

~~FOR OFFICIAL USE ONLY~~

SUPERCONDUCTING COMPUTING

(b)(3)-P.L. 86-36

(b)(6)

Study Leader:

[REDACTED]

Contributors Include:

[REDACTED]

D. Eardley

R. Garwin

J. Goodman

D. Long

D. Meiron

[REDACTED]

W. Press

[REDACTED]

R. Schwitters

[REDACTED]

[REDACTED]

JSR-11-120

February 22, 2012

[REDACTED]

~~FOR OFFICIAL USE ONLY~~

Approved for Release by NSA on 08-07-2020, FOIA Case # 65695


Contents

1 EXECUTIVE SUMMARY	3
1.1 Task Statement	3
1.2 Superconducting Computing Workshop	3
1.3 Findings	3
1.4 Recommendations	5
2 TASK STATEMENTS	6
3 SUPERCONDUCTING COMPUTING WORKSHOP	8
4 SUMMARY OF SUPERCONDUCTING COMPUTING STUDY CONCLUSIONS	9
4.1 Previous Supercomputer Studies	9
4.2 Single Flux Quantum (SFQ) Logic and Memory	9
4.3 Comparison of SFQ Logic with CMOS Technology	10
5 PREVIOUS SUPERCOMPUTER STUDIES	11
5.1 High Technology Multithreaded Technology	11
5.2 Superconducting Technology Assessment	13
5.3 DARPA Exascale Study	15
6 QUESTIONS	18
6.1 Making Comparisons	18
7 INTRODUCTION TO JOSEPHSON JUNCTION DEVICES	20
7.1 Josephson Transmission Line Element	21
7.2 Single Flux Quantum Data Transmission	22
7.3 Resistively Shunted Junction (RSJ) Model	23
7.4 Critical Current Density	25
7.5 Stability against Thermal and Non-Thermal Noise	27
7.6 Lower Limits to the SFQ Switching Energy and Junction Area	29
7.7 Energy-efficient RSFQ (ERSFQ) Logic and Reciprocal Quantum Logic (RQL)	31
7.8 Cooled Memory	32

8	COMPARISON OF SFQ AND CMOS LOGIC	39
8.1	What are Potential Game Changers for Superconducting Computing?	42
8.2	Design Challenge	44
9	FINDINGS AND RECOMMENDATIONS	45
A	APPENDIX: WORKSHOP AGENDA	46
B	APPENDIX: ARCHITECTURAL CHALLENGES	48
B.1	Memory	49
B.2	Processors	50
B.3	Computer System	51
B.4	Shared Memory Multiprocessor	52
B.5	Milestones	53
C	APPENDIX: BASIC JOSEPHSON JUNCTION TUTORIAL	55
C.1	The Resistively-Shunted Josephson Junction	55
C.2	A Flux Memory Element: JJ in a Superconducting Loop	57
C.3	Thermal Fluctuations and the Minimum Critical Current	59
C.4	Scaling Prospects for JJ Logic	60
D	APPENDIX: JOSEPHSON JUNCTION DEVICE MODEL	62
D.1	RSJ Model	63
D.2	Driven Damped Pendulum Model	65
D.3	Tilted Washboard Model	67
D.4	Rapid Single Flux Quantum (RSFQ) Logic	69
E	ACRONYMS	74

1 EXECUTIVE SUMMARY

1.1 Task Statement



1.2 Superconducting Computing Workshop

A three-day workshop on Superconducting Computing was held at the summer study on June 13 to 15, 2011. Leaders in the superconducting field from Hypres, IBM, NASA-JPL, Northrup Grumman, Pacific Northwest National Laboratory, Raytheon-BBN Technologies, Stony Brook University, and the University of Southern California, addressed the Task of our study in a series of presentations, with active discussion.

1.3 Findings

- Single Flux Quantum (SFQ) logic could provide small switching energies ~ 0.1 aJ and switching times ~ 2 ps at a temperature of 4 K.

- Energy-efficient Rapid Single Flux Quantum (ERSFQ) and Reciprocal Quantum Logic (RQL) reduce energy consumption for the 'off' state.
- Avoiding thermal errors sets a fundamental lower limit to the switching energy $E_{sw} > 0.6aJ$ at 4K. Avoiding non-thermal noise would require a larger E_{sw} .
- Avoiding thermal errors determines the minimum critical current and sets the minimum area A and size $A^{1/2}$ of the Josephson junctions, which is also determined by their critical current density J_c .

today $A > 0.3\mu\text{m}^2$; $A^{1/2} > 0.5\mu\text{m}$; $J_c \sim 10\text{kA}/\text{cm}^2$

future $A > 0.03\mu\text{m}^2$; $A^{1/2} > 0.2\mu\text{m}$; $J_c \sim 100\text{kA}/\text{cm}^2$.

- Fundamental lower limits to the switching energy E_{sw} and area A will eventually block scaling to smaller sizes and energies.
- The switching time $\Delta t_{sw} \sim 2ps$ for SFQ logic is comparable to 28 nm CMOS logic. The switching time for $\Delta t_{sw} \sim 10ps$ measured in a ring oscillator.
- The switching energy for CMOS logic is comparable to the 'wall plug' switching energy for SFQ logic that includes to the energy to cool the junctions to liquid He temperatures.

item	28nm CMOS	SFQ Logic	SFQ (wall plug)
switch	100 aJ	0.1 aJ	30 aJ
gate	200 aJ	0.4 aJ	120 aJ

- SFQ advantages - bits travel along superconducting transmission lines with little loss and Josephson Magnetic Random Access Memory (JM-RAM) promises comparatively low read energies.

item	28nm CMOS	SFQ Logic	SFQ (wall plug)
1 bit data, 1 mm wire	100 fJ	< 1 fJ	< 1 fJ
read 32b from 8kB RAM	5000 fJ	0.16 fJ	50 fJ

- Josephson junction memory has the potential advantages of lower read and write energies than CMOS technology.
- JJ memory has large cell areas $\sim 200\mu\text{m}^2$ today for a 250 nm RQL process. Replacing the inductor loop by a multilayer inductor coil or JJ inductors could reduce the cell area to $\sim 1\mu\text{m}^2$.
- The practicality of SFQ technology for a specific application could be assessed by posing a design challenge for a small-scale system compared with CMOS.

1.4 Recommendations

- A petascale superconducting general processor should not be pursued. Single Flux Quantum logic is unlikely to provide an implementation that is superior in speed or energy to CMOS technology in the next decade.
- A diverse set of technologies should be explored today to overcome the bottlenecks of CMOS technology in the future. Examples are: Cooled CMOS technology at 77K or 4K; Low-energy on- and off-chip communication to energies 0.5 pJ/bit; Data transmission over lossless superconducting transmission lines; Electrical signaling over carbon nanotube or graphene interconnects; Low-voltage swing (~ 10 mV) CMOS signaling.

2 TASK STATEMENTS

High performance computing is essential to the mission of the National Security Agency (NSA). After a long period of exponential growth, CMOS technology is beginning to approach fundamental limits to its performance. We need to understand what technology will come next. In the Task Statement presented in the Executive Summary, the [REDACTED]

[REDACTED]

Superconducting computers have existed since the 1950's, and they were actively pursued by IBM through a major R&D program in the 1970's. A new approach to computing was reported by Likharev and Semenov (1991) [1] -- Single Flux Quantum (SFQ) logic. In this approach, each bit of information is represented by one flux quantum inside the computer. A flux quantum is the smallest magnetic flux that can be trapped inside a superconducting loop. Using Josephson junction (JJ) devices, an SFQ computer manipulates the motion of flux quanta to implement the gates needed for Boolean logic. SFQ devices can also trap a flux quantum to act as a memory element. Two new approaches to SFQ logic providing lower switching have been invented recently that greatly reduce the energy consumption: Energy-efficient Rapid SFQ (ERSFQ) and Reciprocal Quantum Logic (RQL).

Single Flux Quantum logic can be very fast, with switching times ~ 1 ps, and it can have very low switching energies ~ 0.1 aJ at liquid-helium temperatures ~ 4 K. SFQ logic has the potential to create computers with higher speed and lower bit energies than conventional CMOS technology. To compare cooled SFQ logic with CMOS switches that operate at room temperature, we need to multiply the switching energy by a factor ≥ 75 to account

for the energy needed to cool the processor to 4K, which is determined by the Carnot efficiency 75 W/W of an ideal refrigerator. Commercial refrigerators achieve cooling factors $\sim 300 \text{ W/W}$ that don't closely approach the Carnot efficiency. Increasing the cooling efficiency is an important goal for a superconducting computers, addressed in Section 8.1.

3 SUPERCONDUCTING COMPUTING WORKSHOP

A three-day workshop on Superconducting Computing was held at the summer study on June 13 to 15, 2011. Leaders in the superconducting field from Hypres, IBM, NASA-JPL, Northrup Grumman, Pacific Northwest National Laboratory, Raytheon-BBN Technologies, Stony Brook University, and the University of Southern California, addressed the Task of our study in a series of presentations, shown in the agenda in Appendix A.

4 SUMMARY OF SUPERCONDUCTING COMPUTING STUDY CONCLUSIONS

In this section we briefly summarize the conclusions of the Superconducting Computing Study.

4.1 Previous Supercomputer Studies

Three studies examined the design of petascale and exascale computers in the past decade. These studies, described in Section 5 below, provide a very useful background for the current study of superconducting computers.

The JASONS agree that conventional CMOS technology will top off, following a long period of exponential growth, since 1970. The growth of the speed of single processors is rolling over, and highly parallel systems are required for high performance. The energy associated with logic, memory, moving bits of information, and communication limits the speed of processor chips and the size of the computer. The cost of ownership of high performance computers is increasing to serious levels, with memory chip areas measured in football fields, and predicted power consumption as high as 0.5 GW, where $1 \text{ GW} \sim \$1 \text{ G/yr}$ in energy cost.

4.2 Single Flux Quantum (SFQ) Logic and Memory

Energy-efficient Rapid Single Flux Quantum (ERSFQ) and Reciprocal Quantum Logic (RQL) logic are fast, low-energy consumption modifications of SFQ logic that promise fast, energy-efficient computation. Just as for

CMOS technology, the need to prevent unwanted thermal switching puts a lower limit on the SFQ switching energy, and it also puts a lower limit to the physical size of a Josephson Junction (JJ). Memory cells can be implemented by trapping a flux quantum with JJ devices, or through new approaches based on a Josephson junction with a magnetic tunneling layer (JMRAM).

4.3 Comparison of SFQ Logic with CMOS Technology

A comparison of the speed and switching energy of SFQ and CMOS logic has been made, including the energy needed to cool JJ technology to 4 K. We find that speed and switching energy (including cooling) are comparable for SFQ and CMOS bits. For data transmission, SFQ logic has a strong advantage that superconducting transmission lines can transfer bits with very little dissipation. The energy per bit for SFQ memory is also favorable. However, both the current size, and projected lower limits to the future size of a SFQ memory cell are quite large compared with DRAM cells. Josephson Magnetic Random Access Memory (JMRAM) cells are being developed to address this problem.

5 PREVIOUS SUPERCOMPUTER STUDIES

Three supercomputer studies in the past decade have explored the possibility of building a petascale or exascale computer using superconducting and CMOS technology. By examining the conclusions of these studies, we can get valuable information about the architecture of high performance computers and the steps that must be taken to produce a working machine.

5.1 High Technology Multithreaded Technology

The Hybrid Technology Multithreaded Technology (HTMT) study sponsored by DARPA, NASA and the NSF in 2003 proposed a hybrid, high performance computer architecture that combined a superconducting parallel processor, with optical coupling to room temperature mass memory. It's helpful to view conceptual drawings of the proposed system shown in Figure 1, because similar challenges would occur for a future high performance computer that uses only superconducting devices, or a hybrid system with both cooled and room temperature components. The overall speed required by the HTMT study requires a large number of CPU chips operating in parallel with a very large mass memory. To achieve acceptable speeds, cooled cache memory is required, located very near the CPUs.

The large surface area required to reach petascale processing rates means that the cooled superconducting CPU chips must be placed in multichip modules, and stacked close to each other to allow cooling to low temperatures, and to reduce the path length for data transfer. A conceptual design for the




Figure 1: Hybrid Technology Multithreaded Technology (HTMT) - Schematic layout of a petaflop general-purpose processor that combines superconducting Rapid Single Flux Quantum (RSFQ) multichip modules cooled to 4K (shown on right) with room temperature memory (shown at left).

4 K cooled multichip RSFQ modules from the HTMT study is shown in Figure 1.

To achieve a sufficiently large mass memory in the HTMT design, a large surface area of room temperature chips was required. At the conventional memory to CPU ratio of 1 byte/flop, a 1 petaflop computer requires ~ 1 million DRAM chips, each with 1 GB capacity, for a total surface area $\sim 100 \text{ m}^2$. The memory for a 100 petaflop computer would need a memory chip area of $10,000 \text{ m}^2$, a surface area of about 2 football fields. In addition to the large area, the memory also requires that read and write data must travel significant distances, adding delays associated with the speed of light $\sim 1 \text{ ft/ns}$. Ten meters gives a delay $> 30 \text{ ns}$. These estimates show that the locations of the processor and memory are important.

Looking at the HTMT Petaflops system layout in Figure 1, we see that the processor core, cooled to 4 K, is contained inside a liquid nitrogen cooled thermal shield at 77 K, located in the center of a circle of room tempera-


ture racks that contain the large-scale memory and network control circuits. Packing the components closely is needed to reduce the communication times, which are significant in a large system. An additional difficulty is data transfer from 4 K to room temperature. Using wires is not attractive, because metals transmit heat as well as electricity. A photonic approach is more attractive, as proposed in the HTMT study.

5.2 Superconducting Technology Assessment


[REDACTED]

[REDACTED] The study included experts from the superconducting community, including some people who also spoke at our 2011 JASON workshop. This study is interesting, because it spells out important steps that are needed to move from the production of small-scale SFQ data processors to a large-scale superconducting computer.

[REDACTED]



SFQ logic transmits a bit of information as a short voltage pulse that contains a single flux quantum. The advantage is that these pulses can be of short duration ~ 1 psec. However the transmission of a bit of information in a pulse also creates difficulties in timing. SFQ gates typically contain registers at their inputs to receive and hold incoming data. If a pulse arrives between two clock pulses, a "1" is registered, for no pulse a "0". Using this approach, one can step data through a pipelined processor. However, for a general-purpose processor, the data paths vary in length with the calculation, creating a large uncertainty in pulse arrival times. The delays can be significant. Pulses travel along a superconducting transmission line at speeds ~ 100 micron/ps, about $1/3$ of the speed of light. Traveling a distance ~ 1 cm across a chip can take ~ 100 ps. To include all of these pulses in the calculation, one would need to reduce the clock rate to < 10 GHz, undoing the potential speed advantage.



because they do not introduce the heat leaks associated with metal wires. The read and write energies of room temperature memory, as well as the associated times, are also important.

5.3 DARPA Exascale Study

DARPA sponsored the Exascale study in 2008 to consider the design of a CMOS computer capable of a 1 exaflops (10^{18} flop/sec) computations by the year 2015. The total power consumed was limited to 20 MW (\$20M/yr in energy cost), to keep the operating expenses reasonable, and the size was limited to 500 conventional server racks. Its interesting to consider the results of the Exascale study, because they lay out the problems faced in constructing a supercomputer of this power using CMOS or an alternative technology.

The record in supercomputer performance is currently the K computer in Japan, which has achieved 10.51 petaflops in LINPACK benchmark testing. Moving to an exaflop would be the next big advance. The Exascale study's conclusions raise a number of points that will be important for any future supercomputer. A summary by the chair, Peter Kogge is presented in "Next Generation Supercomputers", IEEE Spectrum (2011) [2].

The biggest challenge was energy and power. The study was not able to achieve the 20 MW goal. The first estimate reached a total operating power of 67 MW. However, a re-evaluation of the power needed to transfer data between the processors and the large RAM raised the figure to 480 MW, an excessively high operating power that would require $\approx$ \$0.5 G/yr of electricity to operate. Most of the power is associated with data transmission between the processors and the memory.

This design exercise shows us that lowering the energy required to transmit data is very important for an exascale computer.

William Dally, a member of the Exascale study, estimates that the total power of an exaflop computer could be reduced to 50 MW through a careful analysis of locality - the relative placement of processors and memory, as well as the memory hierarchy.

A relatively small memory 3.6 petabytes was used in the Exascale design, in part to lower energy costs of the total system. The byte/flop ≈ 0.004 fraction is much smaller than the 1 byte/flop rule of thumb that is conventionally used for large computers. It appears that there is no fundamental justification for the 1 byte/flop number. Nonetheless, one wonders if the proposed design would have enough memory for rapid calculations. This is likely a complex question to answer that depends on the software and the problem addressed.

Concurrency is important for the Exascale design. Because the processor chip speed has stalled (a 1.5 GHz clock is used in the Exascale design), it is necessary to operate many CPU chips in parallel. This raises important problems for software design. One would like to separate the computation into many small processes that can be done independently with local memory. However, this is often not feasible, and one must expend the energy and time needed to send data across the machine, raising the power and slowing the solution. In a very large parallel machine, a large fraction of the processors may remain idle, decreasing the overall computation rate.

Resiliency becomes important. The number of components in the proposed exascale computer is so large, that it was estimated that an error would occur every 30 to 40 minutes. This problem was brought on, in part, by the

choice to operate the system at 0.5 V, lower than conventional CMOS voltages. The low operating voltage increases the rate of thermal upsets, as well as non-thermal upsets due to cross talk and noise.

6 QUESTIONS

To address the Task Statement of the Superconducting Computing study, we asked a series of questions:

What combination of speed improvements, energy savings, hardware cost and facilities cost would be compelling?

Can SFQ logic be higher speed than CMOS?

Can SFQ logic be lower energy per operation than CMOS?

Can SFQ logic be as low-cost as CMOS?

What could the game-changers be?

Candidate next steps.

6.1 Making Comparisons

To compare the energy consumed by SFQ logic and CMOS technology, it's important to note that power consumed by a conventional CMOS web-server today is split between the CPU (~30 percent), memory (~20 percent), efficiency of the power supply (~20 percent) and other parts of the system. Moving data and communications can also consume a significant fraction of the power, as we have seen in the Exascale Study above. So the switching energy is only part of the picture.

In comparing the performance of different systems, it is also important to specify the application and the nature of the calculation being performed. An application specific integrated circuit (ASIC) can be far more power effi-

cient than a general purpose processor, because it uses constants instead of variables loaded into registers, latches instead of register files, wires instead of multiplexers, and because it has low control overhead. For example, to sum the integers from 1 to 100, an ASIC chip can consume 100X less power than a simple general purpose processor. A large general-purpose CPU can consume 1000X larger power than an ASIC processor with the same capability for a specific problem.

7 INTRODUCTION TO JOSEPHSON JUNCTION DEVICES

A tutorial introduction to Josephson junction devices is presented in Appendices C and D. In this section, we highlight the points that are needed to compare SFQ and CMOS logic.

Josephson discovered that a current I_s can flow through a tunnel barrier between two superconductors with no applied voltage, according to the relation

$$I_s = I_c \sin \gamma \quad (7-1)$$

where I_c is the critical current of the junction and γ is the phase difference between the superconducting order parameter ψ in the two superconducting layers. Josephson also predicted that a voltage V applied across the Josephson junction would create a rapid oscillation in the current

$$\frac{d\gamma}{dt} = \left(\frac{2e}{\hbar} \right) V = \left(\frac{2\pi}{\Phi_0} \right) V \quad (7-2)$$

where e is the charge of an electron and $\hbar$ is Plank's constant. A fundamental physical quantity is the magnetic flux quantum Φ_0 given by

$$\Phi_0 = \frac{h}{2e} = 2.07mV - ps \quad (7-3)$$

The value of the flux quantum follows from the requirement that the order parameter ψ be single valued in a superconductor loop that surrounds a flux quantum - the phase of ψ advances by an integer multiple of 2π . The size of a flux quantum can be expressed as the product of the amplitude of a voltage pulse $V_\phi(t)$ times it's duration Δt ; such a pulse can transport a bit of information through an SFQ circuit.

7.1 Josephson Transmission Line Element

A Josephson transmission line (JTL) element, shown in Figure 2, is the basic building unit of Rapid Single Flux Quantum (RSFQ) gates. As described below, an incoming voltage pulse $V_A(t)$ at the JTL input produces an outgoing pulse $V_B(t)$ at the output, delayed by the switching time Δt_{sw} . A ring oscillator made of JTL elements can be used to test the speed.




Figure 2: A Josephson transmission line (JTL) element, which is the basic building block of Single Flux Quantum (SFQ) logic gates. An incoming voltage pulse (left) produces an outgoing pulse that contains a single flux quantum (right), delayed by the switching time.

The inset to Figure 2 shows the actual layout for a JTL element. Note the large size of the input inductor and the shunt resistor, shown in blue, compared with the Josephson junction. Although the Josephson junctions must have a certain area to avoid thermal upsets (see below) they are typically not the largest components in the circuit.

7.2 Single Flux Quantum Data Transmission

In SFQ logic, a bit of data is transmitted down a superconducting transmission line as a voltage pulse $V_A(t)$, as shown in Figure 3. The time integral of the pulse is equal to the flux quantum:

$$\Phi_0 = \int dt V(t) = V_\Phi \Delta t \quad (7-4)$$

$$V_\Phi \Delta t = 2.07 \text{mV} \cdot \text{ps} \quad (7-5)$$

This relation fixes the product of the pulse amplitude V_Φ and the pulse width Δt .




Figure 3: Voltage pulse that contains a single flux quantum, ringing at the plasma frequency ω_p .

The voltage pulse oscillates at the plasma frequency of the Josephson junction:

$$\omega_p = \left(\frac{2\pi I_c}{\Phi_0 C} \right)^{1/2} \quad (7-6)$$

where C is the capacitance of the junction. For a typical value of the JJ critical current density $J_c = 10 \text{kA/cm}^2$ obtained today, the plasma frequency is quite high $\omega_p \sim 4 \times 10^{12} \text{rad/sec}$, and promises to increase in the future as the critical current reaches values $J_c \sim 100 \text{kA/cm}^2$.

The switching time Δt_{sw} can be found from the input and output voltage pulses of a Josephson transmission line element, as shown in Figure 4. In order to generate the output pulse, the JJ phase γ must rotate by 2π radians, giving

$$\Delta t_{sw} = 2\pi/\omega_p \propto 1/J_c^{1/2} \quad (7-7)$$

For a critical current $J_c = 10kA/cm^2$, we have $\Delta t_{sw} \sim 2ps$. For larger critical currents $J_c = 100kA/cm^2$ in the future, the switching time could be reduced to $\Delta t_{sw} \sim 0.7ps$.




Figure 4: Comparison of incoming and outgoing voltage pulses for a Josephson transmission line element - the delay determines the switching time Δt_{sw} .

7.3 Resistively Shunted Junction (RSJ) Model

The Resistively Shunted Junction (RSJ) model of a Josephson junction, illustrated in Figure 5, and described in detail in Appendix D, presents a simple way to understand the operation of SFQ logic switches. It is helpful to understand that the equation of motion for the RSJ model of a JJ is the same as that for a ball rolling down a tilted washboard.




Figure 5: Tilted washboard model of a resistively shunted Josephson junction. The potential energy is plotted $U(\gamma)$ vs. γ . The phase difference across the junction; the potential oscillates once as each flux quantum passes through.

The potential energy $U(\gamma)$ of the Josephson junction can be found by integrating the Josephson equations:

$$U(\gamma) = -E_J [\cos \gamma + (I_0/I_C)\gamma] \quad (7-8)$$

The potential energy oscillates sinusoidally with the JJ phase difference γ . Each time a flux quantum Φ_0 passes through the junction barrier, the phase γ advances by 2π . The amplitude of the oscillation in U is given by the Josephson energy E_J :

$$E_J = (\Phi_0/2\pi)I_c \sim 3 \times 10^{-20} \text{ J} \quad (7-9)$$

determined by the JJ critical current I_c . The resonant frequency of oscillation near a minimum in U is simply the plasma frequency ω_p . If one dropped the JJ state into one of the minima it would oscillate back and forth at the plasma frequency until damping pulled it to the bottom.

The tilt in the potential energy $U(\gamma)$ is due to a dc drive current I_o . The drive current tries to drive flux quanta through the junction, just as gravity tries to make a ball roll down the tilted washboard.

The total JJ energy is the sum of the potential energy $U(\gamma)$ and the kinetic energy $(1/2)CV^2$ determined by the second Josephson equation, Eq. 7-2. For a resistively shunted junction, damping provided by the shunt resistor acts to dissipate kinetic energy, and pulls the JJ state into one of the minima in U resulting in a steady state in which no flux quanta pass through the junction.

An incoming voltage pulse $V_A(t)$ creates a kinetic energy E_{in} which lifts the JJ state above the minimum. If the amplitude of the incoming pulse is large enough, the state will pass over the barrier and move 'downhill' as a flux quantum passes through the junction, as shown in the figure 5. With little damping, the state would continue to move downhill, and many flux quanta would pass through the junction. However, for SFQ logic, damping associated with the shunt resistor pulls the state into the next minimum in U so that only one flux quantum passes through.

The motion of the state of a JJ from one minimum in the potential energy $U(\gamma)$ to the next quantizes the incoming voltage pulse $V_A(t)$ into either a "1" - emit a flux quantum and move downhill to next minimum - or a "0" - no motion. It also regenerates a small incoming voltage pulse into a full-sized exiting pulse to correct for losses or unwanted reflections in the circuitry. These are analogous to the actions of CMOS buffer to digitize ranges of incoming voltages into either a "1" or a "0", and to regenerate the pulse to drive other gates.

7.4 Critical Current Density

The critical current density J_c of a Josephson junction is an important parameter that determines the JJ switching time Δt_{sw} , through the plasma

frequency ω_p . The critical current density also determines the minimum potential barrier ΔU to avoid thermal and non-thermal upsets. In turn, the barrier ΔU sets the minimum switching energy E_{sw} and the minimum junction area A as described below. The value of J_c is determined by the tunnel barrier material and its thickness.

Larger values of the critical current density J_c are better: they produce shorter switching times, because Δt_{sw} is proportional to $1/J_c^{1/2}$, and smaller area junctions, because A is proportional to $1/J_c$. The ultimate lower limit to the switching time is given by the superconducting energy gap through the Ambegaokar-Baratoff relation D-13 (see Appendix D).

Today's SFQ technology uses a barrier with critical current density $J_c \approx 10 \text{ kA/cm}^2$. In the future, advances in fabrication technology and barrier materials may give larger values $J_c \sim 100 \text{ kA/cm}^2$. Research has achieved critical current densities up to $J_c \sim 400 \text{ kA/cm}^2$ (Miller et al. 1993[4]), so it could be reasonable to consider an ultimate upper limit $J_c \sim 1000 \text{ kA/cm}^2$. The challenge to producing large critical currents is maintaining the uniformity of an insulating barrier that is only a couple of atomic layers thick.

The graph in Figure 6 (Miller et al. 1993 [4]) shows the dependence of the measured critical current density J_c vs. barrier thickness for Nb/ AlO_x /Nb Josephson junctions, the materials currently used for SFQ logic. The barrier is formed by oxidizing a thin Al film to produce AlO_x ; the greater the oxygen exposure, the thicker the AlO_x film.

The results shown in Figure 6 are fit by a planar tunneling model for $J_c < 10 \text{ kA/cm}^2$. The planar model is appropriate for high quality, planar insulating slabs, and it is based on a simple quantum mechanical tunneling. Current production uses $J_c \sim 10 \text{ kA/cm}^2$, because the tunneling process is




Figure 6: Measured critical current density J_c for an aluminum oxide Josephson junction barrier vs. barrier thickness, measured by oxygen growth time (Miller et al. 1993 [4]).

understood and high quality insulating films can be produced.

At higher critical current densities $10 \text{ kA/cm}^2 < J_c < 400 \text{ kA/cm}^2$, the slope of the curve increases. This case is explained by tunneling through a randomly placed collection of quantum point contacts, created by defects in the insulating AlO_x film. Very high critical current densities can be obtained in the lab, but the ability to achieve similarly high values of J_c in a manufacturing process with high uniformity is a challenge. Research to explore other tunneling materials, such as AlN , may help advance the progress toward larger critical current densities.

7.5 Stability against Thermal and Non-Thermal Noise

Preventing errors created by thermal or non-thermal noise is an important topic. In CMOS technology, the number of thermal and non-thermal errors is reduced to zero during the computation. This strict standard requires

a potential energy barrier that is large enough to prevent any energy, and effectively forces the supply voltage to be $V_{DD} \sim 1V$ for room temperature operation. Reducing the supply voltage V_{DD} would lower the energy required to change wires, which is proportional to V_{DD}^2 , as well as the switching energy, but the need to avoid errors sets a lower limit on the reduction.

The thermal switching rate P for a Josephson junction is

$$P = (\omega_p/2\pi) \exp(-\Delta U/k_B T) \quad (7-10)$$

where the attempt rate is set by the plasma frequency ω_p and ΔU is the potential energy barrier to switching, shown in Figure 5. Following the standard CMOS approach, we adapt a no-errors rule that requires a SFQ chip with a million JJs to operate for 30 yrs (10^9 sec) without making a thermal error. This criterion requires $\Delta U > 60k_B T$. The thermal barrier ΔU is determined by the Josephson energy E_J and the tilt of the washboard potential, set by Equation D-10 the bias current I_o ; an expression is given in Appendix D. For a typical bias current $I_o = 0.7I_c$ the thermal barrier is

$$\Delta U = E_J/3 > 60k_B T \quad (7-11)$$

This ‘no errors’ criterion requires that $E_J > 180k_B T$ and sets lower limits to the SFQ switching energy E_{sw} and junction area A , discussed below. By decreasing the barrier height, one can reduce E_{sw} and A in proportion. However the thermal switching rate increases exponentially as the barrier decreases, limiting the range of reduction.

It is interesting to consider the effects of reducing the energy barrier ten times to $\Delta U = 6k_B T$. The error rate per JJ then increases from $P \sim 10^{-15} s^{-1}$ by a factor 3×10^{10} to $P \sim 3 \times 10^{-5} s^{-1}$. At first, this might seem to be unacceptably high. However, for some applications this error rate would

be acceptable. Suppose one is examining a lawn of green grass looking for one white blade, eg. a single “1” in 100 million “0”s. The vast majority of errors would occur on “0”s where they have no effect. And if the white blade of grass is detected, that can be tested by looking at it again. An area where errors should not occur is in the control circuitry of the computer; here the barrier must be kept high.

Non-thermal switching will be introduced in a superconducting computer by crosstalk between different lines and from unwanted reflections caused by changes in line impedance and less-than-perfect terminations. These factors could be important in a complex superconducting chip with many metal layers. The thermal barrier ΔU must also be large enough to protect against non-thermal switching. Tests could be carried out to determine the correct value of the barrier ΔU by deliberately injecting artificial noise into a superconducting computer circuit.

7.6 Lower Limits to the SFQ Switching Energy and Junction Area

The minimum potential energy barrier ΔU that is required to avoid thermal errors also sets lower limits to both the SFQ switching energy E_{sw} and the junction area A .

The SFQ switching energy $E_{sw} = 2\pi E_J$ is proportional to the Josephson energy E_J . Using the no-errors criterion we find $E_J = 3\Delta U > 180k_B T$, which directly determines the lowest switching energy that can be used without inducing thermal errors at $T = 4\text{K}$:

$$E_{sw} > 6 \times 10^{-20} \text{ J} \quad (7-12)$$

No other parameters enter. This means that the barrier ΔU against thermal errors determines a fundamental lower limit to the switching energy E_{sw} that is independent of both the critical current density J_c and the junction area A . This lower limit to the barrier ΔU for SFQ logic is analogous to the lower limit on supply voltage $V_{DD} \sim 1\text{ V}$ for CMOS logic. To avoid unwanted switching caused by non-thermal noise, a somewhat larger barrier ΔU and switching energy E_{sw} may be needed.

The tunnel barrier ΔU also determines the minimum critical current I_c and the minimum junction area A that can be used for error-free operation. In addition to being proportional to the barrier ΔU , the switching energy E_{sw} is given by

$$E_{sw} = \Phi_0 I_c \quad (7-13)$$

where the numerical value follows from the no-errors criteria above. The flux quantum Φ_0 is a fundamental constant, so this sets a lower limit to the JJ critical current

$$I_c = J_c A \quad (7-14)$$

Using the strict criterion $\Delta U > 60k_B T$, we find $I_c > 30\text{ }\mu\text{A}$. However, if we use a barrier $\Delta U > 6k_B T$ of one tenth the height, the critical current drops in proportion to $I_c > 3\text{ }\mu\text{A}$.

Using Eq.7-14 for the minimum critical current I_c , we can find the minimum junction area A and size $A^{1/2}$ that can be used to avoid thermal errors for the critical current densities $J_c \sim 10\text{ kA/cm}^2$ used today:

$$A > 0.3\text{ }\mu\text{m}^2, \quad A^{1/2} > 0.5\text{ }\mu\text{m} \quad (7-15)$$

In the future, for higher values $J_c \sim 100\text{ kA/cm}^2$, the area and size could be reduced to:

$$A > 0.03\text{ }\mu\text{m}^2, \quad A^{1/2} > 0.2\text{ }\mu\text{m} \quad (7-16)$$

The junction area $A > 0.3 \mu\text{m}^2$ and size $A^{1/2} > 0.5 \mu\text{m}$ for today's SFQ technology are large compared with 28 nm CMOS devices currently in production.

However if we allow a larger error rate associated with a lower barrier $\Delta U > 6k_B T$ and a future value of the critical current density $J_c = 100 \text{ kA}/\text{cm}^2$, we find a junction area $A > 0.003 \mu\text{m}^2$ and size $A > 50 \text{ nm}$ that approach 28 nm CMOS devices.

Comparing the area of SFQ and CMOS devices, it seems clear that SFQ technology will always require a greater chip area. The lower limits for JJ area A are larger than the current size of today's CMOS FETs, and the difference will increase in the future, as CMOS moves to smaller size scales. However, JJ chips can be made with very thin substrates, because the power dissipation of JJ processor chips is orders of magnitude less than for CMOS. These factors could greatly reduce the volume required for a processor by stacking many thin chips, as discussed for superconducting memory below.

7.7 Energy-efficient RSFQ (ERSFQ) Logic and Reciprocal Quantum Logic (RQL)

Rapid Single Flux Quantum (RSFQ) logic, proposed by Likharev and Semenov (1991) [1] has the advantages of SFQ logic described above, but the DC bias current I_o consume power at all times, in addition to the switching energy. Two new adaptations of SFQ logic consume much less energy: Energy-efficient Rapid Single Flux Quantum (ERSFQ) logic, and Reciprocal Quantum Logic (RQL).

The circuit diagram for a Josephson Transmission Line (JTL) element for Energy-efficient RSFQ logic is shown in Figure 7. The resistor that supplies DC current bias for RSFQ logic has been replaced by a Josephson junction with a series inductor. The bias JJ is designed so the bias current I_o is set by its critical current. This approach eliminates power dissipation by the bias resistor, and improves the energy efficiency.

RQL uses a different approach illustrated in the JTL element and AND/OR gate circuit diagrams shown in Figure 8. AC power oscillations synchronized with the clock rate are applied to the JTL element and the AND/OR gate through a transformer, so no dc current can flow. In addition, inductive coupling is used to carry out logic operations in the gate, a new approach.

RQL has several advantages compared with EFSFQ logic: The clocked supply provides an overall synchronization of data flow through the computer and helps avoid the timing problems for clock pulses that are caused by different path lengths. Also, by phase shifting the clock at different locations, a directional flow for data can be established in an RQL circuit, like a 3-phase induction motor that spins in only one direction. This approach achieves directional data flow in a circuit composed of only two-terminal devices, a fundamental achievement. Finally RQL gates currently have smaller size, and lower switching energy.

7.8 Cooled Memory

Cooled memory is essential for an SFQ computer, if it is to carry out calculations at high speed with low energy consumption. The mV size of SFQ voltage pulses is not sufficient to switch room-temperature DRAM memory cells, and the energy required is too great. In addition, data transfer between




Figure 7: Energy-efficient Rapid Single Flux Quantum logic Josephson transmission line element.




Figure 8: Reciprocal Quantum Logic (RQL) (left) Josephson transmission line element with clocked inductive power coupling (right) AND/OR logic circuit; first input to OR, second input to AND.

the cooled processor and room-temperature memory poses its own problems through latency and the increased heat leak.

A Superconducting Quantum Interference Device (SQUID) can be used as a memory cell that stores one bit of information in the form of a flux quantum. SQUIDs have the advantages of high speed and low read- and write-energies E_{read} and E_{write} . These devices were created in the 1960's (Ruggiero and Rudman 1990 [5]), and they come in two forms: a dc SQUID composed of two Josephson junctions, and an ac SQUID with one JJ and an inductor. Figure 9 shows the circuit diagram of an ERSFQ D Flip-Flop memory gate based on a dc SQUID. The writing time is $t_{write} \sim 6$ ps and the read and write energies are $E_{read} \sim E_{write} \sim 0.7$ aJ. These are competitive numbers compared with DRAM cells.




Figure 9: Energy-efficient Rapid Flux Quantum (ERSFQ) D flip flop memory gate. The red ring shows a dc SQUID.

However, the inductor needed to trap a flux quantum inside a SQUID memory cell is physically quite large, and it places a lower limit on the area.

To trap a flux quantum, the inductance must be

$$L > 5\left(\frac{\Phi_0}{2\pi I_c}\right) \sim 20pH \quad (7-17)$$

for $I_c = 100 \mu A$. The inductance of a loop of wire with radius r and width w is:

$$L = \mu_0 r \ln(8r/w) \quad (7-18)$$

giving a radius $r > 3\mu m$. The large physical size of the inductor $\sim 3 \times 3 \mu m^2$ leads to memory cells that are currently $\sim 10 \times 20 \mu m^2$ for RQL technology, and larger for ERSFQ. The size, of SQUID-based memory cells is much greater than for DRAM technology, which currently produces memory cells of size $0.1 \times 0.1 \mu m^2$ and density $\sim 10 \text{ Gbit/cm}^2$

In the future, the inductor size for memory cells can be reduced by using a multi-layer inductor in the form of a coil, or by using the linear part of a Josephson junction inductor near the bottom of a minimum in the potential energy $U(x)$. Using these approaches, the ultimate minimum size of an SFQ memory cell could be $\sim 1 \times 1 \mu m^2$ giving a density $\sim 100 \text{ Mbit/cm}^2$ that is larger, but still a factor of 100 below DRAM today.

The area required by SFQ memory appears to make it impractical for large-scale storage. As we have seen from the Exascale study, the large chip area of the CMOS memory required poses serious challenges for the physical size of the computing system, limits processor-memory speed, and increases the total energy consumption.

SFQ memory is naturally suited for use as cache memory, located on the same chip as an SFQ processor, or nearby in a multichip module. In these locations, a small $\sim 100 \text{ kB}$ SFQ memory could provide fast, low energy storage.

It is interesting to consider a 1 petaflop superconducting computer with a fully superconducting memory system. As for the HTMT study referred to above, we'll use the conventional ratio 1 byte/flop of memory capacity to CPU speed, giving a total memory capacity of 1 petabyte. This memory would require one million DRAM chips, each with capacity 1 GB, having a total area $A \sim 100\text{m}^2$. For a future superconducting memory with density $1\text{bit}/\mu\text{m}^2$, the total area for 1 petabyte of JJ memory would be quite large, $A \sim 10,000\text{ m}^2$, posing a challenge to the designers.

Because SFQ electronics does not generate large quantities of heat, one could reduce the volume occupied a JJ memory system considerably by thinning the substrates and stacking the memory chips, with spacing d . It is possible to do this from a thermal point of view, because the heat production for superconducting memory is orders of magnitude lower than for CMOS, because the superconducting transmission lines to memory elements do not consume a significant amount of energy, and because the heat conduction of a metal substrate remains high at low temperatures. The total volume required is then $V \sim Ad$, where A is the total memory area. For a 1 petabyte JJ memory with area $A \sim 10,000\text{m}^2$ and $d \sim 1\mu\text{m}$, we find $V \sim 10\text{m}^3$, a relatively modest volume.

It is possible to cool very large volumes to liquid helium temperatures using current technology. An example is found in the CERN Large Hadron Collider particle accelerator has a circumference of 27 km. The accelerator controls the beam by using a series of 1600 superconducting solenoids cooled to 1.9K. Using similar refrigerator technology, it could be possible to cool a fully superconducting computer system, including processors and memory, to liquid helium temperatures.

To address current problems with JJ memory, new approaches are being

developed that are based on room-temperature Magnetoresistive Random Access Memory (MRAM) technology. MRAM has been developed to provide dense, non-volatile digital storage, as illustrated in Figure 10. For tunneling-based MRAM, electrons can tunnel from a fixed magnetic layer on the bottom through a barrier to a free magnetic layer on the top. A bit of information is stored by the direction of magnetization in the free layer, which is set by passing currents through perpendicular write wires to generate a local peak in magnetic field intensity. The stored bit is read out electrically by measuring the tunnel current. Although MRAM does not compete in density with DRAM, it is attractive for certain applications such as radiation-proof storage on satellites.




Figure 10: Josephson Magnetic RAM (JMRAM) approaches that use a Josephson junction technology to read out a magnetic multilayer, originally developed for Magnetoresistance Random Access Memory (MRAM) at room temperature.

The proposal is to join MRAM and JJ technology by using an MRAM magnetic tunnel layer as the barrier inside a Josephson junction, as shown on the right in Figure 10. Variations on this approach are being pursued at HYPRES, IBM, and Northrup Grumman. The magnetization of the free

layer is written by currents in two perpendicular wires as for conventional MRAM cells, but the readout is done by Josephson junctions. The potential advantages of this approach are compatibility with SFQ logic with rapid low-energy reads, and relatively small cell areas. To write a bit, current pulses are sent through two perpendicular wires that span the memory array, generating magnetic fields that circle both wires, so the write energy will be somewhat larger. This approach looks promising, and it should be explored as a research topic to evaluate its potential.

8 COMPARISON OF SFQ AND CMOS LOGIC

CMOS technology has enjoyed an impressively long period of exponential growth in computing power, through considerable investments of time and money, and it is continuing to progress toward smaller size scales and greater densities. So it is a formidable opponent.

Table 1 presents the energy required for SFQ and CMOS logic, data transfer, and memory. The first column is for 28nm CMOS technology, which is in production today, and the second column for future 7nm CMOS technology taken from the International Technology Roadmap for Semiconductors (ITRS) projections. The third column is for an SFQ switch operating at 4 K, and the fourth column is for an SFQ switch including the energy required to cool it to 4 K, estimated using the factor 300W/W from the efficiency of a commercial cooler.

First, comparing the switching energies for an isolated switch, we find that CMOS and SFQ logic are similar. For SFQ logic we use a switching energy $E_{sw} \sim 0.1aJ$ that is close to the fundamental minimum $E_{sw} > 0.06aJ$ needed to avoid thermal switching at 4 K, discussed above. This number is much smaller than those for CMOS, but adding the cooling energy raises it to $E_{sw} \sim 30aJ$, comparable to the switching energy $E_{sw} \sim 100aJ$ for a 28nm CMOS switch today and the projected value $E_{sw} \sim 18aJ$ for 7 nm CMOS in the future.

To make the energy comparison simple, we used single devices in configurations like a ring oscillator. In a densely packed chip, the generation of heat becomes a major factor for CMOS technology. The allowable switching energy and device density are limited by the chip's ability to carry away heat.

Table 1: Comparison of energies for CMOS and Josephson Junction (JJ) technology for specific actions in a logic circuit; 28nm CMOS is currently available, and the numbers for 7nm CMOS are predicted by the ITRS roadmap. The JJ numbers are shown at low temperatures and including cooling, assuming a refrigerator with 300 W/W reciprocal efficiency.

item	28nm CMOS	7nm CMOS	JJ (4K)	JJ (with cooling)
Switch	100aJ	18 aJ	0.1 aJ [†]	30 aJ [†]
Gate	200 aJ	35 aJ	0.4aJ [†]	120 aJ [†]
1mm Wire	100 fJ	70 fJ	< 1 fJ	< 1 fJ
32b 1mm Bus	3.2 pJ	3.2 pJ	< 1 pJ	< 1 pJ
Read/Write memory bit cell	200 aJ	35 aJ	1 aJ*	300 aJ*
Read 32 bit, 8kB RAM	5000 fJ	875 fJ	0.16 fJ**	50 fJ**
Write 32 bit, 8kB RAM	5000 fJ	875 fJ	5 fJ**	1500 fJ**

[†] $E_{sw} > 0.06$ aJ to avoid thermal switching

* ERSFQ JJ memory cell

** JMRRAM estimate

The heat generation and switching time that one can achieve on a chip can differ, depending on the chip architecture and the nature of the problems it is meant to solve. Superconducting computers have an advantage in that the power dissipated at low temperatures is 1000X lower than CMOS and the total heat dissipated by the chip is far less. As noted above, these advantages could allow one to stack thin superconducting processor chips to reduce the volume occupied.

Comparing the energy required to transfer data, SFQ logic has a strong advantage - the energy required to transmit a bit of information down a superconducting transmission line is very low, essentially zero. Because moving data inside a high performance computer is responsible for a large fraction of the power consumption, as discussed above for the Exascale study, this advantage could be quite important. CMOS logic operates by switching the

output voltage between two levels that correspond to “1” and “0”. Moving a bit from “0” to “1” swings one end of a wire high, and the voltage step runs along the line, charging it up. The resulting charging energy CV^2 is a significant fraction of the power consumption of CMOS chips, and an even larger fraction for a high-performance computer. By contrast, SFQ logic transmit bits of data in pulses. A “1” is encoded by the transmission of a short voltage pulse that contains one flux quantum down a superconducting transmission line. The energy associated with the pulse enters one end of the line and leaves from the other, so very little energy is dissipated, dropping the number to values well below CMOS technology.

For memory, we first compare the read and write energies for CMOS and Energy-efficient Rapid Single Flux Quantum (ERSFQ) memory cells. The ERSFQ numbers including cooling energy are comparable to 28 nm CMOS, but their area is much greater, as discussed above. We then compare the read energies for CMOS and estimates of Josephson Magnetic Random Access Memory (JMRAM), and find a strong advantage for the cooled memory. Much of the energy consumed by CMOS memory is associated with control electronics and data transmission across the chip by changing datalines. These numbers are lower using superconductors. The write energy for JMRAM memory is large compared with the read energy, due to the currents that are needed to magnetically write a bit to a cell.

The energy required by Dynamic Random Access Memory (DRAM) chips decreases with minimum feature size, as shown in Figure 11 (Vogelsang 2010 [6]). DRAM technology currently produces chips with 4 Gbit/chip and cell densities $\sim 10^{10}$ bits/cm² that have read and write energies ~ 10 pJ/bit. The DRAM chip capacity will increase and the energies will drop as the feature size drops from 40 nm today to below 20 nm in the future, as shown.




Figure 11: Dynamic Random Access Memory (DRAM) energy consumption vs. minimum feature size for the recent past, and projected into the future (Vogelsang 2010) [6].

8.1 What are Potential Game Changers for Superconducting Computing?

A key ingredient for energy comparison is the energy needed to cool superconducting electronics to 4 K. Commercial coolers currently have efficiencies ranging from 220 W/W to 1000 W/W. However, the Carnot efficiency between 4 K and 300 K leads to a factor that is only 75 W/W. If a more efficient cooler could be developed with performance ~ 100 W/W, it would drop the SFQ switching energies in the comparison Table 1 by a factor of 3X.

It's attractive to think about using high T_c superconductors to construct SFQ logic that operates at 77K to lower the cost and provide more efficient cooling. However a number of problems occur. To protect against thermal upsets at 77K, a proportionately higher potential barrier ΔU and critical current I_c are required. That means the switching energy will go up by the same factor, undoing the supposed energy advantage. In addition, the higher critical current I_c requires a shunt resistor with very low inductance to avoid trapping a flux quantum, using the analysis above. Numerically, the required inductance is so small, that it cannot be made using the regular fabrication process (Likharev and Semenov 1991 [1]).

If a method is found to make gigabytes of memory that operates at low temperatures with read and write energies ~ 1 pJ/bit, this would undo one of the biggest barriers to superconducting computing. The size must also be small enough to pack the memory near the processor chips. The JMRAM approach currently being explored may lead to small cells with small switching energies. A large body of research on coupled magnetic layers at room temperature creates a range of magnetic alternatives that could be tested. At UC Berkeley, Van Duzer and co-workers have tested cooled hybrid JJ-DRAM memories (Yoshikawa et al. 2005 [7]), but did not achieve low-energy operation.

A potential application that is attractive, is CMOS logic cooled to a liquid nitrogen temperature 77K paired with high- T_c superconducting transmission lines to move data without energy dissipation. Cooling to 77K is relatively cheap and easy. Cooled CMOS will have lower switching energy than the room temperature version, and low-temperature operation may allow switching techniques that would not work at 300K.

If one is willing to use the same degree of flexibility in design and fabrication need for JJ fabrication, there may be attractive opportunities using cooled CMOS.

8.2 Design Challenge

To evaluate the ability of Single Flux Quantum (SFQ) logic to compete with CMOS technology, we suggest a challenge: design a 32-bit processor similar to the Cortex A-9 that operates with high throughput and low latency. Using CMOS as a target, achieve throughputs in SPECInt instructions at rates > 3 billion/sec and pointer traversals at rates > 1 billion/sec. Designing a general purpose processor such as the Cortex A9 would demonstrate the ability of SFQ logic to make processors for standard software such as Linux, and show that SFQ logic can overcome the barriers associated with making a capable general purpose chip.

It is important to note that the Cortex A9 has the computing power of a smart cell phone, so satisfying this challenge would not lead to greater computing power. However the Cortex A9 is an appropriate problem, because it has a well understood architecture that it is designed to be networked into parallel systems.

Also, the challenge is to design, but not to fabricate the chip, because the tools needed to achieve that goal would require substantial work and investment. It is better to explore the feasibility first.

These matters are addressed in Appendix B, which discussed the architectural challenges for superconducting computing.

9 FINDINGS AND RECOMMENDATIONS

The Findings and Recommendations for this report are presented in the Executive Summary.

A APPENDIX: WORKSHOP AGENDA

(b)(3)-P.L. 86-36

~~CONTROLLED~~ UNCLASSIFIED INFORMATION**JASON 2011 Summer Study
Super-C*****Tentative* AGENDA
June 13-15***- All briefs are unclassified -***Monday, June 13**

Time	Title	Speaker	Affiliation
1300-1330	Superconducting Computing – Introduction	Marc Manheimer	NSA
1330-1430	Rapid Single Flux Quantum (RSFQ) Logic	Anna Herr	Northrop Grumman Electronic Systems
1430-1445	Break		
1445-1615	Superconducting Technology Assessment (STA), HTMT	John Spargo	Northrop Grumman Aerospace Systems
1600-1700	Discussion		

Tuesday, June 14

Time	Title	Speaker	Affiliation
0900-1030	Energy-Efficient SFQ Logic & Memory	Oleg Mukhanov	HYPRES
1030-1130	Reciprocal Quantum Logic (RQL)	Quentin Herr	Northrop Grumman Electronic Systems
1130-1145	Break		
1145-1245	Josephson Magnetic Random Access Memory (J-MRAM)	Anna Herr	Northrop Grumman Electronic Systems
1245-1345	Lunch		
1345-1445	Orthogonal Spin-Transfer MRAM	Thomas Ohki/Andy Kent	Raytheon-BBN/NYU
1445-1515	Break		
1515-1615	System Integration	Oleg Mukhanov/ Deep Gupta	HYPRES
1615-1700	Discussion		

~~CONTROLLED~~ UNCLASSIFIED INFORMATION

Wednesday, June 15

Time	Title	Speaker	Affiliation
0900-1000	Architecture Lessons from HTMT	Loring Craymer	USC
1000-1130	Processor Design and insights	Mikhael Dorojevets	Stony Brook University
1130-1145	Break		
1145-1245	Architecture – Advantages and Opportunities	Andres Marques	Pacific Northwest National Laboratory
1245-1345	Lunch		
1345-1445	Fabrication of Superconducting Electronics	Alan Kleinsasser	NASA-JPL
1445-1500	Break		
1500-1600	IBM Vision for Superconducting Computing and Current Work	Bob Wisnieff/ Gerald Gibson/ Mark Ketchen	IBM
1600-1700	Additional questions and Wrap-up		

~~CONTROLLED UNCLASSIFIED INFORMATION~~

B APPENDIX: ARCHITECTURAL CHALLENGES

This Appendix presents an overview of the Architectural Challenges that face the development of a superconducting general processor. It assumes some familiarity with terminology used in the computer architecture community.

JASONS understanding of the current state of the art in Josephson junction circuits is composed of custom designed circuits such as 8-bit ripple carry adders and some small-scale special purpose circuits. We understand that there are some cell libraries, but it is not clear how complete these libraries are and whether they are amenable to use with standard tools.

It is clear that custom designed circuits are no longer possible at the scale that has been suggested. Even a modest microprocessor requires a modern automated CAD flow and will use a licensed processor design in the form of Verilog or VHDL (Very High Speed Integrated Circuit Hardware Description Language).

Building a complete microprocessor is a significantly more complex task than has been attempted, and adds significant challenges due to the large number of interacting components. Due to the requirements for timing the flux pulses, existing design tools cannot be used unmodified. Moving beyond the existing state of the art will require adopting a fully automated CAD flow.

In order to move forward, more complete cell libraries will need to be developed complete with physical, pin, and timing models that can be used

as part of the CAD flow. An additional challenge will be developing routing and placement tools that can be used given the requirements of JJ logic. Existing routing and placement tools that were developed for CMOS are unlikely to be adaptable given the unique properties of JJ logic.

B.1 Memory

Creating a memory of sufficient size and speed to match the performance of the JJ processor will be a difficult task. We were presented with several technologies, but all of them were significantly slower than the cycle time of the JJ processor.

The memory cell size of the JJ-based register memory is approximately 1 micron, so the memory density using this technology will be $\sim 100\times$ less dense than comparable CMOS. Creating L1 and L2 caches that are comparable in size to existing CMOS processors will require significant advances. If the JJ processors turn out to be as fast as claimed, then the sizes of these cache memories will need to be larger than for CMOS processors in order to attain the necessary cache hit rates.

The main memory of the system is an even greater challenge. JASON was presented with three different technologies, all based on magnetic spin. All of these are asymmetric in terms of write speed, so that writes are much slower than reads, and reads are slower than DRAM. The greater the performance gap between the caches and main memory, the larger the caches must be in order to mask this gap.

The asymmetry of the proposed memory technologies needs to be quantified and understood. Depending on the workload, it will have enormous

impact on the size of the caches, the design of the memory subsystem and on the overall performance of the system. It will also impact the classes of algorithms that will run well on the computer system.

Slow writes can be mitigated using interleaving, coupled with a suitably sized cache. This affects the balance of the system, and a memory that is sufficiently fast to service all of the write requests without stalling the processor may provide more read bandwidth than can be utilized.

B.2 Processors

JASON feels that a suitable challenge for moving beyond special-purpose designs is producing a full working general-purpose microprocessor. A suitable example is an embedded processor such as the ARM Cortex M0, the simplest ARM microcontroller. It requires only about 12k gates, but will reveal some of the issues in going from a Digital Signal Processor (DSP) or special purpose circuit to a full microprocessor.

Producing such a microprocessor will require substantial progress in development of the cell library as well as the automated CAD flow.

The goal of this first step is to demonstrate the ability to deal with the integration issues that are present in producing a complete microprocessor.

The measure of success would be to develop a working ARM Cortex M0 processor with 16kB of instruction memory and 16kB of data memory. It should be possible to run some algorithm of interest, say a hashing algorithm on this processor and to compare its performance to the same processor done in CMOS. Comparison would be made in terms of both speed and power consumption.

In order to build a useful general purpose computer system, a microprocessor that is capable of running a modern operating system, tools and application programs is necessary. The next level of challenge is the production of such a processor. JASON recommends the ARM Cortex A9, which has all of the modern features as well as excellent operating system support including a complete tool chain (even Windows is being ported to the ARM architecture). This processor has approximately 1M gates, and will require advances in process and may also require the development of scalable packaging technology to allow multichip modules to be built. The A9 design comes with a 16-64kB L1 cache, which is larger than existing JJ register memories. It also comes with an L2 cache controller. Depending on the developments in main memory technology, the L2 cache is likely going to need to be on the order of 4MB in order to have adequate hit rates and avoid degrading performance when going to the much slower magnetic memory.

The ARM Cortex A9 processor is equivalent to the processor used in modern smart phones. Producing such a processor will require a fully automated CAD flow, including routing and placement tools. At this scale it will no longer be reasonable to rely on significant human involvement in these issues.

Given the expected density of JJ circuits, we expect that a microprocessor such as the Cortex A9 will require multiple chips. As a result, an additional challenge will be the development of scalable packaging technologies.

B.3 Computer System

Building a microprocessor is not sufficient to declare success. A complete

computer is a complex system of interacting components. Some such as the memories are relatively tightly integrated with the processor, while others such as peripherals will operate asynchronously. Many other important issues will come up in integrating the microprocessor with the memory and even simple peripherals.

The ultimate goal of this phase should be to boot Linux on a system based on the Cortex A9 processor with 1 GB of main memory, and with adequate cache memory so that for most workload the processor is not waiting on main memory. The system needs to be balanced in order for the system to have good performance on representative workloads.

B.4 Shared Memory Multiprocessor

The next phase in moving towards a high performance computing system will be a small-scale shared memory multiprocessor. JASON recommends a 16-way Distributed Shared Memory (DSM) processor with directory-based cache coherency. The additional complexity in this phase will be in the cache coherency logic to implement directory-based cache coherency and in the interconnection network. JASON recommends a crossbar or a flattened butterfly switch. The complexity of the interconnect is expected to be intermediate between the Cortex M0 and the Cortex A9 processor. This will also exercise the scalable packaging technology, since it is certain that using existing and projected process technology this will require multiple chips.

Intermodule communication issues will dominate this phase of moving towards a high performance computer system.

B.5 Milestones

1. Develop a sufficiently complete cell library to enable processor, interconnect and memory development.
2. Develop an automated CAD flow.
 - (a) Routing and placement tools will have to be developed to deal with JJ delay dependencies.
3. Gigabit scale main memory (in parallel with processor development).
4. Simple RISC microcontroller processor (ARM Cortex M0, 12k gates)
5. Modern RISC processor (ARM A9, 1M gates, 1664kB L1 cache)
 - (a) Memory technology advance to support L1 cache and operating memory
6. Scalable packaging to enable multichip modules
7. Functional computer system including processor, memory, read only memory (ROM)
 - (a) Should be able to boot Linux and run programs compiled using open source tools
8. 16-way distributed shared memory (DSM)
 - (a) Crossbar or flattened butterfly interconnect
 - (b) Directory-based cache coherency
 - (c) 1GB/per processor

9. Cluster in the dewar, connecting multiple DSM modules using appropriate interconnect such as Dragonfly switch
 - (a) Interconnect technology that is less dependent on precise timing will be required.
10. Cross dewar communication
 - (a) The key issue will be getting bits out of the cryostat.

C APPENDIX: BASIC JOSEPHSON JUNCTION TUTORIAL

This is a simple discussion of Josephson junctions as they relate to superconducting *classical* logic for lower power dissipation and possible route to larger scale supercomputing.

C.1 The Resistively-Shunted Josephson Junction

The basic device used in this technology is a Josephson junction (JJ), formed by a thin tunneling oxide between two superconductors. For the logic applications, the device is shunted by a parallel resistor (except for possibly when the tunnel barrier is very thin, i.e. the current density is very high, and there may be enough single-electron tunneling to make the devices self-shunting). The effective circuit is therefore a parallel combination of the shunt resistor, a parasitic capacitance arising from the junction and the wiring, and an element which passes a non-dissipative current, according to the current-phase relation of the JJ,

$$I(t) = I_C \sin(\delta(t)), \quad (\text{C-1})$$

where $\delta(t) = \frac{2e}{\hbar} \int V(t) dt$ is the phase difference across the junction. The important combination of constants that keeps appearing is the magnetic flux quantum, $\hbar/2e = \Phi_0 = 2 \times 10^{-15} \text{ Wb} = 20 \text{ Gauss} - (\mu m)^2$, which can also be expressed as $\Phi_0 = 2 \text{ mV} - \text{ps} = 2 \text{ mA} - \text{pH}$. There are innumerable opportunities to mess up factors of 2π in the following.

The Josephson tunneling acts as a nonlinear inductor. This is seen by taking the derivative,

$$\frac{dI}{dt} = I_C \dot{\delta} \cos(\delta(t)) = \frac{2eV I_C}{\hbar} \cos(\delta(t)), \quad (\text{C-2})$$

and comparing with the constitutive relation for an inductor, $V = LdI/dt$, we identify the nonlinear Josephson inductance

$$L_J = \frac{\hbar}{2eI_C} \frac{1}{\cos(\delta)}. \quad (\text{C-3})$$

The parallel combination of this Josephson inductor with the capacitance and resistance gives a parallel L-R-C resonator. The characteristic resonance frequency is called the plasma frequency,

$$\omega_P = \frac{1}{\sqrt{L_J C}} = \sqrt{\frac{2eI_C}{\hbar C}} = \sqrt{\frac{2eJ_C}{\hbar \zeta}}. \quad (\text{C-4})$$

Since both the critical current I_C and the capacitance scale with the junction area, this frequency depends only on the current density (i.e. the barrier thickness); for typical Nb/AlOx/Nb junctions at a current density of $J_C \sim 10,000 \text{ A/cm}^2 = 100 \text{ } \mu\text{A}/\mu\text{m}^2$, and the junction specific capacitance of $\zeta \sim 80 \text{ fF}/\mu\text{m}^2$, this plasma frequency is about 300 GHz. For digital applications with single flux pulses, we need to arrange that this resonance is overdamped, $Q < 1$, which means that the I-V curves are non-hysteretic, and the McCumber parameter

$$\beta_C = \frac{2e}{\hbar} I_C R^2 C < 1 \quad (\text{C-5})$$

This parameter is the square of the resonator Q , and means that the shunt resistor and the overall impedance of the device is

$$Z_C = \sqrt{\frac{L_J}{C}} = \sqrt{\frac{\hbar}{2eI_C C}}, \quad (\text{C-6})$$

which scales inversely with the junction area and as the inverse square root of the current density. For a $1 \mu\text{m}^2$ junction that seems typical in use today, it has a critical current of $100 \mu\text{A}$, the critical current and Josephson inductance are $100 \mu\text{A}$ and 3 pH respectively, and the characteristic impedance is 6Ω . Of course, since the system is overdamped, the characteristic timescale can be slower than the plasma frequency, $\omega_c = L/R = 2eI_C R/\hbar$, but this is on the same order. More on scaling of this later.

C.2 A Flux Memory Element: JJ in a Superconducting Loop

All of the building blocks of memory in this technology consist of one or more junctions in a superconducting loop that contributes some ordinary (i.e. linear) inductance, L . (For historical reasons this is known as an RF-SQUID configuration, while a DC-SQUID consists of two junctions in a loop, regardless of how they are used. As another aside, this is also the same circuit, in a different parameter regime, as the "phase" qubit introduced by John Martinis and colleagues for quantum computing.) The flux in this all superconducting circuit must be quantized in units of Φ_0 , in order to maintain the single-valuedness of the order parameter

$$\delta - 2\pi \frac{LI_L}{\Phi_0} = 2\pi n \quad (\text{C-7})$$

where n is an integer, and I_L is the current flowing in the inductor. If one now adds a bias current, the current will divide between the two inductors (the real inductor and the junction), except that, since the junction cannot exceed a certain maximum current, the system is hysteretic - a memory! This happens when the loop inductance is somewhat bigger than the junction's linear inductance, i.e. when the parameter

$$\beta_L = \frac{2e}{\hbar} I_C L = 2\pi I_C L / \Phi_0 \quad (\text{C-8})$$

is greater than one, and it should be typically five or more. This requires a somewhat bulky inductor. Remembering that the inductance of a simple loop is $L = \mu_0 r \ln(8r/w)$ where r is the radius of the loop and w is the width of the wire, and that μ_0 in convenient units is about 1 picohenry per micron, we see that the loop should be five or more microns diameter. This implies that the size of a memory cell is at least 30 square microns. It seems that

current designs have not even pushed on this limit, and the greatest density we heard about was only a few kilobits per square millimeter - not very dense memory. With a multilayer process, one could try to make a smaller, multi-turn inductor. It can be made much smaller using more Josephson junctions, more on this later.

Now if one ramps the current up, eventually the critical current is exceeded, and the phase difference across the junction jumps by 2π , storing a circulating current in the loop, even if the current is then reduced back to zero. Ramping the current in the opposite direction can reset the cell, storing a circulating current in the opposite sense. During the transition between states, the junction phase advances by 2π , and a voltage pulse is developed and dissipated in the shunt resistor. There is no voltage across the junction during the quiescent state, although the conventional RSFQ dissipates a lot of power in the biasing circuitry, if it uses resistors. How much energy was stored in the bit, or used in the transition? The circulating current is on the order of the junction critical current, and so the energy is simply

$$E_{\text{bit}} = (L_J + L)I_C^2 = \frac{L + L_J}{L_J} \frac{\hbar}{2eI_C} I_C^2 = \frac{L + L_J}{L_J} \frac{I_C \Phi_0}{2\pi} \sim I_C \Phi_0. \quad (\text{C-9})$$

In this sense, the superconducting logic is the electromagnetic "dual" of conventional CMOS, where bits are stored as charge on a capacitor. This bit energy is much smaller than in CMOS, for the example of 100 μA critical current, it is only 2×10^{-19} Joules, or about 1 electron volt. This is still much bigger than $k_B T$ (especially at 4 K), but cannot be reduced much, if we want to keep error rates low, as I calculate in next section. It compares favorably with CMOS, where the logic levels are a volt, but the energy to charge the capacitor of about 100 attofarads is about three orders of magnitude higher.

If we could reduce the voltage levels in CMOS in proportion to the temperature (is there a reason why one cannot?), then we could reduce this energy per bit by four orders of magnitude, or T^2 .

Note that the fields and currents induced by storing a bit are still quite small. Since we are storing the bit as one flux quantum ($20 \text{ Gauss-}\mu\text{m}^2$), in a few micron-sized loop, the fields are only a few Gauss, while the critical field of Nb is about 2,000 Gauss. Likewise, the critical current density of a Nb wire is 10^7 to 10^8 A/cm^2 , about three to four orders of magnitude higher than that of a junction. It is tempting to say that we can reduce the size of the junctions, the critical currents, and therefore also the loop sizes to make a memory more dense. However, at a fixed operating temperature we seem to need a fixed and relatively large critical current to prevent thermally-induced phase slips (flux tunneling), and therefore errors.

C.3 Thermal Fluctuations and the Minimum Critical Current

The dynamics of a biased Josephson junction, or a junction in a loop, can be described by the motion of a classical particle in an effective potential, where the phase difference δ is the coordinate. The junction contributes an oscillatory contribution (the energy stored in the junction is periodic in δ). For a current biased-junction, this is the well-known RSJ washboard model, with a potential

$$U(\delta) = -\frac{\Phi_0}{2\pi}(I_b\delta + I_C \cos(\delta)) = -E_J \left(\delta \frac{I}{I_c} + \cos(\delta) \right) \quad (\text{C-10})$$

where $E_J = \hbar I_C / 2e = I_C \Phi_0 / 2\pi$ is the Josephson energy. For a junction in parallel with an inductor (more like the usual RSFQ cell), the potential

includes a parabolic term from the inductor

$$U(\delta) = E_J [\delta^2/2\beta_L + 1 - \cos(\delta) - I_b\delta/I_C] \quad (\text{C-11})$$

For our canonical 100 μA junction, the Josephson energy is about 3×10^{-20} Joules, or 2,400 $k_B T$. Because we want to ensure directionality of the switching, the bias current through the junction is typically about 70% of the critical current. Now the particle can escape from one local minimum by tunneling over the potential barrier created by the sinusoidal term. But this will be about one third of E_J . The error rate from these thermal escapes will be given by

$$\Gamma = \frac{\omega_C}{2\pi} \exp(-\Delta U/kT) \quad (\text{C-12})$$

and the "attempt frequency" is the characteristic timescale calculated above. If we want to ignore all these errors, having a minimum time t_{err} between errors for N gates in a device, we need $\Delta U/kT \sim \ln(\omega_C t_{err} N/2\pi)$. If we assume a billion seconds and a million gates, then $\Delta U \sim 60kT$, or $E_J \sim 180kT$ and the minimum critical current is about $30\mu\text{A}$ at 4 Kelvin.

C.4 Scaling Prospects for JJ Logic

Scaling superconducting computing circuits to smaller sizes and high enough densities to compete with CMOS appears to be a significant but necessary challenge. Due to the possibility of thermal errors, as discussed above, there is a practical limit on the minimum critical current for the junctions. We could reduce junction sizes to scale to more dense circuits, but this requires an increase in critical current density. Densities up to about $100\text{kA}/\text{cm}^2$ have been demonstrated, making the minimum junction size about $0.2\text{ }\mu\text{m}$. However, the inductance will continue to dominate. An alternative is to use other Josephson junctions as the inductors. To keep

the nonlinearity in these junctions small, or equivalently to make sure that only the desired junction in a cell switches, the other junctions need to have about three times larger critical current, making each of them one third of the inductance of the small junction. To get enough total inductance, we need a multi-layer stack of these larger junctions (or a series array) of about ten or fifteen junctions. I think this is not unreasonable to fabricate, especially because this stack of junctions could perhaps be just a single process and patterning step, where a thick multilayer of alternating superconducting layers and thin barriers can be repeated. I also don't think the circuit performance is particularly sensitive to variations of these extra junctions. The critical behavior would be dominated by the single, smaller, junction. So it seems to me reasonable to get a memory cell size down to about one square micron at the smallest. A one megabit memory at this density can fit on one square millimeter, or has a density of 100 megabytes per square centimeter, about 100X lower than current DRAM chips. But a petabyte of memory for a petascale computer requires a daunting 10,000 square meters!

D APPENDIX: JOSEPHSON JUNCTION DEVICE MODEL

The Josephson junction device – We consider a Josephson junction (JJ) constructed from two superconducting layers of area A , separated by a tunnel layer of thickness d . Conventionally, the superconductor is niobium and the insulator is aluminum oxide.

JJ circuit element – The circuit model for a JJ circuit element in a computer is shown in Figure D.1. The nonlinear response of the device is modeled by the JJ junction symbol. The shunting capacitor C is the capacitance between the two superconducting layers. The junction capacitance $C \approx A/d$, dominates the capacitance to ground $C_\Sigma \approx \sqrt{A}$ because $d \ll \sqrt{A}$. The resistance R_n represents the ‘normal’ tunneling resistance of the JJ, *i.e.* the tunneling resistance when the superconducting layers become normal metals. To increase the damping and avoid unwanted ringing, an additional shunt resistor R is added for single flux quantum (SFQ) logic circuits.




Figure D.1: Josephson junction resistively shunted junction (RSJ) circuit model.

D.1 RSJ Model

The equation of motion for the JJ is given by the RSJ model (McCumber 1968 [8], Likharev 1979 [10], Tinkham 2004 [11]). The time-dependent current $I(t)$ through the circuit model shown in Figure D-1 is given by:

$$I(t) = I_c \sin \delta + \frac{V}{R} + C \frac{dV}{dt} \quad (\text{D-1})$$

where I_c is the critical current of the JJ. The Josephson equations for the supercurrent I_s and the voltage V across the JJ are:

$$I_s = I_c \sin \gamma \quad (\text{D-2})$$

$$\frac{d\gamma}{dt} = \left(\frac{2e}{\hbar} \right) V = \left(\frac{2\pi}{\Phi_0} \right) V \quad (\text{D-3})$$

where γ is the phase difference of the superconducting order parameter ψ across the junction.

The magnetic flux quantum Φ_0 for superconductors is given by:

$$\Phi_0 = \frac{h}{2e} \quad (\text{D-4})$$

where h is Planck's constant and $2e$ is the charge of a Cooper pair, giving $\Phi_0 = 20.7 \text{ G } \mu\text{m}^2$. Its value is determined by the fact that the superconducting order parameter ψ must be single-valued in a loop encircling Φ_0 . The line integral $\oint \vec{A} \cdot d\vec{s} = N\Phi_0$ of the vector potential $\vec{A}$ about a closed loop in the superconductor counts the number N of flux quanta enclosed.

Substituting the JJ voltage V from Eq. D-3 into Eq. D-1, we find the equation of motion for the JJ circuit element using the RSJ model (McCumber 1968 [8]):

$$\frac{d^2\gamma}{dt^2} + \left(\frac{1}{RC} \right) \frac{d\gamma}{dt} + \left(\frac{2\pi I_c}{C\Phi_0} \right) \sin \gamma = \left(\frac{2\pi}{C\Phi_0} \right) I(t) \quad (\text{D-5})$$

For small phase differences $\sin \gamma \approx \gamma$, the RSJ model resembles an RLC tank circuit with resonant frequency

$$\omega_P = \sqrt{\frac{2\pi I_C}{C\Phi_0}} = \sqrt{\frac{1}{L_J C}} \quad (\text{D-6})$$

called the ‘plasma frequency,’ where the effective inductance of the Josephson junction is $L_J = \Phi_0/2\pi I_C$. The time constant of the tank circuit is:

$$\tau = RC \quad (\text{D-7})$$

Rewriting the RSJ equation of motion using these quantities we find

$$\frac{d^2\gamma}{dt^2} + \left(\frac{1}{\tau}\right) \frac{d\gamma}{dt} + \omega_P^2 \sin \gamma = \omega_P^2 \left(\frac{I(t)}{I_C}\right). \quad (\text{D-8})$$

For small phase differences, $\sin \gamma \approx \gamma$ and Eq. D-8 is the equation of motion of a simple harmonic oscillator. The quality factor Q is given by

$$Q = \omega_P \tau = \sqrt{\frac{2\pi R^2 C I_C}{\Phi_0}} \quad (\text{D-9})$$

For computer applications, the value of the shunt resistor R is often chosen so the RSJ circuit is critically damped with $Q = \omega_P \tau = 1/2$, and the system most rapidly approaches equilibrium.

The potential energy of the JJ circuit element *vs.* the phase difference γ for a constant current I_0 is:

$$U(\gamma) = -\left(\frac{\Phi_0 I_C}{2\pi}\right) \cos \gamma - \left(\frac{\Phi_0 I_0}{2\pi}\right) \gamma = -E \left[\cos \gamma + \left(\frac{I_0}{I_C}\right) \gamma \right] \quad (\text{D-10})$$

where the Josephson energy E_J is defined by:

$$E_J = \frac{\hbar I_C}{2e} = \left(\frac{\Phi_0}{2\pi}\right) I_C \quad (\text{D-11})$$

The critical current I_C of the Josephson junction is a very important parameter. It’s value is given by:

$$I_C = \frac{V_C}{R_n} \quad (\text{D-12})$$

where V_C the characteristic voltage of the Josephson junction, and is its tunnel resistance in the normal regime. The characteristic voltage V_C does not depend on any parameters of the junction except the superconducting energy gap $\Delta(T)$ and the temperature T . It is given by the Ambegaokar and Baratoff relation (1963 [3]):

$$V_C = \frac{\pi}{2} \left(\frac{\Delta(T)}{e} \right) \tanh \left(\frac{\Delta(T)}{2k_B T} \right) \quad (\text{D-13})$$

At low temperatures $T < T_C$, the energy gap approximates its zero temperature value (Likharev 1979 [10]).

$$\Delta(0) \approx (\pi/1.8)k_B T_C \quad (\text{D-14})$$

Near the critical temperature T_C , the gap $\Delta(T)$ drops rapidly with temperature. The normal tunnel resistance R_n of the JJ drops exponentially with the tunnel barrier thickness d , and is inversely proportional to its area A .

D.2 Driven Damped Pendulum Model

Equation D-8 is mathematically the same as the equation of motion for a driven damped pendulum (D’Humieres et al. 1982 [9]), shown in Figure D.2. For the equation of motion we have:

$$(ml^2) \frac{d^2\theta}{dt^2} + (\kappa) \frac{d\theta}{dt} + (mgl) \sin \theta = \Gamma(t) \quad (\text{D-15})$$

which describes how the angle θ of the pendulum responds to a torque $\Gamma(t)$ applied about the axis. Here ml^2 is the moment of inertia, κ is the viscous damping constant, and mgl is the restoring torque due to the action of gravity on a bob of mass m attached to a rigid, massless rod of length l .

For small angles away from a minimum, the pendulum acts as a simple




Figure D.2: Rigid pendulum with mass m , length ℓ and angle θ . The force of gravity mg creates a torque $mg\ell$ about the axis.

harmonic oscillator with resonant frequency

$$\omega_0 = \sqrt{\frac{g}{l}} \quad (\text{D-16})$$

and time constant

$$\tau = \frac{\kappa}{ml^2} \quad (\text{D-17})$$

Rewriting the Eq. D-15 using these expressions we find

$$\frac{d^2\theta}{dt^2} + \left(\frac{1}{\tau}\right) \frac{d\theta}{dt} + \omega_0^2 \sin \theta = \omega_0^2 \left(\frac{\Gamma(t)}{\Gamma_C}\right) \quad (\text{D-18})$$

where $\Gamma_C = mgl$ is the critical torque needed to lift the bob against gravity.

Equation D-18 has exactly the same form as the JJ equation of motion

Eq. D-8.

The energy of the pendulum for a constant applied torque Γ_0 is given by:

$$U(\theta) = -\Gamma_C \cos \theta - \Gamma_0 \theta = -\Gamma_C \left(\cos \theta + \left(\frac{\Gamma_0}{\Gamma_C}\right) \theta \right) \quad (\text{D-19})$$

where the characteristic potential energy E_P of a pendulum is equal to the critical torque:

$$E_P = mgl = \Gamma_C \quad (\text{D-20})$$

Comparing Eqs. D-8, D-10, and D-11 for the RSJ model with Eqs. D-18, D-19, and D-20, we find that the equation of motion and the potential energy of a JJ device and a driven damped pendulum are given by exactly the same expressions. This means we can use our intuition about driven pendulums to guide our understanding of Josephson junctions.

D.3 Tilted Washboard Model

To understand Josephson junction logic circuits, it is helpful to consider the “Tilted Washboard Model” shown in Figure D.3, which plots the JJ potential energy U vs. the phase γ for a fixed drive current I_0 . The drive current tilts the sinusoidal potential, creating a series of minima in U at lower energies, separated by $\Delta\gamma = 2\pi$ in phase. The JJ moves from one minimum to the next, as a single flux quantum Φ_0 passes through the barrier.




Figure D.3: Josephson junction energy U vs. number N of flux quanta Φ_0 , that have passed through, for $I_0/I_C = 0.6$. The JJ phase difference is $\gamma = 2\pi N$.

One can set up a logic circuit, where a bit of information is represented by Φ_0 , and digital operations are carried out by Josephson junctions – this

is the basis of Rapid Single Flux Quantum (RSFQ) logic (Likharev and Semenov 1991 [1]). At the bottom of each minimum in U , the JJ voltage V falls to zero, the phase γ is constant in time, and the drive current I_0 passes through the junction as a supercurrent. This state blocks the motion of a flux quantum through a JJ, and it can be used to store a flux quantum of information.

The energy ΔU required to escape from a minimum is given by Tinkham (2004 [11]):

$$\Delta U \approx 2E_J (1 - I_0/I_C)^{\frac{3}{2}}, \quad (\text{D-21})$$

where the barrier ΔU has a characteristic size set by the Josephson energy $E_J = \Phi_0 I_C$. The barrier ΔU falls to zero for $I_0 = I_C$ when the minima in U disappear. We can use Eq. D-21 to estimate the error rate of a JJ logic element due to thermal fluctuations. The probabilistic rate of escape P is the product of the attempt frequency $\omega_P/2\pi$ with a Boltzmann factor:

$$P \approx \left(\frac{\omega_P}{2\pi}\right) \exp\left(-\frac{\Delta U}{k_B T}\right) \quad (\text{D-22})$$

where the attempt rate $\omega_P/2\pi$ is determined by the plasma frequency ω_P for oscillations in phase γ near the bottom of a minimum from Eq. D-6.

The total energy E of the JJ is the potential energy U plus the 'kinetic' energy:

$$K = \frac{1}{2} \left(\frac{\Phi_0}{2\pi}\right)^2 \left(\frac{d\gamma}{dt}\right)^2 = \frac{1}{2} C V^2. \quad (\text{D-23})$$

which is proportional to the square of the 'velocity' $d\gamma/dt$, and is equal to the charging energy of the JJ capacitance C . If $V \neq 0$, the JJ phase will change according to the Josephson equation Eq. D-3, a current V/R will flow through the resistor, and the state of the system will change.

One can lift the JJ state over the potential barrier by applying a short voltage pulse that adds a kinetic energy E_{in} , as shown in Figure D.3. The

response of the Josephson junction depends on its damping. For a lightly damped junction ($\omega_P\tau > 1/2$), the state will continue to travel horizontally to the right in Figure D.3 after the exciting voltage pulse has ceased. In the lightly damped regime, the JJ has multiple stable states for drive current I_0 , a zero-voltage state with supercurrent I_0 , and a finite voltage state where part of the drive current flows through R . The system is hysteretic, and the observed state is determined by the history of excitation. These voltage states were used in the original IBM superconducting computer program.

In the more heavily damped regime ($\omega_P\tau \leq 1/2$), the behavior qualitatively changes. In response to a voltage pulse, the JJ state flips out of one minimum and is trapped by the next, as indicated in Figure D.3. During the transfer, a single flux quantum moves through the junction. In its fall to the bottom of the next minimum, the JJ either dissipates the kinetic energy in the shunt resistor, or sends one or more flux quanta out to other JJ devices down a superconducting transmission line – in this case R is the characteristic resistance. The most rapid transitions occur for critical damping ($\omega_P\tau = 1/2$). In this damped regime, the motion of the JJ state is quite simple, providing a natural approach to rapid single flux quantum (RSFQ) digital logic based on the flux quantum.

D.4 Rapid Single Flux Quantum (RSFQ) Logic

Likharev and Semenov (1991 [1]) developed an approach to superconducting logic based on single flux quanta, which they named Rapid Single Flux Quantum (RSFQ) Logic. RSFQ logic is the basis of two leading approaches to superconducting computing, Reciprocal Quantum Logic (RQL) and Energy-efficient Rapid Single Flux Quantum (ERSFQ) logic, that were

presented at the JASON study. We give a brief summary of RSFQ logic characteristics here.

Inside an RSFQ computer, a bit of information is represented by a flux quantum Φ_0 . This provides a natural way to digitize analog signals that is analogous to the classification of a “1” or a “0” in CMOS logic by ranges of voltage. In moving from one minimum to the next, as shown in Figure D.3, one flux quantum passes through the Josephson junction. For the pendulum model shown in Figure D.3, this corresponds to the pendulum bob flipping once over the top.

A flux quantum representing a bit of information is passed from one to another Josephson junction as a voltage pulse $V(t)$ that travels along a superconducting transmission line at the speed of propagation $c_S \sim c/3 \sim 1 \times 10^8 \text{m/sec}$, where c is the speed of light in a vacuum. By integrating the second Josephson equation, Eq. D-3, we find a relation between the peak voltage V_Φ and the pulse duration Δt .

$$\Phi_0 = \int dt V(t) = V_\Phi \Delta t = 2.07 \text{mVpsec}. \quad (\text{D-24})$$

Shorter pulses are obtained for larger peak voltages. One can see that this approach is naturally suited to mV voltage pulses with psec time widths, making is promising for low energy, high speed computing. The upper limit to V_Φ is given by the Josephson characteristic voltage V_C from Eq. D-13:

$$V_\Phi \approx V_C = \frac{\pi}{2} \left(\frac{\Delta}{e} \right) \tanh \left(\frac{\Delta}{2k_B T} \right), \quad (\text{D-25})$$

where 2Δ and T_C are the superconductor’s energy gap and critical temperature.

The duration Δt and energy of a single flux quantum voltage pulse set fundamental limits to the speed and bit energy of an RSFQ computer. For

niobium, the superconducting energy gap is $2\Delta = 3.05\text{meV}$, giving a peak voltage $V_\Phi \approx V_C \approx 2.5\text{mV}$ at temperatures well below T_C . This peak voltage allows a flux quantum to be carried by a picosecond pulse with time duration $\Delta t \approx 0.8\text{psec}$ along a superconducting transmission line. The physical length of the pulse is $\Delta x = c_s \delta t \sim 80\text{ }\mu\text{m}$, using $c_s \sim 1 \times 10^8\text{m/sec}$. These picosecond pulses open the way for very rapid transmission of bits at rates up to $1/\Delta t \approx 1.25\text{THz}$ for niobium.

The energy E_{sw} to switch a Josephson junction from one potential energy minimum to the next, as shown in Figure D.3, is proportional to the Josephson energy E_J , the amplitude of the sinusoidal oscillations in U :

$$E_{sw} = 2\pi E_J \left(\frac{I_0}{I_C} \right) = \Phi_0 I_0 \sim \Phi_0 I_C \quad (\text{D-26})$$

The switching energy E_{sw} is proportional to the junction area A , because the critical current density J_C is a constant determined in the fabrication process by the materials and the barrier thickness. Using representative values $J_C = 10\text{ kA/cm}^2$ and $A = 1\text{ }\mu\text{m}^2$ for Nb/ AlO_x /Nb junctions, we find $E_{sw} \approx 2 \times 10^{-20} J$, a very small value.

The switching energy is a fundamental property, and one would like to design the system to achieve the lowest possible value. In practice, the minimum value of $E_J \propto J_C A$ is determined by the maximum error rate P that is acceptable, where P (Eq. D-22) is exponentially sensitive to E_J . Higher current densities J_C are desirable, because they permit using JJ devices with smaller area A and larger packing density, as well as shorter duration single-flux-quantum voltage pulses (Eq. D-24). To increase the current density J_C , one could use a thinner tunnel barrier to decrease R_n or choose a superconductor with a larger energy gap Δ , as shown in Eqs. D-12 and D-25.

The Josephson junction amplifies the energy $E_{in} \sim (1/2)CV_{in}^{1/2}$ de-

posited by an incoming voltage pulse of height V_{in} , to generate a full rotation of the JJ phase. The energy E_{in} is a fraction α of the Josephson energy – the typical value $\alpha \sim 0.3$ is a compromise between the desire for a small switching energy, and the need to prevent unwanted thermal motion over the potential barrier. During a rotation of the JJ state from one minimum to the next (Figure D.3), the initial potential energy is converted into a kinetic energy $K \sim E_J$. A full flux quantum moves through the junction, creating an outgoing voltage pulse of height V_Φ given by $(1/2)CV_\Phi^2 \sim K$.

In this way, the Josephson junction digitizes the incoming voltage pulse into a “1” or a “0” depending on whether the incoming pulse is large enough to flip the JJ phase over the barrier. In addition, the JJ restores a weak incoming “1” pulse to the voltage V_Φ expected for a full flux quantum. Both properties are required for logic circuits. The amplification factor is $1/\alpha$ in energy, with a typical value $1/\alpha \sim 3$. This provides enough energy to drive two similar JJ devices for a fanout of 2.

The JJ switching energy $E_{in} \sim (1/2)CV_\Phi^2$ is quite small compared with the charging energy CV^2 of a CMOS line, in part because the voltages V_Φ are ~ 1 mV instead of ~ 1 V. However this advantage is offset by the energy needed to cool the superconducting computer to liquid He temperatures, so that the effective switching energy at room temperature is a factor ~ 300 higher for a commercial cooler.

An important advantage of JJ computers is the fact that information is passed along superconducting transmission lines as single-flux-quantum voltage pulses. The energy of a single pulse $E_J \approx 1 \times 10^{-19} J$ is quite small, and it is not dissipated by the line, but enters at one end and leaves at the other. The impedance of the line is chosen to approximately match the JJ devices. A difficulty created by single-flux-quantum pulses, is that timing

become more difficult, because different pulses will travel different distances, and arrive at different times at a gate. This can be managed by adding a register to the gate input that records incoming data during the interval between two clock pulses, and then forward the data to the gate (Likharev and Semenov 1991 [1]). The clock speed must be slow enough to handle the expected range in arrival times for pulses from different parts of the chip.

E ACRONYMS

CAD - Computer Aided Design

CMOS - Complimentary Metal Oxide Semiconductor

DRAM - Dynamic Random Access Memory

DSM - Distributed Shared Memory

DSP - Digital Signal Processor

ERSFQ - Energy-efficient Rapid Single Flux Quantum logic

JJ - Josephson Junction

JMRAM - Josephson Magnetoresistance Random Access Memory

JTL - Josephson Transmission Line

MRAM - Magnetoresistance Random Access Memory

ROM - Read Only Memory

RQL - Reciprocal Quantum Logic

RSFQ - Rapid Single Flux Quantum

SFQ - Single Flux Quantum

SQUID - Superconducting QUantum Interference Device

References

- [1] K. K. Likarev and V. K. Semenov, "RSFQ Logic/Memory Family: A New Josephson Junction Technology for Sub-Terahertz-Clock-Frequency Digital Systems," *IEEE Trans. Appl. Superconductivity* 1, 3 (1991).
- [2] Peter Kogge, "Next-Generation Supercomputers," *IEEE Spectrum* (February 2011).
- [3] Vinay Ambegaokar and Alexis Baratoff, "Tunneling between Superconductors," *Phys. Rev. Lett.* 10, 486 (1963).
- [4] R. E. Miller, W. H. Mallison, A. W. Kleinsasser, K. A. Delin, and E. M. Macedo, "Niobium Trilayer Josephson tunnel junctions with ultrahigh critical current densities," *Appl. Phys. Lett.* 63, 1423 (1993).
- [5] Steven T. Ruggiero and David A. Rudman, Eds. "Superconducting Devices," (Academic Press, London, 1990).
- [6] Thomas Vogelsang, "Understanding the Energy Consumption of Dynamic Random Access Memories," *MICRO '43 Proc. 2010 43rd Ann. IEEE/ACM Int. Sym. on Microarchitecture*, (IEEE Computer Society, Washington, DC, USA (2010).
- [7] N. Yoshikawa, T. Tomida, M. Tokuda, Q. Liu, X. Meng, S. R. Whiteley, and T. Van Duzer, "Characterization of 4K CMOS Devices and Circuits for Hybrid Josephson-CMOS System," *IEEE Trans. Appl. Superconductivity* 15, 267 (2005).
- [8] D. E. McCumber, "Effect of ac Impedance on dc Voltage-Current Characteristics of Superconductor Weak-Link Junctions," *J. Appl. Phys.* 39, 3113 (1968).

- [9] D. D'Humières, M. R. Beasley, B. A. Huberman and A. Libchaber, "Chaotic states and routes to chaos in the forced pendulum," *Phys. Rev. A* 26, 3483 (1982).
- [10] K. K. Likarev, "Superconducting weak links," *Rev. Mod. Phys.* 51, 101 (1979).
- [11] Michael Tinkham, "Introduction to Superconductivity, Second Edition," (Dover Publications, Mineola, NY, 1996).
- [12] P. G. de Gennes, "Superconductivity of Metals and Alloys," (Westview Press, 1999).
- [13] A. W. Kleinsasser, "High Performance Nb Josephson Devices for Petaflop Computing," *IEEE Trans. Appl. Supercond.* 11, 1043 (2000).
- [14] Thomas H. Lee, "Planar Microwave Engineering," (Cambridge University Press, 2004).
- [15] S. M Sze and Kwok K. Ng, "Physics of Semiconductor Devices," (Wiley-Interscience, Hoboken, NJ, 2007).

STANDARD DISTRIBUTION LIST

(b)(6)

DARPA Library
3701 North Fairfax Drive
Arlington, VA 22203-1714

Defense Technical Information Center (DTIC)
8725 John J. Kingman Road
ATTN: DTIC-OA [redacted]
Suite 0944
Fort Belvoir, VA 22060-6218

Director, IDA
Technical Information Services
Room 8701
4850 Mark Center Drive
Alexandria, VA 22311-1882

Director, DTRA
Research Development Office
8725 John Jay Kingman Road
Room 3380, Mail Stop 6201
Fort Belvoir, VA 22060

(b)(3)
(b)(6)

OGA

JASON Library [5]
The MITRE Corporation
3550 General Atomics Court
Building 29
San Diego, CA 92121-1122

Principal Deputy Director
Office of Science, SC-2/Forrestal Building
U.S. Department of Energy
1000 Independence Avenue, SW
Washington, DC 20585

Records Resource
The MITRE Corporation
Mail Stop C025
202 Burlington Road, Rte 62
Bedford, MA 01730-1420

Reports Collection
Los Alamos National Laboratory
Mail Station 5000
MS A150
PO Box 1663
Los Alamos, NM 87545

Superintendent
Code 1424
Attn: Documents Librarian
Naval Postgraduate School
Monterey, CA 93943

[redacted]
Deputy Administrator for
Defense Programs
NA-12
National Nuclear Security Administration
U.S. Department of Energy
1000 Independence Avenue, SW
Washington, DC 20585

[redacted]
Administrator, U.S. Dept of Energy
National Nuclear Security Administration
1000 Independence Avenue, SW
NA-10 FORS Bldg
Washington, DC 20585

[redacted]
Director
Defense Advanced Research Projects Agency
3701 N. Fairfax Drive
Arlington, VA 22203-1714

[redacted]
Deputy Director, Horizon
OSD AT&L/DDR&E/PD
Pentagon, Room 38854
Washington, DC 20520

[redacted]
President & CEO
The MITRE Corporation
Mail Stop N640
7515 Colshire Drive
McLean, VA 22102-7508

(b)(6)

[Redacted]

[Redacted]

Reentry Systems Branch Head, Navy Strategic
Systems Programs
Strategic Systems Programs (Aftn. SP28)
2521 Clark Street, Suite 1000
Arlington, VA 22202-3930

[Redacted]

The MITRE Corporation
Mailstop MDA/ RM 5H305
7515 Colshire Drive
McLean, VA 22102-7508

[Redacted]

3030 Defense Pentagon
Room 3C913A
Washington, DC 20301-3030

[Redacted]

Office of Naval Research
Code 31
800 North Quincy Street
Arlington, VA 22217-5660

[Redacted]

Mailstop T130
7515 Colshire Drive
McLean, VA 22102-7508

[Redacted]

[Redacted]

Office of the Director of National Intelligence
IARPA
Washington, DC 20511

[Redacted]

Assistant Secretary of the Navy
(Research, Development & Acquisition)
1000 Navy Pentagon
Washington, DC 20350-1000

[Redacted]

OSD/DDR&E
Director for Basic Research Laboratories
3030 875 N. Randolph St, Suite 150
Arlington VA 22203

[Redacted]

Principal Deputy Director
DDR&E
3040 Defense Pentagon
Room 3B 854
Washington, DC 20301-3040

[Redacted]

Director, Special Projects Office
Department of Homeland Security
245 Murray Lane SW
Washington, DC 20528

OGA

(b)(3)
(b)(6)



NATIONAL SECURITY AGENCY
CENTRAL SECURITY SERVICE
FORT GEORGE G. MEADE, MARYLAND 20755-6000

FOIA Case: 65695B
6 August 2020

STEVEN AFTERGOOD
FEDERATION OF AMERICAN SCIENTISTS
1112 16TH NW
SUITE 400
WASHINGTON DC 20036

Dear Steven Aftergood:

This responds to your Freedom of Information Act (FOIA) request sent via email to this office on 27 October 2011 for "a copy of the report produced for NSA - JASON Defense Advisory Panel: "Superconducting Computing" JSR-11-120." A copy of your request is enclosed.

Your request has been processed under the FOIA and the document you requested is enclosed. Certain information, however, has been deleted from the enclosure. As stated in the interim response sent to you on 20 June 2012, the report requested was not yet finalized. Rather than have the request denied and a new request submitted, you agreed to wait until the report was finalized. The request was assigned Case Number 65695. For the purpose of fee assessment, you were placed into the "media" fee category for this request and there were no assessable fees.

The document has been reviewed by this Agency as required by the FOIA and this Agency is authorized by various statutes to protect certain information concerning its activities. We have determined that such information exists in this document. Accordingly, those portions are exempt from disclosure pursuant to the third exemption of the FOIA, which provides for the withholding of information specifically protected from disclosure by statutes. The specific statutes applicable in this case is FOIA Exemption (b)(3), Section 6, Public Law 86-36 (50 U.S. Code 3605). Also, personal information regarding other individuals has been deleted from the enclosure in accordance with 5 U.S.C. 552 (b)(6). This exemption protects from disclosure information which would constitute a clearly unwarranted invasion of personal privacy.

The document was also reviewed by the CIA, and they have asked that we protect CIA equities. The specific statutes applicable to CIA are (b)(3) - Section 6 of the Central Intelligence Agency Act of 1949, as amended, and (b)(6) - Section 102A(i)(1) of the National Security Act of 1947, as amended. Any appeal of the denial of CIA information should be directed to that agency.

Since these deletions may be construed as a partial denial of your request, you are hereby advised of this Agency's appeal procedures.

You may appeal this decision. If you decide to appeal, you should do so in the manner outlined below.

- The appeal request must be in writing and addressed to:

NSA/CSS FOIA/PA Appeal Authority (P132)
National Security Agency
9800 Savage Road STE 6932
Fort George G. Meade, MD 20755-6932

- The request must be postmarked no later than 90 calendar days of the date of this letter. Decisions appealed after 90 days will not be addressed.
- Please include the case number provided above.
- Please describe with sufficient detail why you believe the denial of requested information, was unwarranted.
- NSA will endeavor to respond within 20 working days of receiving your appeal, absent any unusual circumstances.

You may also contact our FOIA Public Liaison at foialo@nsa.gov for any further assistance and to discuss any aspect of your request. Additionally, you may contact the Office of Government Information Services (OGIS) at the National Archives and Records Administration to inquire about the FOIA mediation services they offer. The contact information for OGIS is as follows:

Office of Government Information Services
National Archives and Records Administration
8601 Adelphi Rd- OGIS
College Park, MD 20740
ogis@nara.gov
(877) 684-6448
(202) 741-5770
Fax (202) 741-5769

Sincerely,



SHARON C. LINKOUS
Acting Chief, FOIA/PA Office
NSA Initial Denial Authority

Encls:
a/s