

CCC Workshop on Physics & Engineering Issues in Adiabatic/Reversible Classical Computing

TECHNICAL SESSION I—FUNDAMENTAL PHYSICS Fundamental Physics of Reversible Computing—An Introduction



Monday, October 5th, 2020

Michael P. Frank, Center for Computing Research



SAND2020-9778C



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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-NA0003525.

2 Workshop Overview



	Day 1 (Mon. 10/5)	Day 2 (Tue. 10/6)	Day 3 (Wed. 10/7)	Day 4 (Thu. 10/8)	Day 5 (Fri. 10/9)
START TIME (US PDT)	TECHNICAL SESSION I: FUNDAMENTAL PHYSICS	TECHNICAL SESSION II: DEVICE & CIRCUIT TECHNOLOGIES	TECHNICAL SESSION III: ARCHITECTURE & HIGH-LEVEL TOPICS	FIRST DAY OF WORKING MEETINGS	SECOND DAY OF WORKING MEETINGS
8:30 am	Workshop Intro	Keynote: Ed Fredkin	M. Frank, G. Snider, N. Yoshikawa, H. Thapliyal, R. Wille	(9a.) Day 4 Intro.	(9a.) Day 5 Intro.
9:20 am	Mike Frank	Mike Frank		Outbriefs from Breakouts	Outbriefs from Re-Breakouts
10:00 am	Norm Margolus	Sarah Frost-Murphy			
10:20 am	Early Break	Early Break	Early Break	Early Break	Early Break
10:50 am	Neal Anderson	Kevin Osborn	Erik Demaine	Concordance Discussion #1	Concordance Discussion #2
11:10 am	Craig Lent	Ralph Merkle	Robert Glück		
11:30 am	Karpur Shukla	Joe Friedman	Erik DeBenedictis		
11:50 am	Panel / Q&A	Panel / Q&A	Panel / Q&A		
12:10 pm	Late Break	Late Break	Late Break	Late Break	Late Break
12:40 pm until...	Physics Breakout	Techno. Breakout	Arch./HL Breakouts	Re-Breakouts	Final Breakout & Concordance

Abstract Text



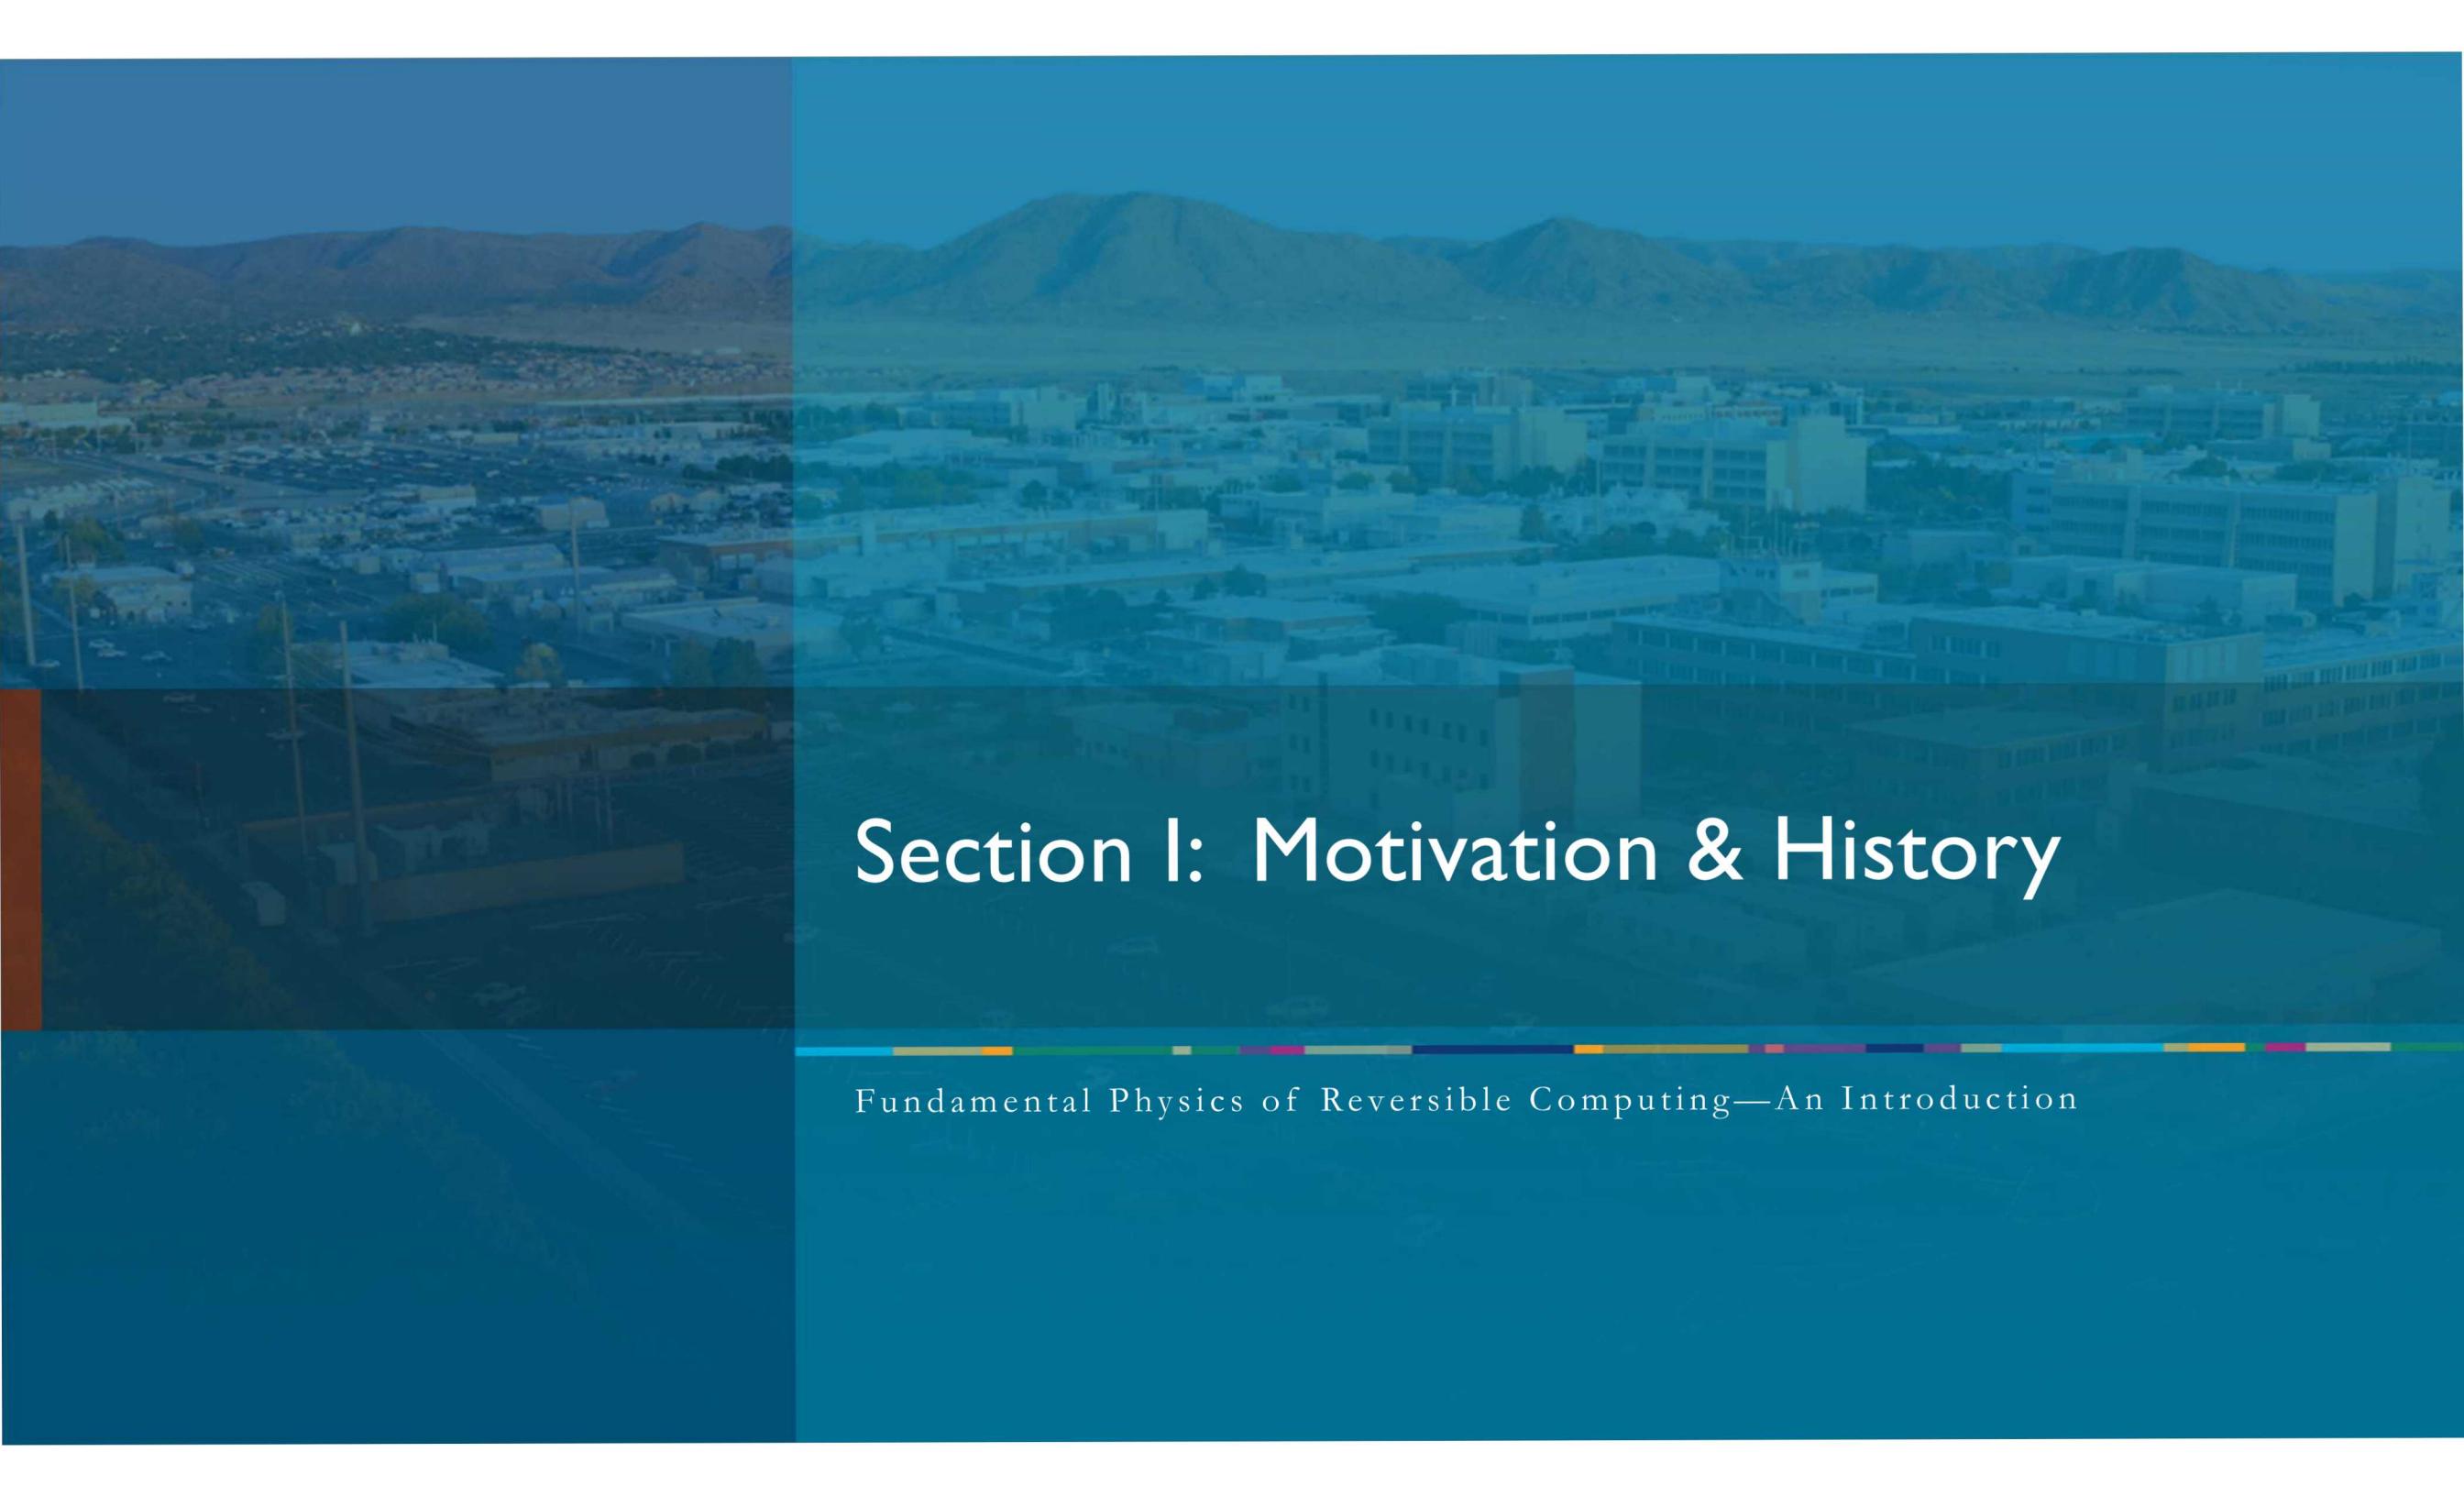
The concept of using reversible computing to circumvent fundamental physical limits on energy efficiency has historical roots going all the way back the 1961 work of Landauer, and was shown to be theoretically workable by Bennett in 1973. But, over the last 59 years, relatively little attention has been paid, from a fundamental physics perspective, to the question of just how energy efficient, as a function of speed, practical physical implementations of reversible computing can be made to be. To finally answer this question in a definitive way is becoming an increasingly important task as the conventional, non-reversible computing paradigm approaches its limits. Recruiting the physics community to turn increased attention to solving this and related problems is one of the major motivations for this workshop. In this talk, we kick off the Fundamental Physics session by giving an overview of what's already known about the fundamental physics of reversible computing, and highlighting some important research challenges in this area.

Outline of Talk

Fundamental Physics of Reversible Computing— An Introduction

- I. Motivation & Brief History
- II. Foundational Topics:
 - Computational and Physical Information
 - Fundamental Theorem of the Thermodynamics of Computing
 - Landauer’s Principle
 - Fundamental Theorem of Traditional Reversible Computing
 - Fundamental Theorem of Generalized Reversible Computing
 - Quantum Dynamics of Classical Computing—a generic formulation
 - Some known fundamental relations between speed and dissipation (Likharev, Feynman)
- III. Looking Ahead:
 - What advances in the fundamental physics of reversible computing are needed to provide a foundation for a robust engineering discipline?





Section I: Motivation & History

Fundamental Physics of Reversible Computing—An Introduction



Why are we here?

- Progress in the energy-efficiency of the conventional (non-reversible) computing paradigm is approaching hard limits, which ultimately trace back to fundamental thermodynamic issues.
 - Industry is already struggling to continue to advance along the traditional scaling path.
- Energy efficiency is a fundamental limiting factor on the economic utility of computing.
 - Without energy efficiency gains, there are diminishing returns from optimizing *every* other aspect of computing.
- Transitioning to the unconventional computing paradigm known as *reversible computing* provides the only physically possible alternative scaling path for allowing the energy efficiency of general digital computing to continue improving indefinitely...
 - And, so far, no fundamental limit to the (even practically) achievable efficiency is known.
- The overall economy is becoming increasingly dependent on computing, as a larger and larger share of economic activity takes place in the cyber realm...
 - Making reversible computing practical thus has the potential to multiply *the total economic value of civilization* (for any given amount of available energy resources) by indefinitely many orders of magnitude.

Semiconductor Roadmap is Ending...



Thermal noise on gate electrodes of minimum-width segments of FET gates leads to significant channel PES fluctuations when $E_g \lesssim 1\text{-}2 \text{ eV}$

- Increases leakage, impairs practical device performance
- Thus, roadmap has minimum gate energy asymptoting to $\sim 2 \text{ eV}$

Also, real logic circuits incur many *compounding* overhead factors *multiplying* this limit:

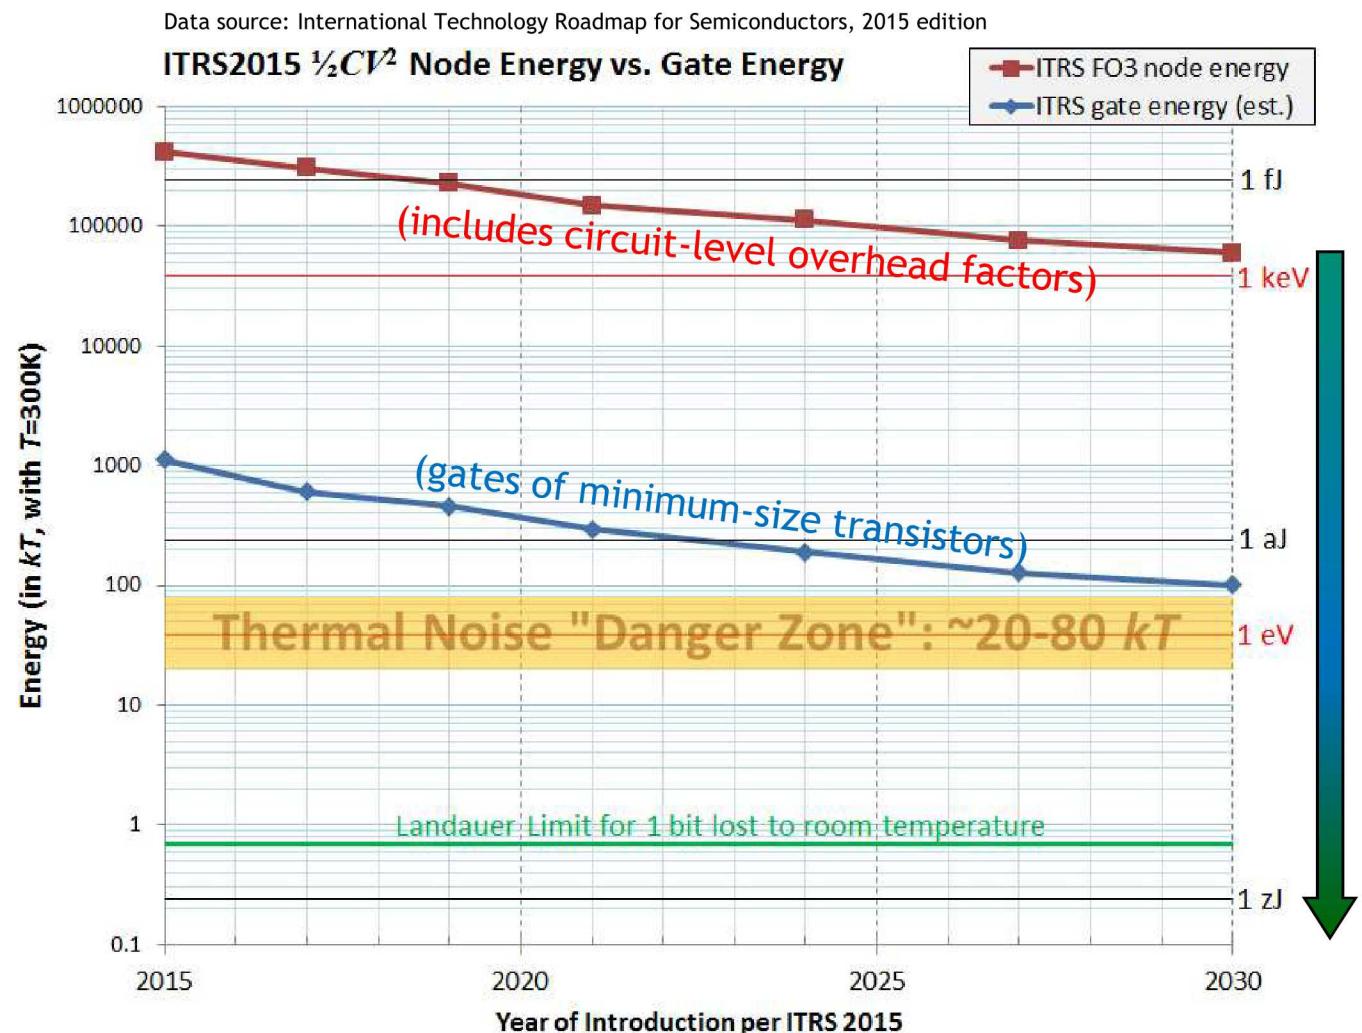
- Transistor width $10\text{-}20\times$ minimum width in fast logic.
- Parasitic (junction, etc.) transistor capacitances ($\sim 2\times$).
- Multiple (~ 2) transistors fed by each input to a given logic gate.
- Fan-out of each gate to a few (~ 3) downstream logic gates.
- Parasitic wire capacitance ($\sim 2\times$).

Due to all these overhead factors, the energy of each logic bit in real logic circuits is many times larger than the minimum-width gate energy!

- $375\text{-}600\times$ (!) larger in ITRS'15.
- ∴ Practical bit energy for irreversible logic asymptotes to $\sim 1 \text{ keV}$!

Practical, real-world logic circuit designs can't just magically cross this $\sim 500\times$ architectural gap!

- ∴ Thermodynamic limits imply much larger practical limits!
- The end is near!

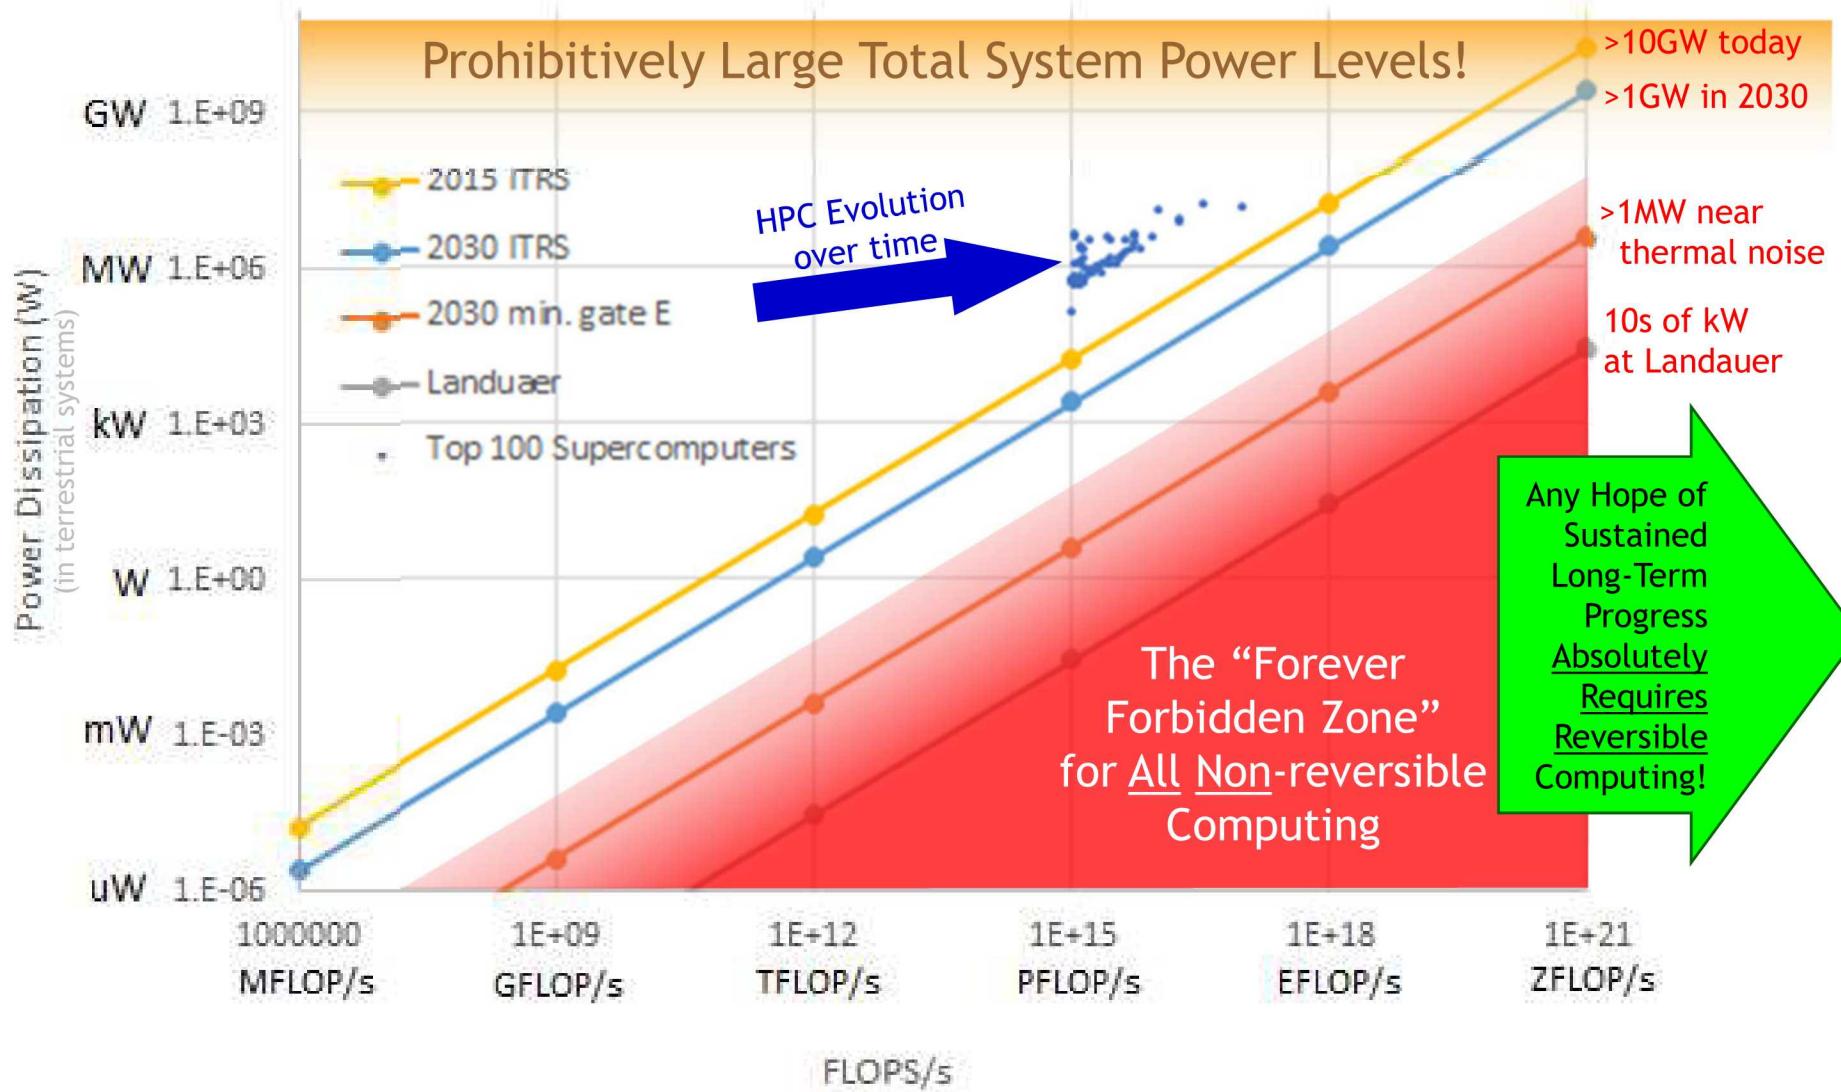


Only reversible computing can take us from $\sim 1 \text{ keV}$ at the end of the CMOS roadmap, all the way down to $\ll kT$.

Implications for FLOPS & power



Note: The limits suggested by the diagonal lines do not even include power overheads for interconnects, memory, or cooling!





Section II: Foundational Topics

Fundamental Physics of Reversible Computing—An Introduction

Section II: Foundational Topics (Outline)



Big Picture: Non-Equilibrium Open Quantum Systems Framework

The (Classical) Computational State Abstraction—And its Quantum Representation

- Proto-computational Basis
- Computational vs. Non-Computational Subspaces
- Time-Dependent Formulation

The Fundamental Theorem of the Thermodynamics of Computation

- Relates computational and physical entropies

Computational Operations

- Logical Reversibility (Unconditional and Conditional)
- Time-Dependent Formulation

Landauer's Principle (properly understood!)

Quantum Model of Classical Computational Operations

- Open vs. Closed System cases
- Quantum Statistical Operating Contexts
- Quantum Contextualized Computations
- What it means to physically implement a (classical) computation

Some known fundamental relations between speed and dissipation (Likharev, Feynman)

Big Picture (Non-Equilibrium Open Quantum System Framework)

Divide the model universe \mathfrak{U} into the “system” (computing system) \mathfrak{S} and its external environment \mathfrak{E} .

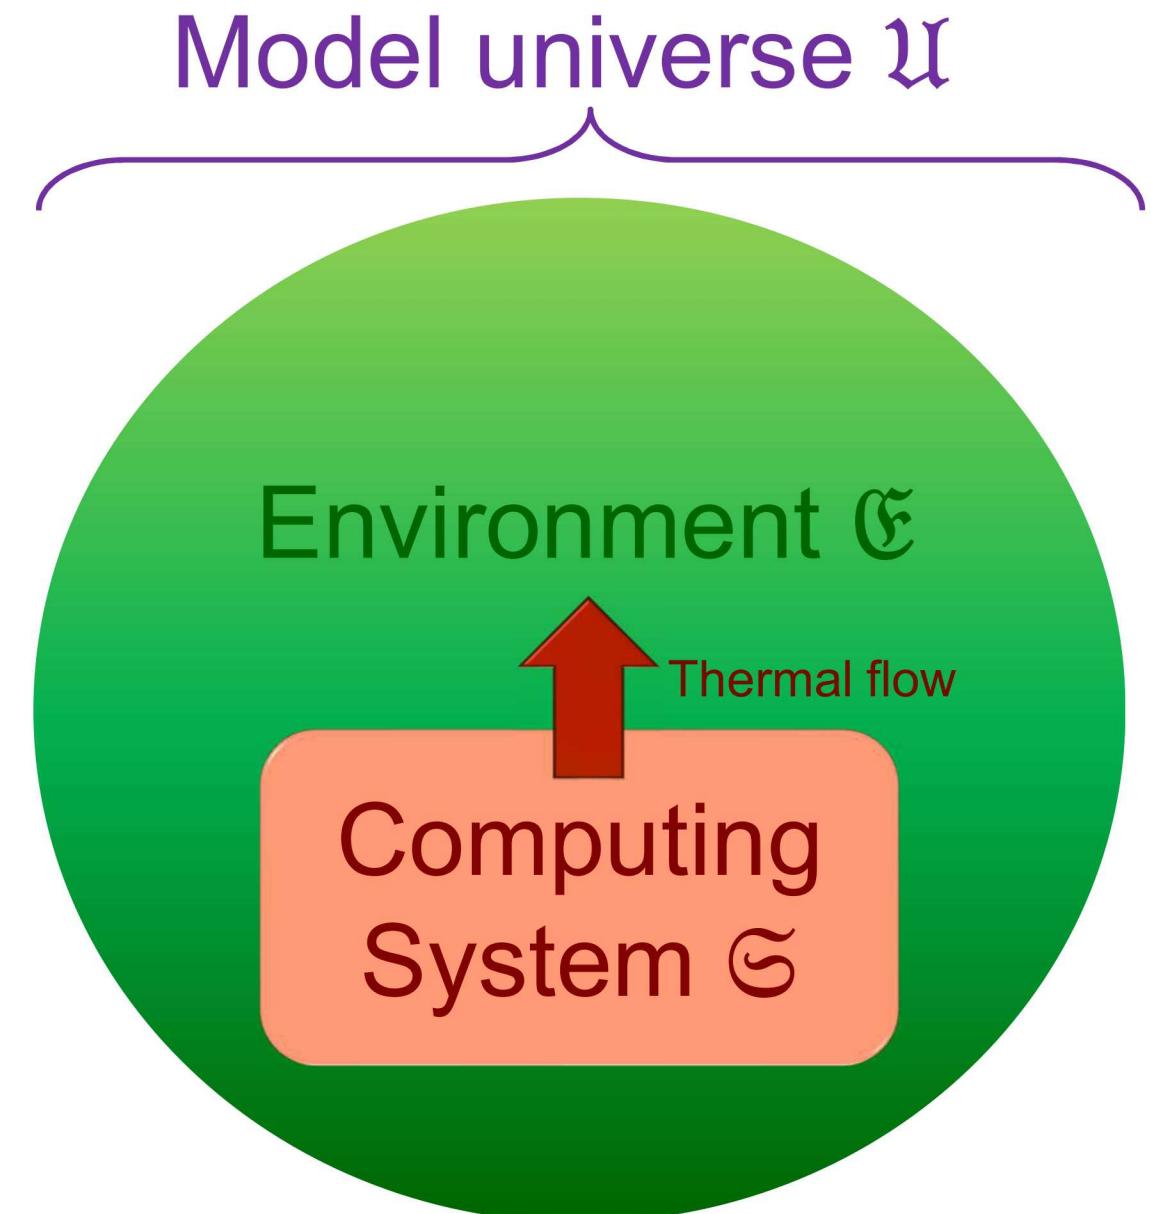
- System \mathfrak{S} can contain its own free energy supply.
 - *E.g.*, a battery
- System \mathfrak{S} can expel waste heat to the environment \mathfrak{E} .
 - Although to some extent, this is all just a formal convenience...
 - A *large enough* closed system can approximate an open one.

Assume the universe’s Hilbert space factorizes:

- $\mathcal{H}_{\mathfrak{U}} = \mathcal{H}_{\mathfrak{E}} \otimes \mathcal{H}_{\mathfrak{S}}$

Also important assumption: Correlations between \mathfrak{S} and \mathfrak{E} aren’t effectively tracked or maintained...

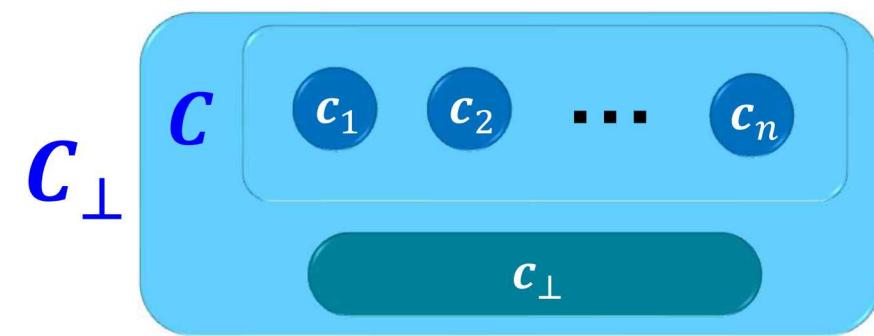
- $\therefore \rho_{\mathfrak{U}} = \rho_{\mathfrak{E}} \otimes \rho_{\mathfrak{S}}$
 - Effectively, after a short thermalization timescale after emitting energy ΔQ .



Computational State Abstraction



There is no need to worry yet about exactly *how* the digital computing machine will be organized, at this early stage of modeling...

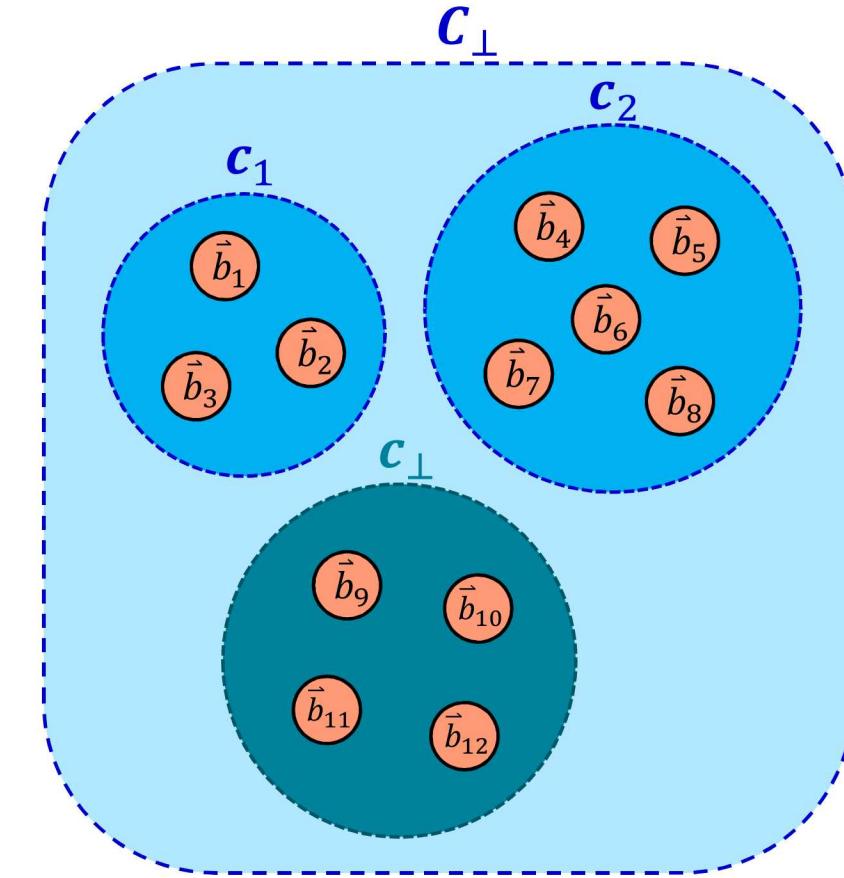


- Instead, just say that, at any given time, we can define some number n of valid discrete *computational states* $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_n$ that the machine could theoretically be in
 - The set of *all* of these defined computational states is called \mathbf{C} .
 - We can add a single extra “dummy computational state” \mathbf{c}_\perp to represent the generic circumstance that the system is not currently strictly occupying *any* of these defined computational states.
 - The computer might be broken, vaporized, or just one of its state bits might just be a *little* bit outside of its defined error margins.
 - The set of *all* “computational states” *including* \mathbf{c}_\perp is called \mathbf{C}_\perp .
 - Call this “the augmented computational state set.”

Proto-Computational Basis

A *proto-computational basis* for the computing system \mathfrak{S} just means *any* appropriate orthonormal basis \mathcal{B} for the system's Hilbert space $\mathcal{H}_{\mathfrak{S}}$ such that the exact computational state $\mathbf{c} \in \mathcal{C}_{\perp}$ is consistently and unambiguously determined by a complete projective measurement of the *quantum* state of the machine onto the basis \mathcal{B} .

- Given such a \mathcal{B} , it follows that \mathcal{C}_{\perp} can then be identified with a set-theoretic *partition* of the set \mathcal{B} .



Example of a computational state space \mathcal{C} consisting of 3 distinct computational states $\mathcal{c}_1, \mathcal{c}_2, \mathcal{c}_{\perp}$, each defined as an equivalence class of basis states in \mathcal{B} . The catch-all state $\mathcal{c}_{\perp} = \mathcal{B} - \bigcup_{i=1}^n \mathcal{c}_i$ is considered computationally invalid.

Computational and Non-Computational Subspaces



In some cases, we may be able to factor the Hilbert space $\mathcal{H}_{\mathfrak{S}}$ of the computing system into separate Hilbert spaces for “computational” and “non-computational” degrees of freedom:

- $\mathcal{H}_{\mathfrak{S}} = \mathcal{H}_{\mathfrak{C}} \otimes \mathcal{H}_{\mathfrak{NC}}$

But, this may not always be possible!

- *E.g.* different states of the computational subsystem \mathfrak{C} may put differing constraints on the \mathfrak{NC} system, such that we can’t properly describe the \mathfrak{NC} system by itself using the very *same* Hilbert space for all cases.

However, in such cases, we can still represent $\mathcal{H}_{\mathfrak{S}}$ via a *subspace sum*:

- $\mathcal{H}_{\mathfrak{S}} = \bigoplus_{i=0}^n \mathcal{H}_{\mathfrak{NC}}^{(j)}$

where $\mathcal{H}_{\mathfrak{NC}}^{(j)}$ is a Hilbert space for the \mathfrak{NC} subsystem that is applicable when the \mathfrak{C} state happens to be \mathbf{c}_j . (Let $\mathbf{c}_0 = \mathbf{c}_{\perp}$.)

Computing System (\mathfrak{S})

Non-Computational Subsystem (\mathfrak{NC})

Computational Subsystem (\mathfrak{C})

Thanks to Karpur for this insight!

Time-Dependent Description

Computational states are *discrete*, but (physical) systems evolve *continuously*; therefore, the important entities in this model have to be treated in a time-dependent way in general:

- $\mathcal{C}(t)$ - The discrete set of computational states defined for the computing system at time t .
 - The cardinality of this set may jump discontinuously at some times.
- $\mathcal{B}(t)$ - A protocomputational basis for $\mathcal{H}_{\mathfrak{S}}$ that could be used to unambiguously measure the computational state at time t .
- $\vec{b}_i(t)$ – The i^{th} particular protocomputational basis state in $\mathcal{B}(t)$.
- $\mathcal{c}_j(t) \subseteq \mathcal{B}(t)$ - The subset of protocomputational basis states that corresponds to the j^{th} computational state in $\mathcal{C}(t)$, in some arbitrary enumeration of computational states.
- $\mathcal{c}_{\perp}(t) = \mathcal{B}(t) - \bigcup_{j=1}^n \mathcal{c}_j$ - The subset of protocomputational basis states for which the computational state is undefined at time t .
- $\mathcal{C}_{\perp}(t) = \mathcal{C}(t) \cup \{\mathcal{c}_{\perp}(t)\}$ – The augmented computational state space at time t ;
 - As before, this is a partition of $\mathcal{B}(t)$. The number and sizes of the subsets may jump discontinuously.
- $\mathbf{c}(t)$ - The actual computational state at time t (if fully decohered).

Fundamental Theorem of the Thermodynamics of Computation

NOTE: I distinguish this from Landauer's Principle proper.



Let $\phi \in \mathcal{H}_{\mathfrak{S}}$ represent a microstate (pure quantum state) of the computing system \mathfrak{S} .

- Let ϕ be hypothetically sampled by applying a complete projective measurement of \mathfrak{S} onto some protocomputational basis \mathcal{B} .
 - Thus ϕ can be identified as ϕ_i , corresponding to some $\bar{b}_i \in \mathcal{B}$.
- The probability distribution $p(\phi_i)$ is given as usual by the Born rule, or equivalently by the diagonal elements of the $\rho_{\mathfrak{S}}$ density matrix in the \mathcal{B} basis.

Note that the distribution $p(\phi_i)$ implies a *derived* distribution over the *computational* states:

$$P(c_j) = \sum_{\phi_i \in c_j} p(\phi_i).$$

And, the total entropy of the physical system (random variable Φ for the state ϕ) can always be written as $S(\Phi) = H(C) + S(\Phi|C)$,

- where C is a random variable for the *computational* state, and S, H are the entropies based on the probability distributions p, P respectively.

Fundamental Theorem Illustrated



The total entropy of any given computing system \mathfrak{S} can always be partitioned as a sum of the entropies associated to its computational vs. non-computational subsystems.

- In this picture, we are implicitly imagining hypothetically sampling \mathfrak{S} by measuring it in an appropriate protocomputational basis \mathcal{B} ...
- When this is not the case, or at times when the system becomes (perhaps briefly) entangled with its environment \mathfrak{E} , we need to be a little bit more careful.
 - Karpur's approach will aim to be a little bit more general here.

Computing System (\mathfrak{S}),
total entropy $S(\Phi) = -\sum p \log p$

Non-Computational
Subsystem ($\mathfrak{N}\mathfrak{C}$)

non-computational /
conditional entropy

$$S_{nc} = S(\Phi|C) = S(\Phi) - H(C)$$

Computational
Subsystem (\mathfrak{C})

$$\text{info. entropy } H(C) = -\sum P \log P$$



For our purposes, a (classical) *computational operation* O on a computational state set \mathbf{C} is a (potentially stochastic) map:

$$O: \mathbf{C} \rightarrow \mathcal{P}(\mathbf{C})$$

- Maps each initial state $c_i \in \mathbf{C}$ to a corresponding probability distribution $P_i \in \mathcal{P}(\mathbf{C})$ over final states.

A computational operation O is called *deterministic* (for our purposes) when the final state entropy $H(P_i) = 0$ for all i .

- Also we can have that O is just *deterministic over a subset* $\mathbf{A} \subset \mathbf{C}$ of initial states, but not the whole set \mathbf{C} .
- If O is not deterministic, we call it *stochastic*.
 - So as not to be confused with the computer science meaning of *nondeterministic*.

Logically Reversible Operations



We say that an operation O is (*unconditionally, logically*) reversible if and only if there is no final state $c_k \in \mathbf{C}$ that is reachable from two different c_i, c_j ($i \neq j$), i.e., where:

$$P_i(c_k) \neq 0 \text{ and } P_j(c_k) \neq 0.$$

- Otherwise, we say that O is *logically irreversible*.

We say that O is *conditionally (logically) reversible under the precondition that the initial state $c \in A$* , for some $A \subseteq \mathbf{C}$, if and only if there is no final state $c_k \in \mathbf{C}$ reachable from two $c_i, c_j \in A$ ($i \neq j$).

- Although it's not very widely known, it's only this weaker, *conditional* form of reversibility (given a context where the precondition is guaranteed to be satisfied) that's required to avoid the information loss that causes necessary dissipation under Landauer's Principle!
- Models of reversible (as well as quantum) computing can be generalized in ways that take advantage of this.

Time-Dependent Case

We can write O_s^t to denote a computational operation being applied over the time interval between starting time $s \in \mathbb{R}$ and terminating time $t \in \mathbb{R}$, with $t > s$:

$$O_s^t: \mathcal{C}(s) \rightarrow \mathcal{P}(\mathcal{C}(t))$$

The remaining definitions (for determinism, reversibility, etc.) change correspondingly in the obvious ways.



Intermission

Fundamental Physics of Reversible Computing—An Introduction

An Important Remark on Landauer's Principle



Some authors identify the *Fundamental Theorem of the Thermodynamics of Computing* (described earlier) with Landauer's Principle, but I would argue that to make that identification is, properly speaking, an misleading misapprehension, which misses certain key points gained through a *proper* understanding of Landauer.

- The *Fundamental Theorem* merely has to do with how we happen to group physical states into computational states, and how total entropy can always be expressed as a sum of computational and non-computational entropies.
 - Merely changing the grouping (e.g. merging or unmerging computational states) does not inherently *increase* total entropy!
 - It only moves it (potentially reversibly) between nominally computational versus non-computational forms.
 - E.g. raise/lower a potential energy barrier separating two degenerate states
 - But I would argue that the statement that is more *properly* called **Landauer's Principle** is actually a very *different* kind of statement, about a somewhat more complicated theoretical setup, namely this:
 - If we take (all, or part of) a (fully or partially) *known* computational state, or (as a special case of this) an unknown state but with *known correlations*, and we *obliviously allow some of the previously-known information to thermalize*, that is, to become *more uncertain*, then this *uncertainty increase* represents (quite immediately, by definition!) a net increase in absolute entropy.
 - The prototypical case, namely, oblivious erasure of a deterministically computed (and therefore, correlated) computational bit, is then just a trivial special case of this—since a correlated bit lost to a thermal environment is quickly thermalized.
 - To make the above statement mathematically precise and rigorous takes just a little bit more work (and some information theory), but is straightforward.

Proof of Landauer's Principle (example for correlated-subsystems case)

Let X, Y be state variables corresponding to *any two disjoint* computational subsystems $\mathfrak{X}, \mathfrak{Y}$ within a larger computer \mathfrak{C} .

- There is a joint probability distribution $P(X, Y)$, and a corresponding joint entropy $H(X, Y)$.
- Reduced entropies $H(X), H(Y)$ of the individual subsystems are defined in the usual way.

The mutual information between \mathfrak{X} and \mathfrak{Y} is defined as:

$$I(X; Y) \stackrel{\text{def}}{=} H(X) + H(Y) - H(X, Y).$$

Now, define the *independent entropy in \mathfrak{Y}* as the *rest of \mathfrak{Y} 's* (reduced subsystem) entropy, *besides* the mutual information I that \mathfrak{Y} has with \mathfrak{X} :

$$S_{\text{ind}}(Y) \stackrel{\text{def}}{=} H(Y) - I(X; Y) = H(Y|X),$$

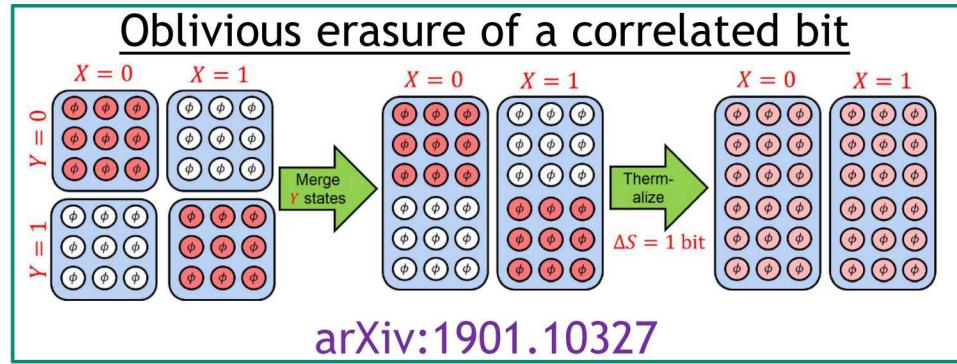
- Same thing as the conditional entropy of \mathfrak{Y} , conditioned on \mathfrak{X} .

Now, consider *erasing \mathfrak{Y}* via *any oblivious* physical mechanism...

- Meaning, set $H(Y) = 0$ *unconditionally*, without reference to X *or to any other information we may have about Y* .
 - E.g., remove an energy barrier separating $Y = 0$ and $Y = 1$ computational states, and call the merged state $\bar{Y} = 0$.

And assume, in general, *non-computational information will fairly rapidly thermalize*. (If not, then why even consider it non-computational?)

- This thermalization process is when/where the absolute entropy increase happens in Landauer!
 - By assumption, environment evolution is not tracked, ergo \mathfrak{C} - \mathfrak{E} correlation is lost.



Note that we could try to “reverse” the whole erasure process to restore the original reduced entropy $H(Y)$ of the \mathfrak{Y} subsystem...

But now, $I(X; Y)_{\text{new}} = 0$ (any correlations have become lost!)

- $\therefore S_{\text{ind}}(Y) = H(Y)$, $\therefore \Delta S_{\text{ind}}(Y) = I(X; Y)_{\text{orig}} = \Delta S_{\text{tot}}$.

If, originally, Y was (*deterministically*) computed from X , then:

- $H(Y|X)_{\text{orig}} = 0$, i.e., $S_{\text{ind}}(Y) = 0$, so $H(Y) = I(X; Y)_{\text{orig}}$.
 - *Apparent* entropy of *all* computed bits is *actually* *entirely* mutual information!
 - a.k.a. “information-bearing entropy” in Anderson’s terminology

Independent entropy (and total universe entropy!) has *increased* by

$$\Delta S_{\text{tot}} = \Delta S_{\text{ind}}(Y) = I(X; Y)_{\text{orig}} = H(Y).$$

\therefore Erasing computed (as opposed to *random!*) bits turns their *digital* information into new physical entropy.

Q.E.D.! ■

Basic Reversible Computing Theory

(For full proofs, see [arxiv.org:1806.10183](https://arxiv.org/abs/1806.10183))

Fundamental theorem of traditional reversible computing:

- A deterministic computational operation is (unconditionally) non-entropy-ejecting if and only if it is *unconditionally* logically reversible (injective over its entire domain).

Fundamental theorem of generalized reversible computing:

- A *specific* (contextualized) deterministic computation is (specifically) non-entropy-ejecting if and only if it is *specifically* logically reversible (injective over the set of *nonzero-probability* initial states).
- Also, for any deterministic computational operation, which is conditionally reversible under some assumed precondition, then the entropy required to be ejected by that operation approaches 0 as the probability that the precondition is satisfied approaches 1.

Bottom line: To avoid requiring Landauer costs, it is *sufficient just to have reversibility when some specified preconditions are satisfied*.

- Basis for practical engineering implementations.
- We'll see some examples in tomorrow's talk.

Traditional Unconditionally Reversible “Gates” (Operations)



NOT
(in-place)



cNOT
(Toffoli)

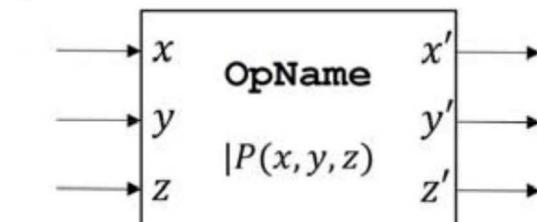


ccNOT
(Toffoli)

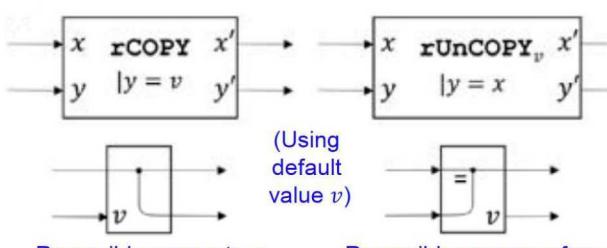


cSWAP
(Fredkin)

Generalized Conditionally Reversible Operations

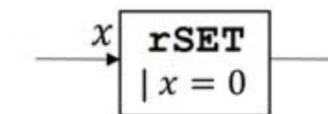


Generic symbol for 3-variable operation

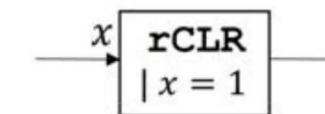


(Using
default
value v)

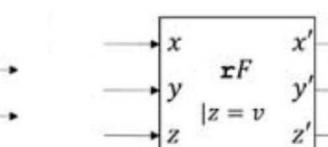
Reversible copy x to y Reversible uncopy y from x



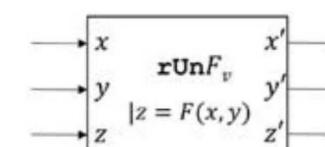
Reversible set-to-one



Reversible clear-to-zero



Reversible
do/undo any
function F ,
w.r.t. default
value of v



Reversible
do/undo any
function F ,
w.r.t. default
value of v

Physical Implementation of Computational Operations



Consider the universe \mathfrak{U} .

- The computer system \mathfrak{S} together with its surrounding environment \mathfrak{E} .
- Let the joint Hilbert space $\mathcal{H}_{\mathfrak{U}} = \mathcal{H}_{\mathfrak{E}} \otimes \mathcal{H}_{\mathfrak{S}}$.

Consider a computational operation O_s^t taking place within \mathfrak{S} .

- Between starting time $s \in \mathbb{R}$ and terminating time $t \in \mathbb{R}$, with $s < t$.

Assuming perfect knowledge of physics, the transformation of $\mathcal{H}_{\mathfrak{U}}$ from time s to t is described by some time evolution operator $U_s^t = U_s^t(\mathfrak{U})$ that applies for \mathfrak{U} between those times.

- In general, the final density matrix $\rho_t = U_s^t \rho_s (U_s^t)^\dagger$.

Note that U_s^t describes the effect of *all* physical processes taking place within \mathfrak{U} , including:

- Dynamical evolution of the physical computational mechanisms in \mathfrak{S} .
- Delivery of needed free energy to the active computing elements in \mathfrak{S} .
- Thermal flows of dissipated energy out into the environment \mathfrak{E} .

We can call this the *open system case*.

Closed-System Variant



Simplified *vs.* open system model, but still physically realistic.

- A real computer *could* actually be operated as an *approximately* closed system for some limited period of time.
 - Until internal energy stores run out, and/or enclosure overheats.

So now, restrict our attention to the subspace of $\mathcal{H}_{\mathfrak{U}}$ that is the Hilbert space $\mathcal{H}_{\mathfrak{S}}$ of the computer system itself.

- Ignore, temporarily, any thermal flows across the \mathfrak{S} - \mathfrak{E} boundary.
 - Imagine that \mathfrak{S} is wrapped in a perfect thermal insulating barrier.
- Now model the effect of the dynamics within \mathfrak{S} as being described by a *local* unitary time-evolution operator $U_s^t(\mathfrak{S})$ operating on $\mathcal{H}_{\mathfrak{S}}$.

Note, the change in the protocomputational basis \mathcal{B} between times s and t can also be modeled by a unitary matrix, $\frac{\mathcal{B}(t)}{\mathcal{B}(s)}U$.

- Then denote a “basis-corrected” version of $U_s^t(\mathfrak{S})$ as:

$$U_s^t(\mathfrak{S}, \mathcal{B}) = \frac{\mathcal{B}(t)}{\mathcal{B}(s)}U \cdot U_s^t(\mathfrak{S}).$$

Quantum Statistical Operating Contexts



This generalizes the concept of a *statistical operating context* or initial probability distribution P that is needed to define a *statistically-contextualized* computational process.

Define as a mixed state ρ_s encompassing all of our uncertainty, as modelers, regarding the initial quantum state of the system at time s , prior to performing the computational operation O_s^t .

Also require that ρ_s is *block-diagonal* in the initial basis $\mathcal{B}(s)$.

- And, the blocks need to correspond to the initial partition $\mathcal{C}(s)$.
 - I.e., no quantum coherences should exist between the different computational states.
 - Formally: ρ_s has no off-diagonal terms between basis states $\vec{b}_1, \vec{b}_2 \in \mathcal{B}(s)$ where $\vec{b}_1 \in \mathcal{C}_i$ and $\vec{b}_2 \in \mathcal{C}_j$ for $\mathcal{C}_i, \mathcal{C}_j \in \mathcal{C}(s)$ where $i \neq j$.
 - This constraint is needed for modeling classical computation.
 - Can weaken when extending this framework to the quantum case.
 - There *can* be coherences *within* a computational state, though...
 - This could correspond, *e.g.*, to physical qubits that may exist within the machine (*e.g.*, long-lived nuclear spins in supporting materials) that are unrelated to the classical digital data being represented.

Quantum Contextualized Computations



This generalizes the concept of a *statistically* contextualized computation $\mathcal{C}(O, P)$ from the Generalized Reversible Computing paper (arxiv.org:1806.10183).

A *(quantum contextualized) computation* $\mathcal{C}_s^t(O_s^t, \rho_s)$ refers to the act of performing the computational operation O_s^t within the computer system \mathfrak{S} when the initial mixed state of \mathfrak{S} at time s is given by a quantum statistical operating context ρ_s .

- Must meet the conditions from the previous slide for $\mathcal{B}(s)$ and $\mathcal{C}(s)$.

What it means to physically implement a given (classical) computation



The basis-adjusted time-evolution operator $U_s^t(\mathfrak{S}, \mathcal{B})$ implements the quantum contextualized computation $\mathcal{C}_s^t(O_s^t, \rho_s)$, written:

$$U_s^t(\mathfrak{S}, \mathcal{B}) \Vdash \mathcal{C}_s^t(O_s^t, \rho_s),$$

- if and only if the density matrix $\rho_t = U_s^t(\mathfrak{S}, \mathcal{B})\rho_s U_s^t(\mathfrak{S}, \mathcal{B})^\dagger$ that results from applying the unitary $U_s^t(\mathfrak{S}, \mathcal{B})$ to the initial mixed state ρ_s has the following property:
 - For any initial computational state $c_i(s) \in \mathcal{C}(s)$ that has nonzero probability under ρ_s , if we zero out all elements of ρ_s outside the set of rows/columns corresponding to $c_i(s)$ and renormalize, and then apply $U_s^t(\mathfrak{C}, \mathcal{B})$ to this restricted ρ'_s , the resulting final mixed state ρ'_t implies the same probability distribution $P_i(t)$ over final computational states in $\mathcal{C}(t)$ as is specified by applying the stochastic map O_s^t to the initial state, $O_s^t(c_i(s))$.

Note: It can OK, under this definition, if small coherences *temporarily* arise between different final computational states in $\mathcal{C}(t)$,

- as long as the subsequent evolution causes them to decay very quickly.
- That is, we don't want these “parasitic” coherences to impact the dynamics of subsequent operations.

Some next steps for this framework looking forwards...



Show some specific examples of time-dependent $\mathbf{B}(\tau)$, $\mathbf{C}(\tau)$ and basis-adjusted unitaries $U_s^t(\mathfrak{S}, \mathbf{B})$ that meet the above definition of the “implements” operator $\mathbb{I}\mathbb{H}$ for the case of desired classical reversible operations O_s^t .

- And, show that some such unitaries can in turn be (at least approximately) implemented by real, buildable physical mechanisms.

Characterize the departures from ideality of theoretically realizable approximate physical implementations of reversible operations O_s^t in terms of the resulting increase ΔS in total entropy.

- Can we derive a general lower bound on ΔS that depends on simple parameters such as the time delay $d = t - s$ to perform the operation?

Generalize the above theoretical treatment as needed to address the (marginally more realistic) open-system case.

Likharev's dissipation limits

Likharev '82 analyzed limits of dissipation for his reversible JJ-based Parametric Quantron (PQ) technology concept.

- Based on analyzing rates of crossing a potential energy barrier through thermal excitation and quantum tunneling.

Main results:

- Limit due to classical thermal excitation over barrier (assuming underdamped junction):

$$W_C \approx \frac{k_B T}{\omega_c \tau} \ln \frac{1}{\omega_A \tau p}.$$

- $\omega_c = \frac{k}{\eta} \approx \frac{2\Delta}{\hbar}$ with elasticity modulus $k = \frac{d^2 U}{dx^2}$ and effective viscosity η ; and 2Δ is the superconducting gap energy;
- ω_A approximates to the JJ plasma frequency $\omega = \sqrt{k/\eta} = \sqrt{2q_e I_c/\hbar C}$, and τ is the cycle period;
- p is the tolerable error probability per operation.

Limit due to quantum-mechanical tunneling through the barrier:

$$W_C \approx \frac{\hbar}{\tau} \ln \frac{1}{\omega_c \tau p}.$$

However! Likharev himself admits the limitations of this analysis:

- It is not a *fundamental*, technology-independent analysis.
- Alternative device concepts might do better!

International Journal of Theoretical Physics, Vol. 21, Nos. 3/4, 1982

Classical and Quantum Limitations on Energy Consumption in Computation

K. K. Likharev

Department of Physics, Moscow State University, Moscow 117234, U.S.S.R.

Received May 6, 1981

Our approach to the problem, of course, leaves open whether it is possible to invent some novel device providing lower power consumption. If we limit ourselves to the quasistatic devices, where the computation can be stopped at any moment, without inducing an error, one can hardly get away from the above estimates. In fact, the only role of the parametric quantron in our discussion has been to demonstrate how a flexible bistable potential well could be physically realized. (Of course, some numerical factors can appear in the estimates if peculiar well shapes are taken into account.)

One can, however, argue that the above-mentioned condition of quasistatics is by no means compulsory, and that the information can be processed by some "dynamical" devices, where the cycle period can be shorter than the relaxation time. This problem is left for further analysis.

Feynman's dissipation limits

In lectures for his 1983-1986 CalTech course, “Potentialities and Limitations of Computing Machines,” Feynman derived a limit on energy dissipation per step for *Brownian machines* (e.g., DNA copying) driven by chemical potentials.

- He concludes that an approximate formula for this is:

$$\text{energy loss/step} = kT \frac{\text{minimum time taken/step}}{\text{time/step actually taken}}$$

However, he mentions in a footnote that a “slight correction” to this expression would be needed for *ballistic* machines, and later argues, quite informally, that in that case, the expression should be:

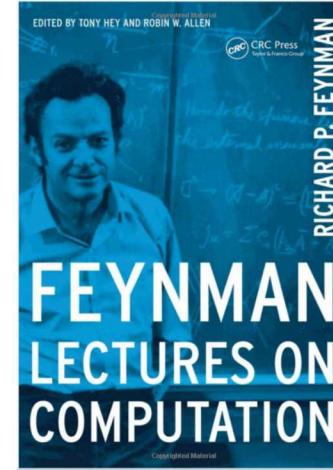
$$kT \frac{\text{time to make collision}}{\text{speed [sic] at which it happens}}$$

$$kT \frac{\text{time to make collision}}{\text{speed at which it happens}}. \quad (5.37)$$

This expression has not been analyzed in any great detail for the billiard ball machine.

- An arguably very similar expression, but:
 - The whole argument in this part of the notes is *extremely* brief and informal (“hand-wavy”)
 - The possible application of e.g. the Landauer-Zener formula for quantum-mechanical scattering processes is not considered at all
 - Modern STA (Shortcuts to Adiabaticity) techniques had not even been developed yet, and so of course are also not considered
 - *Asynchronous* ballistic models (e.g. ABRC) which avoid chaotic instabilities had also not been invented yet

Thus, we must conclude that Feynman’s analysis of this problem is *not definitive*, nor the final word.



An example we gave of reversible computing was that of the chemical process of copying DNA. This involved a machine (if you like) that progressed in fits and starts, going forward a bit, then backwards, but more one than the other because of some driving force, and so ended up doing some computation (in this case, copying). We can take this as a model for more general considerations and will use this “Brownian” concept to derive a formula for the energy dissipation in such processes. This will not be a general formula for energy dissipation during computation but it should show you how we go about calculating these things. However, we will precede this discussion by first giving the general formula⁷, and then what follows can be viewed as illustration.

⁷This rule is pretty general, but there will be exceptions, requiring slight corrections. We will discuss one such, a “ballistic” computer, in §5.5. [RPF]





Section III: Looking Ahead

Fundamental Physics of Reversible Computing—An Introduction

Key Questions for the Physics of Reversible Computing



Are there fundamental (i.e., technology-independent) lower bounds (greater than zero) that follow from general non-equilibrium quantum thermodynamics on energy dissipation per reversible computational operation as a function of, say, the speed of the operation (and/or other fundamental physical parameters such as size, temperature, *etc.*)?

- And, if so, can such bounds be expressed via simple analytic scaling relations?

Can we deduce anything regarding *e.g.* exotic quantum phenomena, materials properties, *etc.*, that would need to be leveraged in order for a technology to saturate the bounds?

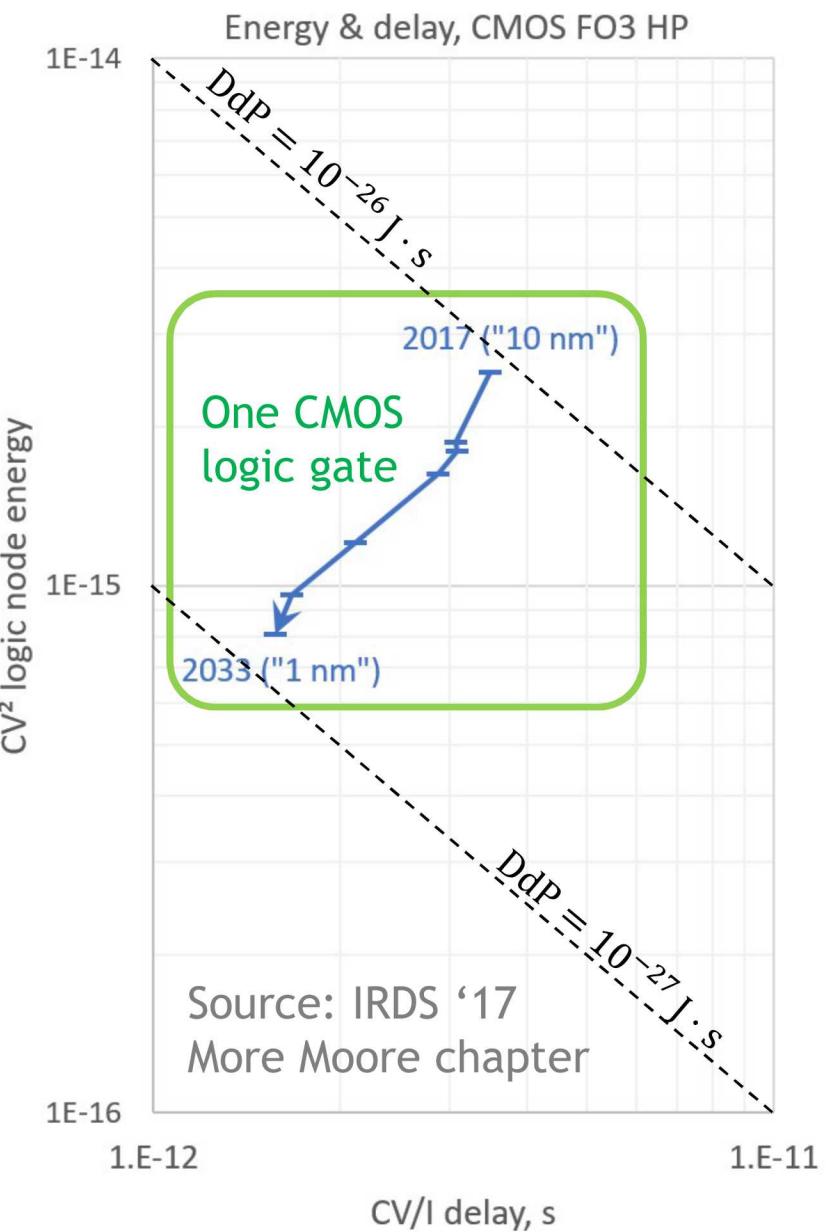
- Examples of quantum phenomena that are (or might be found to be) useful for this:
 - Decoherence-free subspaces (DFSs), Zurek pointer states
 - Topological invariants?
 - *e.g.*, signed flux charge threading a bounded planar superconducting circuit.
 - Dynamical versions of the *quantum Zeno effect* (QZE)?
 - ...*others*???

Answering the above questions can then become a starting point for innovation of breakthrough technologies for reversible computing that exhibit vastly improved engineering characteristics....

Existing Dissipation-Delay Products (DdP) —Non-reversible Semiconductor Circuits

Conventional (non-reversible) CMOS Technology:

- Recent roadmaps (e.g., IRDS '17) show Dissipation-delay Product (DdP) decreasing by only $<\sim 10\times$ from now to the end of the roadmap (~2033).
 - Note the typical dissipation (per logic bit) at end-of-roadmap is projected to be $\sim 0.8 \text{ fJ} = 800 \text{ aJ} = \sim 5,000 \text{ eV}$.
- Optimistically, let's suppose that ways might be found to lower dissipation by an additional $10\times$ beyond even that point.
 - That still puts us at $80 \text{ aJ} = \sim 500 \text{ eV}$ per bit.
- We need at least $\sim 1 \text{ eV} \approx 40 \text{ } kT$ electrostatic energy at a minimum-sized transistor gate to maintain reasonably low leakage despite thermal noise,
 - And, typical *structural* overhead factors *compounding* this within fast random logic circuits are roughly $500\times$,
 - so, $\sim 500 \text{ eV}$ is *indeed* probably about the practical limit.
 - At least, this is a reasonable order-of-magnitude estimate.



Existing Dissipation-Delay Products (DdP)—Adiabatic Reversible Superconducting Circuits

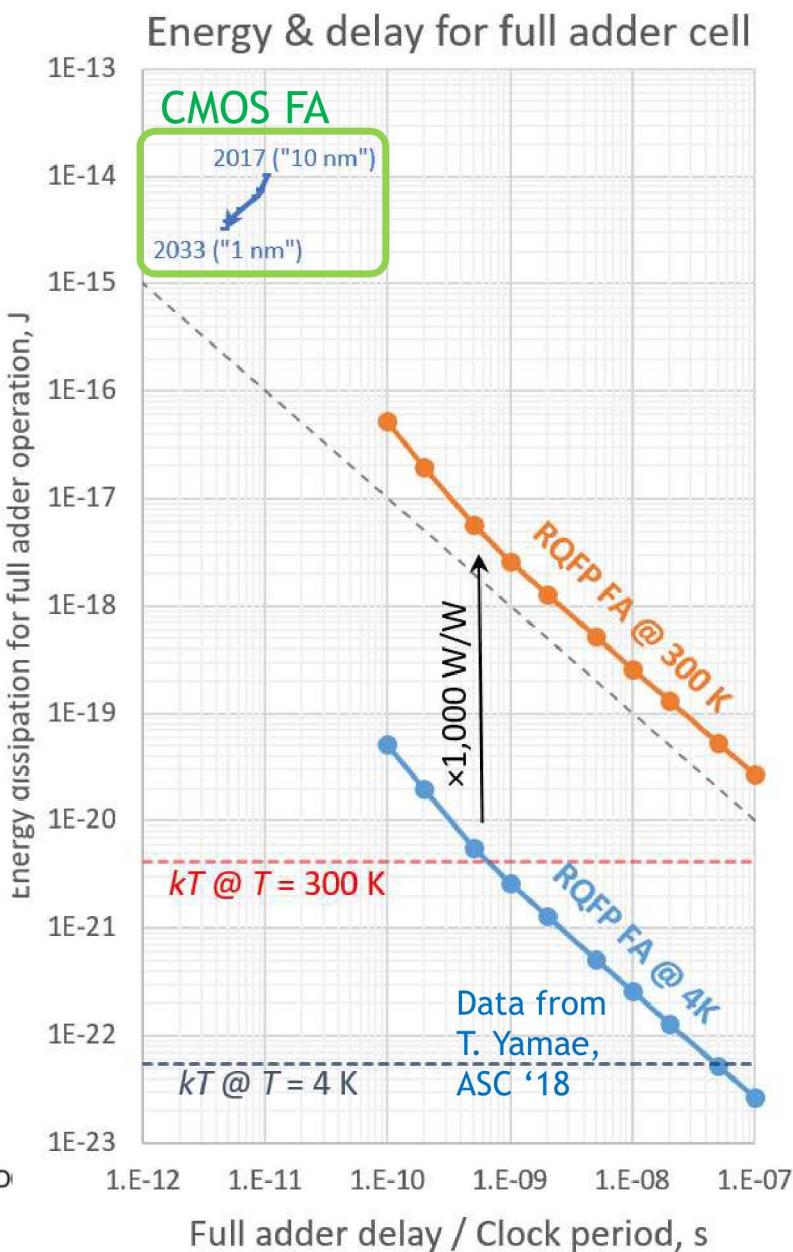
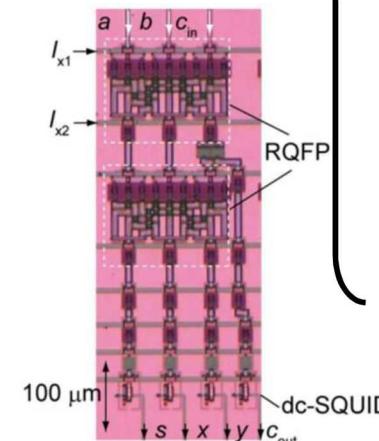
Reversible adiabatic superconductor logic:

- State-of-the-art is the **RQFP** (Reversible Quantum Flux Parametron) technology from Yokohama National University in Japan.
 - Chips were fabricated, function validated.
- Circuit simulations predict DdP is $>1,000\times$ *lower* than even *end-of-roadmap* CMOS.
 - Dissipation extends *far below* the 300K Landauer limit (and even below the Landauer limit at 4K).
 - DdP is *still* better even after adjusting by a conservative factor for large-scale cooling overhead (1,000 \times).

Question: Could some *other* reversible technology do *even better* than this?

- We have a project at Sandia exploring one possible superconductor-based avenue for this...
- But, what are the *fundamental* (technology-independent) limits, if any?

RQFP =
Reversible
Quantum Flux
Parametron
(Yokohama U.)



Can dissipation scale better than linearly with speed?



Some observations from Pidaparthi & Lent (2018) suggest Yes!

- Landau-Zener (1932) formula for quantum transitions in e.g. scattering processes with a missed level crossing...
- Probability of exciting the high-energy state (which then decays dissipatively) scales down *exponentially* as a function of speed...
 - This scaling is commonly seen in many quantum systems!
- Thus, dissipation-delay *product* may have *no lower bound* for quantum adiabatic transitions—*if* this kind of scaling can actually be realized in practice.
 - *I.e.*, in the context of a complete engineered system.
- **Question:** Will unmodeled details (e.g., in the driving system) fundamentally prevent this, or not?

J. Low Power Electron. Appl. 2018, 8(3), 30; <https://doi.org/10.3390/jlpea8030030> Open Access Article

Exponentially Adiabatic Switching in Quantum-Dot Cellular Automata

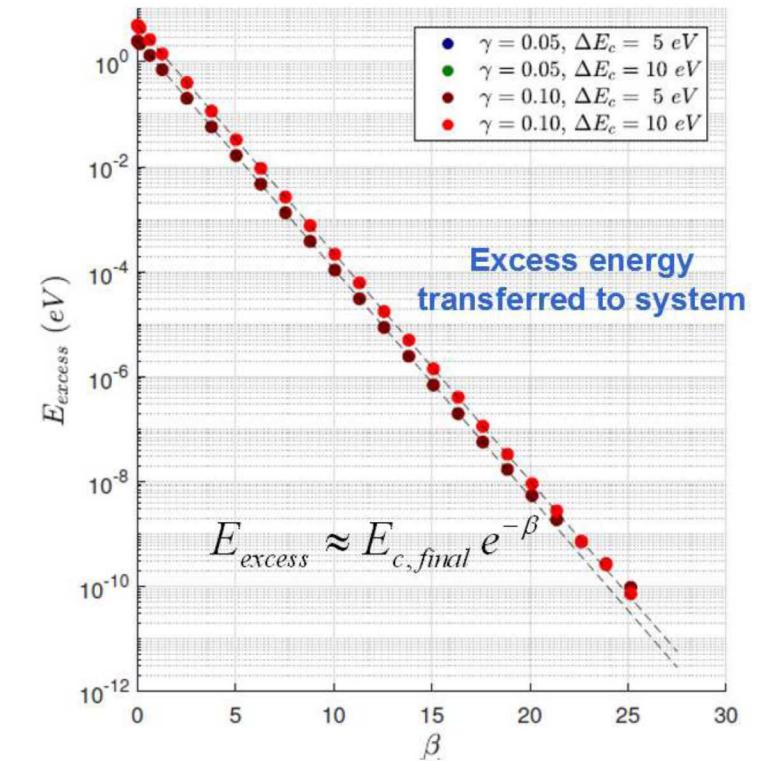
Subhash S. Pidaparthi and Craig S. Lent

Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556, USA

* Author to whom correspondence should be addressed.

Received: 15 August 2018 / Revised: 5 September 2018 / Accepted: 5 September 2018 / Published: 7 September 2018

(This article belongs to the Special Issue Quantum-Dot Cellular Automata (QCA) and Low Power Application)

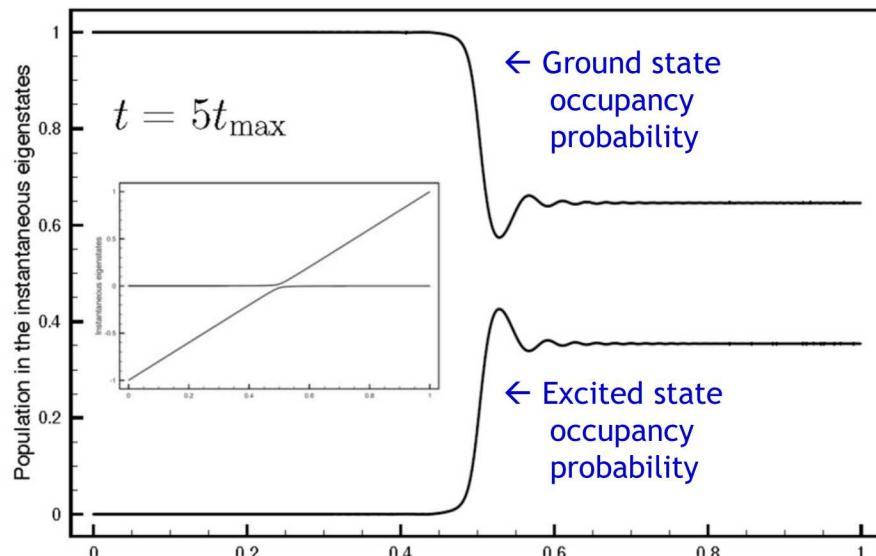


Shortcuts to Adiabaticity (STA)



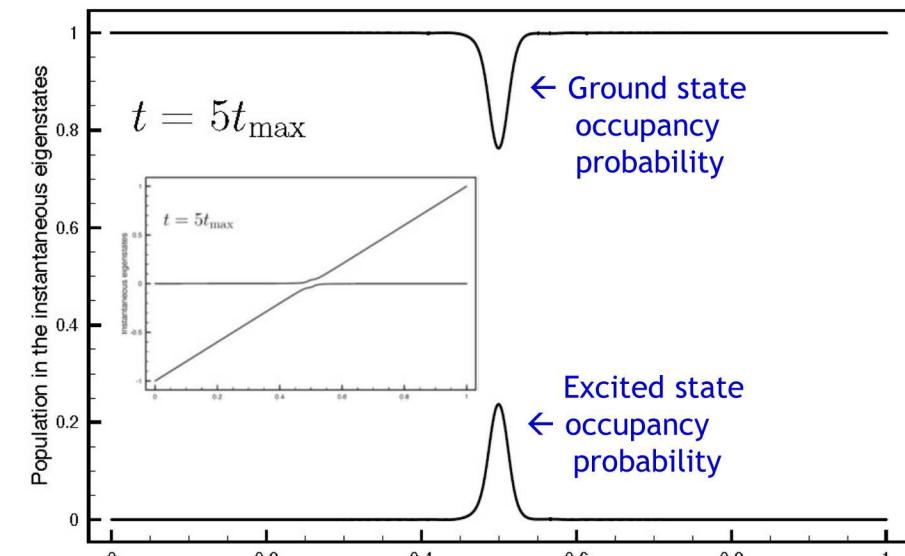
A line of theoretical physics research showing that, *in principle*, quantum state transformations can always be carried out with exactly zero dissipation, even at any given *finite* delay!

- Requires the introduction of a finely-tuned “counterdiabatic” perturbation to the system’s time-dependent Hamiltonian.
- **Again, we ask:** Is this idealized prediction *actually achievable*, if fundamental thermodynamic limits that apply to the complete system are accounted for?



Normal quantum adiabatic process:
Substantial excitation/dissipation

Figure credit:
Collaborator
David Guéry-Odelin
(Université de Toulouse)



Using counterdiabatic protocol:
Zero net excitation/dissipation

Limits to Reversible Computing?

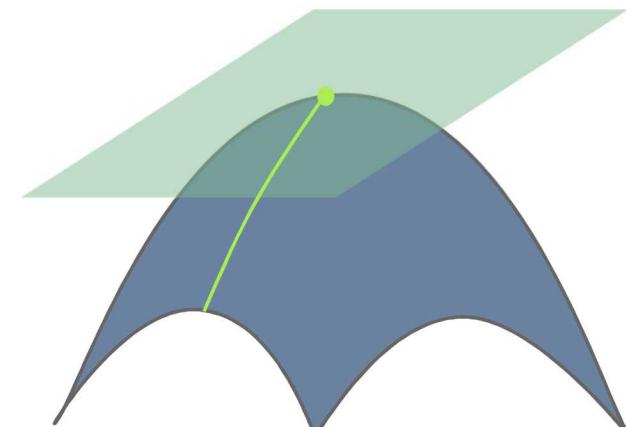
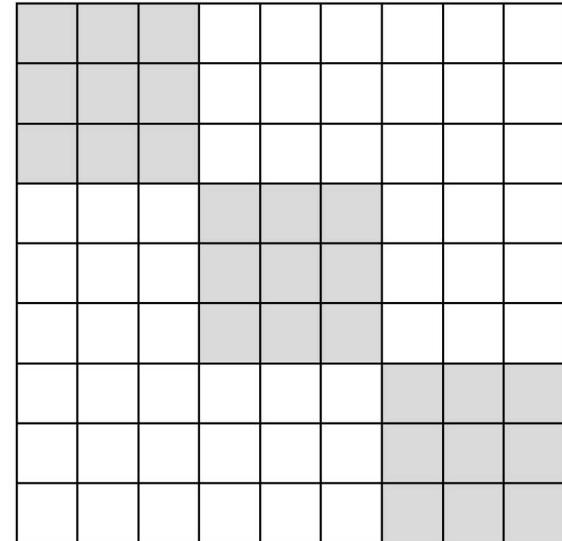
—An approach from the theory of Open Quantum Systems

(Work with Karpur Shukla, Brown University, and Victor V. Albert, CalTech)

- Computational states modelled as *decoherence-free subspace blocks* (DFSB) of overall Hilbert space.
- Quantum Markov equation with multiple asymptotic states: admits subspace dynamics (including DFSB structures) for open systems under Markov evolution.
 - Induces geometric tensor for *manifold of asymptotic states*.
 - Similar to quantum geometric tensor / Berry curvature for closed systems.
 - Current work: use multiple asymptotic state framework to derive thermodynamic quantities...
 - Uncertainty relations, dissipation and dissipation-delay product.



BROWN



Conclusion

Some form of reversible computing is *absolutely required* in order for the energy efficiency (and thus, cost efficiency) of general digital computing to avoid asymptoting against *firm* thermal barriers in the foreseeable future.



- This follows directly from the proper *understanding* of Landauer's Principle (slide 21) that is substantiated by the (rigorous) arguments summarized in slides 22-23.
- Various researchers who have misapprehended Landauer are simply *missing the whole point*.
 - Various critics of Landauer & Bennett have simply failed to appreciate the essential, *unavoidable* role that information-theoretic *correlation* plays in computing, which is the ultimate origin of the *absolute entropy increase* that is rigorously caused by Landauer erasure...
 - ...when the true *meaning* of Landauer's Principle is understood properly!
 - All *conventional* (i.e., non-reversible) digital circuit architectures rely *fundamentally* on frequent *oblivious erasure of correlated bits*, ergo, they can never surpass the Landauer limit (by the elementary proof on slide 22).
 - In contrast, properly-designed *reversible* architectures are designed to *avoid* such erasure, ergo are *not* subject to the Landauer limit.

The fundamental limits of *reversible* computing are still very far from being fully understood...

- There is a significant opportunity for physicists to develop fundamental new results in this area.

Leveraging of exotic quantum phenomena may be required to saturate the fundamental limits.

It seems likely that *breakthrough technologies for reversible computing remain to be discovered*.

- And this, in turn, would lead to incalculable increases in the value of computing, and civilization!