

Superconducting Technology Assessment

Position Papers

Contents:

- Towards a Technology and Architecture Hybrid?
 - Thomas Sterling, Panel Moderator
- Superconductor Technology for High-End Computing System Issues and Technology Roadmap
 - Arnold H. Silver
- Opportunities, Challenges, and Projections for Superconductor RSFQ Microprocessors
 - Mikhail Dorojevets
- Cryogenic Memories for RSFQ Ultra-High-Speed Processor
 - T. Van Duzer
- System Balance and Fast Clocks
 - Burton J. Smith

Towards a Technology and Architecture Hybrid?

Thomas Sterling, Panel Moderator

*Center for Computation and Technology
Department of Computer Science
Louisiana State University
and
Center for Advanced Computational Research
California Institute of Technology*

The challenges to extending the delivered computing capabilities of semiconductor technology through Moore's Law, while manageable in the short term, may prove difficult or possibly impractical in the long term. Even now, the complex interplay of power and performance is resulting in significant changes in previous trends. Clock rates of commodity microprocessors are flattening even as multi-core chips are emerging as the norm for next generation systems. While conventional wisdom has dictated an assumption of continued adherence to the pure CMOS tradition of the last decade and more, the supercomputing community must consider the possibility of alternative technologies, at least in combination with more conventional devices. More than just changing or augmenting the technology base, new architecture structures and programming models may need to be considered and multiple levels to exploit the potential of such advances. This panel addresses these issues with a focus on one possible alternative technology: superconductor devices. Rapid Single Flux Quantum (RSFQ) logic exhibits operational properties in terms of performance and power that now positions it as a potential future leader among alternative digital technologies to augment semiconductor components in hybrid systems. But it is also challenged by lack of maturity and commercial market as well as its reliance on extreme operational temperature regimes. This panel brings together leaders in the fields of technology and computer architecture to consider the possible strategies and potential viability of superconductor based supercomputing.

RSFQ technology may deliver clock rates in excess of an order of magnitude that of the corresponding semiconductor logic and with dramatically reduced power requirements. Further, at least in principle, it is easier to fabricate than heavily doped based semiconductor device fabrication processes. Nonetheless, in spite of decades of research and experience with small fabrication lines, it has not managed to challenge the prevailing semiconductor technologies. However, the increasing difficulties to sustaining current level of growth in density and performance of CMOS within practical power constraints may change this. This panel considers critical issues of technology and architecture and how RSFQ may contribute effectively to future supercomputing next decade. Four major topics will be addressed:

1. superconductor technology
2. micro-architecture using RSFQ
3. hybrid memory systems

4. system architecture incorporating superconductor components

Superconductor technology, its viability and capability, will be discussed by Arnold Silver, the inventor of the original single flux quantum electronic circuits and long-time leader in the field at TRW. Designing micro-architectures using RSFQ devices will be described by Mikhail Dorojevets of SUNY Stony Brook who is the architect of the FLUX test chip. Ted van Duzer of UC Berkeley will discuss a potentially important means of providing superconductor based systems with high capacity, high bandwidth memory which is critical to the success of these systems. And, Burton Smith, Chief Scientist of Cray Inc. will explore the system level implications for computer architecture for future supercomputer systems exploiting superconductor device technology.

Superconductor Technology for High-End Computing System Issues and Technology Roadmap

Arnold H. Silver

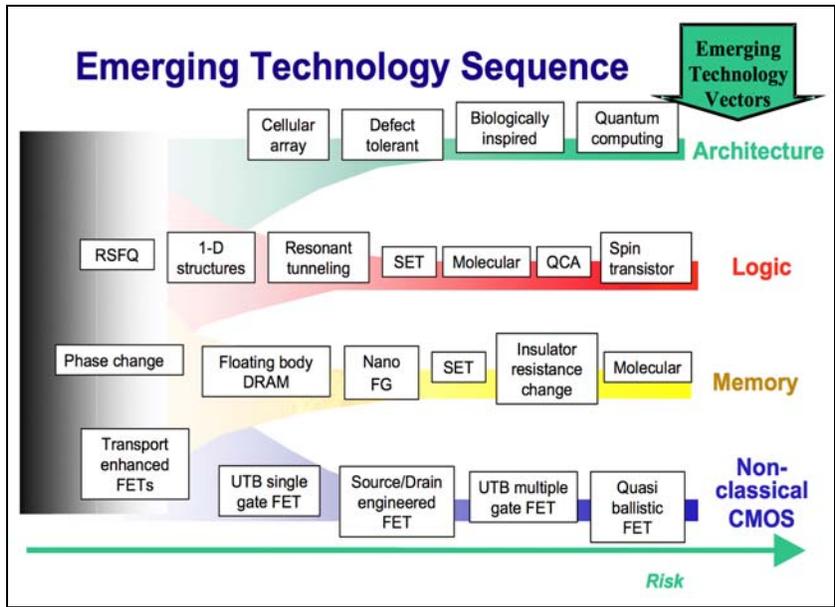


Figure 1. RSFQ was identified the lowest risk (highest maturity) potential emerging technology for processing beyond silicon. (From 2004 ITRS Update.)

under a focused program. They defined a five-year technology Roadmap to meet that goal as illustrated in Fig. 2.

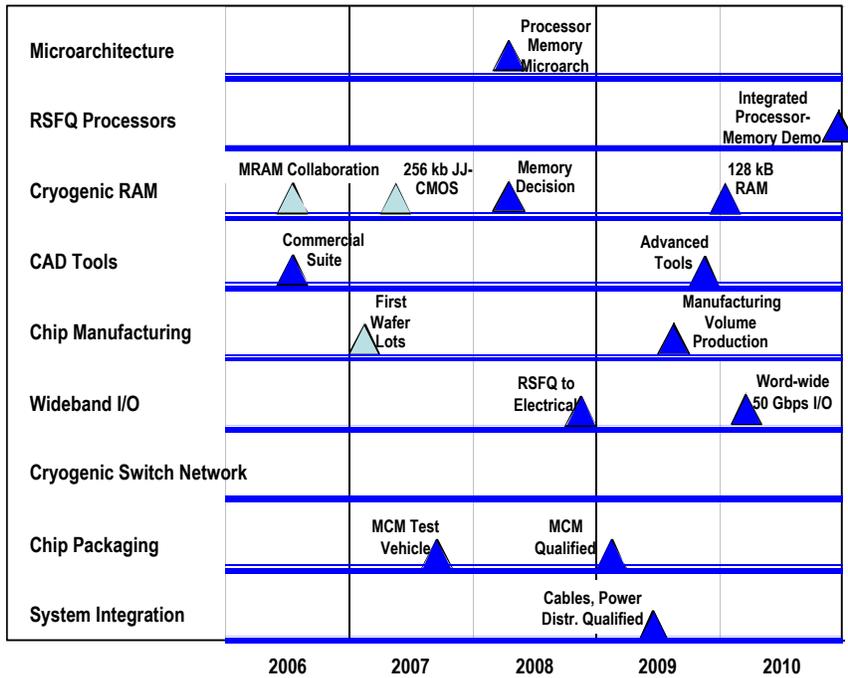


Figure 2. Roadmap for RSFQ technology tools and components.

Introduction

The ITRS 2004 Update on Emerging Research Devices identified superconductor rapid single flux quantum (RSFQ) technology as the most advanced of any of the alternative candidate technologies for extending performance beyond today's semiconductor technology (Fig. 1).

A Superconducting Technology Assessment (STA) Panel assessed the readiness of superconductor RSFQ technology to initiate system development in 2010, including all the elements necessary for implementation of RSFQ processors for high-end computing systems. The Panel concluded that RSFQ VLSI and other necessary technologies could be brought to the required state of readiness by year 2010

The most ambitious milestone is the Integrated Processor-Memory Demo. It requires development of:

1. Processor-cryogenic RAM microarchitecture
2. RSFQ cell library and a suite of CAD tools
3. RSFQ chip manufacturing facility
4. Superconducting MCM
5. Cryogenic RAM
6. 1M-gate RSFQ processor
 - 50 GHz clock
 - Approximately 10 chips on a single MCM
 - Inter-chip communication at the clock frequency.

The total investment was estimated at \$400M over five years. Government investment

will be required to accomplish the Roadmap.

What Is RSFQ Electronics?

RSFQ is the latest generation of high performance superconductor circuits based on Josephson junction devices. Josephson junctions (JJ), the basic superconductor-switching device, can operate in two distinct modes:

- **Voltage latching mode** where junctions switch from the zero-voltage state to the voltage-state of about 2.5 mV.
- The early work (IBM and the Japanese Josephson computer projects of the 1970's and 1980's) used voltage-latching circuits.
- AC power is required to reset the junction to zero voltage.
- **Non-latching mode** where the switching event in a junction generates single magnetic flux quantum pulses.
- RSFQ devices generate, store, and transmit identical magnetic single flux quantum (SFQ) pulses at frequencies approaching 1,000 GHz.
- RSFQ circuits are DC powered.

Table 1 compares CMOS and RSFQ device technologies. Since circuit fabrication is similar to semiconductors, RSFQ leverages VLSI processing technology and CAD tools.

Table 1. Comparison of CMOS and RSFQ Devices

Function	CMOS	RSFQ
Basic Switch	▪ Transistor	▪ Josephson tunnel junction (a 2 terminal device)
Data Format	▪ Voltage levels	▪ Identical picosecond current pulses
Speed Test	▪ Ring oscillator	▪ Asynchronous flip-flop ▪ 770 GHz achieved ▪ 1,000 GHz expected
Data Transfer	▪ Data bus ▪ RC delay and power dissipation	▪ Nearly lossless and dispersion-free superconducting transmission lines that support "ballistic" transfer at ~ 100 μ m/ps
Clock Distribution	▪ Clock bus	▪ Clock pulse regeneration by RSFQ junctions ▪ Nearly lossless and dispersion-free superconducting transmission lines
Logic Switch	▪ Complementary transistor pair	▪ Two-junction comparator
Bit Storage	▪ Charge on a capacitor	▪ Current in an inductor
Power	▪ Volt levels	▪ Millivolt levels
Fan-In, Fan-Out	▪ Large	▪ Small
Power Distribution	▪ Ohmic power bus	▪ Lossless superconducting wiring
Noise	▪ >300 K thermal noise	▪ 4 Kelvin thermal noise that enables low power operation

Is RSFQ Ready for Investment?

Significant development is needed to make RSFQ ready for design and construction of high-end computers. Although RSFQ circuits are still relatively immature, their similarity in function, design, and fabrication to semiconductor circuits permits realistic extrapolations. Progress has been demonstrated on limited budgets by U.S. companies such as Northrop Grumman and HYPRES, and in universities such as Stony Brook University and the University of California, Berkeley. Recent efforts in Japan are making similar progress. Most of the design, test, and fabrication tools are derived from similar semiconductor tools with some modification. Small asynchronous RSFQ circuits have been demonstrated at 770 GHz, and system clocks greater than 50 GHz appear attainable. The extremely low power will enable systems that have greatly increased computational capability and reduced power requirements compared to today's high-end systems.

The Panel concluded that superconductor RSFQ circuit technology is ready for an aggressive, focused investment to meet a 2010 schedule for initiating the development of petaflops-class computing. This judgment was based on:

- An evaluation of progress made in the last decade.
- Projection of an advanced VLSI process for RSFQ in a manufacturing environment.
- A reasonable roadmap for RSFQ circuit development that is coordinated with manufacturing and packaging technologies.

Figure 3 illustrates one possible configuration of the cryogenic system, including processors, RAM, and network switch.

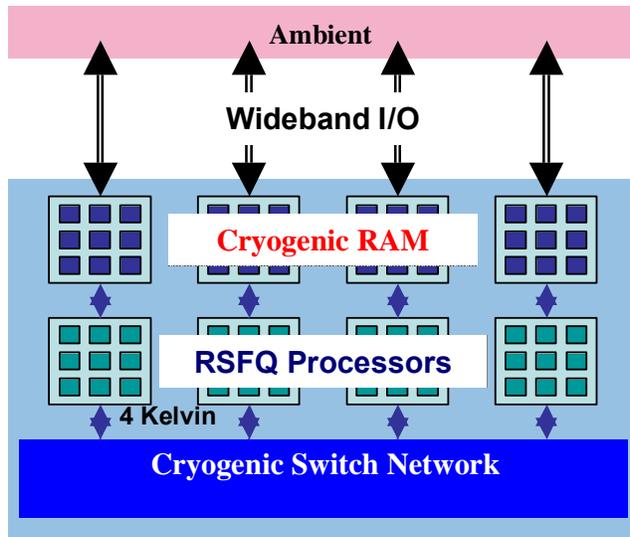


Figure 3. Notional diagram of the cryogenic system. RSFQ processors communicate with local cryogenic RAM and the cryogenic switch network. Cryogenic RAM communicates with ambient electronics via a wideband I/O.

manufacturing technology that can support the volume and cost targets for delivery of known good die for all RSFQ chip types required for a petascale system.

If development continues beyond the 2010 timeframe, a production capability for chips with 250 GHz speeds and densities comparable with CMOS are possible.

Cryogenic RAM

Three attractive candidates for fast, dense cryogenic RAM were identified: hybrid CMOS-JJ RAM, MRAM and hybrid MRAM, and ballistic SFQ RAM. T. Van Duzer discusses RAM.

MCM

MCMs that support 50 GHz communications between chips are necessary. The design of MCMs for RSFQ chips is technically feasible and fairly well understood. However, the design for higher speeds and interface issues need further development. MCMs for processor elements will be much more complex and require more layers of impedance controlled wiring than those built previously, with stringent control of crosstalk and ground-bounce.

The options are to [1] develop a superconducting MCM production capability, [2] find a vendor willing to customize its advanced MCM packaging process to include superconducting wire layers, or [3] procure MCMs with advanced normal metal layers for the bulk of the MCM, then develop an internal process for adding superconducting wiring.

An alternative to planar packaging on MCMs and boards is 3D packaging. Conventional electronic circuits are designed and fabricated using a planar, monolithic approach with only one major active device layer. More compact packaging technologies can bring active devices closer to each other allowing shorter time-of-flight, a critical parameter for systems with higher clock rates. In systems with superconducting components, 3D packaging enables higher component density, smaller vacuum enclosures, and shorter distances between different sections of the system.

Wideband I/O

RSFQ chips dissipate very little power, but the heat load for a petaflops system from heat conduction between the cryostat and room temperature through I/O and power lines will be very significant. Therefore, the I/O design must be a careful balance between electrical SNR and thermal properties. High bandwidth signal I/O requires low-loss, high-density cabling, which translates to high conductivity or large cross-section signal lines. Therefore, the I/O design must find the right balance between thermal and electrical properties.

The challenges imposed by tens of Pb/s bandwidth between the cold and room temperature sections of a petaflops superconducting supercomputer may require novel architectures to best suit optical packet switching, which has the potential to address the shortcomings of electronic switching, especially in the long term. The input data lines can use WDM optical technology, which appears to afford the best electrical-thermal solution. The principal problem is the output circuitry. Since there is not enough power in an SFQ data bit to directly drive

Long Lead Items

While all items in the Roadmap are important, the major long lead development items are RSFQ chip manufacturing, cryogenic RAM, superconducting MCMs, and wideband input/output from 4 Kelvin to ambient electronics.

Chip manufacturing

By 2010 production capability for high-density RSFQ chips should be achievable by application of manufacturing technologies and methods similar to those used in the semiconductor industry. The 2010 capability can be used to produce chips with speeds of 50 GHz or higher and densities of 1-3 million junctions per cm^2 .

The chip manufacturing capability needs to meet the following criteria:

- Earliest possible availability of RSFQ chips for microarchitecture, CAD, and circuit design development efforts. These chips must be fabricated in a process sufficiently advanced to have reliable legacy to the final manufacturing process.
- Firm demonstration of yield and

ambient semiconductor electronics, interface circuits are required to amplify the SFQ voltage pulse. Semiconductor drive circuits consume more power than can be tolerated at the 4-Kelvin stage. One option is to communicate SFQ signals up to an intermediate temperature stage and then optically up to room temperature.

Cryogenic Switch Network

The interconnection network at the core of a supercomputer is a high-bandwidth, low-latency switching fabric with thousands or even tens of thousands of ports to accommodate processors, caches, memory elements and storage devices. The Bedard crossbar switch architecture, with low fan-out requirements and replication of simple cells, is a good candidate for this function.

Power Cables

Superconductor circuits for supercomputing applications are based on DC-powered RSFQ circuits. Due to the low voltage (mV level), the total current to be supplied is in the range of few amps for small-scale systems and can be easily increased to kiloamps for large-scale systems. Serial distribution of DC current to small blocks of logic has been demonstrated, and this will need to be accomplished on a larger scale in order to produce a system with thousands of chips. However, we can expect that the overhead of current-supply reduction techniques on-chip will drive the demand for current supply into the cryostat as high as can be reasonably supported by cabling.

System Integration

System integration is a critical, but historically neglected, part of the overall system design. It is usually undertaken only at later stages of the design. System integration and packaging of RSFQ circuits offer several challenges due to the extremely high clock rates (50-100 GHz) and operation at cryogenic temperatures (4 - 77 K). The design of secondary packaging technologies and interconnects for RSFQ chips is technically feasible and fairly well understood. The lack of a superconducting packaging foundry with matching design and fabrication capabilities could be a major issue.

The design of enclosures and shielding for cryogenic electronic systems is technically feasible and fairly well understood. However, these techniques have never been tested for dimensions on the order of meters. The use of hybrid technologies—superconductor, optical, and conventional electronic components and system interfaces with different physical, electrical and mechanical properties—further complicate the system testing.

Refrigeration

The technology for the refrigeration plant is understood. Commercial cryocoolers are available, but engineering changes may be needed to upscale them for larger scale systems. It may be desirable to consider the trade-off between multiple smaller coolers versus one large cooler. Development toward a 10 W or larger 4 K cooler would be desirable to enable a supercomputer with modular cryogenic units. One key issue is the availability of U.S. manufacturers. Development funding may be needed for U.S. companies to insure that reliable American coolers will be available in the future.

Opportunities, Challenges, and Projections for Superconductor RSFQ Microprocessors

Mikhail Dorojevets
Dept. of Electrical and Computer Engineering
Stony Brook University, Stony Brook, NY 11794-2350
midor@ece.sunysb.edu

Superconductor processors based on Rapid Single Flux Quantum (RSFQ) circuit technology can reach and exceed operating frequencies of 100 GHz, while keeping processor power consumption low. These features provide an opportunity to build compact, multi-petaflops systems with ultra-high-speed 64/128-bit single-chip microprocessors to address the government's critical mission needs for high end-computing (HEC).

The availability of ultra-high-speed, low power superconductor circuit technology is only one of several requirements for successful high-performance system design. In order to be able to initiate the practical design of a superconductor multi-petaflops system, the following critical design challenges need to be addressed:

- Processor microarchitecture;
- Memory;
- Interconnect.

The key characteristics of superconductor processors, such as ultra-high clock frequency and very low power consumption, are due to the following properties:

- Extremely fast (a few-picosecond) switching times of superconductor devices;
- Very low power consumption;
- Ultra-high-speed, superconducting interconnect capable of transmitting signals (picosecond pulses) with negligible attenuation at full processor speed.

Simple sub-micron RSFQ gates (such as toggle flip-flops) have already demonstrated operation frequencies reaching 770 GHz. Currently, the complexity and speed of superconductor chips reached the point when RSFQ chips with tens of thousands Josephson junctions have been demonstrated to operate at ~20 GHz clock frequencies, while less complex chips have reached 50 GHz clock rates. Among those successfully demonstrated chips were small crossbar switches, front-ends for digital signal processing, and experimental microprocessor prototypes.

Another advantage of superconductor circuits is the ballistic transport of pulses over superconducting Nb lines without any RC charge process involved. Transmission rates reaching 60 GHz have been already demonstrated for reliable chip-to-chip communication over lines several centimeters long for picosecond voltage pulses traveling at ~ one third of the speed of light in vacuum (~100 $\mu\text{m}/\text{ps}$).

RSFQ circuits have both dynamic and static power consumption. Each RSFQ gate dissipates static power in its bias resistors that set the operating current for each junction. Currently, a typical junction with 140 μA critical current consumes ~ 200 nW, and a typical clocked gate ~ 2 μW when they are idle. In the meantime, dynamic power dissipation for such gate is ~ 1.4 nW/GHz, i.e., ~ 140 nW/gate at 100 GHz clock frequency. Static power consumption for future

VLSI scale superconductor circuits can be reduced by a factor of 3 by decreasing their bias voltage (currently 2 mV).

As estimated, a 100 GHz RSFQ processor with one million gates and their average junction critical current of 140 μ A would have the total power consumption of \sim 0.8 W at 4.5 K.

While no radical execution paradigm shift is required for superconductor processors, several architectural and design challenges need to be addressed in order to exploit these new processing opportunities.

The issues of RSFQ processor design have been addressed in three projects: the Hybrid Technology Multi-Threaded (HTMT) project, the FLUX project in the U.S., and the Superconductor Network Devices project in Japan (Table I).

Time Frame	Project	Target Clock	Target CPU Performance (peak)	Architecture	Design Status
1997-1999	SPELL processors for the HTMT petaflops system (US)	50-60 GHz	\sim 250 GFLOPS/CPU (est.)	64-bit RISC with dual-level multithreading (\sim 120 instructions)	Feasibility study
2000-2002	8-bit FLUX-1 microprocessor prototype (US)	20 GHz	40 billion 8-bit integer operations per second	Ultrapipelined, multi-ALU, dual-operation synchronous long instruction word with bit-streaming (\sim 25 instructions)	Designed, fabricated
2002-2005	8-bit serial CORE1 microprocessor prototypes (Japan)	16-21 GHz local, 1 GHz system	\sim 250 million 8-bit integer operations per second	Non-pipelined, one serial 1-bit ALU, two 8-bit registers, very small memory (7 instructions)	Designed, fabricated, and demonstrated

The key design challenges at the processor design level are:

- **Microarchitecture**
 - pipelining and clocking for 50-100 GHz RSFQ processors;
 - small area reachable in a single cycle;
 - latency avoidance and tolerance.
- **Memory**
 - wire delay-dominated SFQ RAM;
 - hybrid-technology memory hierarchy.
- **Interconnect**
 - high-bandwidth, low-latency system interconnect;
 - multi-temperature, high-speed, low-power interfaces between the cryogenic core and warm electronics.

Most of the architectural and design challenges are not peculiar to superconductor circuitry but, rather, stem from the processor circuit speed itself. At the same time, some of the unique characteristics of the RSFQ logic will certainly influence the microarchitecture for superconductor processors.

- Conclusions and Goals

The Superconducting Technology Assessment (STA) Panel conducted a thorough evaluation of the status of the superconductor technology in 2005. The STA Panel believes it will be possible to find and demonstrate viable solutions for architectural, design, and fabrication challenges during the 2005-2010 time frame.

The proposed program has two major goals for processor design:

- find viable microarchitectural solutions suitable for 50-100 GHz superconductor RSFQ processors;
- design, fabricate, and demonstrate a 50 GHz, 32-bit, 100 GFLOPS, 1-million gate processor with 128 KB, 200-400 GB/s off-chip local memory integrated on a multi-chip module (MCM).

It is also planned to develop a cell library and a set of CAD tools to allow engineers without deep knowledge of physics of superconductivity to design superconductor circuits of such complexity and speed.

Table 2. Summary of the key opportunities, challenges, and projections for superconductor microprocessors		
Superconductor Technology Opportunities	Architectural and Design Challenges	Projections
<ul style="list-style-type: none"> • Ultra-high processing rates • Very low power consumption in RSFQ processors • Ultra-high-speed superconducting transmission lines with negligible attenuation 	<ul style="list-style-type: none"> • Microarchitecture: <ul style="list-style-type: none"> • pipelining and clocking for 50-100 GHz processors; • small area reachable in a single cycle; • latency tolerance • Memory: <ul style="list-style-type: none"> • wire delay-dominated SFQ RAM; • hybrid-technology hierarchy • Interconnect: <ul style="list-style-type: none"> • high-bandwidth, low-latency system interconnection network; • multi-temperature, high-speed, low-power system interfaces between the cryogenic core and warm electronics 	<ul style="list-style-type: none"> • 100 GHz 64/128-bit processors for HEC • Compact multi-petaflops system core with acceptable power consumption

Cryogenic Memories for RSFQ Ultra-High-Speed Processor

T. Van Duzer.

Electrical Engineering and Computer Sciences, University of Calif., Berkeley, CA 94720-1770
vanduzer@eecs.berkeley.edu

The gap between logic speed and memory access is a growing problem in all computing systems and it is exacerbated for ultra-high speed processors such as the proposed cryogenic Rapid Single Flux Quantum (RSFQ) logic working at 50 –100 GHz. The Superconducting Technology Assessment (STA) Panel considered two levels in a hierarchy of cryogenic memory located off of the processor chip. The first level of off-chip memory would be located on the MCM at 4 Kelvin (4 K) with the processor chip in order to minimize propagation-time delays. We are planning for a 1 Mb memory for this stage. The second-level memory would be much larger and could be located on a more efficient refrigerator stage at 40-77 K.

First-Level Off-Chip RAM

Ideas for the critical first-level 4 K off-chip RAM that could be located on an MCM with the processor are:

- hybrid Josephson-CMOS memory
- single-flux-quantum superconducting memory
- superconducting magnetoresistive RAM (MRAM),

These are listed in the order of their states of development. Since the hybrid memory has already been partially demonstrated and it makes use of the highly developed CMOS processes, we discuss it first. The second one is a single-flux-quantum superconducting memory and the degree of success achieved by several previous projects suggests a high probability of successful development. Such a memory could have speed and/or power advantages over the hybrid memory, which requires amplification of the SFQ pulses to volt levels as inputs to the CMOS parts. The third one is the MRAM, which is the subject of R&D in a number of places for room-temperature applications and should be adaptable to 4 K applications, with the advantage that the word and bit lines could be superconductors, thus eliminating one of the main sources of power dissipation in room-temperature MRAMs. Some studies on the 4 K properties of the magnetic storage devices indicate favorable results. The potential density and adaptability to 4 K operation suggests it should be evaluated as one prospect for the first-level RAM. These are summarized in the table below.

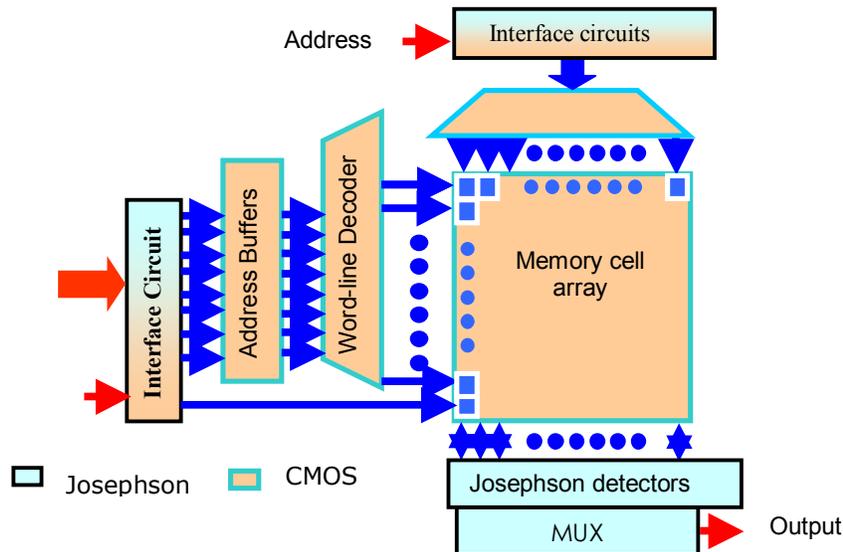
4 Kelvin Off-Chip RAM

Memory type	Readiness for development	Potential density	Potential speed	Potential power dissipation
Hybrid JJ-CMOS	High	High	Medium	Medium
Single Flux Quantum (SFQ)	Medium	Medium	Medium-high	Low-medium
Josephson-MRAM	Low	High	Medium-high	Low-medium

Since the technologies of the three memory concepts are so different from each other, we discuss them separately.

Josephson-CMOS Hybrid RAM

The core of the hybrid Josephson-CMOS RAM is fabricated in a CMOS foundry and can benefit from the existence of a highly developed fabrication process, and the Josephson parts are rather simple. This has the advantage of the high density achievable with CMOS and the speed and low power of Josephson bit-line detection. See the figure below. The entire memory is operated at 4 K so it can serve as the local cryogenic memory for the processor. A 64-kb CMOS memory array made in 0.25 micron CMOS fits in a 2 mm x 2 mm area. As CMOS technology continues to develop, the advances, including density, can be adopted for this hybrid memory. It should be possible to fit a 1 Mb hybrid RAM on a 1 cm² chip. The retention time for charge in a three-transistor DRAM-type memory cell at 4 K is essentially infinite, as has been shown experimentally, so that refreshing is not required; the operation is as though it were an SRAM even though DRAM-type cells are used. If a CMOS fabrication could be made specifically for 4 K operation, the power dissipation could be greatly decreased because of the excellent sub-threshold characteristics of MOS devices at 4 K.



Architecture of the hybrid Josephson-CMOS RAM

Since the hybrid JJ-CMOS RAM has been studied in a university research program for several years, there is a great deal of knowledge derived from extensive simulations and experiments. Computer simulation of Josephson circuits is highly developed and is reliable. A BSIM model (CMOS industry standard at 300 K) has been adapted to 4 K operation and it gives very good agreement with measurements. All components have been simulated for high-speed operation and have been demonstrated experimentally at low speed. According to simulations, the access time for 64 kb should be 500 ps in existing technology, and scaled to 1 Mb it is still subnanosecond. We estimate cycle time to be 300 ps with pipelining. Access time measurements are in progress.

Single-Flux-Quantum Superconducting RAM

A second candidate for the first-level off-chip memory is one that stores single magnetic flux quanta (SFQ) in superconducting loops controlled by Josephson junctions. Such a memory will not require the amplification to volt levels as in the hybrid Josephson-CMOS memory, and this could allow lower power dissipation, and possibly higher speed. More development effort will be required than for the hybrid memory. There have been research projects on several different configurations of SFQ memories. We describe here one that shows a high level of promise.

A 16 kb pipelined SFQ RAM referred to as "CRAM" for cryogenic random access memory design consisting of four 4-kb sub-arrays was estimated to have 400 ps access time and 100 ps cycle time. All components of a 4-kb block were fabricated and tested at low speed. Due to the block-pipeline architecture, the access time and cycle time will scale for a 64-kb RAM as follows: the cycle time is estimated remain the same (~100 ps) with access time somewhat increased to about 600 ps due to an extra decoder. It was projected that with a 20 kA/cm² process, the density would increase, and cycle time would be reduced to 30 ps, with access time in the order to 400 ps. The project was discontinued when work stopped on the HTMT project.

Hybrid Josephson-MRAM

Magnetoresistive random access memory (MRAM) is an alternative memory technology currently under development in the semiconductor industry for high performance, nonvolatile applications. This technology combines a spintronic device with silicon microelectronics to deliver a combination of attributes not found in any other CMOS memory: speed comparable to SRAM, cell size comparable to or smaller than DRAM and inherent nonvolatility independent of operating temperature or device scaling. The memory element has two stable magnetic states measured by a high- or a low-resistance element (bit "1" or "0"), and retains its value without any applied power.

The STA Panel evaluated two types of MRAMs: field switched (FS) tunneling magnetoresistive (TMR) devices and spin momentum transfer (SMT) devices. Both rely on the effect of spin-polarization on the conductivity of a resistive element.

SMT elements are low resistance metals whose magnetoresistive state is set by transferring spin momentum directly from a write current. This current-driven, resistance-based approach provides a unique opportunity to integrate Josephson decoders and read/write circuitry with high speed, high density SMT MRAM cells for cryogenic operation. The details of such a system are under consideration.

Second-Level Cryogenic RAM

We have also considered second-level memories to back up the high-speed first-level memories described above. This memory would be located at a higher-temperature, more efficient stage of the refrigerator. The temperature would be in the 40-77 K range where silicon mobility has its peak value. Two possibilities are:

- CMOS
- MRAM

The purely CMOS memory could take advantage of the extremely low leakage and high mobility existing at cryogenic temperatures. As in the 4 K hybrid, advantage can be taken of the fact that refreshing of the charge in a memory cell is not necessary and compact DRAM-type cells can be used for SRAM-like operation.

For MRAM second-level memory, one could use the field-switched TMR devices. They involve tunnel junctions, read by sensing the resistance of a magnetic tunnel junction with a read-current pulse. The resistance depends on the relative polarization of the magnetic films encasing the tunnel junction. Likewise, the magnetic cell is erased/written with a larger write-current pulse, whose local magnetic field flips the cell polarization and hence bit resistance to the desired state. The control circuits are made in CMOS.

System Balance and Fast Clocks

Burton J. Smith

As clock rates have risen over the years, nearly all aspects of computer implementation from programming model (got caches? got cores?) to component technology have been forced to adapt. The falling cost of transistors has enabled some of this, but does not always help. For example, we have now reached clock rates even in CMOS where skin effect in copper-based transmission lines limits the global bandwidth of large-scale systems so strongly that optical interconnect looks like the only way to retain balance.

Balance is important. Without it, we have systems whose performance is determined solely by their bottlenecks. Amdahl's Law mathematics applies, and most of the investment in the system is fruitless. Also, the more generally applicable the system can be, the more customers it will have and the better the return on both the manufacturer's and the customer's investments. Some of these balance challenges are with us today, but they will become even more numerous and pressing if and as clock rates continue to climb. With disruption of the type under discussion here, the challenges are very severe.

Latency is the most obvious of these challenges, and the faster the clock the worse it becomes even if absolute time-of-flight remains unchanged. Memory latency is the most widely understood problem, but synchronization latency and branch latency are not far behind. Latency tolerance is needed to address these problems because caches, instruction-level parallelism, and branch prediction have already reached or nearly reached the limits of their effectiveness at today's clock rates. Properly implemented, fine-grain multithreading addresses latency in all forms, but few people are familiar with it or its benefits. Nevertheless, it is probably mandatory for systems based on this technology.

Bandwidth is another big challenge. It is coupled to latency by Little's Law, which requires that, in a (conservative) subsystem that transports "things", the product of average latency and average bandwidth equals the average number of things being transported. Because increased bandwidth and increased latency (measured in clocks) demand greatly increased concurrency, some form of parallelism is needed to supply it, which is why multithreading is effective. To generate more concurrency, more processor state is needed and more fast memory (often multi-ported) must be incorporated into the processor. There can be no bottlenecks for concurrency in any subsystem, either, because the composition of subsystems, whether parallel or pipelined, must avoid Amdahl's law in the small as well as in the large. In particular, MPI represents a concurrency bottleneck for global communication whereas the best alternative, shared memory, has been deprecated by most of the HPC community until recently.

The **Temporal Locality** challenge is this: how can data re-use be exploited? The classic consequences of caching remote data as in CC-NUMA systems is poor scaling due to cache miss latency, but this latency can be tolerated; unfortunately, the coherence traffic that results also saps global bandwidth, which is an exceedingly precious resource at large scale. Other techniques need to be explored more fully than they have been; these include how best to exploit streaming locality, what if any remote atomic memory operations are needed, and how to let the compiler direct what data are cached and when. Cache size can probably be reduced at least for a sophisticated multithreading implementation because the statistical averaging that results from

out-of-order thread scheduling means that a few cache misses will not have a very strong impact on performance.

The **Thread Weight** challenge stems from the observation that processor state must grow to manifest the concurrency needed for latency tolerance, and without doing something about this issue the cost of synchronization will increase proportionally. If multithreading is employed for latency tolerance instead of something like vector pipelining, the state per thread remains moderate and synchronization costs need not grow. If temporal locality is to be exploited, multiple threads can be dynamically scheduled and cooperate in their use of shared state so that starting and stopping some threads will not affect performance much. This is largely unexplored architecture and compiler territory.

The **Connectivity** challenge was briefly alluded to previously. It has two aspects. First, long range, high bandwidth connections are more expensive than short slow ones. For copper, skin effect ultimately makes transmission line cost proportional to the cube of the distance for fixed data rate and the square root of the data rate for fixed distance. In addition, interconnections based on exotic materials are always much more expensive than those from conventional technology until the exotic becomes more mundane and the engineering for manufacturability has been done. The cost of optical interconnect is an excellent example.

The **Programmability** challenge has become a colossal problem for nearly all of HPC. Poor programmability has strongly reduced programmer productivity, and discouraged new computational approaches by independent software vendors and by government agencies. It is at least as much a programming language issue as it is an architectural one. It is unclear whether anyone will want to (or even be able to) program systems as parallel as we will need unless something new and different is done.

In summary, many challenges that HPC already faces are exacerbated by very fast clock rates. So far, we have not done very well in addressing these challenges and we probably need to change our ways dramatically even if CMOS somehow proves sufficient for another decade or two. For a technology anything like that considered here, we have no choice but to clean up our act.