Minimizing classical control resources

QRE2022

June 18, 2022

Erik P. DeBenedictis
Zettaflops, LLC

Minimizing dissipation and the million-qubit goal

- Diagram from a major player projecting a million-qubit quantum computer. It looks to me like 10 m of a submarine hull
- Cryo CMOS dissipation now regarded as a scale up limitation
- Could reversible logic allow lower Size Weight and Power (SWaP)?
- Could the same number of qubits be possible in a structure the size of a tank truck?

See 26 mins 1 sec into the youtube https://www.youtube.com/watch?v=mmyq1ubjqO8

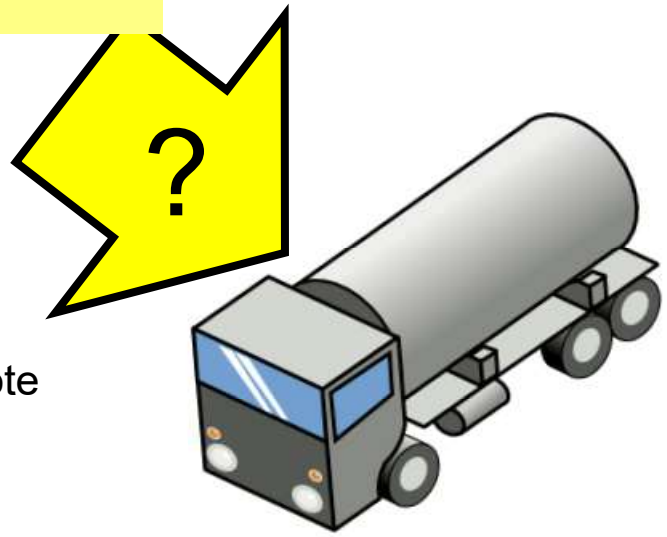

Anthony Megrant, Google, Quantum Week 2021 keynote

Reversible logic reformulated for cryo, summary

- Chips ① developed in 1990s are pretty good demonstrations of reversible logic
- However, the original reversible logic depended on an unmanufacturable component ②
- With knowledge of quantum computer scaling, reversible logic could be reformulated to replace ② with a cryocooler ③

- Energy/heat flow diagram

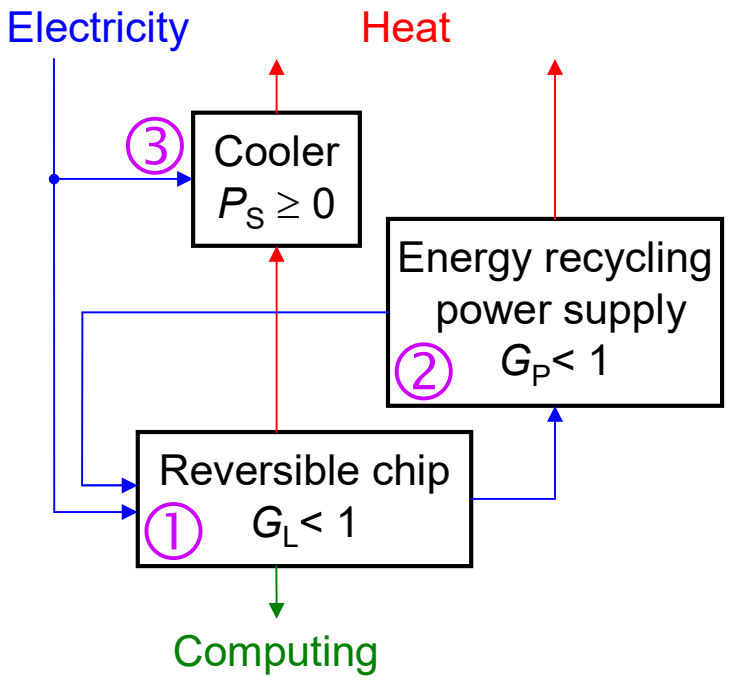

- Detail in backup slides

G_P = the energy recycling efficiency of the power supply
G_L = 1 - $2RC/\tau$, the portion energy not turned into heat
P_S = the cooling over head of the cryo cooler, or 0 if not present (P_S stands for “specific power”)

Outline

- The speaker devised a reversible logic controller (RL controller) design
  - Schematic + spice simulation
  - Plan to use in real quantum computers, but …
  - used as an “existence proof” in this talk
- Existence proof of what?
  - No irreversible gates required in cryostat
  - Transistors required in cryostat $\propto$ source code size (no ∞ stack)
  - Universal architecture in cryostat (.gif/.zip decompressor)
  - Other conclusions in backup material
- Natural metric
  - Irreversible gate-ops + reversible gate-ops + total memory, per quantum operation

Can we make the RL controller with reversible shift registers?

- CMOS ① constant energy/op
- Adiabatic energy per op ② drops with clock period
- Qubit measurement about 1,000× slower than CMOS
- Reversible shift registers have been measured with 10-1,000× lower dissipation than CMOS using the same transistors

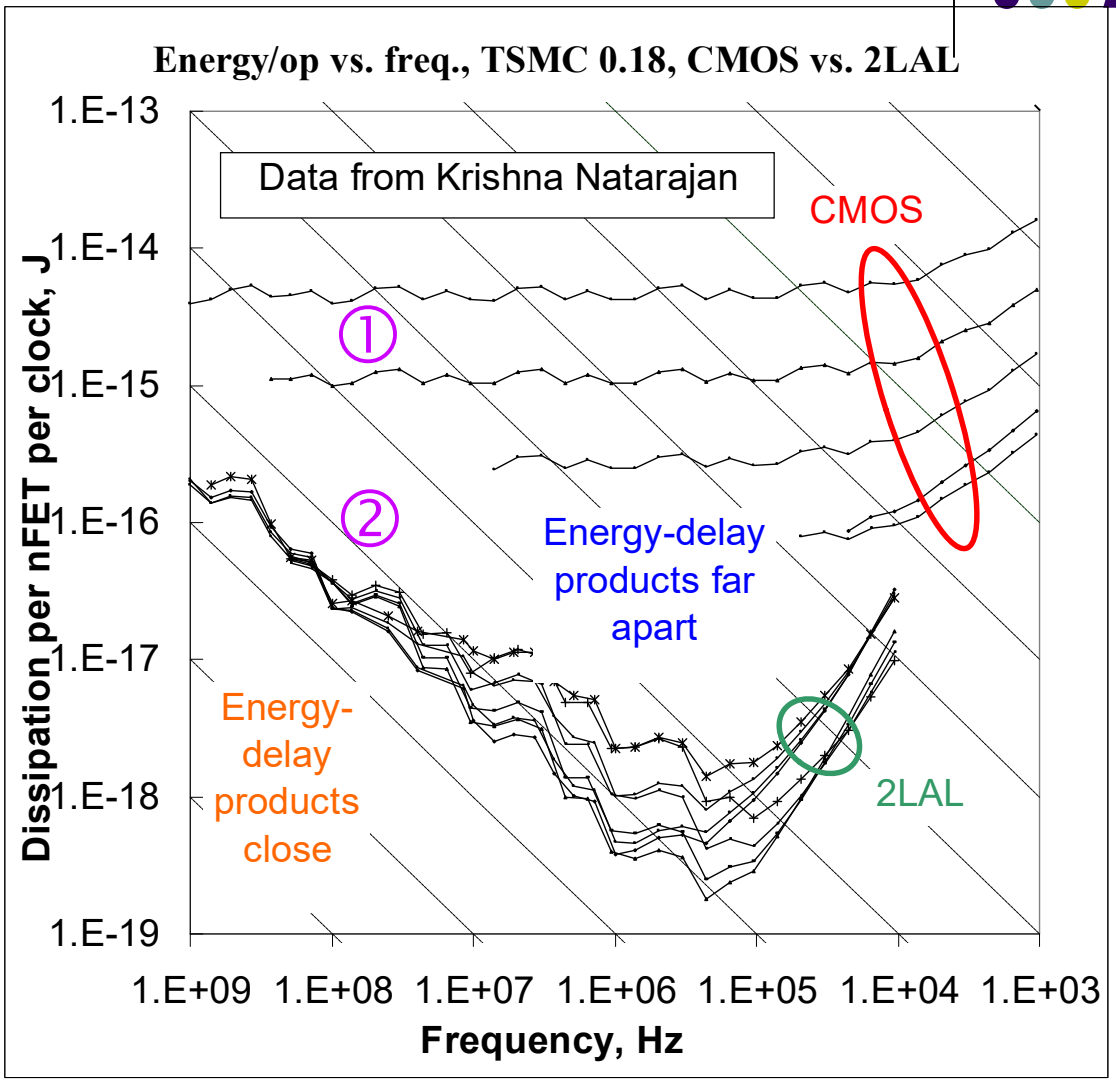

Lowers dissipation by ~131×

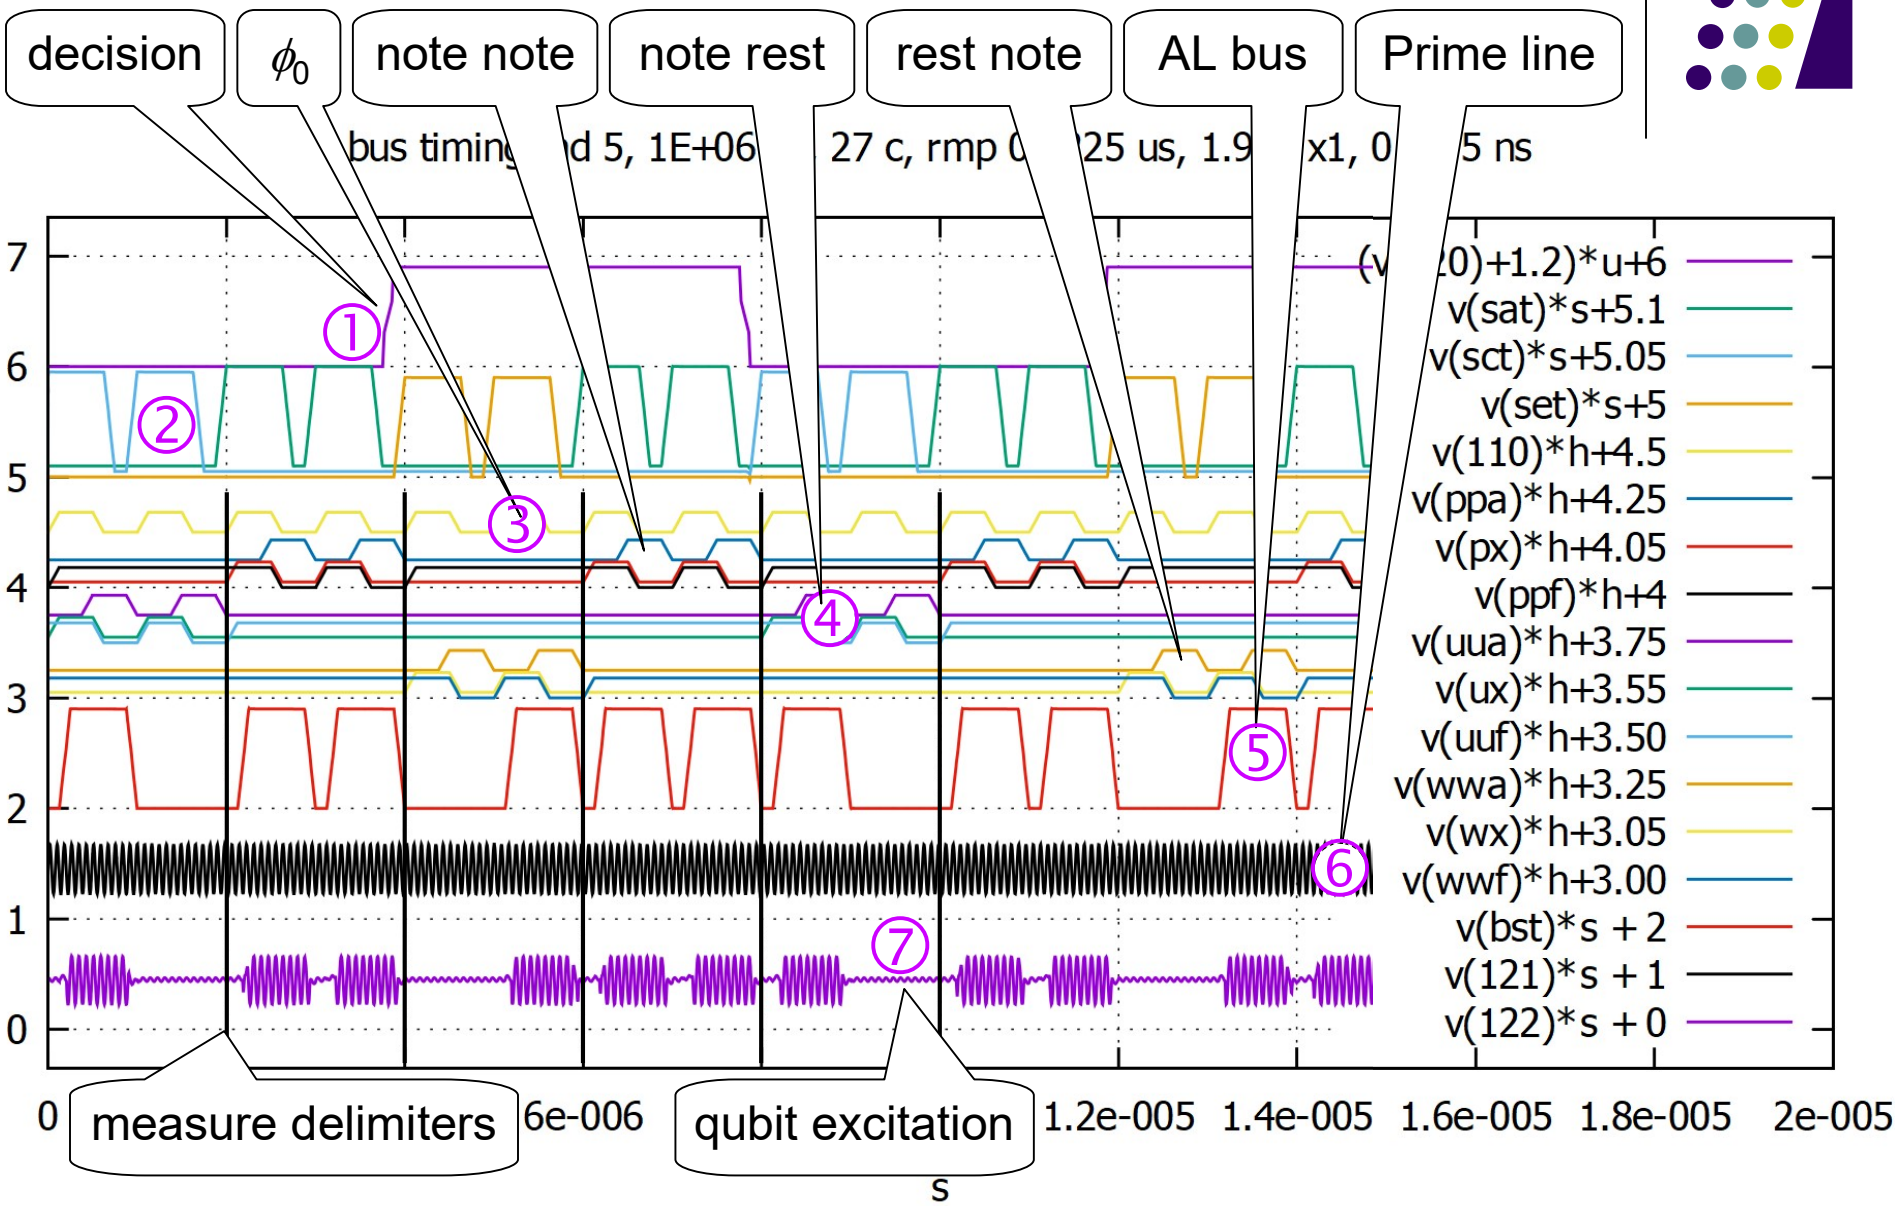

Process: Algorithm-flowchart-schematic-fab

- Error correction procedure

1. Use the circuit of Fig. 2(c) to extract the *XZZXI* syndrome. ①

   (a) If the flag qubit is measured as |->, then use the unflagged circuits analogous to Fig. 2(b) to extract all four syndromes. Finish by applying the corresponding correction from among *IIIII*, *IIZXI*, *IXZXI*,*IYZXI*, *IZZXI*, *IIIXI*, *IIXXI*, *IIYXI*.

   (b) Otherwise, if the syndrome is −1, i.e., the syndrome qubit is measured as |1>, then use unflagged circuits to extract all four syndromes. Finish by applying the corresponding correction of weight ≤ 1.

2. (If the flag was not raised and the syndrome was trivial, then) Similarly extract the IXZZX syndrome. If the flag is raised, then use unflagged circuits to extract the four syndromes, and finish by applying the correction from among *IIIII*, *IIIIX*, *IXXII*, *IIIXX*, *XIIIY*, *IXIII*, *IIIZX*, *IIIYX*…

From: Chao, Rui, et. al "Quantum error correction with only two extra qubits." Physical review letters 121.5 (2018): 050502.

- Create "flowchart"

[[5, 1, 3]] error detection and correction with flags:

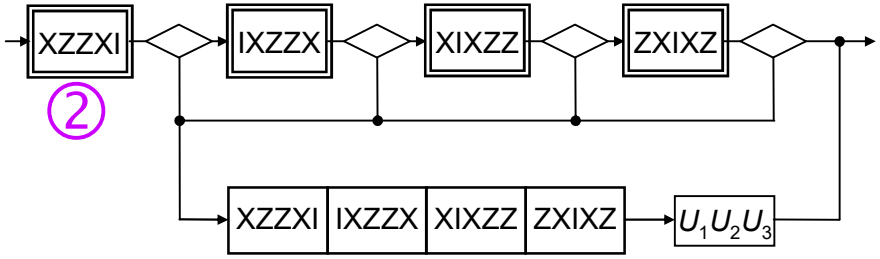

Circuit with flag qubits:

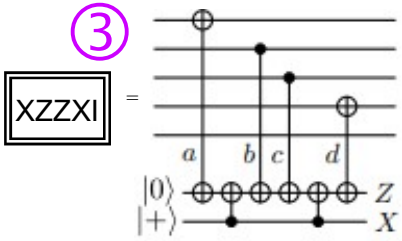

Circuit without flags:

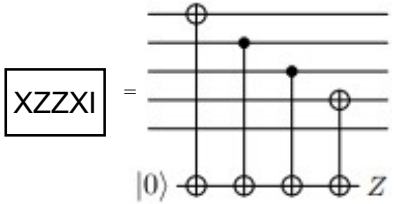

Decision, implemented by a crossover:

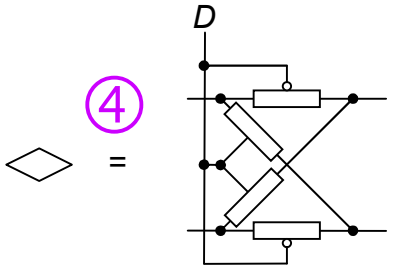

Correction:

$U_1U_2U_3$ = Z, Y, or Z on qubit *n*

Hybrid computing system

- Yields an integrated multi-temperature computing system
  - Function
  - Powertrain
- Only reversible shift registers at cryo stage
- If you think about it, cryo stage is a zip/gif decompressor

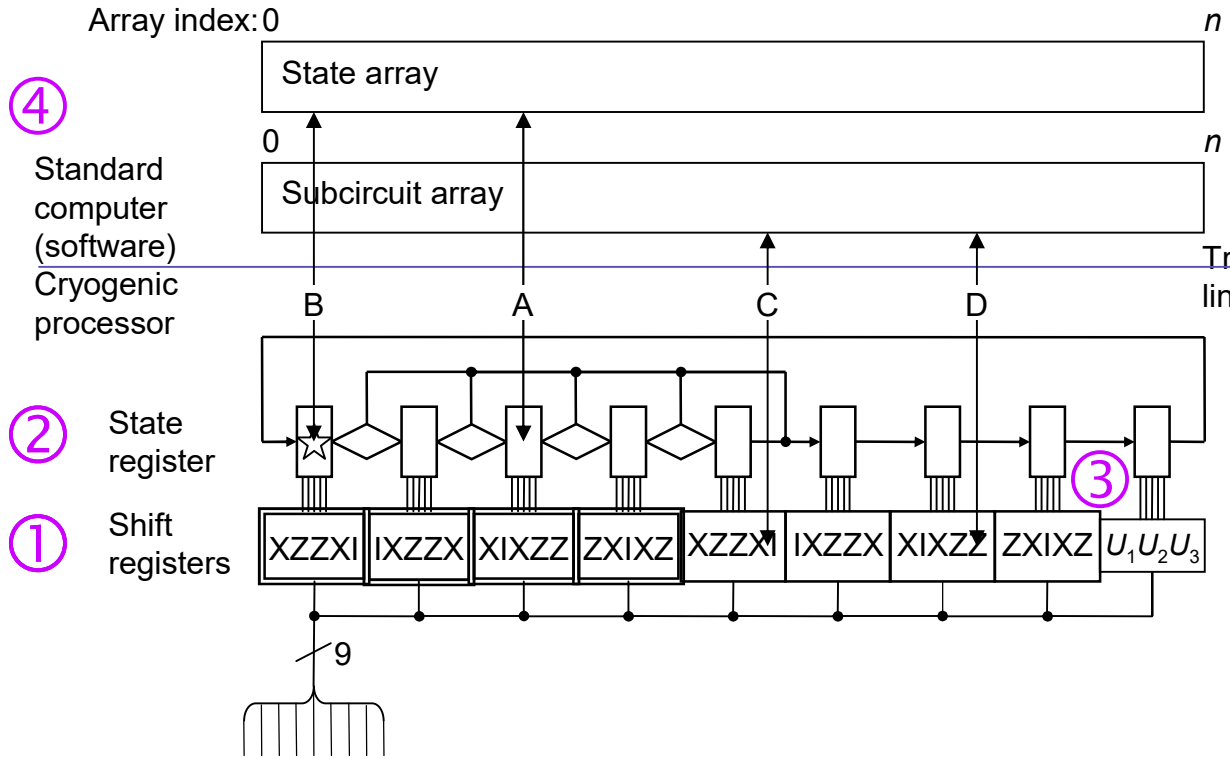

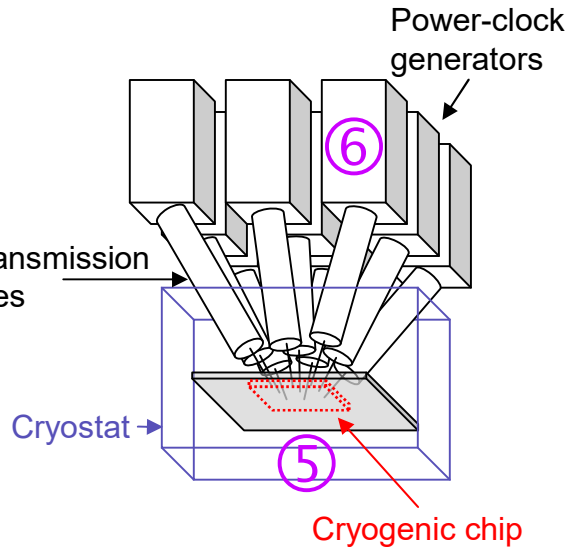

Backup: Data-controlled clock and bus circuits for the boxes

- Substitute these circuits, but keep flowchart interconnect as wires

①

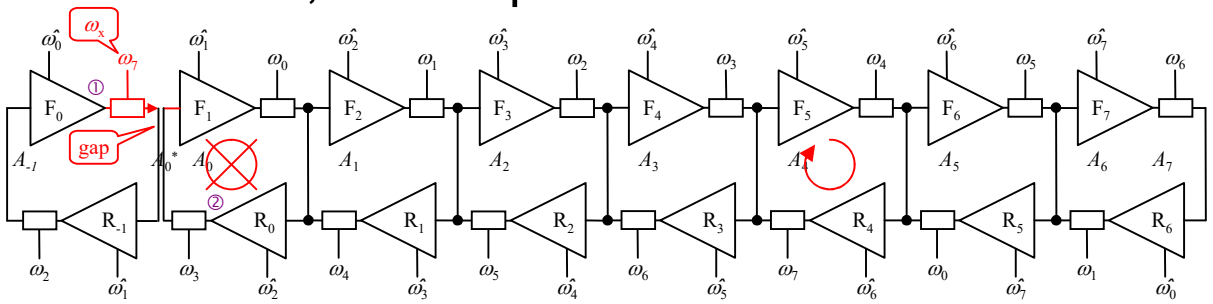

- Substitute crossover for diamonds (previous slide)
- Fab a chip

②

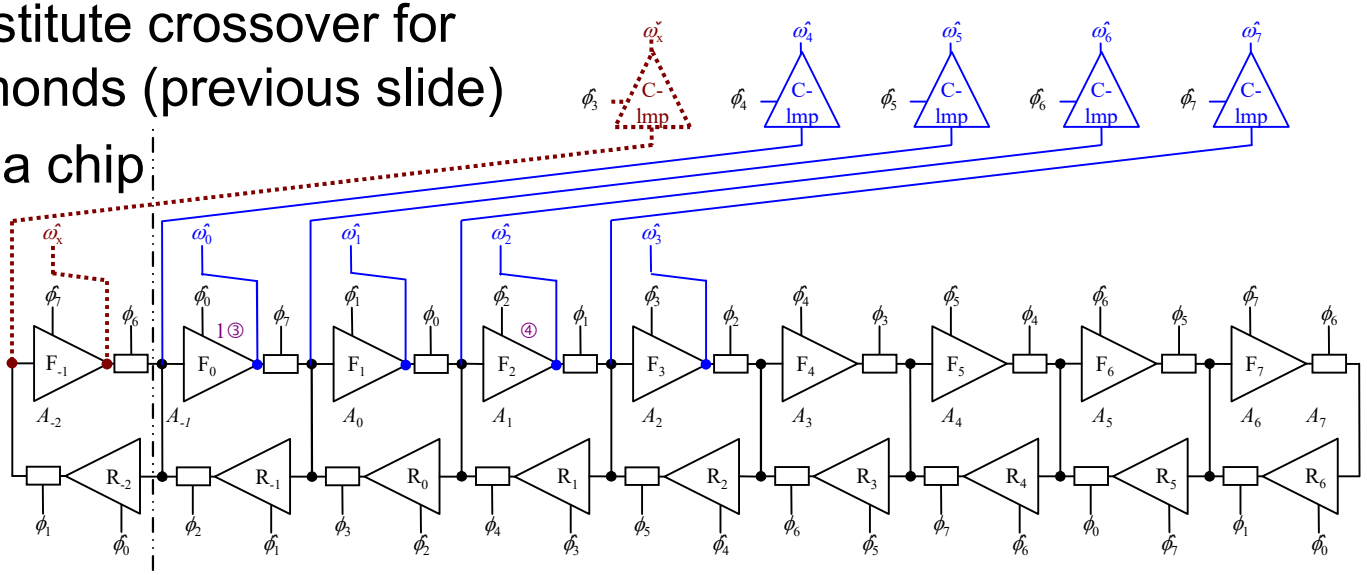

Existence proofs of what? This is what…

- No irreversible gates needed in cryostat
- Number of gates required in the cryostat is proportional to the length of source code rather then depth of recursion stack
- Circuit in cryostat is like a .gif/.zip decompressor

- More interesting results in backup material

Resource estimation I adapt ITRS/IRDS format

①

Table MM-9 Projected Electrical Specifications of Logic Core Device

YEAR OF PRODUCTION	2020	2022	2025	2028	2031	2034
	G48M36	G45M24	G42M20	G40M16	G38M16T2	G38M16T4
Logic industry "Node Range" Labeling (nm)	"5"	"3"	"2.1"	"1.5"	"1.0 eq"	"0.7 eq"
IDM-Foundry node labeling	i7-f5	i5-f3	i3-f2.1	i2.1-f1.5	i1.5e-f1.0e	i1.0e-f0.7e
Logic device structure options	FinFET	finFET LGAA	LGAA	LGAA	LGAA-3D	LGAA-3D
Mainstream device for logic	finFET	finFET	LGAA	LGAA	LGAA-3D	LGAA-3D
DEVICE ELECTRICAL SPECS						
Power Supply Voltage - Vdd (V)	0.70	0.70	0.65	0.65	0.60	0.55
Subthreshold slope (mV/dec) - HP (mV/dec)	78	82	72	75	72	70
Subthreshold slope (mV/dec) - HD (mV/dec)	72	75	63	68	65	65
Inversion layer thickness (nm) [1]	1.10	1.00	1.00	0.90	0.90	0.90
Vt,sat at Ioff=10nA/um - HP (mV)	222	237	212	226	217	211
Vt,sat (mV) at Ioff=100pA/um - HD (mV) [2][3]	345					326
Effective mobility (cm2/V.s)	125					100
Rsd (Ohms.um) [4]	285					221
Ballisticity.Injection velocity (cm/s)	1.39E+07	1				1.46E+07
Vdsat (V) - HP	0.160					0.140
Vdsat (V) - HD	0.178					0.140
Ion (uA/um) at Ioff=10nA/um - HP when Rsd=0 [5]	1737					1504
Ion (uA/um) at Ioff=10nA/um - HP [6]	854					760
Ion (uA/device) at Ioff=100nA/um - HP [7]	91					91
Ion (uA/um) at Ioff=100pA/um - HD when Rsd=0 [5]	1159					861
Ion (uA/um) at Ioff=100pA/um - HD [6]	484	495	546	521	459	347
Ion (uA/device) at Ioff=100pA/um - HD [7]	52	105	105	81	73	42
Cch,total (fF/um2) - HP/HD [8]	31.38	34.52	34.52	38.35	38.35	38.35
Gate height over fin (nm)	25	20	15	15	15	15
Spacer k value	4.0	3.5	3.5	3.0	2.5	2.5
Cch (fF/um) - HP [8]	0.45	0.44	0.39	0.37	0.37	0.37
Cch (fF/um) - HD [8]	0.50	0.50	0.39	0.37	0.37	0.37
CV/I (ps) - FO3 load, HP [9]	1.11	1.02	0.86	0.78	0.78	0.80
I/(CV) (1/ps) - FO3 load, HP [10]	0.90	0.98	1.16	1.29	1.28	1.25
Energy per switching [CV2] (fj/switch) - FO3 load, HP	0.66	0.65	0.49	0.47	0.40	0.33

②

③ These are spreadsheet equations with downward dependencies

④ Ultimate answer: "Energy per switching [CV^2] (fj/switch)"

Resource estimation II: quantum algorithm w/control

① *Algorithm*

Table MM-9

YEAR OF PRODUCTION	Shor	Opti-mize	Alg 3	Alg 4	Alg 5	Alg 6
	G48M36	G45M24	G42M20	G40M16	G38M16T2	G38M16T4
Logic industry "Node Range" Labeling (nm)	"5"	"3"	"2.1"	"1.5"	"1.0 eq"	"0.7 eq"
IDM-Foundry node labeling	i7-f5	i5-f3	i3-f2.1	i2.1-f1.5	i1.5e-f1.0e	i1.0e-f0.7e
Logic device structure options	FinFET	finFET LGAA	LGAA	LGAA	LGAA-3D	LGAA-3D
Mainstream device for logic	finFET	finFET	LGAA	LGAA	LGAA-3D	LGAA-3D
DEVICE ELECTRICAL SPECS						
Po	0.70	0.70	0.65	0.65	0.60	0.55
Su	78	82	72	75	72	70
Su	72	75	63	68	65	65
Inv	1.10	1.00	1.00	0.90	0.90	0.90
Vt,	222	237	212	226	217	211
Vt,	345					326
Effective mobility (cm2/V.s)	125					100
Rs	285					221
Ba	1.39E+07	1				1.46E+07
Vd	0.160					0.140
Vd	0.178					0.140
Ion	1737					1504
Ion	854					760
Ion (uA/device) at Ioff=100nA/um - HP [7]	91					91
Ion	1159					861
Ion	484	495	546	521	459	347
Ion	52	105	105	81	73	42
Cc	31.38	34.52	34.52	38.35	38.35	38.35
Ga	25	20	15	15	15	15
Sp	4.0	3.5	3.5	3.0	2.5	2.5
Cch (fF/um) - HP [8]	0.45	0.44	0.39	0.37	0.37	0.37
Cch (fF/um) - HD [8]	0.50	0.50	0.39	0.37	0.37	0.37
CV/I (ps) - FO3 load, HP [9]	1.11	1.02	0.86	0.78	0.78	0.80
I/(CV) (1/ps) - FO3 load, HP [10]	0.90	0.98	1.16	1.29	1.28	1.25
Energy per switching [CV2] (fj/switch) - FO3 load, HP	0.66	0.65	0.49	0.47	0.40	0.33

② Algorithmic operation counts

③ Operation count per error corrected qubit

④ Classical control parameters

These are spreadsheet equations with downward dependencies

⑤ Ultimate answer: joules for algorithm, e. g. Shor(834737)

Conclusions

- Resource model builds on Landauer's minimum dissipation and physics of computation rather than commercial products (e. g. ST Micro 28 nm, 22FFL)
- Developed RL controller out of 100% reversible shift registers
- Simulations suggest >100× dissipation reduction over a cryo CMOS work-alike (but there are many undetermined factors)
- Can prove various things using RL controller as an existence proof:
  - No irreversible gates needed in cryostat
  - Resources in cryostat are proportional to source code length
  - Universal cryostat architecture (like .gif/.zip controller)

- Additional data at https://zettaflops.org/qre-2022

Problem description

- Quantum computers are a national priority
- The structure on the right does not scale
- For cryogenic qubits, the accepted direction is to compress the data in the cables and use cryogenic electronics to decompress
- *De facto* cryo electronics is cryo CMOS at 4 K, which works about the same as 300 K
- CMOS improving at 2×/decade due to fab, which is not enough for the national priority
- Can we do better?
- Photo by Lars Plougmann, https://www.flickr.com/photos/criminalintent/39660636671, license https://creativecommons.org/licenses/by-sa/2.0/

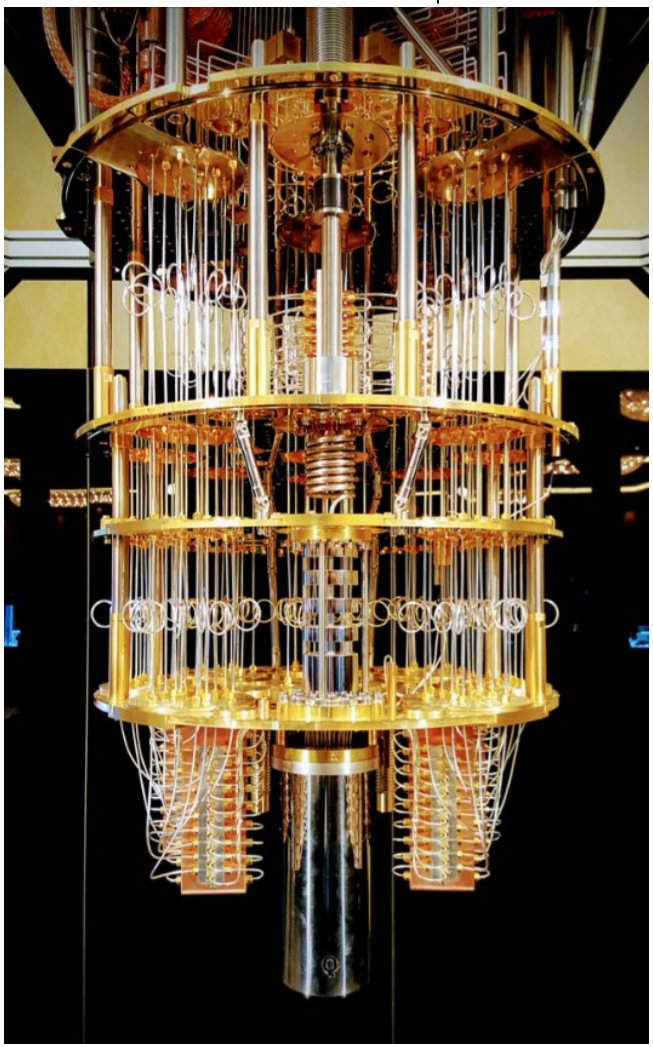

The adiabatic powertrain

- At room temperature, loop ① and ② recycle energy
- At 4 K, cryocooler overhead $P_S \approx$ 1,000, so recycling is no better than avoiding the cryocooler
  - Let ② be a resistor at room temperature $G_P = 0$
- R_{FOM}, Reversible Figure of Merit (FOM),
  - $R_{FOM} = E_{CMOS}/E_{Reversible}$ ④

- Three variable structure

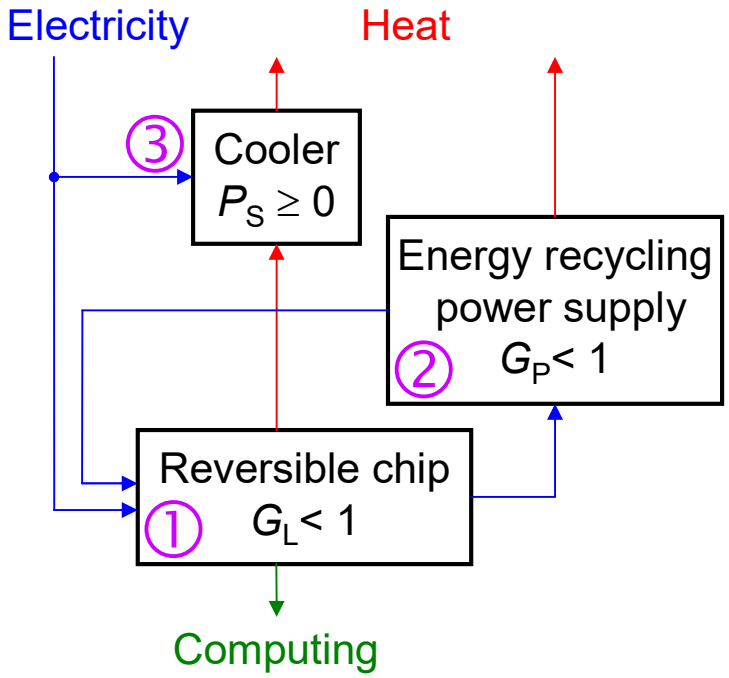

P_S = the cooling overhead of the cryocooler, or 0 if not present (P_S stands for "specific power")
$G_L = 1 - 2RC/\tau$, the portion energy not turned into heat
G_P = the cooling over head of the cryocooler, or 0 if not present
I_{on}/I_{off} is the on/off currents of typical transistors

Ref: Erik DeBenedictis, *Figure of Merit for Reversible Logic Systems*. Zettaflops, LLC technical report ZF011, https://zettaflops.org/wolte-2022/.

R_{FOM} at room temperature

- Reversible Figure of Merit (R_{FOM}) = $E_{reversible}/E_{CMOS}$ with no cryostat

Start with

$$R_{FOM} = \left(1 - G_L \frac{G_P + P_S}{1 + P_S}\right)^{-1} \quad ①\ ②$$

If $P_S = 0$ (heat sink),

$$R_{FOM} = \left(1 - G_L G_P\right)^{-1} \quad ③$$

For reference:

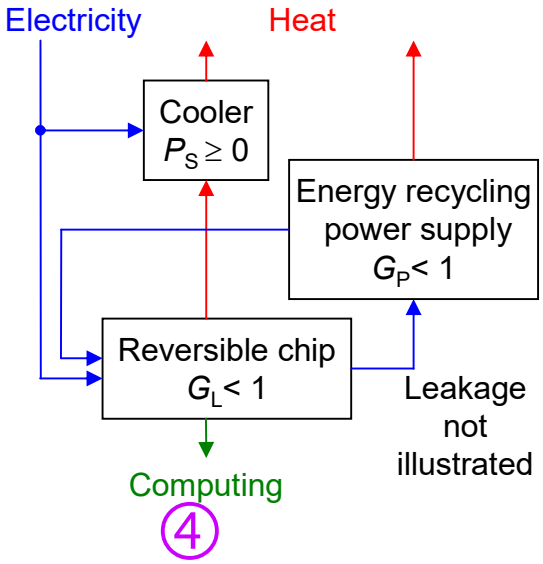

- This is the traditional approach

P_S = the cooling overhead of the cryocooler, or 0 if not present (P_S stands for "specific power")
G_L = 1 - $2RC/\tau$, the portion energy not turned into heat
G_P = the cooling over head of the cryo cooler, or 0 if not present
I_{on}/I_{off} is the on/off currents of typical transistors

R_{FOM} at 4 K

- Reversible Figure of Merit (R_{FOM}) = $E_{reversible}/E_{CMOS}$ with cryostat

Start with

$$R_{FOM} = (1 - G_L \frac{G_P + P_S}{1 + P_S})^{-1}$$ ②

As P_S goes from 0→∞, ①
$(G_P + P_S)/(1 + P_S)$ goes from G_P to 1

$$R_{FOM} = (1 - G_L)^{-1}$$ ③

For reference:

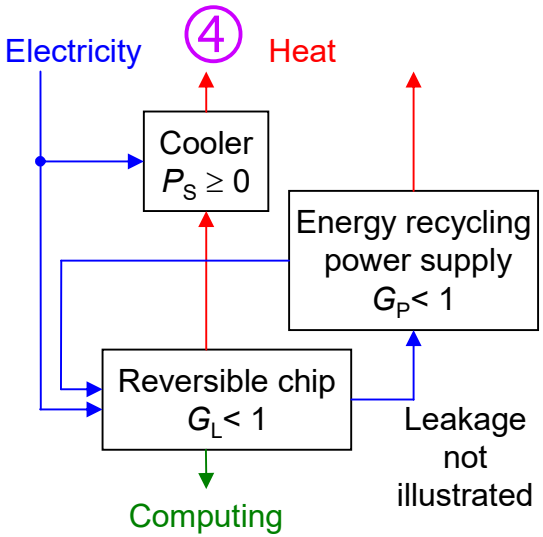

- Chip demos 1995-2008 validated G_L, which is what quantum needs!

P_S = the cooling overhead of the cryocooler, or 0 if not present (P_S stands for "specific power")
G_L = 1 - $2RC/\tau$, the portion energy not turned into heat
G_P = the cooling over head of the cryo cooler, or 0 if not present
I_{on}/I_{off} is the on/off currents of typical transistors

Projects with fewer big risks are more likely to be funded

- A reversible computing system needs a chip and an energy management component
  - Room temperature reversible computing uses an energy recycling power supply—proposed based on resonators, MEMs, and switching circuits
  - If a cryocooler is already present, it will fill the same role
- Reduces risk; suggests two-stage plan

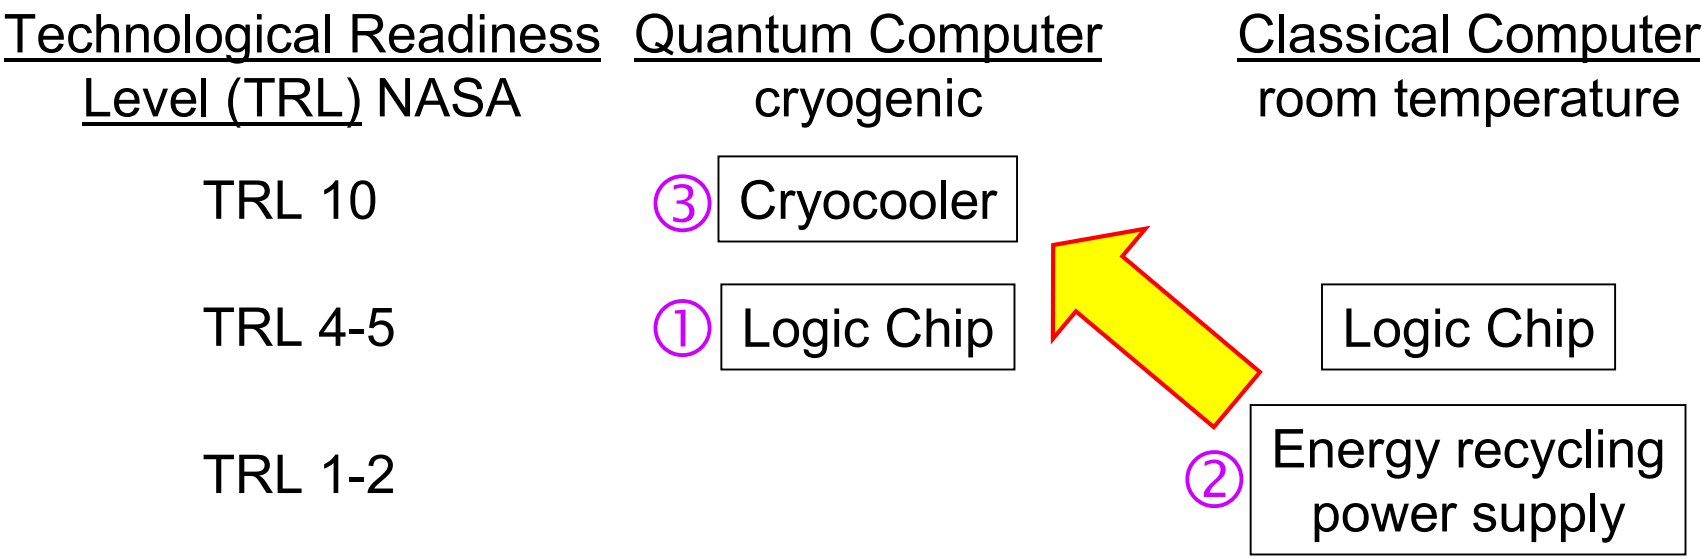

Existence proofs of what? This is what…

No irreversible gates required in cryostat

- Landauer's minimum is "order of *kT* per irreversible function"
- RL controller has no irreversible gates
- Cryo CMOS design tools lay out mostly irreversible gates

- Caveats: One of my papers has a CNOT for optimization; crossover is an externally controlled Fredkin gate; have data-controlled clocks

Reversible circuit is of bounded and reasonable size

- Reversible gate dissipation = $(2\ RC/\tau) \times (\frac{1}{2}CV^2)$
- For 1 μs quantum measurement, $(2\ RC/\tau) \approx 1/1{,}000$
- Bennett showed all-reversible computers are possible, but with overhead [Bennett]
- By construction, RL Controller has no infinite stack and data in shift register proportional to source code size

Flowcharts and Turing completeness

- Circuit is like a data decom-pressor for .gif and .zip files
- Musical measures are “symbols”
- Data stream influences a state machine, steering output of symbols kind of probabilistically based on symbol frequency
- Note: RL controller loads symbols during cryogenic cooldown using irreversible circuits
- Turing complete as a hybrid
- A paper [Böhm 66] showing that a flowchart is only “Turing complete” only if accompanied by a stack
- Righto, the stack is a room temperature and influences a stack-free flowchart through crossovers

Landauer's minimum dissipation in a hybrid system

Physicist vs. computer architect

- Landauer's minimum is kT per irreversible operation
- In a mixed temperature environment, which T do we use?
- Physicist's answer: The T of the environment performing the function
- Computer architect's answer: What are the limits of moving the irreversible operations to an environment where T is most favorable?

Proposed hybrid multi-temperature computing architecture

- Room temperature:
  - Compute the output
  - Compress the output
- Cable:
  - Move compressed string into the cryostat
- Cryo electronics:
  - Decompress the output

Results of Simulation

- ngspice Sky130; CMOS from standard cell
  - R_{FOM} 1 MHz 131 @ 27 C
  - Note: Other orgs chips include an instruction set, raising dissipation
- Simulation output:

```
*** OVERALL ADIABATIC ADVANTAGE
* override J_S in q2.cir to use full circuit and make sure .includes are q2.cir and ar.cir
*.param MD=5 J_S=2 Vp=1.9V     Hz=1e6      T=27  wX=1  ww=500e-9   ll=150e-9   Cw=.01p     Cb=0        xl=1.9      xn=1  xh=1.9      yl=1e6      yn=1  yh=1e6
* 0 Adia , 5 , 7.44291E-13 , Ecyc , 5.92981E-14 , Vp , 1.9 , tw , 1E+06 , Vx , 1.9 , 0.625 , 1E-14 , 27 , 1 , 1 , 960.726
*.param MD=5 J_S=2 Vp=2.15V    Hz=1e6      T=-55 wX=1  ww=500e-9   ll=150e-9
* 0 Adia , 5 , 9.25436E-13 , Ecyc , 5.35942E-14 , Vp , 2.15 , tw , 1E+06 , V
* now switch .includes to c2.cir and cr.cir
*.param MD=5 J_S=2 Vp=1.8V     Hz=1e6      T=27  wX=1  ww=500e-9   ll=150e-
* 0 CMOS , 5 , 7.8273E-12 , Ecyc , 7.77908E-12 , Vp , 1.8 , tw , 1E+06 , Vx
*.param MD=5 J_S=2 Vp=1.8V     Hz=1e6      T=-55 wX=1  ww=500e-9   ll=150e-
* 0 CMOS , 5 , 2.40974E-12 , Ecyc , 2.39103E-12 , Vp , 1.8 , tw , 1E+06 , V
```

③

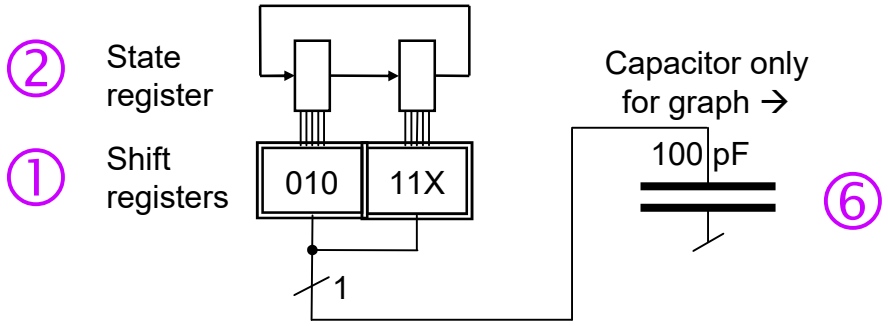

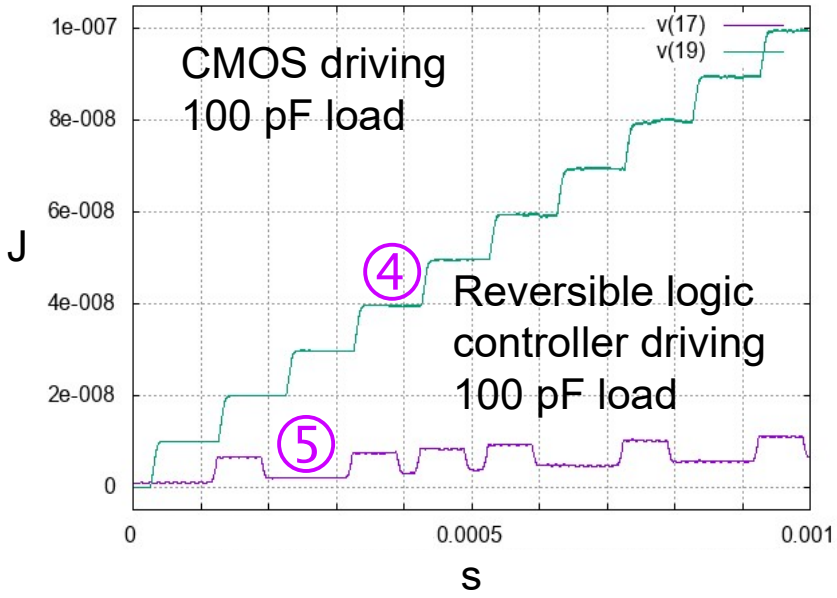

Sky130 validation (in progress)

- Sky130 is an “open” PDK for multi-project wafers, based on a 130 nm process (I have no cryo data)
- Activity
  - “Note note” has been hand-coded in ngspice and various simulation results have been presented at conferences
  - The replicable unit of Q2LAL (circuit family) is shown below. It has been extracted with parsitics and can be incrementally substituted into the hand-coded ngspice
  - Results are in line with predictions
    - 131× advantage over Cryo CMOS from Sky130 standard cells

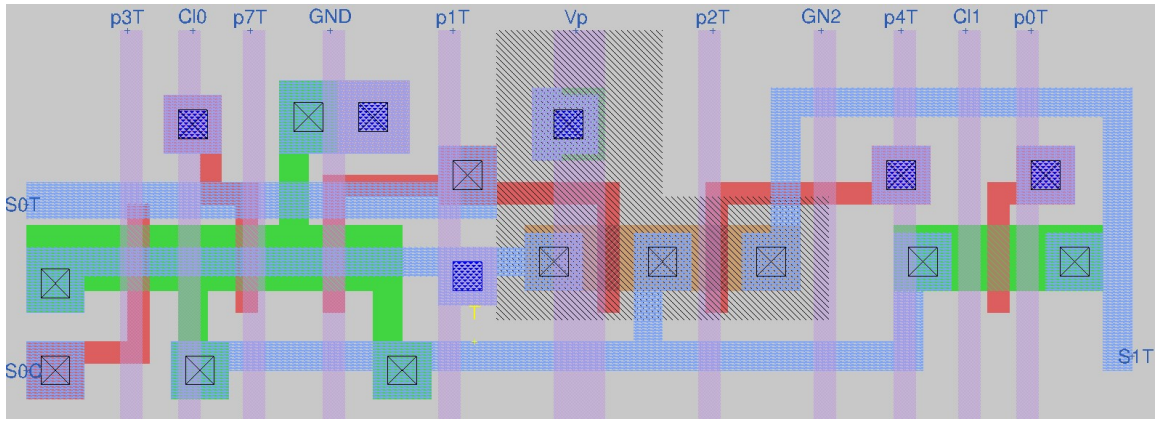

Sky130 full phase

- The cell illustrated is the most common stage
  - Top: The stage
  - 2nd quarter: Vertical flip
  - 3rd quarter: Horizontal flip
  - Bottom: 180° rotation
  - Not optimized
- There are other, less common cells
  - Data-controlled clocks
  - Bus interfaces
  - Crossovers (trivial)