omega\mathfrak{P}}$ is an operation where now the state over Ω starts and ends in the same state. In this decomposition, one of these transformations locally takes the state of Ω away from σ_Ω , and the other returns the state of Ω to σ_Ω . Overall, then, we consider the pair of transformations:

$$\rho_\Omega \otimes \omega_{\text{in},\mathfrak{P}} \mapsto \chi_{\Omega\mathfrak{P}} \mapsto \zeta_{\Omega\mathfrak{P}}. \tag{79}$$

In addition to the properties for $\text{Tr}_{\mathfrak{P}}[\zeta_{\Omega\mathfrak{P}}]$ and $\text{Tr}_\Omega[\zeta_{\Omega\mathfrak{P}}]$ stated above, we have $\text{Tr}_{\mathfrak{P}}[\chi_{\Omega\mathfrak{P}}] = \gamma_\Omega$ for some $\gamma_\Omega \neq \rho_\Omega$.

The composed transition $\rho_\Omega \otimes \omega_{\mathfrak{P}} \mapsto \zeta_{\Omega\mathfrak{P}}$ is the CTO $\rho_\Omega \otimes \omega_{\mathfrak{P}} \mapsto \zeta_{\Omega\mathfrak{P}}$ discussed above, where now it is the state ρ_Ω of Ω that starts and ends in the same state under the composition. Thus, the composed transition must correspond to the constraints in Section 2.2.1.2. This has a direct interpretation in terms of computing if we think of σ_Ω as a reset state. The reset state is conventionally the state we use as our starting point to perform a computation on Ω , which is why it has chosen as the reset state in the first place. Then, the process (79) corresponds to starting with the standard reset state ρ_Ω , using a different subsystem \mathfrak{P} to manipulate the state of Ω (i.e., perform a computation on Ω), and then using \mathfrak{P} once more to perform a conditional Landauer reset on Ω to return it to the reset state. Remarkably, since these compose to yield an overall CTO of the form (11), this process can be achieved with an infinitesimal build-up of mutual information: $I(\Omega : \mathfrak{P}) < \delta_1 + \delta_2$ for $\delta_1, \delta_2 \in \mathbb{R}^+$.

By contrast, we can also consider unconditional Landauer reset as a type of thermal operation, of a form analogous to (16). In this case, we can consider the transformation over $\Omega\mathfrak{P}$ given by:

$$\rho_\Omega \otimes \omega_{\text{in},\mathfrak{P}} \mapsto \chi_{\Omega\mathfrak{P}} \mapsto \sigma_\Omega \otimes \rho_{\text{f},\mathfrak{P}}. \tag{80}$$

Here, Ω and \mathfrak{P} are uncorrelated at the end of the transformation. As before, we can think of σ_Ω as the reset state. We can think of this pair of transformations corresponding to a decomposition of a TO of a form analogous to (16), and, *simultaneously*, the unconditional Landauer protocol in (29). As discussed in Sections 2.2.1.2 and 2.2.1.4, the fact that the final states of Ω and \mathfrak{P} are uncorrelated corresponds to an intrinsic asymmetry between the work of formation F_∞ and the extractable work F_0 ; this loss of energy represents the energy lost as a result of the expulsion of the QMI between Ω and \mathfrak{P} into the environment.

3.4.2. Transformations on Computational States and Catalytic Thermal Operations

For the sake of clarity, we explicitly restate this way of viewing computational operations, now referencing the computational subspace $\mathcal{C} = \mathcal{Q}$ and a control subspace $\mathfrak{K} = \mathfrak{P}$, combining the framework of Section 2.2.1.2 and the current section with the notation and viewpoint of Section 2.2.2.1. As discussed in Section 2.2.2.1, for a subspace \mathcal{C} bearing computational degrees of freedom, we define the reset state $\rho_{r,\mathcal{C}}$ as a standard, known reference state upon which operations can be performed. These operations correspond to known computations, which transform the state of the system from the reset state to one of N known final computational states $\rho_{\ell,\mathcal{C}}$. Then, in the Landauer protocol, we reset the final state back to the reset state, either conditioning the reset protocol on the final state or not conditioning the reset protocol on the final state. These correspond to the conditional and unconditional Landauer protocols, respectively, with the lower bound on the energy cost of each given in Sections 2.2.2.1 and 2.2.2.2. Specifically, we saw that the conditional Landauer protocol was bounded below by zero, whereas the unconditional Landauer protocol was bounded below by the amount of correlated information between the computational state and the subsystem applying the reset potential onto the computational state.

We can understand the process of repeatedly resetting the computational subspace \mathcal{C} to $\rho_{r,\mathcal{C}}$, evolving the state of \mathcal{C} to a final computational state $\rho_{\ell,\mathcal{C}}$, again resetting, again evolving to a final computational state $\rho_{m,\mathcal{C}}$ (which may be the same or different), and continuing in this fashion as a sequence of CTOs as discussed earlier in Section 3.4.1. In particular, since the local state of \mathcal{C} is constantly reset, evolved, and then reset and evolved again and again all under the influence of a secondary operator \mathfrak{K} ; we can consider \mathcal{C} as the catalyst subsystem in the sense of the discussion earlier in this section. A series of computational operations, performed by a subsystem \mathfrak{K} of \mathfrak{S} that's distinct from \mathcal{C} but contained entirely within \mathfrak{S} , transforms the state of \mathcal{C} from $\rho_{r,\mathcal{C}}$ to some $\rho_{\ell,\mathcal{C}}$, which corresponds to our computational operation. \mathfrak{K} then performs the Landauer reset $\rho_{\ell,\mathcal{C}} \mapsto \rho_{r,\mathcal{C}}$ of \mathcal{C} , following either the conditional Landauer protocol (24) or the unconditional Landauer protocol (28), with the corresponding energy costs given by (27) and (35) respectively.

A priori, it may not be clear why we insist that \mathfrak{K} must be the same subsystem that performs the local transformation $\rho_{r,\mathcal{C}} \mapsto \rho_{\ell,\mathcal{C}}$ and the Landauer reset $\rho_{\ell,\mathcal{C}} \mapsto \rho_{r,\mathcal{C}}$. Indeed, these transitions may well be performed by different machines within \mathfrak{S} . However, without loss of generality, we can lump the set of *all* machines which perform operations on \mathcal{C} collectively into a single subspace \mathfrak{K} , and subsequently examine the set of operations that \mathfrak{K} in its entirety performs on \mathcal{C} . In particular, from the decomposition of \mathfrak{S} and \mathcal{C} given in Section 2.1.1, we note that all of these individual machines, as well as their combined collection \mathfrak{K} , *must* correspond to a subspace of \mathfrak{S} .

We can use the techniques of CTOs to examine these transformations, in particular using the argument in Section 3.4.1 to examine the conditional Landauer bound. We start with \mathcal{C} in a reset state, have \mathfrak{K} perform some operations to transform it into a final computational state, and then reset \mathcal{C} to the reset state once again. As in Section 3.4.1, this chain of operations permits us to think of \mathcal{C} as the “catalyst” in a CTO, despite \mathcal{C} being the actual computational system of interest. Then, the means by which we transform from $\rho_{r,\mathcal{C}}$ to $\rho_{\ell,\mathcal{C}}$ and then back to $\rho_{r,\mathcal{C}}$ tells us whether we have a CTO of the form (11) or of the form (16); equivalently, the means by which this pair of transformations takes place tells us whether we have a conditional Landauer reset protocol (22) or an unconditional Landauer reset protocol (29).

In the case of the unconditional Landauer reset protocol, we have the pair of transformations $\rho_{r,\mathcal{C}} \otimes \sigma_{\text{in},\mathfrak{K}} \mapsto \chi_{\mathcal{C},\mathfrak{K}} \mapsto \rho_{r,\mathcal{C}} \otimes \sigma_{f,\mathfrak{K}}$ with $\text{Tr}_{\mathfrak{K}}[\chi_{\mathcal{C},\mathfrak{K}}] = \rho_{\ell,\mathcal{C}}$. This is of the exact same form as the transformation (80). As such, the same conclusion applies: in this case, this transformation corresponds to the CTO described in (16). The final correlation between \mathcal{C} and \mathfrak{K} is ejected into the environment, yielding the irreversible energy difference $F_{\infty} - F_0$. Conversely, in the case of the conditional Landauer reset protocol, we have the pair of transformations $\rho_{r,\mathcal{C}} \otimes \sigma_{\text{in},\mathfrak{K}} \mapsto \chi_{\mathcal{C},\mathfrak{K}} \mapsto \zeta_{\mathcal{C},\mathfrak{K}}$ with $\text{Tr}_{\mathfrak{K}}[\chi_{\mathcal{C},\mathfrak{K}}] = \rho_{\ell,\mathcal{C}}$ and $\text{Tr}_{\mathfrak{K}}[\zeta_{\mathcal{C},\mathfrak{K}}] = \rho_{r,\mathcal{C}}$. Here, we permit the QMI (15) between \mathcal{C} and \mathfrak{K} to build up in both transformations. As before,

the QMI in each transformation can be made as small as possible, but cannot in general be zero. The representation of the conditional Landauer reset protocol as a CTO in which the computational subsystem \mathcal{C} is thought of as the “catalyst” (inasmuch as it returns to the starting state) after two successive operations is given in Figure 13.

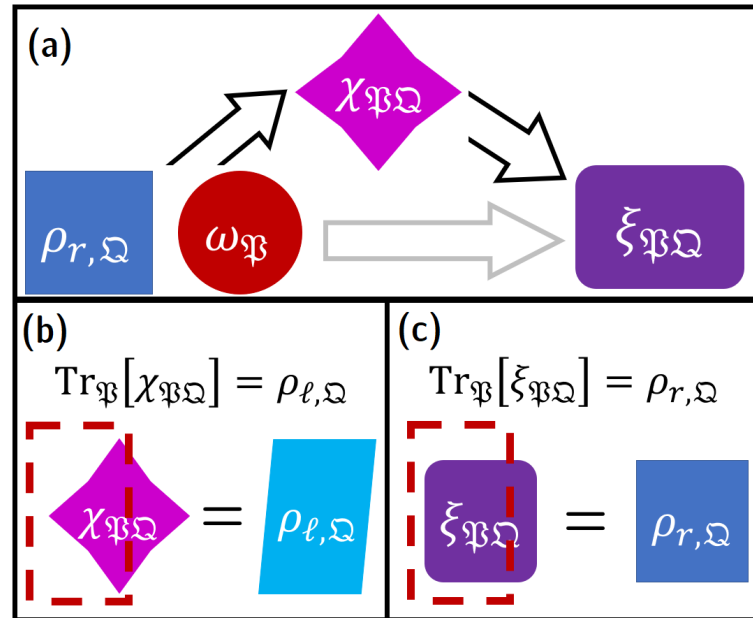


Figure 13. A representation of a computational operation cycle in terms of a CTO, with conditional Landauer reset. (a) The computational system $\mathcal{C} = \Omega$ starts in a standard reset state $\rho_{r, \mathcal{C}}$. An auxiliary system $\mathcal{R} = \Phi$ starting in the state $\sigma_{\mathcal{R}}$ performs a series of operations on $\rho_{r, \mathcal{C}}$, causing $\mathcal{C}\mathcal{R}$ to jointly evolve into a state $\chi_{\mathcal{C}\mathcal{R}}$, with $\text{Tr}_{\mathcal{R}}[\chi_{\mathcal{C}\mathcal{R}}] = \rho_{l, \mathcal{C}}$ as one of the possible known final computational states. (The series of operations that \mathcal{R} performs on \mathcal{C} corresponds to computation.) Then, \mathcal{R} performs a Landauer reset of \mathcal{C} , returning it locally to the standard reset state. In the case of the conditional Landauer reset, pictured here, the reset protocol corresponds to $\mathcal{C}\mathcal{R}$ jointly evolving into the final state $\xi_{\mathcal{C}\mathcal{R}}$, with $\text{Tr}_{\mathcal{R}}[\xi_{\mathcal{C}\mathcal{R}}] = \rho_{r, \mathcal{C}}$. The composition of operations $\rho_{r, \mathcal{C}} \otimes \sigma_{\mathcal{R}} \mapsto \chi_{\mathcal{C}\mathcal{R}} \mapsto \xi_{\mathcal{C}\mathcal{R}}$ is a composition of the form (79), and thus corresponds to a CTO of the form (11) (shown as the gray arrow). (b) Locally in Ω , the state of $\chi_{\Phi, \Omega}$ is given by one of the N final computational states $\{\rho_{l, \Omega}\}_{l=1}^N$; that is, we have $\text{Tr}_{\Phi}[\chi_{\Phi, \Omega}] = \rho_{l, \Omega}$. (c) Locally in Ω , the state of $\xi_{\Phi, \Omega}$ is once again given by the reset state $\rho_{r, \Omega}$; that is, we have $\text{Tr}_{\Phi}[\xi_{\Phi, \Omega}] = \rho_{r, \Omega}$.

3.5. Subspace Representations of Computational and Noncomputational Operations

In Section 2.2.3.2, we discussed some of the basic properties of open quantum systems with multiple asymptotic states evolving under the GKSL approximation. A key aspect of the asymptotic subspace $\text{As}(\mathcal{H}_{\mathcal{G}})$ is that the evolution after GKSL relaxation supports the *full* dynamics available to closed quantum systems. In other words, $\text{As}(\mathcal{H}_{\mathcal{G}})$ supports any dynamics that can be expressed by a Hilbert space of states evolving under a Hamiltonian; in this case, the Hamiltonian \hat{H}_{∞} governs the dynamics of $\text{As}(\mathcal{H}_{\mathcal{G}})$ after the GKSL relaxation. This provides an extremely powerful framework to represent reversible computing operations: if we can represent RC operations for closed system dynamics, we can automatically get a representation for RC operations in GKSL dynamics. As we saw in Section 2.1.2.3, the most general framework for representing a computational subsystem is with the DFS sum $\mathcal{H}_{\mathcal{G}} = \bigoplus_i \mathcal{H}_{\text{DFS}, i}$, where each DFS $\mathcal{H}_{\text{DFS}, i}$ represents a computational state c_i .

When examining the dynamics on $\mathcal{H}_{\mathcal{G}}$, this immediately gives us a way of distinguishing computational and noncomputational operations. In particular, we note that since each DFS corresponds to a specific computational state, a (non-identity) computational operation must *transfer* states from one subspace to another. Conversely, a noncomputational operation *cannot* transfer states from one subspace to another; therefore, it must only be

able at most to rearrange states within each subspace. A direct consequence of this is that noncomputational operations must commute with the DFS structure of the computational system, whereas computational operations in general have no such restrictions and permit coherences between different computational subspaces during the immediate period of computational operation. A visual representation of these two different kinds of operations is provided in Figure 14.

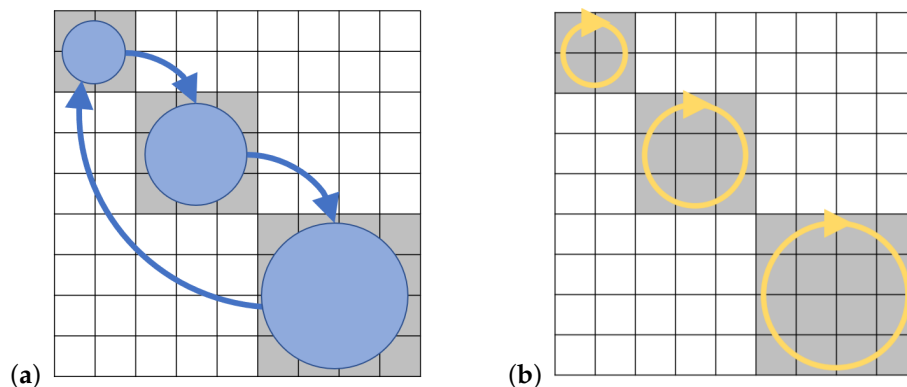


Figure 14. The representation of computational and noncomputational operations in a Hilbert space, such as $\text{As}(\mathcal{H}_{\mathcal{C}})$. (a) Computational operations (blue arrows), which transfer between different computational states (blue circles). As discussed in Section 2.1.2.2, each computational state corresponds to a distinct, orthogonal DFS (gray), with the overall Hilbert space corresponding to the direct sum of these. (b) Noncomputational operations (yellow arrows), which cannot transfer between different computational states and thus can only transfer protocomputational states within the same DFS. Note that a direct consequence of this is that noncomputational operations must commute with the DFS structure.

Our interest here is in classical computing, rather than quantum computing. As a result, we expect no quantum coherences to develop between different computational subspaces; quantum coherences may only exist within a DFS representing a single computational state. However, computational operations of the type shown in Figure 14a necessarily induce coherences; these are, indeed, characteristic of the transfer between one subspace and another. As a result, immediately after a computational operation, $\text{As}(\mathcal{H}_{\mathcal{C}})$ will appear as a single space. For our computer to remain a *classical* computer, then, we require that this space dephase into the DFS sum we expect *faster* than the computer’s ability to resolve distinct times; this is shown in Figure 15. This provides us with a dephasing timescale, which can in fact offer a way to distinguish between classical, quantum, and “approximately classical” computing representations as we tune the dephasing timescale. (Here, by “approximately classical”, we mean those operations where the dephasing timescale is on the same order as the computer resolution timescale.) The relative strength of the dephasing and computer resolution timescales, and the consequences of tuning this relative strength, will be of significant interest for future work.

It is an almost trivial statement to note that any matrix can be represented as the sum of two other matrices. However, the distinction between computational and noncomputational operations discussed above, and the point that a computational operation necessarily mixes the different computational DFSs to temporarily make a single large space, provides us with an interesting decomposition of such operations. In particular, we note that a computational operation which appears as an operator mixing all of the states in $\text{As}(\mathcal{H}_{\mathcal{C}})$ can be decomposed into a ‘noncomputational part’, which commutes with the DFS structure, and a ‘pure computational part’, which contains all of the information involving transfer of states between DFS blocks (and thus, *all* of the information regarding the actual computational content of the operation); this decomposition is shown in Figure 16. A central property of GKSL dynamics with multiple asymptotic states, derived in [18] and discussed there and in [19], is the nontrivial quantum geometric tensor over

$As(\mathcal{H}_{\ominus})$ that emerges, and the dependence that the dynamics exhibits on the QGT. Notably, different shapes of the asymptotic subspace exhibit different geometric signatures; thus, computational operations which mix different DFS states will have a different quantum geometric signature than noncomputational operations. In light of the decomposition of computational operations into noncomputational and pure computational operations, this also means that each of these parts will exhibit distinct quantum geometric signatures which can identify the computational and noncomputational part. As an added benefit, we expect this decomposition to provide additional intuition for the distinction mentioned above between Landauer’s Principle and the Fundamental Theorem. The discussion of this will be provided in much greater detail in the forthcoming work examining the quantum geometric properties of RC operations in general.




Figure 15. The dephasing of $As(\mathcal{H}_{\ominus})$ from a single subspace to a DFS direct sum. For classical reversible operations, this dephasing must occur faster than the computer’s ability to resolve distinct times (i.e., on a timescale faster than the computer can see). Classical and quantum computing operations can be distinguished by the relation between this timescale and the computer resolution timescale.




Figure 16. Decomposition of a computational operation (left) into a noncomputational part (middle), which commutes with the DFS structure that distinguishes different computational states, and a pure computational part (right), which contains *all* of the information regarding the transfer of states between different DFS blocks, and thus all of the information regarding the computational part of the operation. Notably, because the quantum geometric tensor of $As(\mathcal{H}_{\ominus})$ as a single space has a different shape than that of $As(\mathcal{H}_{\ominus})$ as a DFS sum, the QGT of each of these will naturally be distinct as well. As such, the noncomputational part of a computational operation will have a different, measurable quantum geometric signature to the pure computational part of a computational operation.

Beyond simply distinguishing between computational and noncomputational operations, the discussion in Section 2.1.3 highlights how essential it is to distinguish between the different types of computational operations—deterministic irreversible, deterministic reversible, stochastic irreversible, and stochastic reversible. As further discussed in Section 2.1.4, these types of operations are themselves comprised of a set of primitive computational state transitions; namely the bijections, merges, and splits as discussed in Figure 5. As such, in order to understand the representations of the different types of reversible (and, indeed, irreversible) classical operations, we must first find a representation

of these different types of operations. Although the nature of bijections and merges is somewhat self-evident, the case of splits must be handled with slightly more care, since they can be generally expected to result in temporary coherences (which will, typically, quickly decohere). As with the decomposition of computational operations, these will in all likelihood exhibit distinct quantum geometric signatures in lieu of the quantum geometric signatures of different asymptotic subspaces discussed in [18,19]. Along with the previously-mentioned issues regarding quantum geometric signatures of RC operations, the discussion of this will be addressed in the forthcoming work centering on the quantum geometric properties of reversible computing operations more broadly.

4. Discussion

4.1. Essential Consistency of the Classic RC Formulation with NEQT

An important high-level conclusion supported by this paper is that there is no inconsistency between the simple, classic formulation of Landauer's Principle and reversible computing that we reviewed in Sections 3.2 and 3.3, and a more detailed treatment based on NEQT. Indeed, no such inconsistency is possible, since, as we showed, the classic formulation can be presented in a form that makes no equilibrium assumptions whatsoever. The only assumptions we made there were the fundamental unitarity of the underlying quantum evolution, which is also assumed by all of quantum thermodynamics, and the treatment of the environment as immediately thermalizing all ejected information, which is equivalent to the Markov assumption underlying GKSL dynamics as discussed in Section 3.5. Therefore, the more detailed NEQT formulations *cannot* negate the basic results of Sections 3.2 and 3.3. Indeed, we showed how to draw explicit correspondences between a more detailed treatment of classical reversible computing based on generalized CTOs and multiple asymptotic states in GKSL, and the simpler model of Section 2.1 on which the basic results of Sections 3.2 and 3.3 can be based. We next discuss a few specific aspects of this correspondence in more detail.

4.2. CTO Representations of Reversible Computing and System Boundaries

In Section 3.4, we discussed the representation of quantum mechanical models of reversible computing in terms of the general CTO (11). Specifically, we identified the Landauer reset potential(s) $\hat{V}_{lr,\Omega}$ applied to a subsystem Ω as coming from the interaction between Ω and a distinct separate subsystem \mathfrak{P} , where Ω and \mathfrak{P} are both subsystems of the overall system \mathfrak{S} we are examining in our TO. This analysis reinforces the importance of properly drawing the system boundaries, as discussed in [61].

The specifics of where we choose to draw the boundaries between each subsystem, or between the overall system and the environment, plays a vital role in being able to properly identify the reset protocol of a system bearing computational degrees of freedom. Properly drawing the boundaries makes the distinction as to whether a given reset protocol is a *Landauer* reset protocol *per se* or not. In [61], we have the example of the reset of a system \mathfrak{P} coupled to a copy/referent subsystem Ω which stores the same computational state as \mathfrak{P} before reset. The set of transformations on \mathfrak{S} in this case are given by $\{\rho_{in,\mathfrak{P}} \otimes \sigma_{in,\Omega}\}_{i=1}^N \mapsto \{\rho_{r,\mathfrak{P}} \otimes \sigma_{i,\Omega}\}_{i=1}^N$. Comparing these to expressions (22), (23), and (29), we see that identifying \mathfrak{P} as the reset system is commensurate with the Landauer reset definitions, whereas identifying $\mathfrak{P}\Omega$ as the reset system does not count as Landauer reset. (Specifically, identifying \mathfrak{P} as the reset system, this corresponds to conditional Landauer reset.)

In precisely the same way, identifying the system boundaries is essential to representing reset processes as CTOs. The CTO expression (11) requires a specific shape for the starting and ending states (namely, that the overall system \mathfrak{S} be subdivided into a subsystem \mathfrak{P} which locally undergoes the state transformation $\rho_{in,\mathfrak{P}} \mapsto \rho_{f,\mathfrak{P}}$ and a subsystem Ω which locally returns to the same state). Thus, properly identifying the system boundaries plays a vital role in identifying the effect of the global universal CPTP map as a catalytic thermal operation, or some other kind of operation. This dependence on system boundaries

is, indeed, a *general* feature of resource theories [25]. Thus, when applying the results of RTQT to analyzing Landauer reset protocols and formulating bounds on quantities of interest, it is vital to specify the subsystem boundaries properly: these boundaries affect whether or not we can properly classify a reset as a Landauer reset, whether or not we can properly classify a CPTP map as a CTO, and whether or not we can properly relate these two. (Incidentally, we saw a specific example of this earlier: in Section 3.4, this dependence is what constrained the reset potential(s) to be implemented by a separate catalyst subsystem within \mathfrak{S} as a whole.)

4.3. Applicability of the Markov Approximation to Reversible Computing

As mentioned in Section 2.2.3.1, the Markov approximation (and, especially, the Markov approximation for systems with multiple asymptotic states) involves the separation of several different time scales, given by (62) and (72) respectively. As previously mentioned, a direct consequence of this approximation is that fluctuations that take information from \mathfrak{S} to \mathfrak{E} during an intermediate time period and return it to \mathfrak{S} happen much faster than our ability to resolve the dynamics; instead, any information that is ejected from the system to the environment cannot be returned to the system. We might be concerned that this might represent an “artificial” type of information loss in our model that comes from imposing a limitation on the types of processes we consider, which does not reflect computational systems in the real world. In fact, the opposite is true: this assumption matches perfectly with our understanding of the thermalization of ejected correlated information from \mathfrak{S} .

In Section 3.2 and [11,12], we saw that the entropy increase due to Landauer reset was a consequence of the thermalization of mutual information ejected into the environment. In particular, the thermalization of entropy ejected into the environment occurs at a time frame much faster than our ability to capture the dynamics of the system. Thus, in our model, the role of fluctuations between \mathfrak{S} and \mathfrak{E} originating from \mathfrak{S} is suppressed: the effect of perturbative fluctuations looks identical to the effect of environment perturbations. This is exactly the kind of behavior we expect for a system \mathfrak{S} bearing computational degrees of freedom in a larger open system evolution: the system has no practical way of tracking mutual information that is transferred to the environment at a given time and then is transferred from the environment back to the system at a later time. Any perturbations arising from information that originated from \mathfrak{S} at an earlier time appears to \mathfrak{S} as indistinguishable from perturbations due to \mathfrak{E} . This is also true for models of computation where the computational degrees of freedom reside in a subsystem \mathfrak{P} being acted on by one or more orthogonal subsystems \mathfrak{Q}_i of \mathfrak{S} : a crucial part of the model is that neither \mathfrak{P} nor any of the \mathfrak{Q}_i s are able to track information after it has been ejected into \mathfrak{E} ; and, indeed, this is a completely physically realistic framework. Thus, the Markov assumption, in addition to being a vital calculational tool to retrieve closed form expressions, also represents the real-world dynamics of the system.

4.4. Relationship to the Stochastic Thermodynamics of Computation

The *stochastic thermodynamics of computation* [10,130] is a framework for examining the entropy cost of classical computational systems that has gained substantial prominence in recent years. As such, an important question is the relationship between this framework and the NEQT framework and results given in Sections 2.2.1.2–2.2.2.2, which we discuss here.

In the stochastic thermodynamics of information, the thermodynamic properties of classical computational systems are examined [10] from the point of view of purely classical information theory, relying on the properties of the continuous probability distributions of classical random variables that arise therein. In this approach, classical computation is represented as a continuous-time Markov chain (CTMC), represented by a directed acyclic graph (DAG) which in turn represents a set of functions over a given Boolean string $s \in \{0,1\}^n$ of length n . Then, in the stochastic thermodynamics of information framework, the entropy cost of transitions is calculated by examining the difference in the classical relative entropies (i.e., the Kullback-Leibler divergences) of the CTMC distribution

relative to an arbitrary distribution before and after a single set of state evolutions along the graph. This framework explicitly does not consider the perspective from quantum information theory,³⁸ in contrast with the representation of classical computation in terms of quantum channels (e.g., as in [1–6,8,31–34]). Instead, the thermodynamic properties of computation are derived solely from examining the entropy production and entropy flow rates of probability distributions which evolve under this CTMC representation.

Despite avoiding the quantum information representation for classical computation (and indeed for nonequilibrium dynamics), we would expect *no* differences whatsoever in the conclusions we get from this framework compared to the framework focusing on quantum information theory. At its core, the technique of correlation engineering [11,15] relies on performing operations on correlated systems in such a way as to minimize the mutual information build-up and to make the Helmholtz free energy cost of the transition arbitrarily small, as discussed in Section 2.2.1. Although [11,15] and the discussion in Section 2.2.1 have focused on quantities such as the quantum mutual information, the principles of correlation engineering are in fact completely oblivious to whether the correlations and mutual information quantities are classical or quantum in nature. Indeed, since the examination in these works of the effect of correlations on overall entropic cost do *not* rely on whether or not the system state commutes with the thermal reference state or the Hamiltonian, we can freely substitute classical information quantities (such as the Kullback-Leibler divergence) and retrieve valid statements for correlation engineering of classical thermodynamic systems.³⁹

For individual systems in isolation, this is precisely what we find: the stochastic thermodynamics of computation is completely able to reproduce the results regarding the distinctions between logical and thermodynamic reversibility found earlier in [2,3,7,132]. Specifically, the now-famous result in those works that logically irreversible operations (such as erasure of a single, isolated bit or computational system) can be nevertheless thermodynamically reversible has also been reproduced in [10,130]. This reflects our exact point about the proper interpretation of Landauer’s principle that we elaborated upon in Section 3.2, which is that erasing bits in isolation when they contain mutual information with other bits results in thermodynamic irreversibility, and an amount of entropy increase corresponding to the loss of mutual information.

Nevertheless, the subject of extending the framework of the stochastic thermodynamics of computation to *correlated* systems appears to remain a matter of controversy. In particular, the example of the thermodynamic reversibility of erasure of a single bit *in isolation* has been used to argue that a *correlated* system cannot realize operations that are both logically and thermodynamically reversible to within a mutual information difference of $\delta \in \mathbb{R}^+$, as given in (15). Here, it appears that the misunderstanding has persisted due to confusing the matter of isolated versus correlated systems: the example in [2,3,132] and reproduced in [10,130] continues to focus on the case of an *isolated* system, whereas [11,13,15,61] highlight that correlation engineering relies entirely on reducing the correlation buildup between systems and applying operations in a way that does not increase the net entropy flow from the overall system.⁴⁰

It must be emphasized again that there is no fundamental disagreement between the results calculated in [11,13,15,61]. Indeed, there *cannot* be, since all of these paths start with precisely the same set of fundamental assumptions about the underlying dynamics of the system; that is, that the dynamics are described by some CPTP map which we can express in terms of a larger subspace as per the Stinespring dilation theorem. Thus, although some

³⁸ In the words of [10] directly, “[this paper] will not consider the thermodynamics of quantum mechanical information processing, or indeed any aspect of quantum computation.”

³⁹ Indeed, this precise property is the insight underlying the β -ordering and thermomajorization curve technique in [47].

⁴⁰ Somewhat more confusingly, §IX-B in [10] appears to use this example to claim that correlated systems cannot realize such operations; however, Example 6 in §V-B in [10] appears to use a general argument to verify that logically irreversible operations on *correlated* systems cannot be made thermodynamically reversible when each irreversible operation is at the subsystem level. This is precisely the same principle used in [11,13,15,61]. The source of the internal disagreement in [10] remains unclear.

of the *conclusions* in [10,130] seem to disagree with the consensus of [11,13,15,61,133], this is only a matter of mistaken identity: [10,130] has erroneously drawn conclusions about *correlated* systems from a valid calculation about the thermodynamic reversibility of a logically irreversible process applied to an *independent* system (see also footnote 37). Furthermore, although some of the conclusions of [10,130] may be drawn from a mistaken conflation between the properties of independent and correlated systems, it is worth re-emphasizing that the underlying techniques in these works remain completely valid. Indeed, the stochastic thermodynamics of computation may nevertheless serve as a useful additional tool to examine correlation engineering of correlated systems. (This further reinforces how delicate the issue of system boundaries is, as discussed in [61].)

4.5. Thermodynamically Reversible Transformations of Extended Systems

An appropriate caveat to Landauer's principle could be to mention that even a *correlated* state in a computer could in principle be erased in a thermodynamically reversible way if, rather than erasing bits in isolation, we instead considered the full space of *all possible* thermodynamic transformations, but what, in detail, would a concrete protocol for accomplishing such a transformation look like? One such protocol would be to simply *unwind the correlations* by running, in reverse, a reversible computation that could have computed those correlations in the first place, thereby leaving us, potentially, with an array of independently random input bits, that could then each be thermodynamically reversibly erased in isolation from each other. However, note that the existence of protocols such as this one does not refute the need for reversible computing, since this protocol itself uses reversible computing to begin with.

Fundamentally, if we wish to be able to construct complex computational processes by composing them out of primitive transformations which operate *locally* on part of the full computational state, then by definition, those primitives *cannot* operate monolithically on the state of the entire extended system, and thus they *must* be logically reversible if we wish to avoid entropy increase in the face of the kind of non-local correlations which naturally arise, as a matter of course, in typical large computations. This is not to say that some radically different alternative basis for computation that was *not* based on local computational primitives *at all* would necessarily be impossible to develop, but until that has been done, reversible computing remains the most promising and well-developed available avenue towards performing general digital computing in a thermodynamically efficient way.

That being said, exploration of alternative protocols for thermodynamically efficient computing may someday also be worthwhile. One conceptual example of such an alternative is illustrated by [134],⁴¹ in which a chaotic dynamical system whose strange attractor is engineered to encode the state of an extended Boolean circuit is *monolithically* transformed adiabatically from an old state to a new state of the *entire circuit* all at once. This example illustrates that thermodynamically efficient alternatives to performing computations via a sequence of *local* logically reversible transformations *do* in fact exist. This example even still preserves a kind of compositionality, albeit at a different level, in the sense that the *structure* of the extended system is still composed out of local interaction Hamiltonian terms representing individual Boolean logical constraints, even though the *transformation* of the system is done monolithically.

So, the existence of the example of [134] could be cited as partial vindication of assertions that reversible computing via *local* transformations is not *necessarily* the only way to accomplish digital computation in a thermodynamically efficient way. However, the results of [134] appear to suggest that the penalty for doing *non-local* thermodynamically reversible transformations of digital machines is generally to incur an exponential increase in time complexity—which, in retrospect, is not surprising, since otherwise, one could use methods similar [134] to solve NP-complete problems using only polynomial resources

⁴¹ Detailed results of the study previewed in [134] may be found in the presentation notes available at <https://tinyurl.com/Frank-DeBenedictis-16>.

(by constraining the output of the computation, rather than the input). So, it may still be the case that reversible computing remains the only physically possible way to achieve thermodynamically efficient digital computation with only modest (i.e., polynomial) resource overheads.

We should also note in passing that, apart from [134], many other, more “analog” approaches to physical computing with dynamical systems also exist; see Section 4.2 of [135] for a survey. As with [134], the time-evolution of other, more analog kinds of conservative dynamical systems that may be useful for computing could also conceivably be engineered to approach thermodynamic reversibility, although most existing analog computing schemes have not been specifically designed to do so.

4.6. Future Directions

Although providing a valuable framework both theoretically and for our purposes in modeling reversible computation, the structure of GKSL systems with multiple asymptotic states [17–19] leaves several questions remaining. These questions are of great interest to the theory of open quantum systems generally, and also offer substantial insights for open quantum systems models of reversible computing. One important question is the question of energy dissipation for generic nonequilibrium protocols. The notion of thermodynamic length has been developed [136] for GKSL systems with *single* asymptotic states. Minimization of this length provides a characterization of the minimal dissipation of the time evolution of a Hamiltonian in an open quantum system. Meanwhile, general thermodynamic uncertainty relations have also been developed [137] for single asymptotic states. These provide a general set of uncertainty relations between the currents of a system in an asymptotic state and the entropy production rate of the system, in terms of the information geometric metric (i.e., the Fisher information) on nonequilibrium asymptotic states.

Extending these notions to multiple asymptotic states can help us develop expressions of energy dissipation for time evolving systems for multiple asymptotic states. In order to characterize the efficiency of reversible information processing operations, a key figure of merit that we are interested in is the dissipation as a function of delay, $D(d)$. For reversible computing in an open quantum system, multiple asymptotic state framework, we want to be able to characterize the minimal dissipation for any Hamiltonian we might want to write that can represent the model we are using for a reversible computer.⁴² This will, in general, be a function of the amount of time (i.e., the *delay*) of the operation.

Fundamentally, both the dissipation and delay will depend quite fundamentally and intrinsically on the underlying geometry of the asymptotic states. From the derivation of an adiabatic (i.e., Berry-like) curvature on the space induced by asymptotic states [18,19], we can immediately anticipate that the expression for thermodynamic length in open quantum systems will depend intricately on this induced curvature. Given that quantum speeds for GKSL systems are expressed through a suitable metric on the information geometry of states [29], we can expect that the delay can be derived similarly.⁴³ By relating the information geometry metric tensor with the QGT, we can derive expressions for the delay and in turn the dissipation in terms of the QGT. This can then provide us with dissipation as a function of delay.

The development of $D(d)$ will involve as intermediate steps several quantities which will be of interest to the open quantum systems community as a whole, such as the thermodynamic length and the thermodynamic uncertainty relations for GKSL systems with multiple asymptotic states. These will be essential for understanding properties

⁴² This is not the only energetic cost of interest when examining models of computation: we may also be interested in the minimum energy required to perform a computation [138] or the *maximum* information cost that an operation can take [139]. Since we intend to develop our expression for $D(d)$ for classical RC operations using NEQT (and in particular its quantum information formulation), the expression for $D(d)$ will serve in concert with these other energetic costs to provide a strong characterization of the energetic constraints of classical reversible and quantum computations.

⁴³ To elaborate slightly, note that the dissipation-delay relation is not necessarily directly bounded by the quantum speed limit, which is defined in terms of dynamical energy *invested*, not energy *dissipated*; however, we can expect that the derivation of the delay will still involve considerations of the speed limit, to the extent that dissipation can be bounded as a fraction of the dynamical energy.

of open quantum systems which have direct bearing on classical reversible computing operations. The bearing that these properties will have on RC operations *per se* also depends on the particulars of how RC operations are represented in open quantum systems. We have here provided some of the initial groundwork for these representations, but several specific details which are vital for understanding RC operations and the application of open quantum system properties to RC remain to be developed in full. Here, we have identified several key remaining details—specifically, the quantum geometric signature that different RC operations will have, the representations of merges and splits in the framework of GKSL dynamics with multiple asymptotic states, and the specific timescale under which the result of a computational operation dephases to retrieve the DFS sum structure characteristic of classical RC operations. These details are currently being developed in a forthcoming work, but they are likely not an exhaustive list of the specific characteristics of open quantum systems representations of RC operations.

5. Conclusions

At this time, much work remains to complete the task of fully fleshing out a useful physical theory of classical reversible computing based on the tools of modern quantum thermodynamics and quantum information. Our goal, in this article, was to lay some key conceptual foundations for that effort, point the way towards further progress that can be made, and review some important preliminary results.

Our primary conclusion thus far, from this line of work, is that the core insights from the classic theory of the thermodynamics of computing which originally motivated the field of reversible computing, rather than being contradicted by the modern non-equilibrium quantum thermodynamics perspective, are, to the contrary, supported by it. We argued that the most appropriate understanding of Landauer's Principle is as a statement about the absolute *increase* in total entropy that is required when correlated information is lost, and reviewed the theorems showing that *only* computational operations that are (fully or conditionally) logically reversible can avoid such increases, when applied to subsystems that exhibit correlations to other subsystems, as is normal for subsystems that contain computed information. In addition, we provided a complementary way of seeing conditional and unconditional Landauer reset in terms of catalytic thermal operations, which helps shed some light on the underlying nonequilibrium quantum thermodynamic principles at play distinguishing the reset processes.

Even in its early stages of development, the GKSL dynamical perspective on computational operations in open quantum systems is already showing quite surprising implications, of interest both on a purely theoretical level as well as for applications to reversible computing models. In particular, we see that the quantum geometric properties of the space supporting RC operations play a central role in governing these operations and the dynamics of systems supporting reversible computation. This offers a tantalizing glimpse into the rich geometrical structure which underpins and can help support RC operations, and suggests that the discovery of these signatures *generally* can support reversible computational operations in exotic structures. Much more work remains to be done in teasing out the geometric structure of RC operations, with implications of substantial interest both to reversible computer engineering and the theory of open quantum systems generally.

We should note, in passing, that there is a useful concept of *effective* physical entropy that can be considered to include not just the statistical form of entropy considered here, but also measures of information complexity or *algorithmic randomness* [140]. One can summarize such concepts by saying that, for pragmatic purposes, physical entropy effectively includes both *unknown* and *known but incompressible* information. However, such expanded conceptions of entropy do not affect the particular concerns of this paper, since algorithmically random, incompressible information remains equally incompressible regardless of whether the use of logically reversible algorithms is considered.

In conclusion, we see that there is a potentially enormous long-term practical value to be gained from seriously studying the limits and potentialities of physical mechanisms

designed to efficiently implement reversible computational processes, with an eye towards making technologies for general digital computing far more efficient. In this article, we have reviewed a number of key theoretical tools from modern non-equilibrium quantum thermodynamics which we believe will be useful for continuing this line of work and investigating the physics of reversible computing in more depth. We intend to continue this effort in future papers, and we invite other interested researchers to join us.

Author Contributions: Conceptualization, M.P.F. and K.S.; Formal analysis, M.P.F. and K.S.; Funding acquisition, M.P.F. and K.S.; Investigation, M.P.F. and K.S.; Methodology, M.P.F. and K.S.; Project administration, M.P.F.; Supervision, M.P.F.; Visualization, M.P.F. and K.S.; Writing—original draft, M.P.F. and K.S.; Writing—review & editing, M.P.F. and K.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded in part by the Laboratory Directed Research and Development (LDRD) and Advanced Simulation and Computing (ASC) programs at Sandia National Laboratories, a multimission laboratory managed and operation by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. It was also supported in part by the U.S. Army Research Office (ARO) under cooperative agreement W911NF-14-2-0075 and BAA W911NF-19-S-0007, and in part by the U.S. Air Force Office of Scientific Research (AFOSR) under grant number FA9550-19-1-0355. This document describes objective technical results and analysis. Any subjective views or opinions that might be expressed in this document do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Approved for public release, SAND2021-6489 J.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: The authors would like to thank Victor Albert, Neal Anderson, Gavin Crooks, Ed Fredkin, John Goold, Giacomo Guarnieri, David Guéry-Odelin, Norm Margolus, Markus Müller, Kevin Osborn, Subhash Pidaparathi, Greg Snider, David Wolpert, and Noboyuki Yoshikawa for helpful discussions. M.P.F. would also like to thank Rudro Biswas, Robert Brocato, Erik DeBenedictis, Rupert Lewis, Nancy Missert, and Brian Tierney for their contributions to the reversible computer engineering efforts at Sandia, and Gladys Eden for her love and encouragement. K.S. would also like to thank Hannah Watson for her boundless love and emotional support, and Jimmy Xu for his exceptional moral support.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AQFP	Adiabatic quantum flux parametron
ASIC	Application-specific integrated circuit
BARC(S)	Ballistic asynchronous reversible computing (in superconductors)
CMOS	Complementary metal-oxide-semiconductor (circuit/technology)
CPTP	Completely positive trace-preserving (map/channel)
CTMC	Continuous-time Markov chain
CTO	Catalytic thermal operation(s)
DAG	Directed acyclic graph
DPI	Data processing inequality
FET	Field-effect transistor
GKSL	Gorini-Kossakowski-Sudarshan-Lindblad (operator/theory)
LvN	Liouville-von Neumann (equation)
NEQT	Non-equilibrium quantum thermodynamics
nFET	n-type FET

PES	Potential energy surface
QC	Quantum computation
QCA	Quantum-dot cellular automaton
QRD	Quantum relative divergence
QRT	Quantum resource theory
QGT	Quantum geometric tensor
RA-CMOS	Reversible adiabatic CMOS
RC	(Classical) reversible computing
RNRL	Reversible nanomechanical rod logic
R-QCA	Reversible QCA
RRE	Relative Rényi entropy
RTQT	Resource theory of quantum thermodynamics
RQFP	Reversible quantum flux parametron
SPICE	Simulation Program with Integrated Circuit Emphasis
TO	Thermal operation(s)

Appendix A. Minimum-Energy Scaling for Classical Adiabatic Technologies

In this appendix, we briefly present the derivation for the scaling of minimum energy dissipation for reversible technologies such as RA-CMOS (Section 2.3.1) that obey classic adiabatic scaling and that can be characterized in terms of relaxation and equilibration timescales.⁴⁴

First, we assume (as is the case for “perfectly adiabatic” technologies such as [48]) that the total energy dissipation per clock cycle E_{diss} in a reversible circuit can be expressed as a sum of *switching losses* and *leakage losses*,

$$E_{\text{diss}} = E_{\text{sw}} + E_{\text{lk}}, \quad (\text{A1})$$

and further, that switching and leakage losses depend on the signal energy E_{sig} and transition time t_{tr} approximately as follows:

$$E_{\text{sw}} \simeq E_{\text{sig}} \cdot c_{\text{sw}} \cdot \frac{\tau_r}{t_{\text{tr}}}, \quad (\text{A2})$$

$$E_{\text{lk}} \simeq E_{\text{sig}} \cdot c_{\text{lk}} \cdot \frac{t_{\text{tr}}}{\tau_e}, \quad (\text{A3})$$

where τ_r, τ_e are the relaxation and equilibration timescales, respectively, and $c_{\text{sw}}, c_{\text{lk}}$ are small dimensionless constants characteristic of a particular reversible circuit in a specific family of technologies, such as [48]. In practice, although these specific formulas are only approximate, they approach exactness in the regime $\tau_r \ll t_{\text{tr}} \ll \tau_e$.

Then, now treating (A2), (A3) as exact, we can write:

$$E_{\text{diss}} = E_{\text{sig}} \left(c_{\text{sw}} \tau_r \cdot \frac{1}{t_{\text{tr}}} + \frac{c_{\text{lk}}}{\tau_e} \cdot t_{\text{tr}} \right). \quad (\text{A4})$$

We can collect the constants, absorbing them into adjusted timescales $\tau_r' = c_{\text{sw}} \tau_r$ and $\tau_e' = \tau_e / c_{\text{lk}}$, so

$$E_{\text{diss}} = E_{\text{sig}} \left(\tau_r' \cdot \frac{1}{t_{\text{tr}}} + \frac{1}{\tau_e'} \cdot t_{\text{tr}} \right). \quad (\text{A5})$$

Setting the derivative of (A5) with respect to t_{tr} equal to zero, we find that E_{diss} is minimized when

$$\tau_r' \frac{1}{t_{\text{tr}}^2} = \frac{1}{\tau_e'}, \quad (\text{A6})$$

⁴⁴ Note that this particular scaling analysis does not extend to families of technologies that may potentially offer some approximation to a Landau-Zener type of *exponential* quantum adiabatic scaling, such as R-QCA (see Section 2.3.3).

or in other words, when

$$t_{tr} = \sqrt{\tau_r' \tau_e'} \tag{A7}$$

at which point E_{sw} and E_{lk} are equal. The minimum energy dissipation per cycle is then

$$E_{diss} = 2E_{sig} \sqrt{\frac{\tau_r'}{\tau_e'}} \tag{A8}$$

Thus, for any given reversible circuit design in a family of technologies with given values of the constants c_{sw}, c_{lk} , in order for E_{diss} to approach 0 as the technology develops, we must have that the ratio of equilibration/relaxation timescales $\tau_e/\tau_r \rightarrow \infty$, and, if the relaxation timescale τ_r is fixed, this implies that also the (minimum-energy) value of the transition time $t_{tr} \rightarrow \infty$. These requirements were mentioned in Section 2.3.1.

More specifically, in order to increase the peak energy efficiency of a reversible circuit by a factor of $N\times$, in a given family of technologies obeying classic adiabatic scaling, this requires that the timescale ratio τ_e/τ_r must be increased by $N^2\times$, and (assuming τ_r is fixed) the transition time t_{tr} for minimum energy will increase by $N\times$.

Appendix B. Vectorization of the Operator Algebra on Quantum States

In ordinary quantum mechanics, expressions such as (57), (61), and (66) can be easily solved using operator algebra techniques. This same principle holds for operators which operate on other operators (e.g., $\hat{\mathcal{L}}$ operates on density matrices $\rho \in \mathcal{D}(\mathcal{H}_{\mathcal{G}})$), known as *superoperators*. Since the space of L^2 -bounded operators $\mathcal{B}(\mathcal{H}_{\mathcal{G}})$ forms a Hilbert space in its own right under the Hilbert-Schmidt inner product $\langle \hat{A}, \hat{B} \rangle := \text{Tr}[\hat{A}^\dagger \hat{B}]$, the exact same operator algebra techniques can be applied to examining superoperators.

A simple way of explicitly writing down these techniques for the L^2 -bounded operators is using the process of *vectorization*, which we very briefly discuss here following the excellent presentation in [19]. Succinctly, vectorization is the process of rewriting matrices in $\mathbb{F}^{m \times n}$ as vectors in \mathbb{F}^{mn} via the mapping $|v\rangle\langle w| \mapsto |v\rangle \otimes |w\rangle$. (Here, \mathbb{F} is the field (\mathbb{R} or \mathbb{C}) that the matrices live over.) In terms of a basis $\{|b_i\rangle\}_{i=1}^N$ of $\mathcal{H}_{\mathcal{G}}$, the vectorization of an operator $\hat{A} \in \mathcal{B}(\mathcal{H}_{\mathcal{G}})$ appears as:

$$\hat{A} = \sum_i c_i |b_i\rangle\langle b_i| \mapsto |A\rangle\rangle := \bigotimes_i c_i |b_i\rangle. \tag{A9}$$

As a concrete example, the vectorization of a 2×2 matrix appears as:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \mapsto \begin{pmatrix} a \\ c \\ b \\ d \end{pmatrix}.$$

The “double-ket” notation for vectorized matrices acts largely the same as the familiar Dirac notation:⁴⁵

- Superoperators $\hat{\mathcal{O}}$ act on operators \hat{A} as $\hat{\mathcal{O}}|\hat{A}\rangle\rangle = |\mathcal{O}[\hat{A}]\rangle\rangle$.
- The Hermitian adjoint of $|\hat{A}\rangle\rangle$ is $\langle\langle\hat{A}|$.
- The Hermitian adjoint of $\hat{\mathcal{O}}|\hat{A}\rangle\rangle$ is given by $\langle\langle\hat{A}|\hat{\mathcal{O}}^\dagger = \langle\langle\hat{\mathcal{O}}^\dagger[\hat{A}]|$.
- The Hilbert-Schmidt inner product is given by $\langle\langle\hat{A}|\hat{B}\rangle\rangle := \text{Tr}[\hat{A}^\dagger \hat{B}]$.

⁴⁵ Although our intuitions from Dirac notation directly carry over to the “double-ket” notation, translating back and forth from the vectorized expressions to the operator expressions can be somewhat nontrivial. Care must be taken when doing so, although discussing these difficulties is beyond the scope of this paper.

- Thus, the trace of \hat{A} is given by $\langle\langle \mathbb{1} | \hat{A} \rangle\rangle$.
- The basis $\{|b_i\rangle\}_{i=1}^N$ of $\mathcal{B}(\mathcal{H}_{\mathfrak{E}})$ gives a corresponding basis of $\mathcal{B}(\mathcal{H}_{\mathfrak{E}})$: $|b_{ij}\rangle\rangle = |b_i\rangle\langle b_j|$.
- From this structure, changing the basis of $\mathcal{H}_{\mathfrak{E}}$ changes the basis of $\mathcal{B}(\mathcal{H}_{\mathfrak{E}})$, and thus, the explicit decompositions of the vectorized operators $|\hat{A}\rangle\rangle \in \mathcal{B}(\mathcal{H}_{\mathfrak{E}})$ and the superoperators $\hat{\mathcal{O}}$.

However, the basis change in $\mathcal{H}_{\mathfrak{E}}$ directly reflects a basis change in $\mathcal{B}(\mathcal{H}_{\mathfrak{E}})$: transforming $|b_i\rangle \mapsto |c_i\rangle$ directly corresponds to the transformation $|b_{ij}\rangle\rangle = |b_i\rangle\langle b_j| \mapsto |c_{ij}\rangle\rangle = |c_i\rangle\langle c_j|$. Thus, we don't need any "extra" information in the transformation: everything can be expressed entirely in terms of what lives in $\mathcal{B}(\mathcal{H}_{\mathfrak{E}})$, without needing to further reference $\mathcal{H}_{\mathfrak{E}}$.

- An additional complication that does not appear with ordinary Hilbert spaces is the operator algebra structure $\mathcal{B}(\mathcal{H}_{\mathfrak{E}}) \times \mathcal{B}(\mathcal{H}_{\mathfrak{E}}) \rightarrow \mathcal{B}(\mathcal{H}_{\mathfrak{E}})$; thus, we need to describe the vectorized version of matrix multiplication. Explicitly, the vectorized product of the operators \hat{A} , \hat{B} , and \hat{C} is $|\hat{A}\hat{B}\hat{C}\rangle\rangle = (\hat{C}^\top \otimes \hat{B})|\hat{A}\rangle\rangle$.

In vectorized form, the GKSL Equation (61) appears as:

$$\hat{\mathcal{L}}|\rho_{\mathfrak{E}}\rangle\rangle = \left[-i(\mathbb{1} \otimes \hat{H}_{\mathfrak{E}} - \hat{H}_{\mathfrak{E}}^\top \otimes \mathbb{1}) + \frac{1}{2} \sum_{a,b>0} \kappa_{ab} \left(2\hat{F}_{ab}^* \otimes \hat{F}_{ab} - \mathbb{1} \otimes \hat{F}_{ab}^\top \hat{F}_{ab} - (\hat{F}_{ab}^\top \hat{F}_{ab})^\top \otimes \mathbb{1} \right) \right] |\rho_{\mathfrak{E}}\rangle\rangle. \tag{A10}$$

Meanwhile, the vectorized form of the adjoint GKSL Equation (64) is given by:

$$\hat{\mathcal{L}}^\dagger|\hat{A}\rangle\rangle = \left[i(\mathbb{1} \otimes \hat{H}_{\mathfrak{E}} - \hat{H}_{\mathfrak{E}}^\top \otimes \mathbb{1}) + \frac{1}{2} \sum_{a,b>0} \kappa_{ab} \left(2\hat{F}_{ab} \otimes \hat{F}_{ab}^* - \mathbb{1} \otimes \hat{F}_{ab}^\top \hat{F}_{ab} - (\hat{F}_{ab}^\top \hat{F}_{ab})^\top \otimes \mathbb{1} \right) \right] |\hat{A}\rangle\rangle. \tag{A11}$$

We can also derive the formal solution to the adjoint evolution equation by examining $\text{Tr}[\hat{A}^\dagger(0)\rho_{\mathfrak{E}}(t)] = \langle\langle \hat{A}^\dagger(0) | \rho_{\mathfrak{E}}(t) \rangle\rangle$:

$$\begin{aligned} \text{Tr}[\hat{A}^\dagger(0)\rho_{\mathfrak{E}}(t)] &= \langle\langle \hat{A}^\dagger(0) | \rho_{\mathfrak{E}}(t) \rangle\rangle = \langle\langle \hat{A}^\dagger(0) | e^{t\hat{\mathcal{L}}} | \rho_{\mathfrak{E}}(t) \rangle\rangle \\ &= \langle\langle e^{t\hat{\mathcal{L}}^\dagger} \hat{A}^\dagger(0) | \rho_{\mathfrak{E}}(t) \rangle\rangle = \langle\langle \hat{A}^\dagger(t) | \rho_{\mathfrak{E}}(t) \rangle\rangle. \end{aligned} \tag{A12}$$

Here, since we have $\langle\langle \hat{A}^\dagger(t) | = \langle\langle \hat{A}^\dagger(0) | e^{t\hat{\mathcal{L}}}$ and $e^{t\hat{\mathcal{L}}}|\hat{A}(0)\rangle\rangle = |\hat{A}(t)\rangle\rangle$, we have $|\hat{A}(t)\rangle\rangle$ satisfying the differential equation:

$$\frac{d|\hat{A}(t)\rangle\rangle}{dt} = \hat{\mathcal{L}}|\hat{A}(t)\rangle\rangle \tag{A13}$$

Finally, the vectorized solutions to the differential Equation (61) for time-dependent and time-independent GKSL superoperators are, respectively:

$$|\rho_{\mathfrak{E}}(t)\rangle\rangle = \mathcal{T} \exp \left\{ \int_{t_0}^t dt' \hat{\mathcal{L}}(t') \right\} |\rho_{\mathfrak{E}}(t_0)\rangle\rangle \tag{A14a}$$

$$|\rho_{\mathfrak{E}}(t)\rangle\rangle = e^{t\hat{\mathcal{L}}}|\rho_{\mathfrak{E}}(t_0)\rangle\rangle. \tag{A14b}$$

These expressions are the same as (57) and $\rho(t) = e^{t\hat{\mathcal{L}}}\rho(t_0)$, but crucially they now allow us to examine the spectral decomposition of (and analytic functions of) $\hat{\mathcal{L}}$.

References

1. Landauer, R. Irreversibility and Heat Generation in the Computing Process. *IBM J. Res. Dev.* **1961**, *5*, 183–191. [CrossRef]
2. Bennett, C.H. Logical Reversibility of Computation. *IBM J. Res. Dev.* **1973**, *17*, 525–532. [CrossRef]
3. Bennett, C.H. The Thermodynamics of Computation—A Review. *Int. J. Theor. Phys.* **1982**, *21*, 905–940. [CrossRef]
4. Bennett, C.H.; Landauer, R. The Fundamental Physical Limits of Computation. *Sci. Am.* **1985**, *253*, 48–57. [CrossRef]
5. Landauer, R. Computation: A Fundamental Physical View. *Phys. Scr.* **1987**, *35*, 88–95. [CrossRef]
6. Bennett, C.H. Notes on the History of Reversible Computation. *IBM J. Res. Dev.* **1988**, *32*, 16–23. [CrossRef]
7. Bennett, C.H. Notes on Landauer’s Principle, Reversible Computation, and Maxwell’s Demon. *Stud. Hist. Phil. Mod. Phys.* **2003**, *34*, 501–510. [CrossRef]
8. Nielsen, M.A.; Chuang, I.L. *Quantum Computation and Quantum Information*; Cambridge University Press: Cambridge, UK, 2000.
9. Frank, M.P. The Indefinite Logarithm, Logarithmic Units, and the Nature of Entropy. *arXiv* **2005**, arXiv:0506128.
10. Wolpert, D.H. The Stochastic Thermodynamics of Computation. *J. Phys. A Math. Theor.* **2019**, *52*, 193001. [CrossRef]
11. Frank, M.P. Physical Foundations of Landauer’s Principle. In Proceedings of the 10th International Conference RC 2018: Reversible Computation, Leicester, UK, 12–14 September 2018; Kari, J., Ulidowski, I., Eds.; Lecture Notes in Computer Science 11106; Springer: Cham, Switzerland, 2018; pp. 3–33. [CrossRef]
12. Frank, M.P. Physical Foundations of Landauer’s Principle. *arXiv* **2019**, arXiv:1901.10327.
13. Goold, J.; Paternostro, M.; Modi, K. Nonequilibrium Quantum Landauer Principle. *Phys. Rev. Lett.* **2015**, *114*, 060602. [CrossRef]
14. Guarnieri, G.; Campbell, S.; Goold, J.; Pigeon, S.; Vacchini, B.; Paternostro, M. Full Counting Statistics Approach to the Quantum Non-Equilibrium Landauer Bound. *New J. Phys.* **2017**, *19*, 103038. [CrossRef]
15. Müller, M. Correlating Thermal Machines and the Second Law at the Nanoscale. *Phys. Rev. X* **2018**, *8*, 041051. [CrossRef]
16. Funo, K.; Ueda, M.; Sagawa, T. Quantum Fluctuation Theorems. In *Thermodynamics in the Quantum Regime*; Binder, F., Correa, L.A., Gogolin, C., Anders, J., Adesso, G., Eds.; Fundamental Theories of Physics 195; Springer Nature: Cham, Switzerland, 2018; pp. 249–273. [CrossRef]
17. Albert, V.V.; Jiang, L. Symmetries and Conserved Quantities in Lindblad Master Equations. *Phys. Rev. A* **2014**, *89*, 022118. [CrossRef]
18. Albert, V.V.; Bradlyn, B.; Fraas, M.; Jiang, L. Geometry and Response of Lindbladians. *Phys. Rev. X* **2016**, *6*, 041031. [CrossRef]
19. Albert, V.V. Lindbladians with Multiple Steady States. Ph.D. Thesis, Yale University, New Haven, CT, USA, January 2018.
20. Deffner, S.; Campbell, S. *Quantum Thermodynamics*; Morgan & Claypool: San Rafael, CA, USA, 2019.
21. Binder, F.; Correa, L.A.; Gogolin, C.; Anders, J.; Adesso, G. (Eds.) *Thermodynamics in the Quantum Regime*; Fundamental Theories of Physics 195; Springer Nature: Cham, Switzerland, 2018. [CrossRef]
22. Goold, J.; Huber, M.; Riera, A.; del Rio, L.; Skrzypczyk, P. The Role of Quantum Information in Thermodynamics—A Topical Review. *J. Phys. Rev. A Math. Theor.* **2017**, *49*, 143001. [CrossRef]
23. Ng, N.H.Y.; Woods, M. Resource Theory of Quantum Thermodynamics: Thermal Operations and Second Laws. In *Thermodynamics in the Quantum Regime*; Binder, F., Correa, L.A., Gogolin, C., Anders, J., Adesso, G., Eds.; Fundamental Theories of Physics 195; Springer Nature: Cham, Switzerland, 2018; pp. 625–650. [CrossRef]
24. Lostaglio, M. An Introductory Review of the Resource Theory Approach to Thermodynamics. *Rep. Prog. Phys.* **2019**, *82*, 114001. [CrossRef]
25. Chitambar, E.; Gour, G. Quantum Resource Theories. *Rev. Mod. Phys.* **2019**, *91*, 025001. [CrossRef]
26. Alicki, R.; Lendi, K. *Quantum Dynamical Semigroups and Applications*; Lecture Notes in Physics 717; Springer: Berlin/Heidelberg, Germany, 2007. [CrossRef]
27. Breuer, H.-P.; Petruccione, F. *The Theory of Open Quantum Systems*; Oxford University Press: Oxford, UK, 2007.
28. Banerjee, S. *Open Quantum Systems*; Texts and Readings in Physical Sciences 20; Hindustan Book Agency: New Delhi, India; Springer Nature: Singapore, 2018. [CrossRef]
29. Deffner, S.; Campbell, S. Quantum Speed Limits: From Heisenberg’s Uncertainty Principle to Optimal Quantum Control. *J. Phys. Rev. A Math. Theor.* **2017**, *50*, 453001. [CrossRef]
30. Guéry-Odelin, D.; Ruschhaupt, A.; Kiely, A.; Torrontegui, E.; Martínez-Garaot, S.; Muga, J.G. Shortcuts to Adiabaticity: Concepts, Methods, and Applications. *Rev. Mod. Phys.* **2019**, *91*, 045001. [CrossRef]
31. Nakahara, M.; Rahimi, R.; SaiToh, A. *Mathematical Aspects of Quantum Computing 2007*; Kinki University Series on Quantum Computing 1; World Scientific: Singapore, 2007.
32. Wolf, M.M. Quantum Channels and Operations Guided Tour. Unpublished. Available online: <https://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MichaelWolf/QChannelLecture.pdf> (accessed on 27 May 2021).
33. Attal, S. Lectures in Quantum Noise Theory. Unpublished. Available online: <http://math.univ-lyon1.fr/~attal/chapters.html> (accessed on 27 May 2021).
34. Wilde, M.M. *Quantum Information Theory*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2017.
35. Preskill, J. Lecture Notes: Quantum Computation. Unpublished. Available online: <http://theory.caltech.edu/~preskill/ph219/> (accessed on 27 May 2021).
36. Deffner, S.; Jarzynski, C. Information Processing and the Second Law of Thermodynamics: An Inclusive, Hamiltonian Approach. *Phys. Rev. X* **2013**, *3*, 041003. [CrossRef]

37. Barato, A.C.; Seifert, U. Stochastic thermodynamics with information reservoirs. *Phys. Rev. E* **2014**, *90*, 042150. [[CrossRef](#)] [[PubMed](#)]
38. Strasberg, P.; Schaller, G.; Brandes, T.; Esposito, M. Quantum and Information Thermodynamics: A Unifying Framework Based on Repeated Interactions. *Phys. Rev. X* **2017**, *7*, 021003. [[CrossRef](#)]
39. Frank, M.P. Foundations of Generalized Reversible Computing. In Proceedings of the 9th International Conference RC 2017: Reversible Computation, Kolkata, India, 6–7 July 2017; Phillips, I., Rahaman, H., Eds.; Lecture Notes in Computer Science 10301; Springer: Cham, Switzerland, 2017; pp. 19–34. [[CrossRef](#)]
40. Frank, M.P. Generalized Reversible Computing. *arXiv* **2018**, arXiv:1806.10183.
41. Zurek, W.H. Decoherence, Einselection, and the Quantum Origins of the Classical. *Rev. Mod. Phys.* **2003**, *75*, 715. [[CrossRef](#)]
42. Bekenstein, J.D. Holographic Bound from Second Law of Thermodynamics. *Phys. Lett. B* **2000**, *481*, 339–345. [[CrossRef](#)]
43. Zurek, W.H. Pointer Basis of Quantum Apparatus: Into What Mixture Does the Wave Packet Collapse? *Phys. Rev. D* **1981**, *24*, 1516–1525. [[CrossRef](#)]
44. Stinespring, W.F. Positive Functions on C^* Algebras. *Proc. Am. Math. Soc.* **1955**, *6*, 211. [[CrossRef](#)]
45. Pechukas, P. Reduced Dynamics Need Not Be Completely Positive. *Phys. Rev. Lett.* **1994**, *73*, 1060. [[CrossRef](#)]
46. Brandão, F.; Horodecki, M.; Ng, N.H.Y.; Oppenheim, J.; Woods, S. The Second Laws of Quantum Thermodynamics. *Proc. Natl. Acad. Sci. USA* **2015**, *112*, 3275–3279. [[CrossRef](#)]
47. Horodecki, M.; Oppenheim, J. Fundamental Limitations for Quantum and Nanoscale Thermodynamics. *Nat. Commun.* **2013**, *4*, 2059. [[CrossRef](#)]
48. Frank, M.P.; Brocato, R.W.; Tierney, B.D.; Missert, N.A.; Hsia, A. Reversible Computing with Fast, Fully Static, Fully Adiabatic CMOS. In Proceedings of the 2020 International Conference on Rebooting Computing (ICRC), Atlanta, GA, USA, 1–3 December 2020. [[CrossRef](#)]
49. Bergmann, P.G.; Lebowitz, J.L. New Approach to Nonequilibrium Processes. *Phys. Rev.* **1955**, *99*, 578. [[CrossRef](#)]
50. Parrondo, J.M.R.; Horowitz, J.M.; Sagawa, T. Thermodynamics of Information. *Nat. Phys.* **2015**, *11*, 131–139. [[CrossRef](#)]
51. Beaudry, N.J.; Renner, R. An Intuitive Proof of the Data Processing Inequality. *Quant. Inf. Comp.* **2011**, *12*, 432–441. [[CrossRef](#)]
52. Renithasamy, S.; Wilde, M.M. Relative Entropy and Catalytic Relative Majorization. *Phys. Rev. Res.* **2020**, *2*, 033455. [[CrossRef](#)]
53. Müller-Lennert, M.; Dupuis, F.; Szehr, O.; Fehr, S.; Tomamichel, M. On Quantum Rényi Entropies: A New Generalization and Some Properties. *J. Math. Phys.* **2013**, *54*, 122203. [[CrossRef](#)]
54. Van Erven, T.; Harremoës, P. Rényi Divergence and Kullback-Leibler Divergence. *IEEE Trans. Inf. Theory* **2014**, *60*, 3797–3820. [[CrossRef](#)]
55. Rényi, A. On Measures of Dependence. *Acta Math. Acad. Sci. Hung.* **1955**, *10*, 441–451. [[CrossRef](#)]
56. Audenaert, K.M.R.; Datta, N. α -z-Rényi Relative Entropies. *J. Math. Phys.* **2015**, *56*, 022202. [[CrossRef](#)]
57. Klimesh, M. Entropy Measures and Catalysis of Bipartite Quantum State Transformations. In Proceedings of the 2004 IEEE International Symposium on Information Theory (ISIT), Chicago, IL, USA, 27 June–2 July 2004. [[CrossRef](#)]
58. Klimesh, M. Inequalities that Collectively Completely Characterize the Catalytic Majorization Relation. *arXiv* **2007**, arXiv:0709.3680.
59. Turgut, S. Catalytic Transformations for Bipartite Pure States. *J. Phys. A Math. Theor.* **2007**, *40*, 12185. [[CrossRef](#)]
60. Wilming, H.; Gallego, R.; Eisert, J. Axiomatic Characterization of the Quantum Relative Entropy and Free Energy. *Entropy* **2017**, *19*, 241. [[CrossRef](#)]
61. Anderson, N. Conditional Erasure and the Landauer Limit. In *Energy Limits in Computation*; Lent, C.S., Orlov, A.O., Prod, W., Snider, G., Eds.; Springer Nature: Cham, Switzerland, 2019; pp. 65–100. [[CrossRef](#)]
62. Partovi, M.H. Quantum Thermodynamics. *Phys. Lett. A* **1989**, *137*, 440–444. [[CrossRef](#)]
63. Perarnau-Llobet, M.; Riera, A.; Gallego, R.; Wilming, H.; Eisert, J. Work and Entropy Production in Generalised Gibbs Ensembles. *New J. Phys.* **2016**, *18*, 123035. [[CrossRef](#)]
64. Anderson, N. Landauer’s Limit and the Physicality of Information. *Eur. Phys. J. B* **2018**, *91*, 156. [[CrossRef](#)]
65. Kraus, K. General State Changes in Quantum Theory. *Ann. Phys.* **1971**, *64*, 311–335. [[CrossRef](#)]
66. Talkner, P.; Lutz, E.; Hänggi, P. Fluctuation Theorems: Work is Not an Observable. *Phys. Rev. E* **2007**, *75*, 050102(R). [[CrossRef](#)] [[PubMed](#)]
67. Esposito, M.; Harbola, U.; Mukamel, S. Nonequilibrium Fluctuations, Fluctuation Theorems, and Counting Statistics in Quantum Systems. *Rev. Mod. Phys.* **2009**, *81*, 1665. [[CrossRef](#)]
68. Reeb, D.; Wolf, M.M. An Improved Landauer Principle with Finite-Size Corrections. *New J. Phys.* **2014**, *16*, 103011. [[CrossRef](#)]
69. Touil, A.; Deffner, S. Information Scrambling versus Decoherence—Two Competing Sinks for Entropy. *PRX Quantum* **2021**, *2*, 010306. [[CrossRef](#)]
70. Born, M. Quantenmechanik der Stoßvorgänge. *Z. Phys.* **1926**, *37*, 863–867. [[CrossRef](#)]
71. von Neumann, J. *Mathematische Grundlagen der Quantenmechanik*; Springer: Berlin/Heidelberg, Germany, 1932.
72. Guryanova, Y.; Friis, N.; Huber, M. Ideal Projective Measurements Have Infinite Resource Costs. *Quantum* **2020**, *4*, 222. [[CrossRef](#)]
73. Deffner, S.; Paz, J.P.; Zurek, W.H. Quantum Work and the Thermodynamic Cost of Quantum Measurements. *Phys. Rev. E* **2009**, *94*, 010103. [[CrossRef](#)]
74. Breuer, H.-P.; Bugarth, D.; Petruccione, F. Non-Markovian Dynamics in a Spin Star System: Exact Solution and Approximation Techniques. *Phys. Rev. B* **2004**, *70*, 045323. [[CrossRef](#)]

75. Breuer, H.-P.; Gemmer, J.; Michel, M. Non-Markovian Quantum Dynamics: Correlated Projection Superoperators and Hilbert Space Averaging. *Phys. Rev. E* **2006**, *73*, 016139. [CrossRef]
76. Ivanov, A.; Breuer, H.-P. Extension of the Nakajima-Zwanzig Approach to Multitime Correlation Functions of Open Systems. *Phys. Rev. A* **2015**, *92*, 032113. [CrossRef]
77. Lindblad, G. On the Generators of Quantum Dynamical Semigroups. *Commun. Math. Phys.* **1976**, *48*, 119–130. [CrossRef]
78. Gorini, V.; Kossakowski, A.; Sudarshan, E.C.G. Completely Positive Dynamical Semigroups of N -Level Systems. *J. Math. Phys.* **1976**, *17*, 821. [CrossRef]
79. Erdős, L. Lecture Notes on Quantum Brownian Motion. In *Quantum Theory from Small to Large Scales*; Fröhlich, J., Salmhofer, M., Mastropietro, V., De Roeck, W., Cugliandolo, L.F., Eds.; Lecture Notes of the Les Houches Summer School 95; Oxford University Press: Oxford, UK, 2012; pp. 3–98. [CrossRef]
80. Caldeira, A.O.; Leggett, A.J. Path Integral Approach to Quantum Brownian Motion. *Physica A* **1983**, *121*, 587–616. [CrossRef]
81. Kossakowski, A. On Quantum Statistical Mechanics of Non-Hamiltonian Systems. *Rep. Math. Phys.* **1972**, *3*, 247–274. [CrossRef]
82. Ingarden, R.S.; Kossakowski, A. On the Connection of Nonequilibrium Information Thermodynamics with Non-Hamiltonian Quantum Mechanics of Open Systems. *Ann. Phys.* **1975**, *89*, 451–485. [CrossRef]
83. Wolf, M.M.; Cirac, J.I. Dividing Quantum Channels. *Commun. Math. Phys.* **2008**, *279*, 147–168. [CrossRef]
84. Wolf, M.M.; Eisert, J.; Cubitt, T.S.; Cirac, J.I. Assessing Non-Markovian Quantum Dynamics. *Phys. Rev. Lett.* **2008**, *101*, 150402. [CrossRef]
85. Woit, P. *Quantum Theory, Groups, and Representations*; Springer International Publishing: New York, NY, USA, 2017. [CrossRef]
86. Baumgartner, B.; Narnhofer, H. Analysis of Quantum Semigroups with GKS-Lindblad Generators: II. General. *J. Phys. A Math. Theor.* **2008**, *41*, 395303. [CrossRef]
87. Ticozzi, F.; Viola, L. Quantum Markovian Subsystems: Invariance, Attractivity, and Control. *IEEE Trans. Aut. Control* **2008**, *53*, 2048–2063. [CrossRef]
88. Blume-Kohout, R.; Ng, H.K.; Poulin, D.; Viola, L. Information-Preserving Structures: A General Framework for Quantum Zero-Error Information. *Phys. Rev. A* **2010**, *82*, 062306. [CrossRef]
89. Deschamps, J.; Fagnola, F.; Sasso, E.; Unamità, V. Structure of Uniformly Continuous Quantum Markov Semigroups. *Rev. Math. Phys.* **2016**, *28*, 1650003. [CrossRef]
90. Pastawski, F.; Preskill, P. Code Properties from Holographic Geometries. *Phys. Rev. X* **2017**, *7*, 021022. [CrossRef]
91. Younis, S.G.; Knight, T.F., Jr. Practical Implementation of Charge Recovering Asymptotically Zero Power CMOS. In Proceedings of the 1993 Symposium Research on Integrated Systems, Seattle, WA, USA, February 1993; Ebeling, C., Borriello, G., Eds.; MIT Press: Cambridge, UK, 1993; pp. 234–250. Available online: <ftp://publications.ai.mit.edu/ai-publications/pdf/AITR-1500.pdf> (accessed on 27 May 2021).
92. Younis, S.G. Asymptotically Zero Energy Computing Using Split-Level Charge Recovery Logic. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, June 1994. Available online: <http://hdl.handle.net/1721.1/11620> (accessed on 27 May 2021).
93. Frank, M.P. Reversibility for Efficient Computing. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, June 1999. Available online: <http://hdl.handle.net/1721.1/9464> (accessed on 27 May 2021).
94. Venkiteswaran, A.; He, M.; Natarajan, K.; Xie, H.; Frank, M.P. Driving Fully-Adiabatic Logic Circuits Using Custom High-Q MEMS Resonators. In Proceedings of the International Conference on Embedded Systems and Applications, ESA '04 and VLSI, VLSI '04 (ESA/VLSI 2004), Las Vegas, NA, USA, 21–24 June 2004; Arabia, H.R., Guo, M., Yang, L.T., Eds.; CSREA Press: Las Vegas, NA, USA, 2004; pp. 5–11. Available online: <http://revcomp.info/legacy/revcomp/AdiaMEMS/MLPD-04.pdf> (accessed on 27 May 2021).
95. Zulehner, A.; Frank, M.P.; Wille, R. Design Automation for Adiabatic Circuits. In Proceedings of the 24th Asia and South Pacific Design Automation Conference, ASPDAC '19, Tokyo, Japan, 21–24 January 2019; ACM: New York, NY, USA, 2019; pp. 669–674. [CrossRef]
96. Frank, M.P.; Brocato, R.W.; Conte, T.M.; Hsia, A.; Jain, A.; Missert, N.A.; Shukla, K.; Tierney, B.D. Special Session: Exploring the Ultimate Limits of Adiabatic Circuits. In Proceedings of the 2020 IEEE 38th International Conference on Computer Design (ICCD), Hartford, CT, USA, 18–21 October 2020. [CrossRef]
97. Feynman, R.P. *Feynman Lectures on Computation*; CRC Press: Boca Raton, FL, USA, 2000.
98. Frank, M.P. Common Mistakes in Adiabatic Logic Design and How to Avoid Them. In Proceedings of the International Conference on Embedded Systems and Applications, ESA '03, Las Vegas, NA, USA, 23–26 June 2003; Arabia, H.R., Yang, L.T., Eds.; CSREA Press: Las Vegas, NA, USA, 2003; pp. 216–222. Available online: <http://revcomp.info/legacy/revcomp/MLPD03-Mistakes-paper.pdf> (accessed on 27 May 2021).
99. Takeuchi, N.; Yamanashi, Y.; Yoshikawa, N. Reversible Logic Gate Using Adiabatic Superconducting Devices. *Sci. Rep.* **2014**, *4*, 6354. [CrossRef]
100. Takeuchi, N.; Yamanashi, Y.; Yoshikawa, N. Recent Progress on Reversible Quantum-Flux-Parametron for Superconductor Reversible Computing. *IEICE Trans. Electr.* **2018**, *101*, 352–358. [CrossRef]
101. Yamae, T.; Takeuchi, N.; Yoshikawa, N. A Reversible Full Adder using Adiabatic Superconductor Logic. *Superconductor Sci. Tech.* **2019**, *32*, 035005. [CrossRef]
102. Lent, C.S.; Tougaw, P.D.; Porod, W.; Bernstein, G.H. Quantum Cellular Automata. *Nanotechnology* **1993**, *4*, 49–57. [CrossRef]

103. Amlani, I.; Orlov, A.O.; Toth, G.; Bernstein, G.H.; Lent, C.S.; Snider, G.L. Digital Logic Gate Using Quantum-Dot Cellular Automata. *Science* **1999**, *284*, 289–291. [[CrossRef](#)]
104. Lent, C.S.; Isaken, B. Clocked Molecular Quantum-Dot Cellular Automata. *IEEE Trans. Electr. Dev.* **2003**, *50*, 1890–1896. [[CrossRef](#)]
105. Pidaparathi, S.S.; Lent, C.S. Exponentially Adiabatic Switching in Quantum-Dot Cellular Automata. *J. Low Power Electr. Appl.* **2018**, *8*, 30. [[CrossRef](#)]
106. Pidaparathi, S.S.; Lent, C.S. Energy Dissipation During Two-State Switching for Quantum-Dot Cellular Automata. *J. Appl. Phys.* **2021**, *129*, 024304. [[CrossRef](#)]
107. Drexler, K.E. Molecular Engineering: An Approach to the Development of General Capabilities for Molecular Manipulation. *Proc. Natl. Acad. Sci. USA* **1981**, *78*, 5275–5278. [[CrossRef](#)] [[PubMed](#)]
108. Drexler, K.E. Molecular Machinery and Manufacturing with Applications to Computation. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, 1991. Available online: <http://hdl.handle.net/1721.1/27999> (accessed on 27 May 2021).
109. Drexler, K.E. *Nanosystems*; John Wiley & Sons: Hoboken, NJ, USA, 1992.
110. Merkle, R.C.; Freitas, R.A., Jr.; Hogg, T.; Moore, T.E.; Moses, M.S.; Ryley, J. *Molecular Mechanical Computing Systems*; Tech. Report #046; Institute for Molecular Manufacturing: Palo Alto, CA, USA, 2016. Available online: <http://www.imm.org/Reports/rep046.pdf> (accessed on 27 May 2021).
111. Merkle, R.C.; Freitas, R.A., Jr.; Hogg, T.; Moore, T.E.; Moses, M.S.; Ryley, J. Mechanical Computing Systems Using Only Links and Rotary Joints. *J. Mech. Robot.* **2018**, *10*, 061006. [[CrossRef](#)]
112. Hogg, T.; Moses, M.S.; Allis, D.G. Evaluating the Friction of Rotary Joints in Molecular Machines. *Mol. Syst. Des. Eng.* **2017**, *2*, 235–252. [[CrossRef](#)]
113. Frank, M.P. Asynchronous Ballistic Reversible Computing. In Proceedings of the 2017 IEEE International Conference on Rebooting Computing (ICRC), Washington, DC, USA, 8–9 November 2017. [[CrossRef](#)]
114. Frank, M.P.; Lewis, R.M.; Missert, N.A.; Wolak, M.A.; Henry, M.D. Asynchronous Ballistic Reversible Fluxon Logic. *IEEE Trans. Appl. Supercond.* **2019**, *29*. [[CrossRef](#)]
115. Frank, M.P.; Lewis, R.M.; Missert, N.A.; Henry, M.D.; Wolak, M.A.; DeBenedictis, E.P. Semi-Automated Design of Functional Elements for a New Approach to Digital Superconducting Electronics: Methodology and Preliminary Results. In Proceedings of the 2019 IEEE International Superconductive Electronics Conference (ISEC), Riverside, CA, USA, 28 July–1 August 2019. [[CrossRef](#)]
116. Wustmann, W.; Osborn, K.D.; Osborn Team. Autonomous Reversible Fluxon Logic Gates. APS March Meeting 2017 Abstract. Available online: <https://ui.adsabs.harvard.edu/#abs/2017APS..MARC46006W/abstract> (accessed on 27 May 2021).
117. Osborn, K.D.; Wustmann, W. Ballistic Reversible Gates Matched to Bit Storage: Plans for an Efficient CNOT Gate Using Fluxons. In Proceedings of the 10th International Conference RC 2018: Reversible Computation, Leicester, UK, 12–14 September 2018; Kari, J., Ulidowski, I., Eds.; Lecture Notes in Computer Science 11106; Springer: Cham, Switzerland, 2018; pp. 189–204. [[CrossRef](#)]
118. Wustmann, W.; Osborn, K.D. Reversible Fluxon Logic: Topological Particles Enable Gates Beyond the Standard Adiabatic Limit. APS March Meeting 2018 Abstract. Available online: <https://ui.adsabs.harvard.edu/#abs/2018APS..MARK15004W/abstract> (accessed on 27 May 2021).
119. Yu, L.; Wustmann, W.; Osborn, K.D. Experimental Designs of Ballistic Reversible Logic Gates Using Fluxons. In Proceedings of the 2019 IEEE International Superconductive Electronics Conference (ISEC), Riverside, CA, USA, 28 July–1 August 2019. [[CrossRef](#)]
120. Osborn, K.D.; Wustmann, W. Reversible Fluxon Logic for Future Computing. In Proceedings of the 2019 IEEE International Superconductive Electronics Conference (ISEC), Riverside, CA, USA, 28 July–1 August 2019. [[CrossRef](#)]
121. Osborn, K.D.; Wustmann, W. Reversible Fluxon Logic With Optimized CNOT Gate Components. *IEEE Trans. Appl. Supercond.* **2021**, *31*, 1300213. [[CrossRef](#)]
122. Wustmann, W.; Osborn, K.D. Reversible Fluxon Logic: Topological Particles Allow Ballistic Gates Along One-Dimensional Paths. *Phys. Rev. B* **2020**, *101*, 014516. [[CrossRef](#)]
123. Wustmann, W.; Osborn, K.D. Reversible Fluxon Logic with Shift Registers. APS March Meeting 2020 Abstract. Available online: <https://meetings.aps.org/Meeting/MAR20/Session/A36.3> (accessed on 27 May 2021).
124. Strasberg, P.; Esposito, M. Non-Markovianity and Negative Entropy Production Rates. *Phys. Rev. E* **2019**, *99*, 012120. [[CrossRef](#)]
125. Bonança, M.V.S.; Nazé, P.; Deffner, S. Negative Entropy Production Rates in Drude-Sommerfeld Metals. *Phys. Rev. E* **2021**, *103*, 012109. [[CrossRef](#)]
126. De Haas, W.J.; Wiersma, E.C.; Kramers, H.A. Experiments on Adiabatic Cooling of Paramagnetic Salts in Magnetic Fields. *Physica* **1934**, *1*, 1–13. [[CrossRef](#)]
127. Kunzler, J.E.; Walker, L.R.; Galt, J.K. Adiabatic Demagnetization and Specific Heat in Ferrimagnets. *Phys. Rev.* **1960**, *119*, 1609. [[CrossRef](#)]
128. Pecharsky, V.K.; Gschneidner, K.A., Jr. Magnetocaloric Effect and Magnetic Refrigeration. *J. Magnet. Magn. Mater.* **1999**, *200*, 44–56. [[CrossRef](#)]
129. Sagawa, T. Thermodynamic and Logical Reversibilities Revisited. *J. Stat. Mech. Theor. Exp.* **2014**, *2014*, P03025. [[CrossRef](#)]
130. Wolpert, D. Overview of Information Theory, Computer Science Theory, and Stochastic Thermodynamics for Thermodynamics of Computation. In *The Energetics of Computation in Life and Machines*; Wolpert, D., Kempes C., Stadler, P., Grochow, J., Eds.; Santa Fe Institute Press: Santa Fe, NM, USA, 2019; pp. 1–36.

131. Church, A. An Unsolvable Problem of Elementary Number Theory. *Am. J. Math.* **1936**, *58*, 345–363. [[CrossRef](#)]
132. Sagawa, T. Thermodynamics of Information Processing in Small Systems. Ph.D. Thesis, University of Tokyo, Tokyo, Japan, 2013. [[CrossRef](#)]
133. Sagawa, T. Second Law, Entropy Production, and Reversibility in Thermodynamics of Information. In *Energy Limits in Computation*; Lent, C.S., Orlov, A.O., Porod, W., Snider, G., Eds.; Springer Nature: Cham, Switzerland, 2019; pp. 101–139. [[CrossRef](#)]
134. Frank, M.P.; DeBenedictis, E.P. A Novel Operational Paradigm for Thermodynamically Reversible Logic: Adiabatic Transformation of Chaotic Nonlinear Dynamical Circuits. In Proceedings of the 2016 International Conference on Rebooting Computing (ICRC), San Diego, CA, USA, 17–19 October 2016; pp. 1–8. [[CrossRef](#)]
135. IEEE. *Beyond CMOS Chapter, International Roadmap for Devices and Systems*, 2020th ed.; IEEE: Piscataway, NJ, USA, 2020. Available online: <https://irds.ieee.org/editions/2020/beyond-cmos> (accessed on 27 May 2021).
136. Scandi, M.; Perarnau-Llobet, M. Thermodynamic Length in Open Quantum Systems. *Quantum* **2019**, *3*, 197. [[CrossRef](#)]
137. Guarnieri, G.; Landi, G.T.; Clark, S.R.; Goold, J. Thermodynamics of Precision in Quantum Nonequilibrium Steady States. *Phys. Rev. Res.* **2019**, *1*, 033021. [[CrossRef](#)]
138. Gea-Banacloche, J. Minimum Energy Requirements for Quantum Computation. *Phys. Rev. Lett.* **2002**, *89*, 217901. [[CrossRef](#)]
139. Deffner, S. Energetic Cost of Hamiltonian Quantum Gates. *arXiv* **2021**, arXiv:2102.05118.
140. Zurek, W.H. Algorithmic Randomness and Physical Entropy. *Phys. Rev. A* **1989**, *40*, 4731. [[CrossRef](#)]