Quantum Plus Classical Computation

CCC Reversible Computing Workshop,
October 5-9, 2020

Erik P. DeBenedictis
Zettaflops, LLC

Quantum Plus Classical Computation

- Quantum information technology matured a lot in the last year, leading to the long-term prospect that general, or quantum plus classical, computers will become the most capable and energy efficient. However, the technology that interfaces between the quantum and classical information domains has different requirements from both qubits and CMOS. A careful analysis indicates that reversible technology is ideally suited for this purpose. Thus, reversible technology may be destined to enter the mainstream as a enabler for processing quantum information in lieu of the often cited objective of its restarting Moore's law.

This is the “title and abstract” given to CCC.

Thesis of Talk

- Fully general computers may include qubits
- Reversible logic may be a missing link for the control electronics of a quantum computer
  - Cryogenic adiabatic transistors circuits (CATC)
  - Goal is ultra-low power and slow is fine
  - May allow greater scale up of the quantum computer
- Restarting Moore's Law may be harder

The General Computer

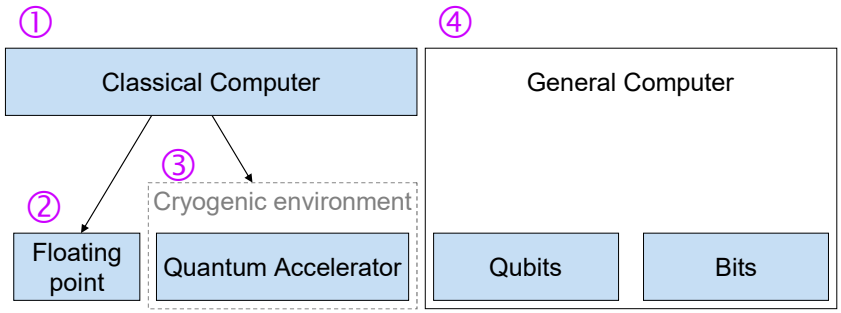

- A one-chip quantum computer outran the fastest supercomputer.
- In general, the most capable computer may include qubits and bits
- However, some problem classes do not benefit from qubits

There is a common view today that computer are classical Turing machines, but are permitted to have “accelerators.” Examples of accelerators are the 8087 floating point unit that accompanied the Intel 8086 processor. It was an accelerator because floating point is more complex than everything else in a microprocessor put together and it warrants a special chip. A quantum accelerator like the Sycamore quantum chip associated with ORNL’s Summit supercomputer fits the same model.

I disagree with this view. The idea of a floating point coprocessor lasted only a few years in the 1980s after which floating point became just another data type in (what is not) out of order processors.

I think the same will happen to quantum computing. There is no benefit to a hard dividing line between classical and quantum information. Humans can show great ingenuity, and I’m expecting that some people will figure out how to integrate quantum and classical information in a processor before long.

This will yield advantages for some problems. In cases where quantum information does not offer an advantage, it will still be possible to use a microprocessor.

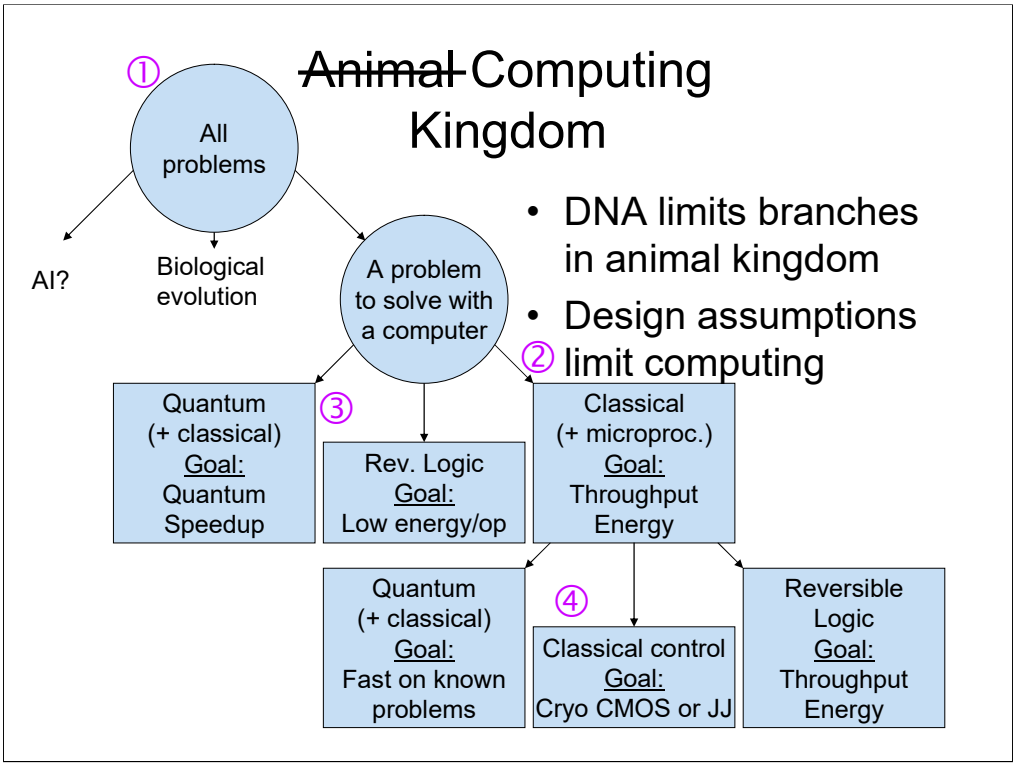

Our goals for computers depend on how we got engaged with them.

DNA is a computing substrate responsible one of the greatest advances of all time: life.

However, we do not expect a microprocessor to have an instruction that causes it to grow additional DRAM like our fingers grow fingernails. Instead, microprocessors are expected to be Turing machines that have high throughput and reasonable energy efficiency when measured at electric power grid rates.

While quantum and reversible computing originated decades ago as exotic ideas, there are many people that view both quantum and reversible computing as a "post-Moore's law computing approach." With this positioning, we might want quantum and reversible computing to propel supercomputers to Zettaflops and beyond. Yet in the last year, quantum computers have become interesting in their own right due to demonstrations of "quantum supremacy" on different applications.

The point of this talk is that reversible computing may be ideally suited for the cryogenic classical portion of a quantum computer, which may be key to quantum computer scale up. Reversible computing would be appropriate for its ultra low energy consumption and low noise generation. However, a quantum computer's throughput comes from the qubits not the control electronics, so reversible computing and/or logic would not be expected to have high throughput.

Future R&D for reversible computing/logic needs will diverge depending on whether the objective is high throughput to restart Moore's law or ultra low energy for cryogenic operation.

Let's explore the difference in these objectives, which will conclude that the shortest term opportunity for reversible computing/logic will be in support of quantum computers.

Quantum-Classical Functions

- Quantum Error Correction
- There is an argument that quantum operations create no heat
- Sometimes you need a local classical operation
  - Diagram shows a loop
  - Apply a microwave waveform until the measurement returns false
  - Generating the waveform takes power, but blocking it does not
- "Quantum floating point"

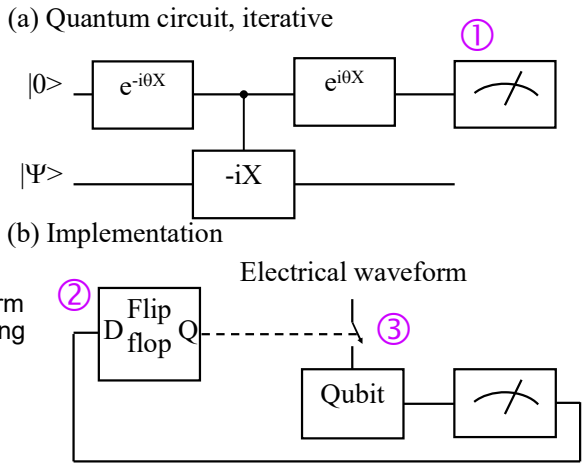

This slide is an example of an intrinsically mixed quantum-classical function.

In quantum error correction, for example, error correction codes are checked on may qubits. The qubits with errors need to have their errors corrected, but the qubits without error are left alone.

It is possible to have an external microwave signal perform operations on may qubits at near zero energy, such as assessing an error correction code or making a correction.

However, deciding whether to enable a microwave signal on a specific qubit requires a classical decision about whether the qubit was in error. The classical decision is governed by Landauer's minimum energy.

Thus, the mixed quantum-classical function does indeed consume energy—in the classical electronics.

There are two ways to control the error correction:

1. Run a wire from the cryogenic environment to room temperature and store the error state in a flip flop. Then use the flip flop to gate a microwave signal onto a cable that runs from room temperature to the cryogenic environment. However, the cabling will occupy space and leak heat, both of which will limit scalability.

2, Alternatively, put a flip flop in the cold environment. The flip flop will dissipate heat and limit scaling due to heat production in the cold environment and cooling overhead.

However, what if we used reversible logic? Reversible logic can operate at arbitrarily low energy levels.

This is a direction to explore.

Technology Objectives for Classical Control System

	Microprocessor	Quantum Computer (cryogenic part)
Throughput	$N_{Gates}\ f_{Clk}$ ①	$N_{qubits}\ S(N_{qubits})\ f_{Clk}$ ②
Cost of Ownership	$N_{Gates}\ (\$_{Gate} + \$_{energy})$ ③	$N_{qubits}\ k\ (\$_{Gate} + \$_{energy}\ 300\ K/T_q) + N_{qubits}\ \$_q$ ④

- Microprocessors
  - ① – maximize clock rate
  - ③ – make energy efficiency trades based on municipal power rates
- Quantum computer
  - ② – maximize quantum speedup, sacrifice clock rate
  - ④ – energy efficiency trades include cooling overhead (1,000× at 4 K)

N_{Gates} = number of gates
N_{qubits} = number of qubits
$\$_{Gate}$ = cost of gate
$\$_q$ = cost of qubit
$\$_{energy}$ = lifetime cost of energy

$S(N_q)$ = quantum speedup
T_q = qubit temp
f_{Clk} = clock rate
k = number of classical gates to support a qubit

Cost equations:

Microprocessors work better with high throughput electronics. However, quantum computers have an additional quantum speedup component. Increasing clock rate increases throughput linearly, but increases in quantum speedup may be exponential. Therefore, the wise engineer would be willing to run the logic slowly if it results in more quantum speedup. This direction has not explored (at all?) for reversible computing.

Optimizing the cost of a microprocessor system involved the cost of buying a gate (or chip) $gate versus the cost of powering it over its lifetime $energy. Reversible computing fails this tradeoff analysis as evidenced by the fact that nobody uses reversible computing. However, the equivalent optimization equation for reversible computing put an refrigeration overhead factor in front of the energy cost. This is 300 K / T at 100% Carnot, but is typically 1,000× at 4 K and could be higher. The bottom line is that an energy savings technology like reversible computing/logic can be a mediocre tradeoff at room temperature but a no-brainer at cryogenic temperatures.

Impact of Throughput as a Goal

- Adiabatic circuits dissipate less heat, but only at low frequencies
- Low frequency reduces microprocessor throughput, but
- quantum computer throughput comes from the qubits

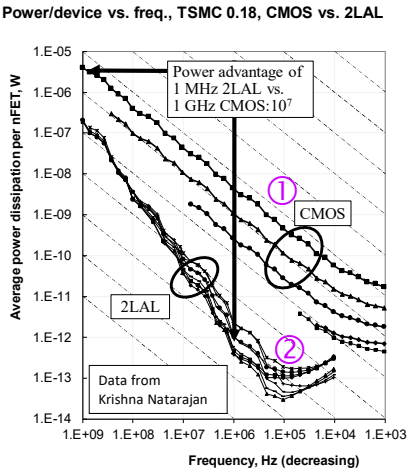

This is the first of five slides explaining how an electrical circuit can decrease the load on a cryocooler.

Voltage-based electrical circuits store $\frac{1}{2}CV^2$ energy on capacitors, the capacitors being a combination of transistor gates and wiring capacitance. CATCs do not change this, but save energy. The energy is saved in the cooling system.

(left) Beyond Moore's law technologies have been extensively studied at room temperature. The blue lasso shows devices studied by the Nanotechnology Research Initiative (NRI) lining up on one side of a frontier of constant energy-delay product.

(middle) On log-log axes (of different scales): The horizontal lines labeled "CMOS" have constant energy per operation as a function of delay—although the level of the line changes with supply voltage. The left end of each horizontal CMOS line is the maximum speed given drive current, and defines the position of the dot on the left hand chart.

However, both theory and simulation show that adiabatic circuits dissipate less power than a CMOS circuit. Energy is conserved, so it must go someplace. We'll explain where it goes in a few slides.

(right). If memory is of interest, the same data can be plotted on power-delay axes. These are the correct curves for the power required to hold a bit in a memory.

Cryo Makes Reversible Viable

- CMOS dissipates energy in the chip
- Adiabatic circuits dissipate energy in the waveform generator
  - For class B circuits, resonators can be better
- Not much help at room temperature
- If there is a temperature difference, adiabatic will typically bipass the refrigerator

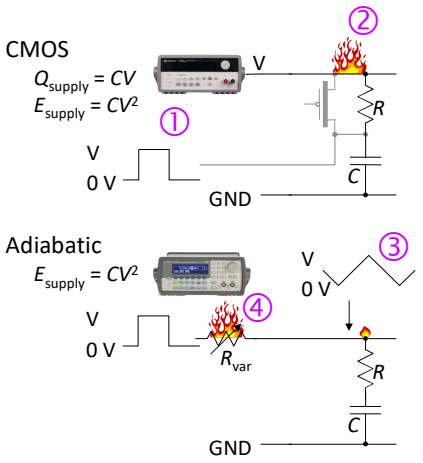

While we have reduced the heat produced in the gates, we've really just moved the location where energy is turned into heat. If the waveform generator uses a "class A" output amplifier to cut supply voltage V until it looks like a ramp, we are still charging capacitors from a fixed voltage – so the resistors must now dissipate ½CV^2, as before. So the total energy is the same, we've merely changed the location where the energy turns into heat from the circuit to the power supply.

In a cryogenic setup, the gates are at a cold temperature, such as 4 K, but signal generator is at room temperature. Thus by moving energy from the cold environment (a) to room temperature (b), we reduce the amount of energy going through the cooling apparatus. At 4 K, this eliminates about 1,000x.

Inversion

- As strictly defined, 2LAL cannot invert
- Quad rail solution: Create two copies of a circuit, the second with complementary logic. Inversion comprises swapping signals between copies
- Competitive circuit SCRL can only invert; need a non-inverting buffer.
- New solution: Extend clock

- Modified clocks

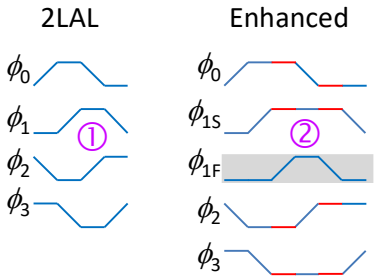

- Slower because there are more ramps

2LAL has four clocks, as shown. Note that the inverse of each clock is another clock.

Enhancement is to add time on either side of the flat top of ϕ_1. This lets us create a ϕ_{1S}, S for slow, and ϕ_{1F}, F for fast.

Existing circuits work fine with either ϕ_{1S} of ϕ_{1F} for ϕ_1.

2LAL inverter

- Spawns/uncomputes a stream with inverted data
- Can generate the third and fourth rails
- Red components generate an inverted signal
- Blue components uncompute an inverted signal. A. k. a. absorbs signal energy

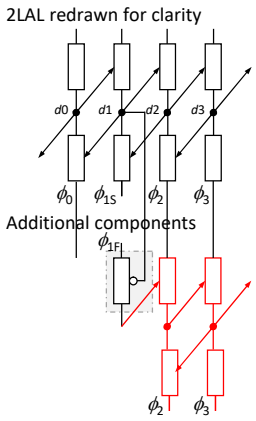

This is the inverter. See report for details.

Essentially, a data stream *d* can spawn a data stream with the inverse of *d* (red components). Signal energy has to be absorbed to maintain reversibility, which is the blue components.

nFET-only 2LAL stages

- The reversible logic families have an awful lot of transistors
- Cutting quad rail to dual rail helps a lot
- Could we turn pass gates into single transistors?
  - Sort of, can do for alternate stages; savings ~25%.

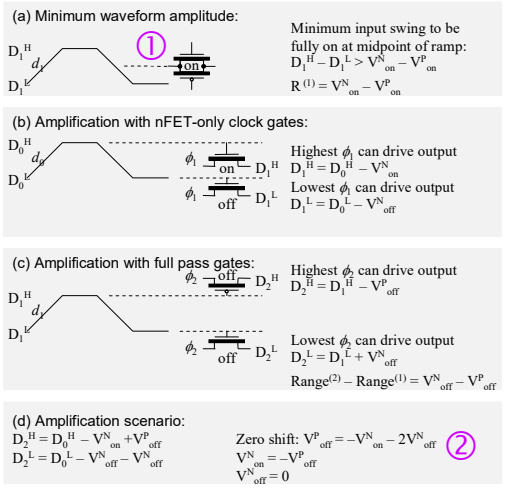

This slide show the advantage of new transistor types that have been discovered since the popular (transistor-based) reversible circuits were discovered (1990s).

See Erik DeBenedictis, Enhancements to Adiabatic Logic for Quantum Computer Control Electronics, report ZF002, https://zettaflops.org/CATC.

Role of cryo FDSOI

- Without FDSOI, pFET threshold needs to be much larger in magnitude than nFET. That would require a custom CMOS process
- New FDSOI for IoT has a fourth terminal that adjusts threshold

- FDSOI overview
  - 4th terminal VB (backbias)
  - Shifts threshold

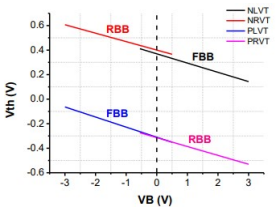

Fig. 13. From http://soiconsortium.eu/wp-content/uploads/2017/08/AnalogRF_28FDSOI_ACathelin_19092017.pdf

If we set clock voltages and transistor thresholds appropriately, we can eliminate the pFETs from pass gates in up to half the clock phases (or stages).

This requires pFET to have a threshold several times that of nFETs (in magnitude; pFET thresholds are negative by convention). However, modern FDSOI transistors can have a dynamically adjustable threshold.

This idea has not been tested and is unlikely to work very well at room temperature. The reason is that subthreshold slope increases dramatically at cryo, reducing the voltage swings. There may not be enough control over threshold to meet requirements for this approach at room temperature.

Conclusions

The interface between quantum and classical information is unique

- The quantum control environment follows a different path through the computing kingdom
- The environment alters current notions of
  - throughput (its from the qubits not the control electronics)
  - the cost of energy (much higher due to cooling overhead)
- Changes design choices, leading to more freedom of action.
  - New cooling strategy, circuits, devices (4-terminal transistors), architectures, more
- Could increase scaling of quantum computers, make them viable

Priority Research Direction

Physical demonstrations: We can project better results than CMOS, such as Horse Ridge. But we don't know the limits

- Fab and test circuits and test energy efficicncy, device spread, speed
- Make quantum computer test chips using reversible logic
- Assess scale up potential

Design and software: There were no new reversible transistor circuits after 1999, until we started exploring the cryo area 2020

- Look for adiabatic circuits that implement important quantum architecture primitives, testing by simulation
- Design tools will be needed, such as synthesis

Backup: Vision of a Hybrid Quantum Computer

- Additional information
  - See https://www.zettaflops.org/CATC
  - See https://arxiv.org/pdf/1912.11532
- A quantum computer would comprise
  - Qubits
  - Josephson junction classical electronics
  - Cryogenic adiabatic transistor circuitry
- Each has advantages

Quantum Computer Vision including CATC-based Control

- Clock-power signals generated at room temperature
- Low-level functions processed by cryogenic adiabatic transistor circuits, such as waveform generation, quantum error correction.
- High level control via a microprocessor at room temperature
- Cryo-specific advances include enhanced circuits, electrical improvements, …

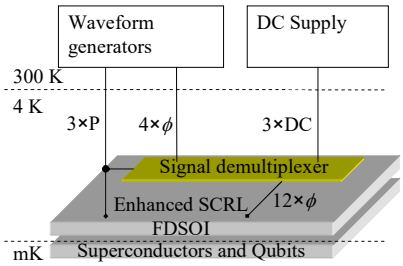

The slide illustrates reversible computing in a quantum computer control context.

A cryogenic module would contain both qubits and classical electronics, both transistorized and Josephson junction based.

Power supplies would be at room temperature where energy is cheaper by 1,000× or more.

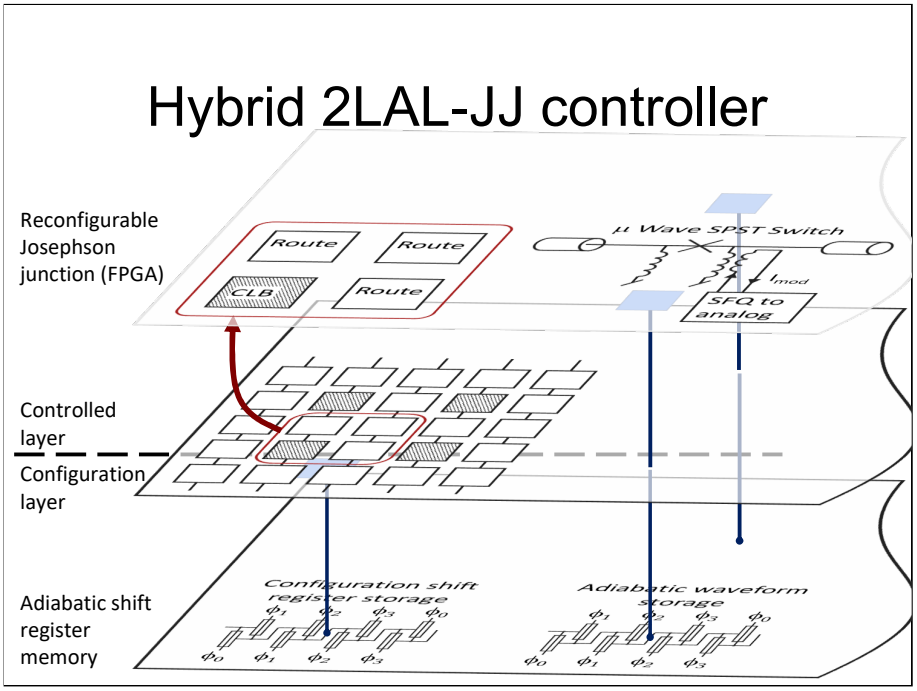

Today, most JJ circuits are fabricated by evaporating niobium on to a blank silicon wafer, which is needed for mechanical stability. The suggestion is to use a completed silicon wafer instead of a blank one. This has been done before physically, but this talk is presenting new design principles.

The lower layer is the semiconductor layer with adiabatic transistor circuits, which communicates (unidirectionally) with the JJ layer via voltages. Circuits for converting (field effect) voltages to forms compatible with JJs are not shown, but could be a superconducting FET or just a large semiconductor FET that interrupts or doesn't interrupt a SFQ pulse.

The adiabatic layer is illustrated as holding the FPGA configuration(s) or a digitized waveform.

See https://arxiv.org/pdf/1912.11532 for further explanation.