Exceptional service in the national interest

Generalized Reversible Computing and the Unconventional Computing Landscape

Michael P. Frank
Center for Computing Research
Sandia National Laboratories

Presented at the EE Computer Systems Colloquium
Stanford University, Oct. 18, 2017

Approved for Public Release
SAND2017-11167 PE

Sandia National Laboratories is a multi-mission laboratory managed and operated 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-NA0003525.

Talk Abstract (hidden when presenting)

With the end of transistor scaling now in sight, the raw energy efficiency (and thus, practical performance) of conventional digital computing is expected to soon plateau. Thus, there is presently a growing interest in exploring various unconventional types of computing that may have the potential to take us beyond the limits of conventional CMOS technology. In this talk, I survey a range of unconventional computing approaches, with an emphasis on reversible computing (defined in an appropriately generalized way), which fundamental physical arguments indicate is the only possible approach that can potentially increase energy efficiency and affordable performance of arbitrary computations by unboundedly large factors as the technology is further developed.

2

Talk Outline

- Motivation:
  - The end of transistor scaling…
    - and implications for supercomputing performance
- Breaking down the causes of energy dissipation in computing…
  - And some candidate technological approaches for addressing them
    - We can call the range of ideas the "Unconventional Computing Landscape"
- **Reversible computing** as the *only* long-term sustainable solution…
  - Landauer's Principle from fundamental physics
  - Fundamental Theorems of Reversible Computing Theory
    - Traditional & General theorems
  - Example technological implementations of reversible computing
    - Adiabatically-switched voltage-coded logic
    - Superconducting logics (adiabatic & ballistic)
    - Nanomechanical logics
- Conclusion

3

Semiconductor Roadmap is Ending…

- Thermal noise on gates of minimum-width segments of FET gates leads to channel PES fluctuations when $E_g \lesssim$ 1-2 eV
  - Increases leakage, impairs practical device performance
    - Thus, ITRS has minimum gate energy asymptoting to ~2 eV
- Also, real logic circuits incur many further overhead factors:
  - Transistor width 10-20× min.
  - Parasitic (junction, etc.) transistor capacitances (~2×)
  - Multiple (~2) transistors fed by each input to a given logic gate
  - Fan-out to a few (~3) logic gates
  - Parasitic wire capacitance (~2×)
- Due to all these overheads, the energy of each bit in real logic circuits is many times larger than the min.-width gate energy
  - 375-600× (!) larger in ITRS'15
    - ∴ Practical bit energy for irreversible logic asymptotes to ~1 keV!
- Practical, real-world logic circuit designs can't just magically cross this ~500× architectural gap!
  - ∴ Thermodynamic limits imply much *larger* practical limits!
    - The end is near!

Data source: International Technology Roadmap for Semiconductors, 2015 edition

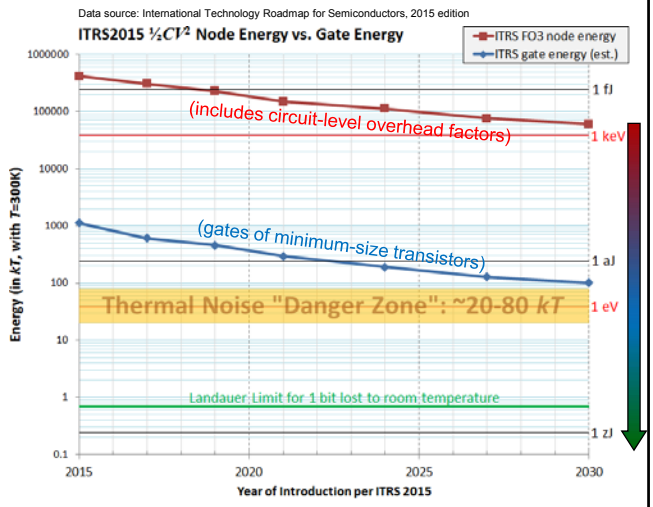

Only reversible computing can take us from ~1 keV at the end of the CMOS roadmap, all the way down to $\ll kT$.

4

Implications for FLOPS & power

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

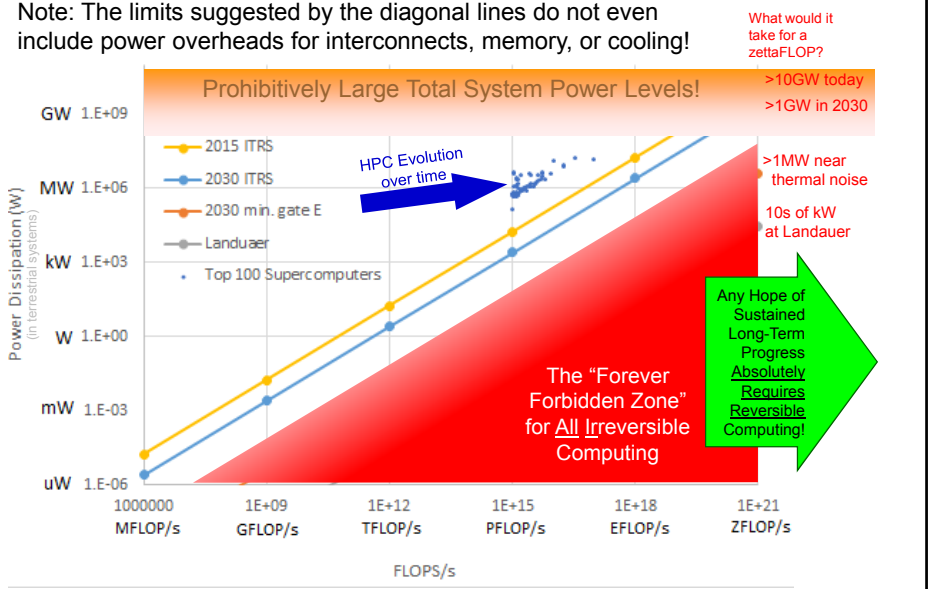

5

Major Sources of Dissipation in Real High-Performance Computing Systems

and some ways to address them…

- Cooling in data centers / machine rooms
  - Upscaling; and/or, site near rivers, in Iceland… or in a crater at a Lunar pole?
- In-cluster networking (between/within cabinets)
  - Optical interconnects will likely become increasingly widely used for this
- Memory systems, and associated communication
  - Research is ongoing on future fast/nonvolatile memory technologies
  - PIM-type architectures can vastly reduce communication overheads
- Board-level, MCM, and chip-level interconnects
  - Optical is being investigated for use at these levels as well
- On-chip memories/caches, SRAM/DRAM – new devices
- Architectural overheads (interpreting ISAs etc.)
  - Application logic migrates to FPGAs, specialized architectures (e.g., neural), ASICs
- Wiring within and between functional units
  - "Steep" low-V switches are being investigated to lower CV^2 energy
- Parasitic overheads of logic gate structures
  - Novel logic families → Mostly played out, except for adiabatic/beyond CMOS
- Gate energy of minimum-sized devices
  - This one is becoming limited by thermal noise → Reason the roadmap is ending

A severely under-appreciated fact: Thoroughly applying reversible / adiabatic principles allows us to scale down power requirements at *all* of these levels simultaneously!

6

Sandia Si Photonics Integrated Circuits

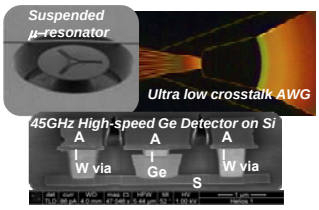

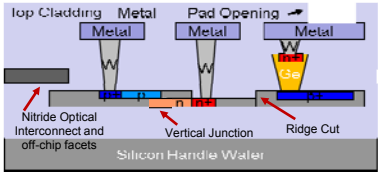

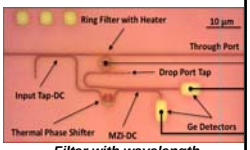

Filter with wavelength stabilization

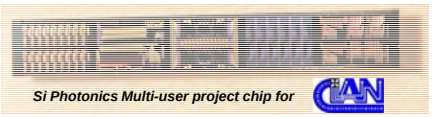

Si Photonics Multi-user project chip for CIAN

Tx/Rx Si Photonics chip

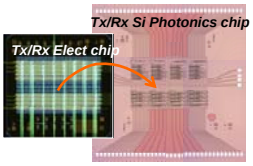

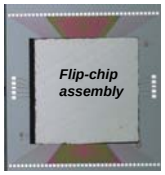

Silicon-based Integrated Devices

- **Components**: Low energy modulators, detectors and electrically controlled optical switches.
- **Heterogeneous integration** with electronics, III-V integrated photonics for lasers, amplifiers, and electro-optic materials
- **Applications**: transmitters, receivers for data links, optical active beam steering, optical logic (matrix multiply), low noise oscillators, optical network add-drop node (CIAN), optical channel monitor (spectrum analyzer) (CIAN)

7

Neural computing at Sandia Labs leverages a large research foundation

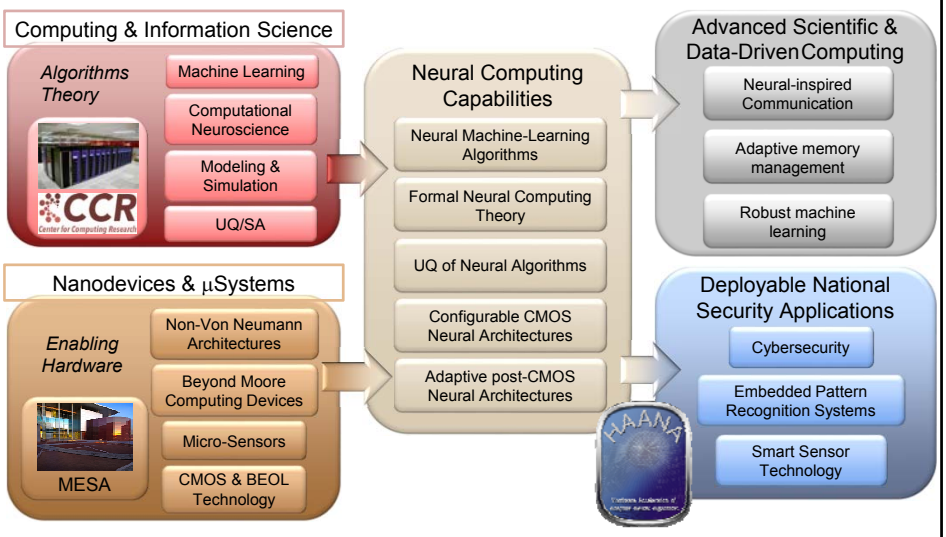

8

Hardware Acceleration of Adaptive Neural Algorithms (HAANA)

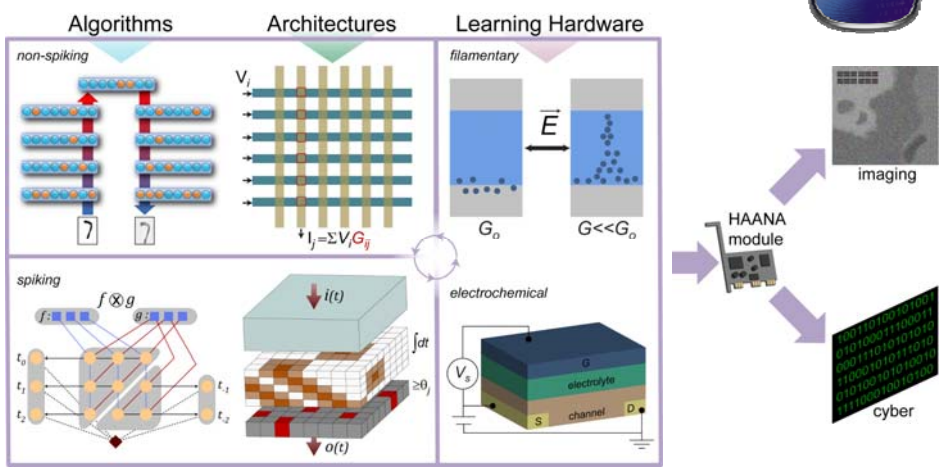

Neurogenesis deep learning: Draelos et al, IJCNN 2017
Spiking network algorithms: Severa et al., ICRC 2016
Digital neuromorphic architecture: Smith et al., IJCNN 2017

Resistive switching model: Mickel et al, Adv Mater 2014
Electrochemical transistor: Fuller et al., Adv Mater 2016
Resistive crossbar accelerator: Agarwal et al., IJCNN 2016

9

Thermodynamics and Information

- Physical ***entropy*** quantifies *uncertainty about the detailed microstate of a system*.
  - First postulated by Boltzmann (H-theorem)
  - Integral to modern physics (Von Neumann entropy)
  - Depends on modeler's state of knowledge (Jaynes)
- The *reversibility (injectivity)* of microphysics <u>underlies</u> the Second Law of Thermodynamics.
  - Entropy of a closed system cannot decrease!
    - Conserved by unitary quantum time-evolution.
  - But, entropy can *increase* if we have any uncertainty about the dynamics, or do not track it in detail
- At the most fundamental level, physical information *cannot be destroyed*.
  - Only *reversibly* transformed, and/or transferred between different subsystems…

$$S[p] = \mathrm{E}_p[\log p^{-1}]$$

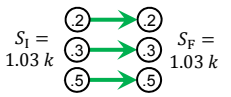

Bijective microphysics → No "true" entropy change

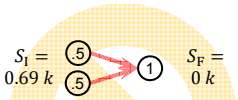

Irreversible microphysics → Entropy would decrease (Second Law of Thermo. would be violated)

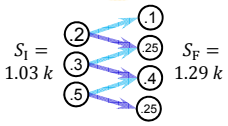

True dynamics uncertain (or not tracked in detail) → Entropy increases

10

Visualizing Entropy of Grouped States

- Can represent a hierarchy of events in a tree structure…
  - Branch ***thickness*** = event probability p.
  - Branch ***length*** = *incremental surprise* Δs associated w. event,
    - relative to whatever base event it's branching off from.
  - Branch ***area*** = event's *incremental heaviness* $\Delta h = p\Delta s$, *i.e.*,
    - its contribution to total entropy, in addition to its base event's.
- *Grouping* events into larger events has these effects:
  - Thicknesses (probs.) of branches combine in parent branch
  - A corresponding part of the total length (surprise) of each branch is reassociated to parent (stem) branch.
  - Note: The total heaviness $\boldsymbol{H}$ of all branches and stems (total entropy $\boldsymbol{S}$) is not changed at all by any grouping/ungrouping!!

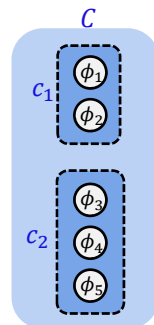

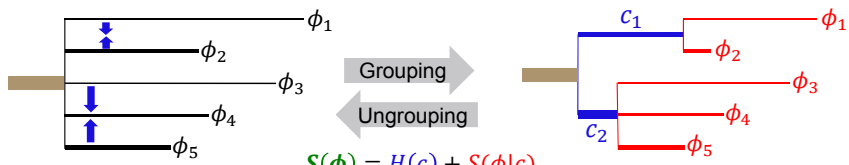

$$\boldsymbol{S}(\boldsymbol{\phi}) = H(c) + S(\phi|c)$$

cf. chain rule of conditional entropy

Total system entropy = computational entropy + non-computational entropy

Proving Landauer's Limit

- Follows directly from thermodynamics!
- A *computational state* c_j is just an equivalence class of physical states ϕ_i
  - On the left we see two computational states c_0, c_1, each with probability 0.5
- The computational subsystem has an induced information entropy $H(c)$.
  - Here, it is $H(c) = \log 2 = 1 \text{ bit} = k \ln 2$.
- The *non*-computational subsystem (everything else) has expected entropy

$$S_{\text{nc}} = S(\phi|c) = S(\phi) - H(c) = S - H$$

  - The conditional entropy of the physical state ϕ, given the computational state c.

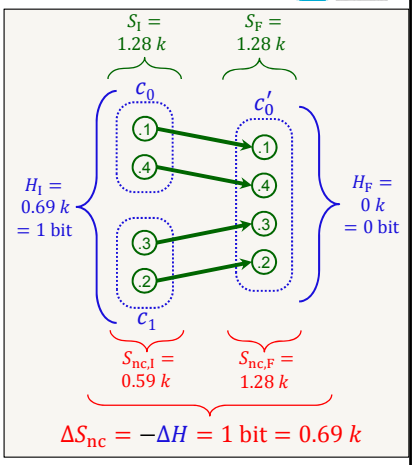

- Thus, if the computational entropy <u>*decreases*</u> (note here $\Delta H = -1 \text{ bit}$),
  - The *non*-computational entropy *must* <u>*increase*</u> by $\Delta S_{\text{nc}} = -\Delta H$ (here, $k \ln 2$).
- Thus, to lose a computational bit, we must eventually dissipate energy $\Delta E_{\text{diss}} = kT \ln 2$ as heat to some environment at some temperature T.

Logically Irreversible, Oblivious Erasure of a Correlated Bit

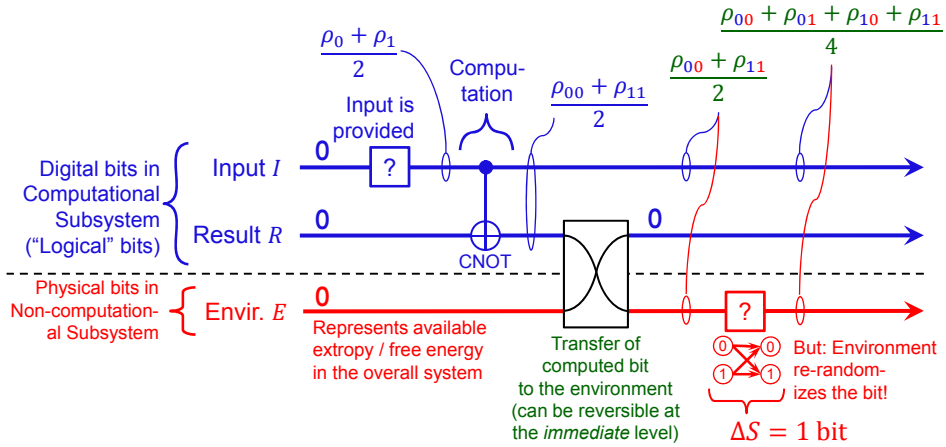

Moving a computed, correlated bit to an (*unpredictable!*) thermal environment *necessarily, inevitably* loses its correlations, and thus increases entropy!

13

Logically Reversible, Non-oblivious Decomputation of a Correlated Bit

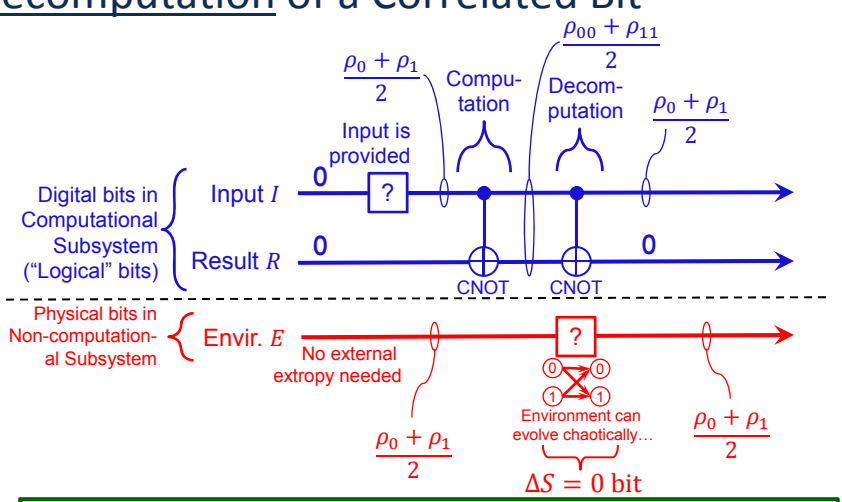

Decomputing correlated bits, instead of ejecting them to the thermal environment, avoids losing correlations & increasing entropy!

14

Reinventing Reversible Computing Theory

(Frank, "Foundations of Generalized Reversible Computing," RC17.)

- Our very first task, in developing the theory of reversible computing, can be characterized as answering this question:
  - What are the necessary and sufficient conditions that must be met, *at the logical level*, in order for a computational process to avoid ejecting entropy from the computational state? (*I.e.*, for $\Delta H \geq 0$.)
    - Or more generally, to approach 0 entropy ejection, $\Delta H \to^{-} 0$.
- Landauer attempted to answer this question in his 1961 definition of ***logical reversibility***, and the entire traditional theory of classical reversible computing has been based on this definition…
  - Only problem: **Landauer's definition is wrong!**
    - In the sense that, in fact, it is provably a *sufficient*, but **not** a *necessary* logical-level condition to avoid entropy ejection from a computation.
      - Let's see why…

… RESTORE TO ONE is an example of a logical truth function which we shall call *irreversible*. We shall call a device *logically irreversible* if the output of a device does not uniquely define the inputs.

… Now assume that the computer is logically reversible. Then the machine cycle maps the 2^N possible initial states of the machine onto the same space of 2^N states, rather than just a subspace thereof.

Devices, Operations, Computations

- We'll distinguish several different concepts:
  - ***Device*** – Can perform one (or more) types of *operations*.
    - A given device has some associated local state information
      - Includes states of I/O terminals, any internal states of device
  - ***Operation*** – a *(computational) operation* is a map O from *initial (computational) states* to *final states* (locally)
    - The terms "input" and "output" are too ambiguous – avoid!
    - We can also consider partial maps (undefined = don't-care)
    - More generally, the operation O could even be a stochastic map…
      - A probabilistic transition rule $r_{ij} = \Pr[c_F = c_{Fj} \mid c_I = c_{Ii}]$.
        - » However, that case is not our main focus at present
  - ***Computation*** – for us, a *computation* is a computational operation performed within a specific operating context
    - Context specifies/constrains the initial state probabilities
      - These are essential for a meaningful thermodynamic analysis!

Avoiding Entropy Ejection: The Traditional (CS theorist's) Way

- Considering what's required of an operation if it is to be non-entropy-ejecting in *all* definable operating contexts leads to the traditional definition of logical reversibility for computational *operations*:
  - **Fundamental Theorem of Traditional Reversible Computing Theory:** A total deterministic operation O is not even *potentially* entropy-ejecting if and only if O is a bijective map (permutation) of the full space C consisting of all of the device's describable computational states.
    - **Definition:** We call such an O an *(unconditionally) logically reversible operation*.

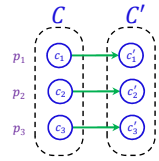

- This theorem leads to the traditional view that classical reversible computations must consist of sequences of complete bijective operations such as "Fredkin gates," "Toffoli gates," "Controlled-NOT" gates, *etc.*, which you may have heard of…
  - This is a very widely-cited model, but, contrary to widespread belief, it is *NOT* the most general, most useful model of reversible computing!!

17

Avoiding Entropy Ejection: The More General (Engineers'!) Way

- Considering requirements for an operation to be non-entropy ejecting in an actual, *specific* operating context P_1 gives a *different* logical reversibility concept suited for *actual* (contextualized) *computations*:
  - **Fundamental Theorem of *Generalized* Reversible Computing Theory:** A total deterministic operation O is not *specifically* entropy-ejecting in a *given* operating context P_1 if and only if O is injective over at least (*i.e.*, when restricting its domain to) the *active subset*

$$A = \{c_i \mid p_i \equiv P_1(c_i) > 0\} \subseteq C$$

consisting of the initial computational states that are assigned nonzero probability by P_1.

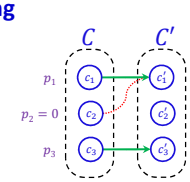

    - **Definition:** Here, we can call $\mathcal{C} = (O, P_1)$ a *logically reversible computation*.
- Given that $\mathcal{C} = (O, P_1)$ is logically reversible, we can also say that:
  - O is a *conditionally (logically) reversible operation* (this is always true),
  - O_A is a *conditioned (logically) reversible operation* with *assumed set* A,
  - A is an *adequate precondition for the (logical) reversibility* of O.
    - Any subset of A is also an adequate precondition for the reversibility of O.

18

Reversible Computing with Adiabatic Circuits

Some early history (pre-2000):

- Fredkin and Toffoli, 1978 (DOI:10.1007/978-1-4471-0129-1_2)
- Seitz *et al.*, 1985 (CaltechCSTR:1985.5177-tr-85)
- Several papers at PhysComp '92:
  - Koller and Athas (DOI:10.1109/PHYCMP.1992.615554)
  - Hall; Merkle (DOIs:10.1109/PHYCMP.1992.615549; 10.1109/PHYCMP.1992.615546)
    - General-purpose reversible methods, but for combinational logic only
- Younis & Knight, 1993 (http://dl.acm.org/citation.cfm?id=163468)
  - First fully-reversible, fully-adiabatic *sequential* circuit technique (CRL)
- Younis & Knight, 1994
  - Simplified 3-level adiabatic CMOS design family (SCRL) – Buggy though.
- Subsequent work at MIT by myself and other students, 1995-99
  - Several reversible and/or adiabatic demonstration chips

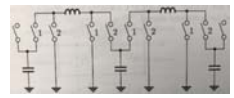

Figure reproduced with permission

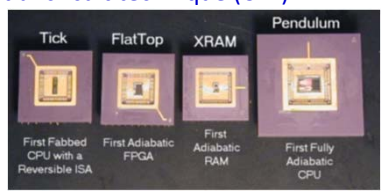

Conventional *vs.* Adiabatic Charging

For charging a capacitive load C through a voltage swing V

- Conventional charging:
  - Constant *voltage* source

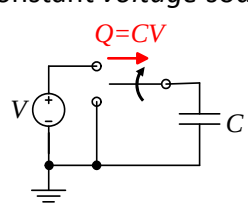

  - Energy dissipated:

$$E_{\text{diss}}^{\text{conv}} = \frac{1}{2}CV^2$$

- Ideal adiabatic charging:
  - Constant *current* source

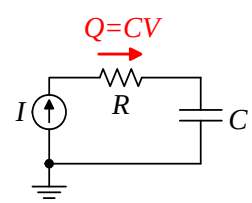

  - Energy dissipated:

$$E_{\text{diss}}^{\text{adia}} = I^2Rt = \frac{Q^2R}{t} = CV^2\frac{RC}{t}$$

Note: Adiabatic charging beats the energy efficiency of conventional by advantage factor: $A = \frac{E_{\text{diss}}^{\text{conv}}}{E_{\text{diss}}^{\text{adia}}} = \frac{1}{2}\frac{t}{RC}$

Adiabatic Switching using MOSFETs –
Inherently reversible, but requires GRC to model this properly

Switching is, inherently, only *conditionally* reversible!

- E.g.: Even *a single MOSFET* can implement a certain (conditioned) reversible COPY operation…
  - Operation sequence is as follows:
    0. Driving node D is initially statically held at 0, input *A* also 0.
    1. Input A is externally supplied (D&B connected iff *A* is high)
    2. Externally transition driver D from 0 to (weak) logic high 1′
    3. Voltage level on node B follows D iff A is strong logic high (**1**)
  - Note: Given a (loose) assumed precondition of $\bar{B}$,
    - *i.e.*, if all initial states with $B = 1$ have probability 0,
    - this indeed performs a reversible COPY operation, $\mathrm{rCOPY}(A, B \mid \bar{B})$.
  - Note: The output in this case is not full-swing,
    - In this diagram, primes (′) denote reduced-voltage logic high signals
  - A notation precisely describing this operation's semantics is:
    - $[\overline{AB}]$ if $B = 0$ then $B := A$ *(else, leave state unchanged)*
- **Note:** Traditional reversible computing theory based on unconditionally reversible operations is insufficient to model the logical/physical reversibility of this operation!
  - Illustrates why GRC is needed for real circuit engineering

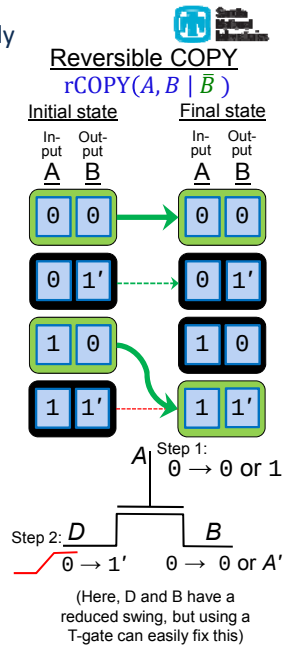

(Here, D and B have a reduced swing, but using a T-gate can easily fix this)

Conditionally-Reversible Boolean Logic in Adiabatic Circuits

- This simple CMOS structure can be used to do/undo latched reversible **rOR** operations
  - Example of 2LAL logic family (Frank '00)
    - Based on CMOS transmission gates
    - Uses dual-rail complementary signals (PN pairs)
- Computation sequence:

  Precondition: Output signal *Q* is initially at logic 0

  By design, driving signal *D* is also initially logic 0

  1. At time 1 (@1), inputs *A*, *B* transition to new levels
     - Connecting *D* to *Q* if and only if *A* or *B* is logic 1
  2. At time 2 (@2), driver *D* transitions from 0 to 1
     - *Q* follows it to 1 if and only if *A* or *B* is logic 1
     - Now *Q* is the logical OR of inputs *A*,*B*
- Reversible things that we can do afterwards:
  - Restore both *A*, *B* to 0 (latching *Q* in place), or,
  - Undo above sequence (decomputing *Q* back to 0)

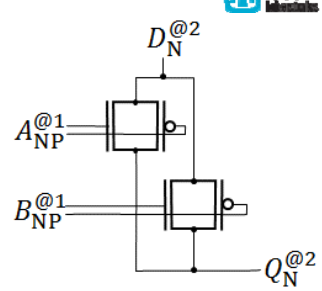

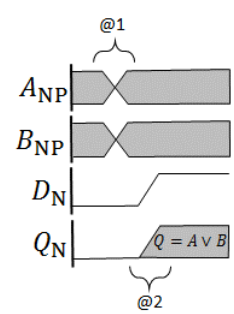

2LAL Shift Register Structure

- 1-tick delay per logic stage: Animation: http://y2u.be/c18mDIOq1IQ

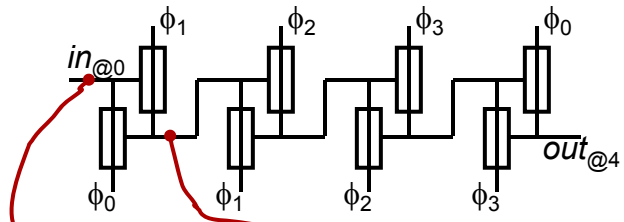

- Logic pulse timing and signal propagation:

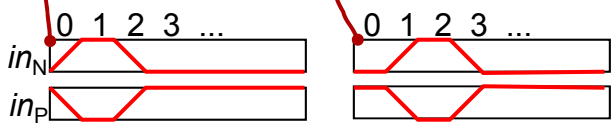

10/17/2017

Simulation Results (Cadence/Spectre)

Power vs. freq., TSMC 0.18, Std. CMOS vs. 2LAL

2LAL = Two-level adiabatic logic (invented at UF, '00)

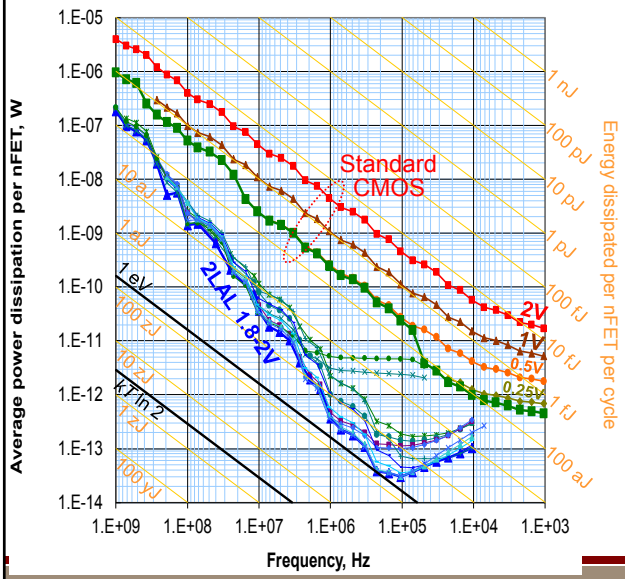

- Graph shows per-FET power dissipation vs. frequency
  - in an 8-stage shift register.
- At moderate freqs. (1 MHz),
  - Reversible uses < 1/100th the power of irreversible!
- At ultra-low power levels (1 pW/transistor)
  - Reversible is 100× faster than irreversible!
- Minimum energy dissipation per nFET is **< 1 electron volt**!
  - 500× *lower* dissipation than best irreversible CMOS!
    - 500× *higher* computational energy efficiency!
- Energy transferred per nFET per cycle is still on the order of 10 fJ (100 keV)
  - So, energy recovery efficiency is on the order of 99.999%!
    - Quality factor $Q = 100{,}000$!
      - Note this does not include any of the parasitic losses associated with power supply and clock distribution yet, though

24

Spectrum of Trapezoidal Wave

- Relative to mid-level crossing, waveform is an odd function
  - Spectrum includes only odd harmonics $f, 3f, 5f, \ldots$
- Six-component Fourier series expansion is shown below
  - Maximum offset with $11f$ frequency cutoff is $< 1.7\%$ of V_{dd}

$$v_{f6}(t) = V_{dd}\left[\frac{1}{2} + \frac{4\sqrt{2}}{\pi^2}\left(\sin\theta + \frac{\sin 3\theta}{3^2} - \frac{\sin 5\theta}{5^2} - \frac{\sin 7\theta}{7^2} + \frac{\sin 9\theta}{9^2} + \frac{\sin 11\theta}{11^2}\right)\right]$$

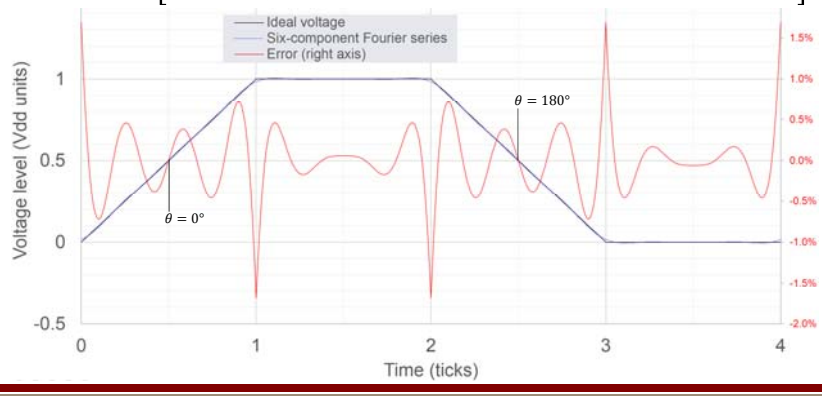

Ladder Resonator Structure

- Can build trapezoidal resonator w. a ladder circuit made of parallel passive bandpass filters, each a sinusoidal LC resonator
  - Each "rung" of ladder passes a different odd multiple of the fundamental clock frequency f
  - Adjust L/C ratio to obtain a target Q value on each path, given parasitic R,C values
- Excite the circuit with a driving signal containing the right distribution of frequency component amplitudes
  - Each frequency component gets amplified by the Q value of its corresponding rung
  - If all rungs are designed to the same target Q, we can just use a trapezoidal driver
- For high Q, clock period must be long compared to the total parasitic RC…
  - Max. possible $Q_n = \frac{1}{2\pi} \cdot \frac{t_{\mathrm{period},n}}{(RC)_{\mathrm{parasitic}}}$

(for $V_{\mathrm{dd}} \approx 1.75$ V ↓)

Example values:

harmonic mode (n)	frequency f	component amplitude Va	inductance L	capacitance C
1	230 kHz	1000.00 mV	691.98 nH	691.98 nF
3	690 kHz	111.11 mV	230.66 nH	230.66 nF
5	1150 kHz	-40.00 mV	138.40 nH	138.40 nF
7	1610 kHz	-20.41 mV	98.85 nH	98.85 nF
9	2070 kHz	12.35 mV	76.89 nH	76.89 nF
11	2530 kHz	8.26 mV	62.91 nH	62.91 nF

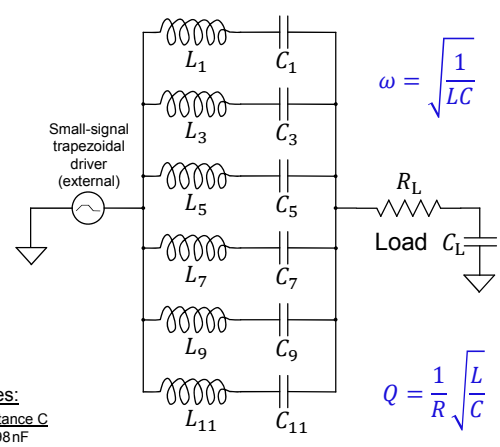

$$\omega = \sqrt{\frac{1}{LC}}$$

$$Q = \frac{1}{R}\sqrt{\frac{L}{C}}$$

Ladder Resonator for Odd Harmonics

Reversible Superconducting Circuits

- Early history of such technologies:
  - Likharev '77 (DOI: 10.1109/TMAG.1977.1059351)
    - ***Parametric Quantron*** (PQ) – Allows switching between bistable states of a SQUID under control of an externally supplied current I_C.
  - Hosoya *et al.* '91 (DOI: 10.1109/77.84613)
    - ***Quantum Flux Parametron*** (QFP) – A fluxon can be shuttled between two adjacent SQUID loops adiabatically under the influence of external controls.
- Some more recent works of interest:
  - Ren & Semenov '11 (DOI: 10.1109/TASC.2011.2104352)
    - Demonstrates empirical dissipation < kT per JJ per cycle
  - Takeuchi *et al.* papers (DOIs: 10.1088/0953-2048/26/3/035010, 10.1038/srep06354, 10.1063/1.4817974, 10.1038/s41598-017-00089-9)
    - AQFP (adiabatic QFP)
- All of the above schemes use adiabatic transformations that are driven by external control signals
  - Conceptually similar, at a high level, to adiabatic CMOS
- Next we'll see a different, *ballistic* approach to reversible computing in superconducting circuits…

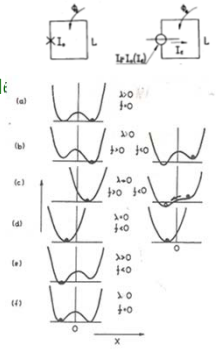

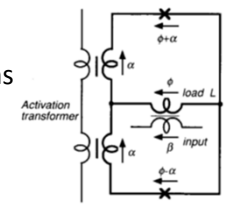

Asynchronous Ballistic Reversible Computing

- Some problems with all of the existing *adiabatic* schemes for reversible computing:
  - In general, numerous power/clock signals are needed to drive adiabatic logic transitions
  - Distributing these signals adds substantial complexity overheads and parasitic power losses
- Ballistic logic schemes can eliminate the clocks!
  - Devices simply operate whenever data pulses arrive
  - The operation energy is carried by the pulse itself
    - Most of the energy is preserved in outgoing pulses
    - Signal restoration can be carried out incrementally
- But, *synchronous* ballistic logic has some issues:
  - Unrealistically precise timing alignment required
  - Chaotic amplification of timing uncertainties when signals interact
- Benefits of *asynchronous* ballistic logic:
  - Much looser timing constraints
  - Linear instead of exponential increase in timing uncertainty per logic stage
  - Potentially simpler device designs
- A new effort just started at Sandia to try to implement ABRC in superconducting circuits

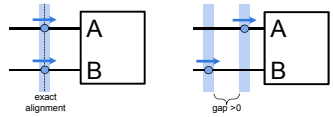

Synchronous Ballistic Asynchronous Ballistic

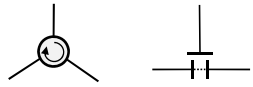

Rotary (Circulator) Toggled Barrier

Example ABR device functions

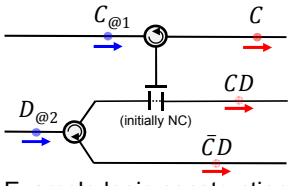

Example logic construction

Physical realizations of ABRC

- To be useful, the ABRC model needs to be realized in a specific physical implementation technology that actually provides near-thermodynamically-reversible operation.
  - Need some kind of soliton-like, near-ballistically-propagating pulse,
    - or some sort of fundamental particle or quasiparticle state.
  - Need a stationary physical state variable that can stably maintain at least two distinct states
    - for use in the toggling devices
  - Need a means of physically interacting the pulses with the states
    - in ways that can reliably, and almost physically-reversibly, implement at least a universal subset of the 2- and 3-terminal primitive devices.
- One intriguing possible candidate implementation technology is to use superconducting circuits…
  - SFQ (single flux quantum, or fluxon) pulses on appropriately constructed superconducting transmission lines can carry info. with relatively low dispersion and high propagation velocity (approx. 2/3 c)
    - Fluxons are naturally quantized by the SQUID-like circuits that produce them, and are naturally polarized (carry 1 bit's worth of +/– polarization state information per pulse)
      - Need to select suitable ABRC primitives operating on arity-2 signals
  - Fluxons trapped in loops (SQUID-like structures) can hold data quiescently
    - Generally, loops hold integer numbers of fluxons in some range: …, –2, –1, 0, +1, +2, …
  - How exactly to implement the reversible interactions?
    - A 3-year, internally-funded project at Sandia is just starting to investigate this…

29

A Very Recent Advance!

Wustman (LPS) & Osborn (JQI) '17 (preprint), "Efficient reversible logic gates without adiabatic constraint: Fluxon resonant scattering with polarity changes"

- The circuit shown at right can be considered as a 2-terminal ABRC device for binary pulses (fluxons)
  - The specified function is to preserve or flip the polarity of a fluxon passing through, depending on device parameters
- Here, the "wires" are LJJ transmission lines
  - Major loss mechanism is resonant plasmon emission
    - With lattice spacing $0.4\lambda_J$, fluxon decay time is ~10^7 junction switching times given initial $v = 0.6c$.

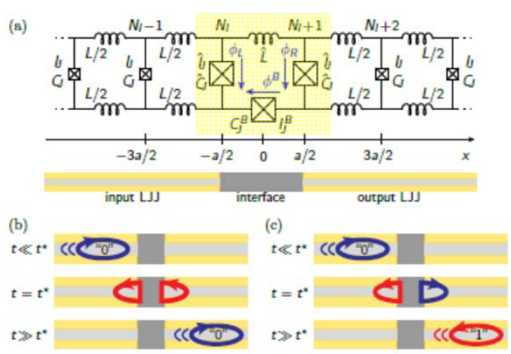

- W&O's paper also describes some more complex (4-terminal) devices
  - But they haven't yet explored asynchronous operation…

30

W&O's simulation of identity/NOT

- Direct numerical integration of JJ circuit's equations of motion
  - Lagrangian:

$$\mathcal{L} = \left(\frac{\Phi_0}{2\pi}\right)^2 \left[\sum_n \frac{C_{J,n}}{2}(\dot{\phi}_n)^2 + \frac{C_J^B}{2}(\dot{\phi}^B)^2\right] - \left(\frac{\Phi_0}{2\pi}\right)\left[\sum_n I_{J,n}(1-\cos\phi_n) + I_J^B(1-\cos\phi^B)\right] - \frac{1}{2}\sum_n \left[L_n^A(I_n^A)^2 + L_n^B(I_n^B)^2\right]$$

  - Gives a discrete approximation to the sine-Gordon equation:

$$\ddot{\phi} - c^2\phi'' + \omega_J^2 \sin\phi = 0$$

- Scattering interaction at interface is nearly elastic
  - Loss in fluxon velocity of only 4%
    - Loss in energy of 2.1-2.5%

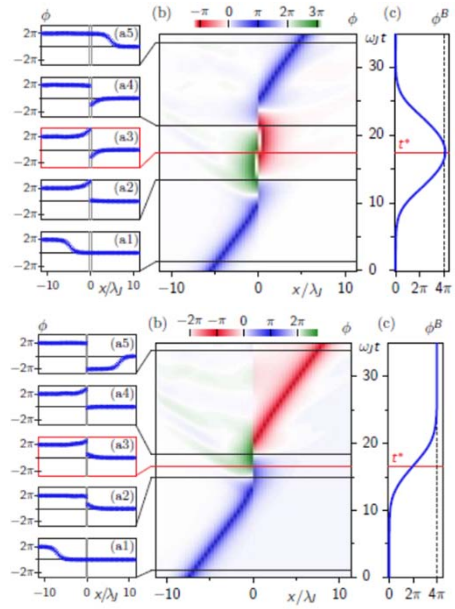

Nanomechanical Rotary Logic

Merkle et al., IMM Report 46 and Hogg et al., DOI: 10.1039/C7ME00021A

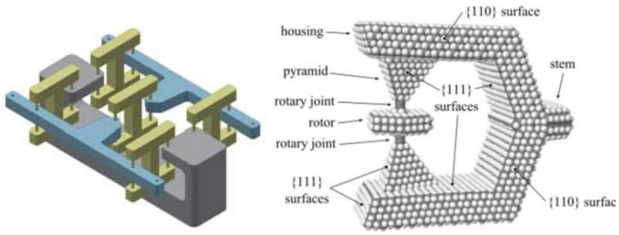

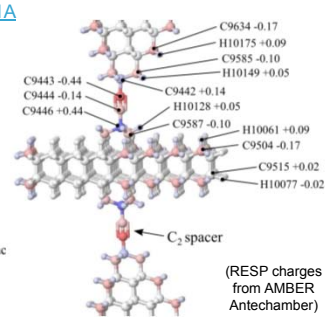

(RESP charges from AMBER Antechamber)

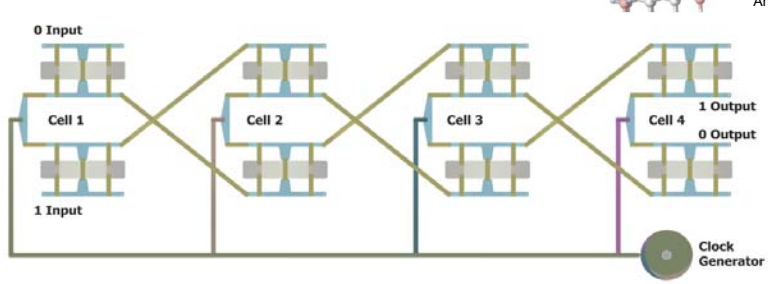

Rotary Logic Lock Operation

- Videos animate schematic geometry of a pair of locks in a shift register
- Molecular Dynamics modeling/simulation tools used for analysis include:
  - LAMMPS, GROMACS, AMBER Antechamber
- Simulated dissipation:
  - ~4×10^{-26} J/cycle at 100 MHz
    - 74,000× below the Landauer limit for irreversible ops!
- Speeds up into GHz range should also be achievable

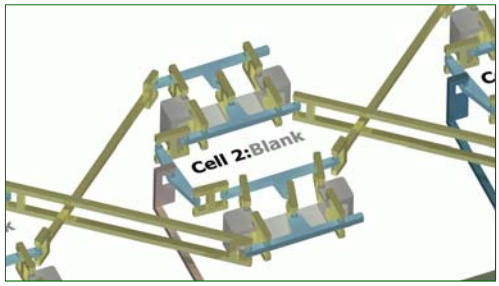

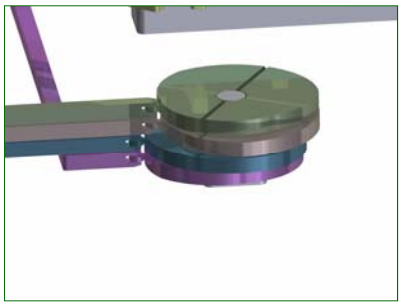

33

Conclusion

- Cost-efficient (therefore, power-limited) raw performance of conventional computing technology faces imminent limits…
  - Due, in practice, to energy efficiency limits that are ultimately rooted in extremely fundamental thermodynamic issues…
- Various "unconventional" computing concepts aim to attack the remaining energy-inefficiencies that exist at various levels of the computer design stack at present…
  - However, most of these concepts only address a few of the levels, and so the efficiency gains that they can offer are ultimately limited
- *Only one* unconventional computing paradigm consistent with fundamental physics can potentially increase energy efficiency of *all* computations by *indefinitely large* factors *throughout* a computer system, at *all* levels, from devices to networks: **Reversible computing**
  - We presented the general principles of reversible computing, and illustrated a few candidate implementation technologies.
- **Reversible computing is the only possible long-term path forward, towards an unbounded future for computing.**
  - **It deserves greatly increased attention from the larger R&D community.**

34