Rebooting Quantum Computing

Erik P. DeBenedictis
Principal
Zettaflops, LLC
Albuquerque, NM, USA
erikdebenedictis@zettaflops.org

Elie K. Track
CTO
Hypres , Inc.
Elmsford, NY, USA
etrack@hypres.com

***Abstract*—Hardware R&D projects in quantum computing are mostly divided into the sub areas of qubits, cryo-CMOS control, and control based on Josephson junction electronics. This paper makes the case that hybrid solutions are possible and should give better results. Specifically, there is recent work demonstrating how to trade the performance of qubits for the energy dissipation of the control electronics. In addition, reversible logic can have properties that lie in between those of cryo-CMOS and Josephson junction electronics, forming an "impedance matcher" that allows all three technologies to be used together. The idea of hybridization also applies to Rebooting Computing's goal of picking the successor to Moore's law from a series of options. Instead, this paper suggests a second reboot that considers hybrids of several of the originally proposed approaches.**

Keywords—quantum computer, energy efficiency, cryo-CMOS, reversible logic, Josephson junction, SFQ, quantum error correction, scale up, Metric-Noise-Resource, roadmap, IRDS

I. Introduction

Many parties, including the IEEE Rebooting Computing Initiative, redefined computing-related R&D priorities in response to the end of Moore's law. The new priorities directed attention toward reversible logic, Josephson junction (JJ) electronics, quantum computing, and other approaches. As shown in Fig. 1, reversible logic and JJs did not thrive for the reasons shown by the × symbols, but quantum computing has attracted substantial investment and is now undergoing scale up [1-2] in preparation for commercialization.

Yet new data this year [3-4] shows that scale up of quantum computers will be dependent on reducing the dissipation of classical electronics, which is or was a topic of Moore's law. There would be a "research dependency deadlock" if quantum computing, a successor to Moore's law, could only succeed if Moore's law continued.

The point of this paper is that there is a gap in the reasoning above. The first pass through Fig. 1 was driven by consumer electronics so there was an unstated assumption everything would operate at room temperature. Calling the second pass through Fig. 1 "rebooting quantum computing," operating temperatures will be specified by optimization algorithms to as low as 100 mK without regard to consumers.

While reversible logic and JJ electronics did not thrive for room temperature applications, the point of this paper is that these technologies should work, and even have important roles, in a cryogenic quantum computer. This paper uses [5-6]'s model of a quantum computer, but goes beyond that paper by discussing specific new technologies.

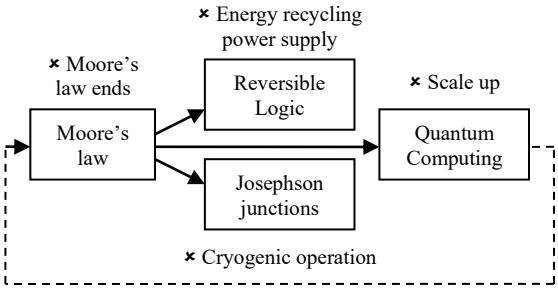

Fig. 1. Rebooting computing started when Moore's law ended. Three of the proposed "post Moore" directions were reversible logic, Josephson junction logic and quantum computing. Each has encountered an obstacle indicated by an ×. The obstacle to quantum computing scale up can be addressed by additional generations of Moore's law, but this creates a cycle. This paper shows that obstacles facing reversible logic and Josephson junctions will not apply on the second pass around the loop.

Vendor roadmaps [1-2] show qubit count growing exponentially to a million raw superconducting qubits—although many applications identified by analysts require vast multiples of a million qubits. Recent cryo-CMOS controllers demonstrated by Intel [3] and IBM [4] dissipate about 25 mW at 4 K, which is 5 W per qubit at room temperature (based on 200× cryocooler overhead). This translates to 5 MW per million qubits or $5M per year given a rough electricity cost of a dollar per watt-year. So, the power bill will be an obstacle to scaling up to a vast multiple of a million qubits.

The situation is analogous to Exascale supercomputers confronting and addressing exponentially growing power requirements. "Research papers that came out in 2008 predicted that an exaflop system would consume between 150 to up to 500 megawatts of energy" [7], or $150-$500M/year. This was too pricey for the US Government so it funded special-purpose, energy efficient computational units ("accelerators") that provide a 10× reduction—and Exascale supercomputers are coming online now at 20 MW. The community may need the quantum equivalent of what was called the Exascale Initiative.

A recent paper [6] presented a power estimation method called Metric-Noise-Resource (MNR). MNR assesses subcomponents contributing to a quantum computer's dissipation and concludes that control electronics is key and, when multiplied by a cryocooler overhead factor, becomes the main determiner of operating cost. However, a quantum computer's performance compared to a classical computer is primarily determined by qubit lifetime γ^{-1}, i.e. the quality of the qubit, and of course implicitly on the number of qubits and run time.

Refs. [5, p. 185] and [6, Fig. 11] identify a tradeoff between qubit lifetime and the dissipation of control electronics that can be used as a tool for designing a system constrained by a budget: raising qubit lifetime γ^{-1} by S will lower dissipation by S^2. Thus, the energy consumption E_{new} of a new quantum computer is

$$E_{new} = (\gamma_{new}/\gamma_{existing})^{-2}\ (A_{BC,new}/A_{BC,existing})\ E_{existing}, \qquad (1)$$

where $E_{existing}$ is the energy consumption of an existing quantum computer, $\gamma_{existing}^{-1}$ and γ_{new}^{-1} are qubit lifetimes of the existing and new computers, and A_{BC} is the control electronics dissipation. The initial set of parameters for the existing computer can be obtained from published measurements [3-4]. So, (1) allows what-if tests on design improvements. Once a new design is constructed, its measurements can define the "existing" parameters for the next generation.

The remainder of this paper will explain how reversible logic and JJ electronics will fit into the emerging framework for quantum computer scale up, revealing top-level properties about needed functions, the environment in which they must be performed, and where their metrics fit into system-level performance metrics.

II. Applying Reversible Logic to Quantum Computing

This section gives two reasons reversible logic did not thrive in the first pass through Fig. 1, but should thrive on the second pass. Quantum and classical reversible computing were both introduced at the Physics of Computation Conference in 1981, where Richard Feynman launched quantum computing with his now-famous paper "Simulating Physics with Computers" [8] and Edward Fredkin and Tommaso Toffoli launched classical reversible computing with "Conservative Logic" [9]. It should be no surprise that these two approaches with a common background work well together. Readers can consult [10] for an introduction to reversible logic.

A. Transistor-based reversible logic works better at cryo

Fig. 2 develops an expression for the reversible logic-to-CMOS energy factor $\mathscr{E}_{RC} = E_R/E_C$, where E_R and E_C are the reversible and CMOS energies for performing the same function, both parameterized by refrigeration overhead and hence applicable at all temperatures.

Fig. 2 derives

$$\mathscr{E}_{RC} = E_R/E_C = 1 - G_L(G_P + P_S)/(1 + P_S), \qquad (2)$$

where $G_L \lesssim 1$ and $G_P \lesssim 1$ are efficiency metrics for the reversible logic circuit and the energy recycling power supply. P_S is specific power, the number of watts required by the cryocooler to remove one watt from the cryostat to room temperature.

Reversible logic did not thrive at room temperature because the G_P parameter of energy recycling power could not be made high enough without physically large resonators. In a room temperature environment, $P_S = 0$ and $\mathscr{E}_{RC}$ simplifies to

$$\mathscr{E}_{RC}^{(PS=0)} = 1 - G_L G_P, \qquad (3)$$

which shows how the G_P parameter connects to the metric $\mathscr{E}_{RC}$. Since G_L and G_P cannot generate energy, $G_L \le 1$ and $G_P \le 1$. With these constraints, the value of $1 - G_L G_P$ will be dominated by whichever of G_L or G_P is furthest below 1. Measurements in Fig. 2c and [11, Table III] show G_P of the energy recycling power supply to be a number of times further from 1 than the G_L of the logic.

While $P_S \approx 1{,}000$ for a 4 K laboratory cryocooler, cryocooler efficiency depends on scale. $P_S \approx 200$ is typical of cryocoolers in particle accelerators and a good guess for a

(a) Figure of merit (lower is better):

$\mathscr{E}_{RC} = E_R/E_C = G_L\ (G_P + P_S) / (1 + P_S)$ ❶ ❼

(b) Energy flow diagram:

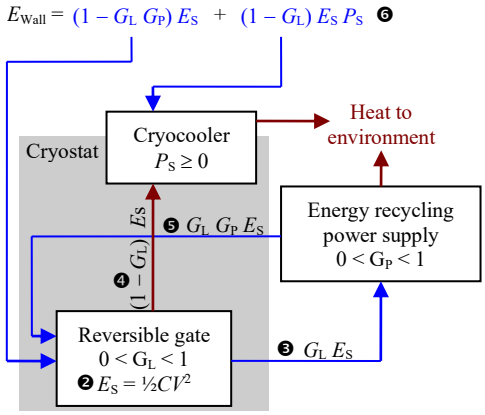

(c) Parameter values for circuit types:
CMOS: $G_L = 0$ by definition
Reversible logic: $G_L < 0.998$, $G_P < 0.993$ by measurement
Room temperature (heat sink cooling): $P_S = 0$

Fig. 2: (a) This figure computes the wall power consumption of reversible logic compared to cryo-CMOS, $\mathscr{E}_{RC} = E_R/E_C$ ❶, a lower-is-better figure of merit. (b) Signal energy is $E_S = ½CV^2$ ❷. A logic gate ejects the portion G_L ❸ of the signal energy electrically for energy recycling. The remaining portion $(1 - G_L)$ ❹ is turned into heat. An energy recycling power supply stores energy temporarily and then recycles it into a clock waveform, but this process has losses and only recycles fraction G_P of its input energy, for a total of $G_L G_P E_S$ ❺. Removing heat ❹ from the cryostat requires the cryocooler to draw a P_S times as much power from the wall, or $(1 - G_L) E_S P_S$ ❻. (c) The previous discussion yields E_{Wall} for reversible logic, or $E_R = E_{Wall}$. If $G_L = 0$, the previous explanation applies to cryo-CMOS, so $E_C = E_{Wall}^{(GL=0)}$ and $\mathscr{E}_{RC} = E_{Wall}/E_{Wall}^{(GL=0)}$ ❼.

future quantum supercomputer. Since all the other terms in the expression for $\mathscr{E}_{RC}$ are close to 1, P_S values of 1,000 and 200 are effectively infinite. If we simplify $\mathscr{E}_{RC}$ for $P_S \to \infty$, we get

$$\mathscr{E}_{RC}^{(P_S \to \infty)} = 1 - G_{L,} \quad (4)$$

where the large P_S made the G_P term disappear.

Substituting a resistor for the energy recycling power supply eliminates the large resonator, but results in a $G_P = 0$. However, this incredibly poor recycling efficiency does not create a problem because there is no G_P term in (4).

A reversible chip of the nRERL logic family was tested in 2000 and "the nRERL circuit consumed 0.19% of the dissipated energy of a static CMOS circuit" [11, p. 873]. This figure is from the chip in isolation (i.e. G_L but no G_P), yielding or $\mathscr{E}_{RC} = 0.0019 = 1/526$ or $G_L = 0.9981$. This test shows plenty of upside, but needs to be repeated at 4 K.

B. How fast does control electronics need to be?

Another reason reversible logic did not thrive for room temperature applications is that reversible logic's greatest energy efficiency advantage occurs when the circuits run more slowly than CMOS, but this section shows that control electronics should run more slowly than CMOS anyway.

A microprocessor's throughput and hence its value is proportional to clock rate, but a quantum computer's throughput is determined by the qubits, not the classical control system. The classical control system needs to keep up with the qubits, but running faster than the qubits is not helpful.

The connection between the speed and throughput of reversible logic is in the circuit equation $1 - G_L = 2RC/\tau$ [10, p. 11], where R is the resistance of a transistor in the 'on' state, C is the capacitance of the electrical node comprising transistor terminals and interconnect wire, and τ is ¼ of the period of the AC clock. The CMOS gate delay is close to RC, so the term $2RC/\tau$ is essentially a slowdown factor compared to CMOS.

The actions of researchers provide data points on the speed needed by control electronics. The three cryo-CMOS quantum computer control chips considered in this paper include digital controllers at 4 K [3-4], and 100 mK [12]. These controllers were fabricated with CMOS processes suitable for microprocessors with 1-4 GHz clock rates. The circuits at 4 K [3-4] include mixed-signal digital-to-analog converters (DACs), analog-to-digital (ADCs), and signal processors that run at clock rates around 1 GHz. However, the digital electronics in one chip has a 1/16 clock (62.5 MHz) [4, p. 360] and another was tested at 8 frequencies from 10 MHz to 1 GHz [12, Fig. 4e]. Furthermore, the researchers in [6, p. 8] use a time step τ_{step} = 100 ns (10 MHz) for the quantum computer, saying it was "determined by its slowest gate." These clock periods are several orders of magnitude longer than a CMOS gate delay and support the idea that designers are willing to trade the speed of the digital electronics for energy efficiency.

The empirical data from [3-4, 12] follow a pattern based on complexity of behavior. The portions of the controller chips with the slower clocks include something akin to an instruction set. Instruction sets can create a wide range of behaviors based on software, yet any specific behavior can be hard coded for greater efficiency.

Ref. [5, p. 186] then considered the possibility that the DACs, ADCs, etc. could be "turned off" when idle. If CMOS is only powered up 10% of the time because it is not needed at other times, dissipation drops to 10% while still getting the job done. This familiar power-savings method was used in the most scalable solution discussed in [5], but the next section shows how to do better.

III. An all-semiconductor improvement

The preceding section sets the stage for a solution that exploits the novel properties of reversible logic.

The previous section indicated that designers choose clock periods for the purely digital cryo-CMOS that is more than an order of magnitude longer than the natural RC time constant of the transistors. Even so, the clock can be turned off most of the time and still get the job done. So, how about using reversible logic where the clock would be "on" all the time, yet at a lower frequency? This would allow the equation $1 - G_L = 2RC/\tau$ to yield a G_L closer to 1 where there is a higher degree of energy recycling. This is the implementation of scenario B in [5, footnote 7 on p. 169] referencing [13].

The authors have included an example of the approach above in other publications [14], yet these other publications use a small number of externally generated signals in lieu of on-chip DACs—a variant of the PL/AL architecture [15]. There is also a publication [16] that discloses both the semiconductor approach above and hybrid approach discussed in the next section.

A. Combining improvements

The large number of potential technology hybrids will require refining the term A_{BC}. This paper has been written to be consistent with scenarios A, B, and C in [6], with values of A_{BC} = 1, 10^{-2}, and 10^{-4}, but there are at least four improvements for A_{BC} that can be used in combination with each other. These improvements are defined below and include the authors' opinion about the maximum upside of each in parenthesis:

$A_{process}$ (5×): Both cryo-CMOS and reversible logic implementations should update naturally to newer semiconductor processes over time. For example, current devices and published papers use semiconductors processes like 22 nm FinFET Low Power (22FFL) where a successor might use a future 3 nm Gate-All-Around (GAA) process.

$A_{reversible}$ (100×): The circuit is changed from CMOS to reversible logic.

$A_{architecture}$ (10×): Various design optimizations are applied to both cryo-CMOS and reversible logic. This would include powering down subsystems when they are not being used [5, p. 186] and replacing current instruction set processors with optimized state machines.

A_{hybrid} (5×): Improvement through selective use of JJs, discussed below.

IV. Hybrid Semi- Superconductor Improvement

From the perspective of Rebooting Computing, a semi-superconducting hybrid creates many new opportunities. The novelty in this paper is to connect these new directions to a quantum computer scale up model.

A. Physical structure

JJ chips are created using an augmented semiconductor process. Fabrication starts with a blank silicon wafer as a mechanical substrate for evaporative deposition of patterned superconductor material, insulator, and normal metal.

In lieu of a blank silicon wafer, a semi- superconductor hybrid starts with a semiconductor wafer containing circuits that will perform parts of the ultimate hybrid chip's function. The physical interface between transistors and JJs could be through standard layer-to-layer vias and interface circuits would convert between voltage-based signals on the semiconductor layer and Single Flux Quantum (SFQ) pulses. DC voltages and currents could also be passed across the interface.

B. Architectural issues

If one looks at control electronics as a function, how should the function be divided into circuits that are placed on the semiconductor and superconductor layers?

Table 1 illustrates the properties of the devices in the hybrid. The partitioning strategy is to allocate complex circuits, such as memory and complex logic, to the semiconductor layer. Transistors are small, so complex circuits will not consume too much chip area. While transistors can run about as fast as JJs, the power consumption of each reversible logic circuit grows quadratically with the speed of the circuit, so the designer will want to run most of the transistor circuits at MHz speeds. JJs are physically large but are energy efficient even at high speed. So, the designer's motivation should be to reserve JJs for simple circuits that require high speed.

TABLE I. Properties of Devices in the Hybrid

	Semiconductor reversible	**Josephson Junction**
Speed	High power above 1 MHz	GHz
Power	High power above 1 MHz	Always very low
Density	10^9 or more devices/chip	10^6 devices/chip
Signal type	Voltage based with typical 1 V swing	Current signals of 10-100 μA

The published literature on cryo-CMOS quantum controllers [3-4, 12] assign digital logic, DACs, ADCs, and various analog components to clock domains based on similar criteria. Engineering a semi- superconductor hybrid could start by looking at existing cryo-CMOS controllers and assigning circuits with digital clock rates or analog bandwidth of 1 GHz or higher to the superconductor layer and the other circuits to the semiconductor layer.

C. Emergent properties of the hybrid

The combination of semi- and superconductors opens up new design opportunities, which will be explained with an example. Cryo-CMOS qubit controller chips [3-4] use an analog mixer circuit to combine a stable multi-GHz sine wave with highly specialized envelope waveform produced by a DAC from bits stored in local (cryogenic) memory. JJs may be the best choice for the DAC and mixer, but JJs' large physical size limits the amount of the memory available to hold envelope waveforms. The approach in Fig. 3 exploits the new opportunity.

Waveform storage has an intrinsically sequential access pattern, and semiconductor reversible logic is quite efficient for implementing shift registers—as long as the speed is not so high that dissipation becomes excessive. Several data streams can be combined into one stream that is several times as fast. However, reversible logic becomes energy inefficient at high speeds, creating a disincentive for using the combining process in semiconductors too many times. So, the idea is to store digitized envelopes in wide shift registers, illustrated as 128 bits wide in Fig. 3c. The wide stream would be narrowed and sped up through an interleaving multiplexer circuit [13, Fig. 6] until it is the correct width for a JJ-based DAC. However, the width and speed of the stream when it crosses the interface from semi- to superconductors will affect the effectiveness of the resulting system. If the signal is too fast, the semiconductors will dissipate too much power, but if the signal is too wide, the number of JJs required will consume too much chip area. So, Fig. 3 illustrates two combining stages on the

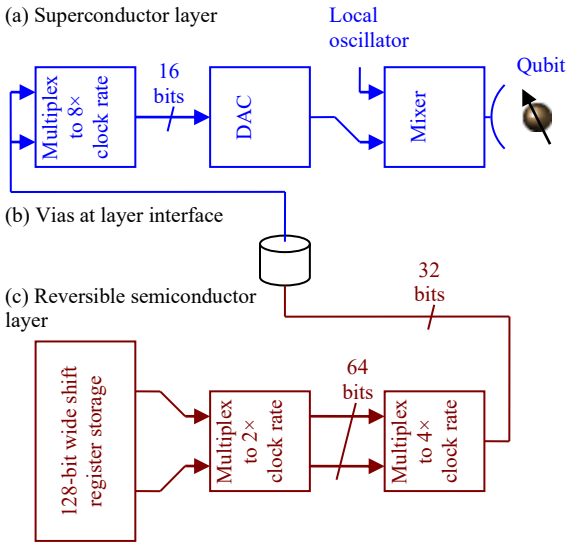

Fig. 3. A hybrid structure for waveform generation. (c) Digitized waveforms are stored semiconductor shift registers devices due to transistors' small size and high storage density. The shift registers are wide (128 bits) and slow to minimize the number transistors operating at high speed and hence high power. However, the stream can be narrowed to 32 bits and sped up to 4× clock rate. (b) The data is transferred to between technologies using standard vias, resulting in single flux quantum (SFQ) pulses. (a) The faster but physically larger JJs further narrow the stream to 16 bits at 8× clock rate, driving a SFQ to analog DAC, which then drives a fully analog mixer. The resulting waveform drives qubits.

semiconductor layer and one on the superconductor layer.

Rebooting Computing has considered reversible logic and JJs to be in scope, but engineering the interface between such technologies in a hybrid system has not been considered in scope to date—but perhaps could be in the future.

D. Hybrids enable new architectures

Clearly, Moore's law will not continue to reduce CMOS dissipation to $A_{BC} = 10^{-4}$ of current levels, but [6, p. 27] includes "scenario C" for SFQ that shows that a reduction of this magnitude could be key to quantum computer scale up.

DACs are the component with the highest dissipation, so they should benefit most from a technology hybrid. JJ DACs capable of generating qubit control signals have been built [17], yet it is not known at present how to scale their approach to the 100 nW level needed to reach $A_{BC} = 10^{-4}$. If future innovations can reach 100 nW, Fig. 3 could become an improvement over the cryo-CMOS controller approach [3-4].

However, there is an approach for controlling qubits with a stream of SFQ pulses. For background, the size of SFQ pulses is very uniform because it is determined by physical constants rather than geometry, where geometry inevitably varies during manufacturing. Exposing a qubit to a specific timed stream of SFQ pulses will apply a gate operation to the qubit [18].

It will be important to know whether envelope modulation or an SFQ pulse stream yields the longest qubit lifetime γ^{-1} and gate fidelity, because qubit lifetime combines with A_{BC} to determine the viability of scaled up quantum computers. Yet, hybrid technologies may be needed merely to answer this question. Lack of dense memory is mentioned in [18] by architects working on scalability, who apparently feel handicapped by the size of JJs and the dissipation of transistors. This provides an opportunity for Fig. 3 to help.

In other documents [16], the authors consider hybrid architectures like Fig. 3 but with more than just DACs. In general, these documents describe a semiconductor layer that creates control words at MHz rates. The control words move across an interface to the superconductor layer, which decodes the control words and uses JJs to apply specified functions to GHz waveforms originating at room temperature and shared by multiple qubits. These hybrid architectures can use a broad variety of JJ-based mixed-signal circuits, including switches, mixers, phase shifter, filters, phase-locked loops, etc.

V. Conclusions

This paper assembles ideas that could lead to R&D plans, a roadmap, benchmarks, or a second pass at Rebooting Computing.

Ref. [6] created an analytical framework for estimating and optimizing the combined quantum, classical, and algorithmic aspects of a quantum computer. Recent cryo-CMOS chips [3-4, 12] calibrate the framework to the current state of the art and [6] deals with the future as scenarios A, B, and C (where A is about 10× lower energy than [3-4] and B and C are additional steps of 100×).

A contribution of this paper is to explain how scenarios B and C could be reached by improvements that might come from an R&D program—improvements that include but go beyond the authors' work [14, 16]. Plugging plans for improvements into the framework of [6] would yield a roadmap.

Improvements could be rated by benchmarks. For example, Fig. 3 shows a waveform playback subsystem. A benchmark could be created that assesses the performance of different implementations of waveform playback such as, for example, chip area and dissipation for playing back a standard waveform. The benchmark results could be plugged into the framework [6] to show how much waveform playback subsystems from different R&D groups would help a roadmap.

At a higher level, this work argues that a second pass at Rebooting Computing should be based on technology teamwork rather than "winner takes all."

Quantum computer R&D has been like a horse race. Investors have horses with names like "cryo-CMOS," "SFQ," "reversible," "transmons," and "ions" that compete in races to see which one will win the purse.

In contrast, consumer electronics uses different varieties of transistors, with names like "CMOS," "DRAM," "Flash," and "RF." It is counterintuitive, to the authors anyway, that Apple would run a competition to determine whether RF transistors are better than DRAM transistors, or vice versa, and then redesign iPhones with just the winning type.

An idea in this paper is that quantum computing has now supplied Rebooting Computing with a goal, a goal that would be best satisfied by putting several familiar horses into a team and seeing how much the team can pull when working together.

The authors' past work [14, 16] had shown how reversible logic could act as mortar to bind semi- and superconductor circuits, taken as bricks, into a hybrid that has additional properties, like a brick wall. This work shows how the additional properties address important issues in quantum computer scale up.

There is another aspect to this story: It turns out that the CNOT gate was originally called the Feynman gate. As mentioned earlier, two of the horses, which could be named "qubit" [8] and "reversible logic" [9], were born at the same conference in 1981. Investigating a little more deeply reveals that the horses have a common ancestor, namely both approaches are based on a new set of "reversible" logic gates that have special significance in physics—rather than being based on coordinating conjunctions in English like AND-OR-NOT. So, the team has a brother and sister with a unique family talent. This paper tries to exploit this talent. Readers will of course note that CNOT (Feynman), Fredkin, and Toffoli are principal multi-input gates in both quantum and classical reversible computing.

This paper's novelty in the quantum area is to put qubits together with classical reversible computing. In a sense, the novelty is to put Humpty Dumpty together again—yet we must acknowledge that Humpty Dumpty was originally whole.

REFERENCES

[1] Neven, Harmut, "Quantum AI Update," Google Symposium 2022, offset 15:00. https://eventsonair.withgoogle.com/events/qss-2022 (note: this link requires registration).

[2] Expanding the IBM Quantum roadmap to anticipate the future of quantum-centric supercomputing, https://research.ibm.com/blog/ibm-quantum-roadmap-2025.

[3] Pellerano, Stefano, et al. "Cryogenic CMOS for Qubit Control and Readout." 2022 *IEEE Custom Integrated Circuits Conference (CICC)*. IEEE, 2022. DOI: https://doi.org/10.1109/CICC53496.2022.9772841 https://pure.tudelft.nl/ws/portalfiles/portal/122719163/Cryogenic_CMOS_for_Qubit_Control_and_ReadoutTaverne.pdf.

[4] Frank, David J., et al. "A Cryo-CMOS Low-Power Semi-Autonomous Qubit State Controller in 14nm FinFET Technology." *2022 IEEE International Solid-State Circuits Conference (ISSCC)*. Vol. 65. IEEE, 2022. DOI: https://doi.org/10.1109/ISSCC42614.2022.9731538.

[5] Fellous-Asiani, Marco. *The resource cost of large scale quantum computing*. Diss. Université Grenoble Alpes 2022. https://arxiv.org/abs/2112.04022v2.

[6] Fellous-Asiani, Marco, et al. "Optimizing resource efficiencies for scalable full-stack quantum computers." *arXiv preprint arXiv*:2209.05469 (2022). https://arxiv.org/abs/2209.05469v1.

[7] Frontier to Meet 20MW Exascale Power Target Set by DARPA in 2008 https://www.hpcwire.com/2021/07/14/frontier-to-meet-20mw-exascale-power-target-set-by-darpa-in-2008/.

[8] Feynman, Richard P. "Simulating physics with computers." *International Journal of theoretical physics* 21.6/7 (1982).

[9] Fredkin, Edward, and Tommaso Toffoli. "Conservative logic." *International Journal of theoretical physics* 21.3-4 (1982): 219-253.

[10] Younis, Saed G. *Asymptotically Zero Energy Computing Using Split-Level Charge Recovery Logic*. Diss. Massachusetts Institute of Technology, 1994. https://dspace.mit.edu/handle/1721.1/7058.

[11] Lim, Joonho, Dong-Gyu Kim, and Soo-Ik Chae. "nMOS reversible energy recovery logic for ultra-low-energy applications." *IEEE Journal of Solid-State Circuits* 35.6 (2000): 865-875. DOI: https://doi.org/10.1109/4.845190 https://s-space.snu.ac.kr/bitstream/10371/16274/1/nMOS%20reversible%20energy%20recovery%20logic%20for%20ultra-low-energy%20applications.pdf.

[12] Pauka, S. J., et al. "A cryogenic CMOS chip for generating control signals for multiple qubits." *Nature Electronics* 4.1 (2021): 64-70. DOI: https://doi.org/10.1038/s41928-020-00528-y.

[13] DeBenedictis, Erik P. "Adiabatic circuits for quantum computer control." *2020 International Conference on Rebooting Computing (ICRC)*. IEEE, 2020. https://doi.org/10.1109/ICRC2020.2020.00004 https://debenedictis.org/erik/CATC/Log_Shift_Reg_v1.02.pdf.

[14] DeBenedictis, Erik P. "*Classical Reversible Logic Circuits for Quantum Computer Control*," Zettaflops, LLC Technical Report ZF010. https://debenedictis.org/erik/CATC/MgtECompZF010v3.pdf.

[15] Hornibrook, J. M., et al. "Cryogenic control architecture for large-scale quantum computing." *Physical Review Applied* 3.2 (2015): 024010 DOI: https://doi.org/10.1103/PhysRevApplied.3.024010 https://arxiv.org/abs/1409.2202.

[16] DeBenedictis, Erik P. Managing Energy in Computation with Reversible Circuits. Patent Application No. WO2022197556. September, 2022. https://patentscope.wipo.int/search/en/detail.jsf?docId=WO2022197556.

[17] Sirois, Adam J., et al. "Josephson microwave sources applied to quantum information systems." IEEE Transactions on Quantum Engineering 1 (2020): 1-7. DOI: https://doi.org/10.1109/TQE.2020.3045682.

[18] Jokar, Mohammad Reza, et al. "DigiQ: A Scalable Digital Controller for Quantum Computers Using SFQ Logic." *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 2022. DOI: https://doi.org/10.1109/HPCA53966.2022.00037 https://arxiv.org/abs/2202.01407.