

Document Type: Preliminary Proposal

Title: **Adiabatic Logic for High-Bandwidth  
Networking Equipment: A Proposed  
Feasibility Study**

Sponsoring organization: Nortel Networks  
Program: Global External Research

Proposer: Dr. Michael P. Frank (Principal Investigator)  
Proposing organization: University of Florida  
College of Engineering  
Departments of Computer & Information Science & Engineering (CISE)  
and possibly also Electrical & Computer Engineering

Date written: Nov. 5, 2000

## **Abstract**

Due to the approach of various fundamental limiting factors to logic technology scaling, future generations of digital semiconductor technology are anticipated to eventually offer greater potential improvements in raw performance per unit circuit size (and cost) than in performance per unit power, meaning that the performance of many future digital systems, including networking equipment, may become heavily power- and cooling-limited, rather than circuit-size or circuit-cost limited, if, that is, the logic circuits continue to be based on traditional, irreversible logic styles.

Adiabatic, reversible logic techniques offer a new dimension of scalability in energy efficiency, which can reduce a system's power requirement *per unit throughput* to almost arbitrarily low levels, perhaps thousands of times lower than with conventional irreversible logic. The primary drawback of adiabatic techniques is that they require increasing the minimum number, and therefore cost, of a system's logic elements in order to maintain or improve performance, as individual devices are slowed down in order to operate them in an adiabatic mode which reduces the energy dissipated per operation.

Therefore, the degree to which adiabatic techniques can be cost-effective within the context of a specific application depends not only on the raw characteristics of the underlying semiconductor technology, but also on system-level issues, such as the amount of space available within packages and enclosures for additional circuitry, and the relative cost of adiabatic circuits versus (alternatively) more aggressive power and cooling systems.

In this document, we propose a two-year feasibility study, with an option for follow-on research, to take place at UF to determine if and when adiabatic techniques will be cost-effective in the design of systems typical of networking and telecommunications equipment such as is produced by Nortel Networks, and to assess the size of the benefits that the use of adiabatic techniques may confer.

This study will be based on (1) a close examination of the most up-to-date predictions of the low-level performance and power characteristics of future logic device technologies, (2) on adiabatic circuit designs, power supplies, and accurate power/performance scaling models currently under development at UF, and (3) on a detailed, multidisciplinary examination of the system design issues relevant to overall cost-efficiency in the context of typical networking appliances.

# 1. Background

## 1.1. Performance and Power Dissipation in Traditional Irreversible Digital Electronics.

The technology of digital electronics, in networking switches as well as in general purpose computers, is traditionally completely reliant on *irreversible* physical mechanisms for its operation. Irreversible switching dissipates roughly  $1/2 CV^2$  energy per low-level bit-operation, where  $C$  is circuit node capacitance and  $V$  is the logic signal voltage. As a result, in a given semiconductor technology generation, the rate at which bit-operations can be performed in traditional irreversible systems is limited, in a roughly proportional way, to the power consumption (and consequent heat generation) of the system, thus making it difficult to scale the up the performance of systems while respecting constraints on system power and size (since more densely-packed enclosures are more difficult to keep cool).

Historically, the semiconductor industry has enabled continuing increases in system performance per unit power by decreasing the physical size of circuit elements in all dimensions, which decreases node capacitance, while simultaneously enabling the devices to function with lower power supply voltages. Thus over the years, the characteristic bit energy has decreased, roughly in concert with the increased rate at which bit-operations can be performed per unit of die surface area, typically maintaining the power dissipation of typical high-performance logic chips within an order of magnitude of the 10 W level. Meanwhile, improvements in low-power processor architectures have reduced the number of raw bit-operations that are typically performed per unit of useful computational work. So overall, in the past, the power required per unit of die area has not increased as quickly as has performance.

However, for continuing this trend in the future, several problems loom. The most energy-efficient logic architectures are nearing the natural limit of their efficiency - namely, the situation where practically every bit-transition is doing some useful computational work. And, at a lower level, bit energies themselves will eventually approach their fundamental thermodynamic minimum of  $kT \ln 2$  - the energy corresponding to 1 bit's worth of entropy at temperature  $T$  - where  $k$  is Boltzmann's constant [Lan61, MeiDav00]. Lowering the operating temperature can reduce this bit energy, but not the total power requirements, since the entropy removed from the system via cooling ultimately has to be deposited in some external reservoir near room temperature (300 K), with a corresponding  $kT$  energy cost. Additionally, *reliable* bit-operations will require bit energies several times *larger* than  $kT$  to avoid errors due to thermal noise - for example, an error rate of  $10^{-27}$  (i.e., roughly 1 thermal error every 30 years in a billion-logic-gate, high-activity-factor, GHz system) requires a bit energy  $> \ln(10^{27}) kT = 62 kT$  [Fra99a].

## 1.2. Reversible adiabatic logic.

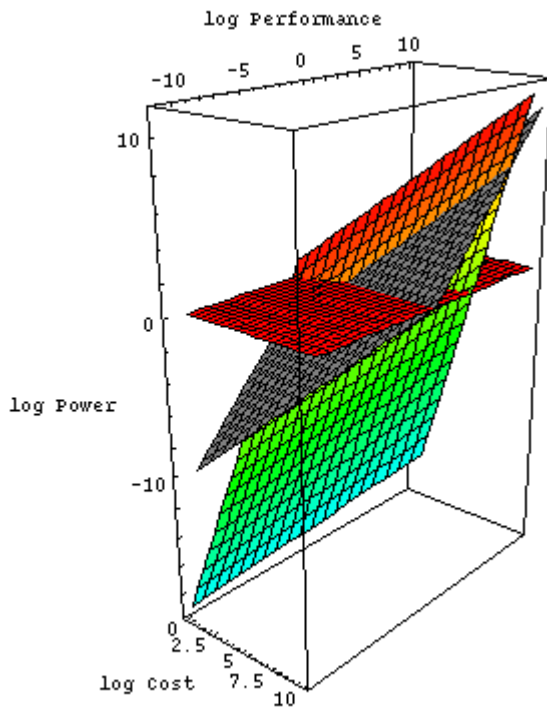
Fortunately, it seems that, in principle at least, there is a computing technique that sidesteps these bit energy limits, and allows performance to scaled up faster than power even within a fixed technology generation. That technique is called *reversible computing* or *adiabatic logic* (among other names), and the essence of it is to abandon the traditional reliance on *irreversible* bit-operations, which dissipate the entire bit-energy of a circuit node every time a new value for the node is computed. Instead, *reversible* operations are used, which execute a smooth, physically *adiabatic* transition from old values to new values. Bit energies of  $1/2 CV^2$  (and above the  $kT$  limit) are still required to be present, but rather than being dissipated on every cycle, these energies are merely transferred back and forth between logic and power supply as needed, mostly recycling their energy, except for frictional losses which can be made almost as small as desired.

Utilizing adiabatic transitions for computing is not a trivial endeavor. It turns out that the complete application of the reversible approach requires changes in the circuit style, logic design, processing architecture, and, in general purpose systems, to the instruction set architecture, programming languages, and software algorithms as well. Fortunately, most of the gains can be realized through changes to only the lowest of these levels, and in any case, the changes needed at all levels are fairly straightforward to apply

when needed, as the author and his colleagues demonstrated in their work at MIT [Fra99a]. Several partially- to fully-adiabatic processors have been demonstrated [Ath+97, Fra+98, Vie99].

The primary drawback to the adiabatic approach is that it requires any given low-level state transition to be performed gradually, relative to the maximum speed that the device's electrical characteristics would allow. This is not as much of a drawback as it may at first seem, because in many real applications, overall performance is actually limited by power and cooling concerns rather than by the electronics itself, so that if, for example, we were to pack a given enclosure full of as many circuits as we could afford, we could not run all those circuits at their top speed *anyway*, due to the resulting power requirement and problem with overheating.

Given a traditional irreversible circuit running at its top speed, if we have to reduce its power by a factor of  $N$ , we can accomplish this by reducing either the clock speed or the number of logic gates by a factor of  $N$ . But, either method applied to traditional circuits generally reduces performance by a factor of  $N$  as well. But, with adiabatic techniques, we can instead reduce the clock speed by a factor of only order  $N^{1/2}$ , reducing performance by only that factor, but reducing power by a full factor of  $N$ , since both the rate of bit-operations, and the frictional loss per bit operation go down by an  $N^{1/2}$  factor. Or, if we have room for extra circuits (and their attendant cost), we can even *increase* the number of gates acting simultaneously by a factor of  $N$ , then reduce the clock speed by a factor of  $N$ , and achieve roughly the same performance as originally, but with  $N$  times lower power. Power scales down with the *square* of clock frequency in adiabatic systems, removing the linearity constraint present in the traditional design-tradeoff space, effectively adding a new dimension to the design space.



**Figure 1.** Cost-performance-power surfaces for conventional (dark gray) and adiabatic (shaded color) computing. The coloring of the shaded plane from red to blue indicates progressively lower power per unit performance. All axes are logarithmic, with an arbitrary base & units. The horizontal plane shows a hypothetical power "ceiling" that might be imposed by application constraints. Note that for conventional systems, the power ceiling translates into a maximum performance that can be achieved below the ceiling, independent of circuit cost. The conventional surface can be reduced to its upper-left edge, where the circuits are running at their maximum frequency and the cost per performance is minimal. But, adiabatic systems permit scaling to higher performance within the power ceiling, if circuit costs somewhat greater than the minimum conventional cost can be tolerated. The figure takes into account the constant-factor overheads that are associated with adiabatic computing, but not the limits to power scaling due to leakage, power supply efficiency, and other limiting factors.

Whether increases in circuit size (compared to a maximum-speed traditional circuit) are feasible or cost-effective depends entirely on the context of the circuit's application within some larger system. For example, how much physical space is available within an IC package or a board enclosure for squeezing in more and more layers of circuitry? (Normally one would not even consider adding more circuits because of heat dissipation concerns, but recall that with adiabatics, by introducing *more* circuitry we can, perhaps somewhat counter-intuitively, *lower* the total system power.) Also, exactly how much additional circuit cost

can we afford, before the cost begins to dominate the total cost of the system? Fortunately, the baseline cost for the electronics shrinks over time, which gives us more room for cost increases. But beyond some point, would it be cheaper to use conventional high-power irreversible electronics, but buy a more aggressive cooling system? These are the issues we will be addressing in this study.

These issues are complex and interacting, and putting them all together, along with accurate projections of the raw characteristics of future logic device technologies (which also bear on the results), and developing an accurate model of the power/performance/cost scaling of highly optimized adiabatic circuit design methodologies, is altogether a substantial research issue. Also, in order for the potential benefits of adiabatic technology to be fully realized, research advances are needed in the area of power supply design; I am currently researching power supply issues in collaboration with my colleagues Khai Ngo, Ken O, and Rob Fox in the Electrical and Computer Engineering department at UF. One avenue being actively pursued is the use of MEMS electromechanical oscillators to provide the high-quality, resonant power supply signal that is needed to drive adiabatic circuits [Fra00a].

### **1.3. High-bandwidth Networking Equipment.**

Vukovic *et al.* at Nortel Networks have estimated that by 2004, the bandwidth requirements for a typical PoP (Point of Presence, *e.g.*, an ISP facility) on the Internet will exceed 100 Tb/s, and that furthermore, if the Moore's Law trends for computational power efficiency are projected forward to this year (and bandwidth figure), total PoP power requirements will approach 1-2 MW, which exceeds by a factor of 20 what is considered to be a more reasonable PoP power level of 50-100 kW, which itself corresponds to an electrical power cost around \$5000/month, at typical present rates around \$0.10/kWh.

The question of interest for this proposal is whether adiabatics can be used to decrease power requirements for this sort of networking system by a factor of 20 (or greater) without substantially increasing total system cost or reducing throughput. The networking application seems to be an especially challenging one for adiabatics, because in high-bandwidth networks, the physical layer will typically *require* very high frequencies to be present at the periphery of the system, where A/D transceivers interface to a raw physical (electronic or optical) signaling medium. A successful adiabatic design for the digital core of the system will need to fan out signals from the relatively serial but high-frequency periphery to a large, massively parallel (but low-frequency) digital core of the system. The low frequencies in the core parts of the system may unacceptably increase data latencies through the system, unless care is taken to parallelize the system design (at all layers) more aggressively than in conventional networking equipment.

### **1.4. Limiting factors for adiabatic power savings.**

Even assuming that the system design can be successfully and thoroughly parallelized to the extent necessary to achieve the desired level of adiabatic power savings, other limiting factors must be taken into consideration before we can know the true extent to which adiabatics can reduce a system's power requirements:

1. Leakage characteristics of the underlying logic device technology. In any ordered system above absolute zero temperature, there is some non-zero rate at which entropy is produced within the system, thereby costing free energy. Transistor circuits exhibit this phenomenon most prominently in the form of thermally-excited leakage of electrons across the potential energy barrier that is presented by a nominally "off" transistor. This leakage or *subthreshold conduction* current depends in magnitude on the height of the threshold barrier, and thus is increasing, as a fraction of on-currents, as transistors shrink and their characteristic voltages decrease in succeeding technology generations [ITRS]. Leakage power sets a lower bound on the reduction in dissipation achievable through adiabatics. The use of variable-threshold transistors may alleviate this problem (*cf.* [Raj00]), but in the long run (>10 years) probably new device mechanisms will be needed. But for this decade, threshold voltages are projected to remain high enough to enable significant power reductions.

2. Power supply efficiency, and scaling of  $Q$  with frequency. Adiabatic circuits require a resonant, energy-recovering AC power supply. Conventional approaches for delivering the needed signal fail to scale power dissipation down as frequency decreases as effectively as power is scaled in the adiabatic logic circuits themselves [Sve95]. Some elements, such as the inductors in LC oscillators, may tend to scale up in size and thus cost as the resonant frequency is decreased. Furthermore, even at a given frequency, the  $Q$  of the primary oscillator in the power supply (which roughly gives the characteristic damping time, in cycles) sets the maximum factor of reduction in energy dissipated per cycle. We are currently investigating the possibility of using micromechanical MEMS oscillators (which have  $Q$ s in the tens of thousands, at reasonable frequencies) as the basis of an adiabatic power supply which would permit significant reductions in power per unit performance [Fra00a].
3. Capability of system design to absorb the extra costs of adiabatic hardware. As we discussed earlier, this depends on the extent to which power dominates cost (or limits performance). We characterize this in terms of a quantity we call the *power premium*, which is the ratio between power-related costs and raw computing hardware costs (per unit of operational lifetime) in a hypothetical non-adiabatic system meeting performance requirements.
4. Fraction of system power that is amenable to adiabatic conversion. Some low-power computing applications, particularly mobile systems, may have components of system power (transmitters, displays, propulsion) that cannot be reduced through adiabatics. In the PoP application, however, most system power is projected to result from data processing, although care should be taken to ensure that other system components such as signal transceivers and disk arrays do not also contribute unreasonable levels of power. (Although adiabatic principles of operation could in theory be applied to those subsystems as well.)

## 2. Proposed Research Program

The focus of this program will be to conduct a detailed feasibility study on the possible use of adiabatic technology to alleviate the power dissipation concerns for future data processing systems of interest to Nortel (such as PoP/networking systems). But since adiabatics is not yet an off-the-shelf technology, some background research and design work needs to be included, in order to facilitate a more realistic and accurate feasibility analysis.

### 2.1. Goals.

Since the present document is but a preliminary proposal, many details of the scope of the proposed work remain to be defined, through further interaction between UF and Nortel. However, some suggested goals for the work include the following options:

1. Improve present designs for adiabatic logic systems and their corresponding power supplies; this design work will enable more accurate assessment of the capabilities of adiabatic designs.
2. Prototype the designs from goal 1, to validate the theoretical analysis of their power/performance characteristics.
3. Complete theoretical work (currently underway) on general analytical models for comparing the cost/power/performance tradeoffs of adiabatic versus conventional data-processing systems.
4. Compile realistic expected parameters for upcoming generations of semiconductor technology and possible competing logic technologies that may arise on the horizon. Similarly for cooling system technologies.
5. With assistance from Nortel systems engineers, compile realistic estimates for the important power-analysis parameters that are specific to the expected future PoP/networking applications of interest to Nortel. Example parameters: Cost parameters, subsystem mix, performance and power requirements.
6. Combine the technology- and application-specific parameters (from goals 4+5) to form an realistic roadmap of the potential effectiveness of adiabatic technologies for solving power dissipation concerns in future systems of interest to Nortel.

## 2.2. Resources

Resources available at UF for this effort include:

- *Proposer's time:* The proposed research is my primary current research area of interest; *e.g.*, I have a career proposal in to NSF to pursue work in this area. I would therefore be happy to commit a large portion (*e.g.* 50%) of my time to this project. Any other funding I might receive overlapping the scope of Nortel's will serve to complement & enhance this research program. I am well qualified to work in this area due to my extensive recent Ph.D. work [Fra99] on reversible computing and my present work on building up a research program in this area at the University of Florida [Fra00b].
- *Proposer's colleagues:* I already have initiated collaborations with several faculty members in the Electrical & Computer Engineering Department at UF to do research in this area, of adiabatic circuits. Those faculty members include experts in VLSI design, power systems design, and semiconductor device modeling. The final version of this proposal may include some of these colleagues if Nortel wishes. There is also an active nanotechnology research collaboration at UF which I am involved with, which impacts this area. And, I am presently collaborating with a large group of colleagues at several universities in the area of future nanoelectronics technology [Raj00].
- *Graduate and undergraduate students:* There are many skilled and well-qualified graduate students in the ECE and CISE departments at UF who would be suited to contribute to this project under research assistantships. Presently, there are many more graduate students who need support than there are positions, so it is a buyer's market. Additionally, there are many bright undergrads at UF who could contribute helpful work inexpensively as research assistants, or even for free for their senior design projects. I have been invited to apply for a courtesy appointment in the Electrical and Computer Engineering department, which will make my research projects more visible to the student labor pool over there (though even without this, I have already recruited some students from that department).
- *Hardware, software & prototyping:* The CISE department provides the usual general computing infrastructure; a small number of additional machines would be helpful for this research. Software needed for this project includes mathematical analysis tools. The CISE department owns licenses to the three major packages: Matlab, Maple, and Mathematica (the latter of which was used to produce figure 1). We may also need use of commercial VLSI design tools, which ECE owns licenses to, though Nortel may wish to contribute to the renewal of those licenses. For prototyping of MEMS-based power supply oscillators, I have contacts with the MEMS fabrication facility at Sandia National Labs [Fra00a].

## 2.3. Schedule and Work Plan Summary

With the idea of a 2-year study in mind, the work for each of the 6 goals in section 2.2 could be quantified as roughly a 2-semester project supported by myself, 1 colleague, and 1 graduate student (or 2 undergraduates). In any given semester (including summer) the team would be working on 2 of the above goals, on average. Each goal would require a 25% time commitment from myself plus at least 10% from a collaborating colleague, for a total of 4 grad student years and 1.4 faculty-years to cover the 6 goals over 2 years. More students would of course be welcome (up to perhaps twice as many) and would allow us to improve the quality and/or thoroughness with which we cover each of the six goals.

Suggested start dates for the project include Jan. 2001 (Spring semester), May 2001 (Summer semester), or September 2001 (Fall semester). However, other start dates could be accommodated as required by Nortel's funding cycle.

Deliverables for the project would include a detailed, final written feasibility study, status reports twice yearly, and write-ups for all the designs and design tools and analytical methods produced in the course of the study.

A full proposal would refine the above schedule and work plan to the level of detail desired by Nortel.

## **2.4. Budget Estimate**

I estimate that in our departments (CISE and ECE) at UF there presently is a rough cost level of \$150K per faculty-year (including benefits and overhead) and \$50K per grad-student-year (including tuition and overhead). For 4 grad-student-years and 1.4 faculty years, this comes to \$410K. I suggest an additional \$100K over the two years to cover prototyping costs, design software licenses, travel to relevant conferences, computer equipment for the project, and the attendant overhead for these costs, for a total of roughly **\$510K** to achieve all six goals.

If Nortel wishes to invest more than that amount on this research, I would estimate that up to another \$500K could be productively utilized for research in this area under my direction based here at UF over 2 years, by means of harnessing more faculty commitment and more student man-months. In the past, I have budgeted a detailed 2-year proposal totaling **\$1M** in this research area [Fra99b]; a proposal with similar scope could be written for Nortel.

Alternatively, the proposal could be de-scoped to produce a less-accurate feasibility estimate more quickly (in 1 year) for about **\$250K**. If the results from that initial study are positive (i.e., that adiabatic computing looks beneficial within a strategic-planning time horizon), a follow-on project would produce more accurate estimates, and begin more detailed design work to support the development of adiabatic technologies.

A full proposal would refine the above (very rough) cost estimates to produce a detailed budget.

## **3. Conclusion.**

This document constitutes an unofficial, exploratory proposal by UF to Nortel to produce a detailed feasibility study for the use of adiabatic circuit technology to provide a "disruptive solution" to the projected problem of power consumption in future networking centers operating at the 100 Tb/s level and beyond. Possible proposal cost levels range from \$250K for a rough 1-year study, to \$500K-\$1M for a detailed 2-year study that includes significant research & prototyping of new design concepts.

This proposal is only a preliminary suggestion, and it is open to modification and refocusing by all parties. The author is open to feedback and may be contacted at [mpf@cise.ufl.edu](mailto:mpf@cise.ufl.edu) or (352) 392-6888.

## **4. References**

- [Ath+97] William C. Athas, N. Tzartzanis, Lars Svensson, and L. Peterson. "A low-power microprocessor based on resonant energy." *IEEE Journal of Solid-State Circuits*, **32**(11):1693-1701, Nov. 1997.
- [Fra+98] Michael P. Frank, Carlin Vieri, M. Josephine Ammer, Nicole Love, Norman H. Margolus, and Thomas F. Knight, Jr. "A scalable reversible computer in silicon." Pages 183-200 of C. S. Calude *et al.*, editors, *Unconventional Models of Computation*, Springer, 1998.
- [Fra99a] Michael P. Frank, *Reversibility for Efficient Computing*, MIT Ph.D. thesis, 1999. Derived manuscript available at <http://www.cise.ufl.edu/~mpf/manuscript>.
- [Fra99b] Michael P. Frank, "Dynamic Optimization of Semi-Adiabatic Power-Managed Architectures," preproposal to DARPA Power Aware Computing and Communications program, 1999. <http://www.cise.ufl.edu/~mpf/proposal.ps>.

- [Fra00a] Michael P. Frank, "MEMS-based resonant power supplies for driving ultra-low-power adiabatic logic circuits," draft abstract for NSF SGER program proposal, <http://www.cise.ufl.edu/~mpf/proposals/mems.html>, 2000.
- [Fra00b] Michael P. Frank, "Mike Frank's Funding Proposals," web page, <http://www.cise.ufl.edu/~mpf/proposals.html>.
- [ITRS] SEMATECH, *International Technology Roadmap for Semiconductors* (1999 Edition), <http://www.sematech.org>.
- [Lan61] Rolf W. Landauer, "Irreversibility and heat generation in the computing process," *IBM Journal of Research and Development*, **5**:183-191, 1961.
- [MeiDav00] James D. Meindl and Jeffrey A. Davis, "The Fundamental Limit on Binary Switching Energy for Terascale Integration (TSI)", *IEEE Journal of Solid-State Circuits*, **35**(10):1515-1516, Oct. 2000.
- [Raj00] Rajendra Singh *et al.*, "Nanoelectronics Science and Engineering Center," Proposal to NSF, Nov. 2000.
- [Sve95] Lars Svensson, "Adiabatic switching," chapter 6 of A. Chandrakasan and R. Brodersen, *Low Power Digital CMOS Design*, Kluwer, 1995.
- [Vie99] Carlin J. Vieri, *Reversible Computer Engineering and Architecture*. MIT Ph.D. Thesis, 1999.
- [Vuk00] Alex Vukovic, John Watkins, Mirjana Vukovic, and Michal Tencer, "Emerging Technologies of Next Generation Networks: Power Density Paradigm," presentation given at 50<sup>th</sup> Electronic Components & Technology Conference, May 21, 2000.