

1. Project Name: **Design & feasibility study for practical adiabatic logic driven by custom high-Q MEMS/NEMS clock/power resonators**

2. **Principal Investigators:** Dr. Michael P. Frank, CISE/ECE Depts.; Dr. Huikai Xie, ECE Dept.

Collaborators: Dr. William R. Eisenstadt, ECE Dept.; Dr. Khai Ngo, ECE Dept.; Dr. Toshikazu Nishida, ECE Dept.; Dr. Zhiyu Pan, Research & Technology Center, Robert Bosch Corp. (subcontractor for MEMS services)

3. **University:** University of Florida

4. **Mailing address:** CSE Bldg., Rm. 301, P. O. Box 116120, Gainesville, FL 32611

5. **Telephone & email:** Primary contact: Michael Frank, (352) 392-6888, mpf@cise.ufl.edu

6. Problem to be addressed. A critical challenge facing the semiconductor industry is to maintain the trend of past improvements in power-performance efficiency until and beyond the end of the present roadmap. The traditional approach to low power, centered around voltage scaling, is facing severe and imminent practical limitations, due to the high sub-threshold conduction across low- V_T devices, and consequent high standby power consumption. Mixed- V_T and power-managed architectures can help alleviate this problem, but only to a limited extent. Even if voltage scaling could itself be continued, by the time the end of the present-day roadmap is reached, ITRS targets have the min-size transistor $\frac{1}{2}CV^2$ gate energies attaining a level at which the quantum information content (number of electron states) in a gate is only a few hundred k_B ($1 k_B \approx 1.44$ bits); see fig. 1. Not far beyond this point, thermal noise may become a significant concern. Even if the noise problem can be dealt with by error correction techniques or improved device isolation, the irrevocable $k_B \ln 2$ limit for bit storage would still apply and be reached by ~ 2035 .

Since the irreversible erasure of any circuit node having a quantum information content of I bits *unavoidably* requires that a quantity of energy $I \times (k_B \ln 2) \times (300 \text{ K})$ must be dissipated into the room-temperature environment, we emphasize that at *no matter what point* the physical information per bit levels off (whether for practically-limited reasons in only a few years, or for more fundamental reasons in a few decades), *further improvements in performance per power* beyond that point will become impossible, *if* we continue to rely on ordinary irreversible logic (that is, logic that destructively modifies its output node value on each gate operation). Therefore, the industry is urged to turn increased attention towards the only physically possible alternative: Reversible Computing.

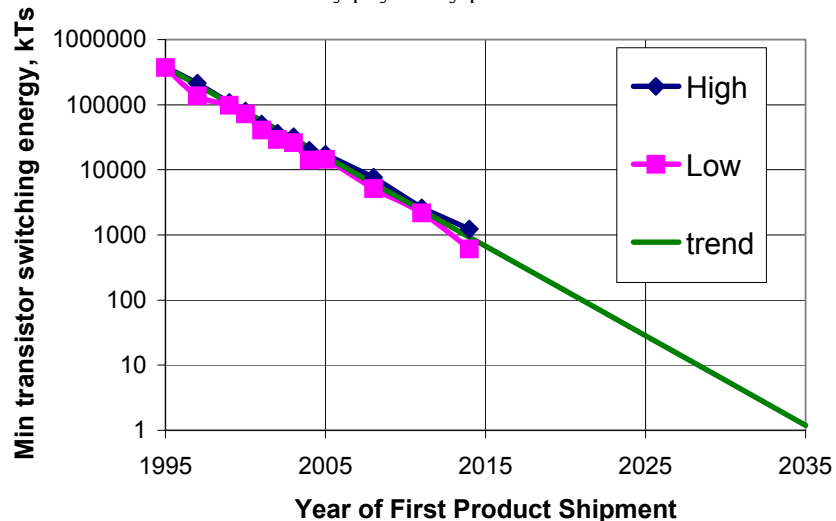


Figure 1. Minimum-sized transistors' $\frac{1}{2}CV^2$ gate energies, in units of room-temperature ($T=300\text{K}$) thermal energy $k_B T$, based on “high-performance” and “low-power” device geometry & supply voltage figures from the 1999 edition of the roadmap. This also roughly measures the physical information per minimum-sized bit, in units of Boltzmann’s constant k_B . (The 2002 update trendline is slightly lower.) The minimum physical information per bit of $0.69 k_B$ would be reached circa 2035, but some higher practically attainable minimum may arrive much sooner.

7. Objective. We aim to carry out detailed design and simulation studies to validate our expectations of the feasibility (in terms of commercially relevant cost-efficiency) of optimized *adiabatic circuits* based on *reversible logic*, which we have previously prototyped in CMOS [1,2], as a viable means to recover an increasingly-large fraction of circuit node energies. We emphasize that in the long run, the use of adiabatic techniques can be rigorously proven to be the *only* way (consistent with fundamental physics) to circumvent the power-performance limits faced by all non-adiabatic technologies. Our previous high-level analytical studies [3] have indicated that, due to the increasing extent to which power dissipation (rather than CV/I gate delay) is becoming the dominant limiting factor on performance, the cost-efficiency of adiabatic solutions is in fact presently increasing relative to non-adiabatic ones. To illustrate the situation, if all of the present trends in device size, energy, cost, *etc.* are

hypothetically extrapolated exponentially until fundamental limits are reached, we obtain the graph shown in Figure 2, below. This projection indicates that adiabatics actually begins winning in cost-efficiency even for a high-performance 100W desktop scenario within the scope of the roadmap, and it potentially achieves several orders of magnitude further cost-efficiency boost over the course of subsequent decades. For more highly power-constrained scenarios, the advantages of adiabatics would be even greater. We emphasize that this analysis already accounts for the nontrivial circuit-complexity and cycle-time overheads that are associated with using adiabatic reversible logic.

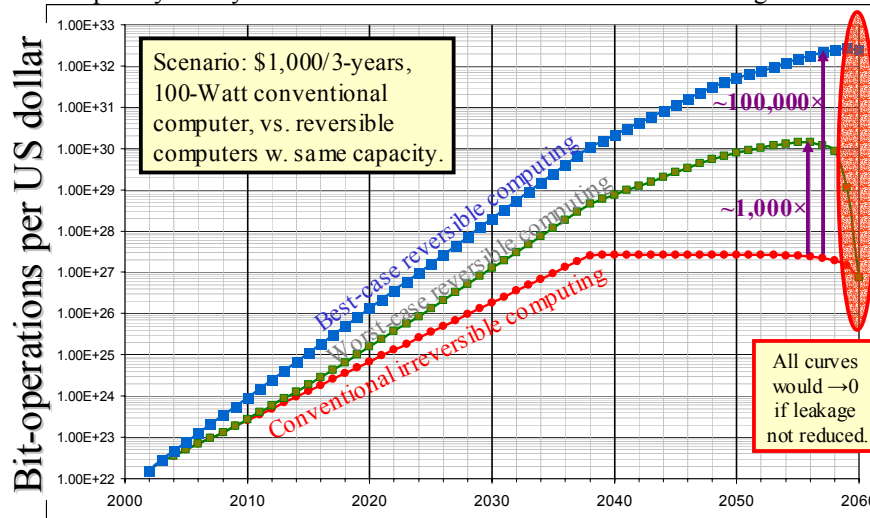


Figure 2. Results from a detailed numerical model [3] of cost-efficiency of optimized reversible (adiabatic) versus irreversible (conventional) computers in future generations of technology. The cost-efficiency of irreversible computing eventually hits a technology-independent thermodynamic brick wall, and cannot improve further so long as the cost of energy (or the heat flux limit of the cooling technology) is fixed. In contrast, the cost-efficiency of reversible computing can continue to improve far beyond this point, limited only by achievable minimum energy leakage rates, which have no known fundamental lower limit. (However, for generating this particular graph, an arbitrary minimum leakage rate of $1 k_B T/\text{ms}/\text{device}$ was assumed.) Note that the advantages of reversible computing are projected to rise to as much 1,000-100,000 \times by the 2050s. This model even takes into account the algorithmic overheads of reversibility, and the proportionality of energy dissipation to speed in adiabatic processes.

8. Novelty. Adiabatic circuits by themselves are not particularly new—quasi-adiabatic logic styles were first investigated by Boyd Watkins at Philco-Ford as early as 1967 [4], as well as in the late 70’s by Fredkin and Toffoli at MIT [5] and in the 80’s by Seitz at Caltech [6]. In the early 90’s, the circuit concepts were further refined, and the connections with reversible computing theory were solidified, by groups at several places including ISI [7], PI Frank’s old group at MIT [8], and IBM [9], with a rush of papers by numerous groups following since then. However, the deal-killer (from a commercial perspective) of all of the adiabatic circuit concepts presented to date has been the energy dissipation in the clock/power supply. Adiabatic circuits require a quasi-trapezoidal voltage waveform. This signal is used for clocking, power delivery, and energy recovery, and it must be generated resonantly with high Q and a tightly-constrained wave shape to avoid severely limiting the overall system-level power savings that can be attained through adiabatic design. Unfortunately, conventional all-electronic resonators based on LC circuits—particularly if clamping power MOSFETs or low coil-count integrated inductors are used—suffer from a low Q , typically only in the tens, a level that prevents high enough energy efficiency to compensate for the significant circuit complexity overheads that are imposed by the use of a mostly-adiabatic logic design style.

In contrast, state-of-the-art MEMS/NEMS technology offers Q ’s in the thousands and even tens of thousands, with frequencies that have recently been extended even up into the GHz range required for high-performance computing [10]. Due to these characteristics, MEMS has recently even become commercially viable as an element in analog RF circuits for communications.

Inspired by this situation, we propose to use custom-designed MEMS resonators to solve the power-supply problem of adiabatic circuits. Simple off-the-shelf MEMS components will not work, because of adiabatics’ special requirement for a trapezoidal voltage waveform, rather than the usual sinusoidal oscillator curve. However, a carefully-tailored custom resonator geometry ought to be able to produce the exact wave shape that we need.

Therefore, this project requires cooperation between experts in all of these areas: (1) mechanical design (Xie), (2) MEMS processes (Xie/Robert Bosch Corp.), (3) VLSI design (Eisenstadt/Ngo), (4) adiabatic circuits & reversible computing (Frank), and (5) systems engineering (Frank).

To our knowledge, this particular concept, of using MEMS oscillators to remove the remaining barrier to the practicality of adiabatics, has never been previously explored. Yet, based on our preliminary analyses, it appears to be a near-ideal silver bullet that will finally transform adiabatics into a cost-effective approach for enhancing

digital system power-performance well beyond conventional limits. Confirming this expectation with a detailed design study would be highly worthwhile endeavor.

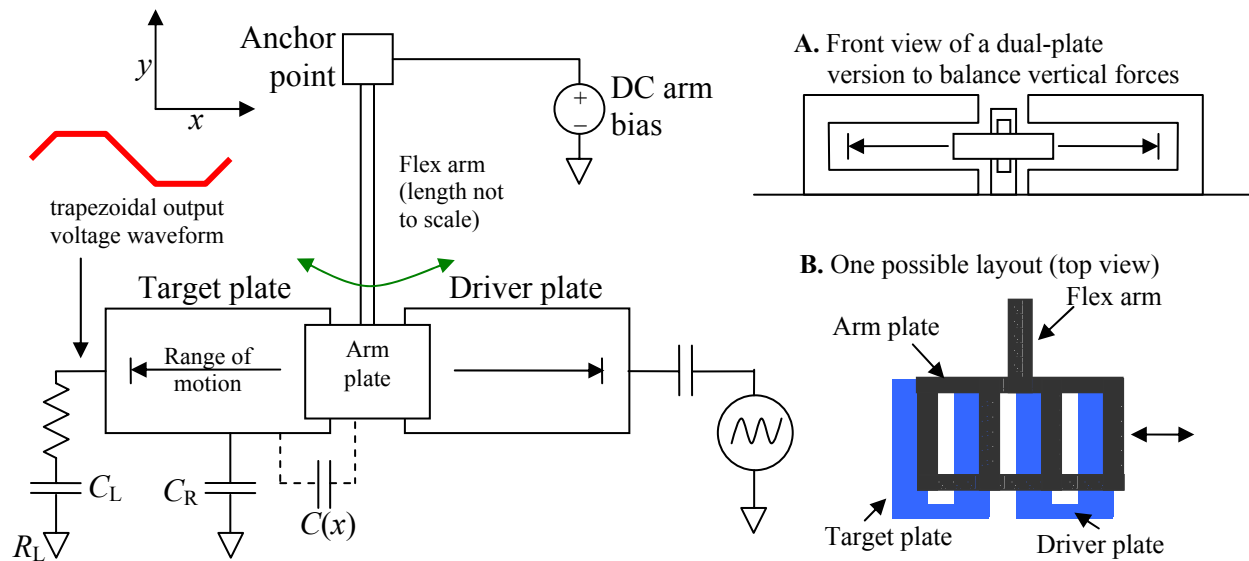


Figure 3. MEMS concept for adiabatic clock/power supply. An electromechanical resonator will be designed with a custom geometry that will yield the desired waveform shape. C_L is the input load gate capacitance of a logic circuit. C_R and $C(x)$ are respectively the reference capacitance and variable capacitance, and form a voltage divider, where x is the horizontal position of the center of the arm plate. The output voltage waveform is determined by $C(x)$. The arm plate has full or zero overlap with the target plate when x falls in certain ranges of values, while $C(x)$ changes linearly (to first order) between these ranges. Thus roughly trapezoidal output voltage waveforms are generated. Further tuning of the electrode shape will be used to compensate for fringe capacitances and more closely approximate the desired waveform. High Q is essential for reducing the drive voltage as well as minimizing the energy dissipation. The horizontal plate electrodes may be replaced by interdigitated comb fingers (vertical fins). In addition to cantilever beams, fixed-fixed or free-free flexures may also be used. A free-free beam MEMS polysilicon resonator with resonance frequency of 90 MHz and Q of 8,400 has been demonstrated [11]. Even higher Q and resonance frequency can be expected by using single-crystal silicon (SCS) microstructures. In this project, we will use SCS resonators packaged in vacuum for high Q . Our fabrication process is also CMOS-compatible, so that signal attenuation and noise resulting from chip-to-chip wire bonding can be avoided. The resonant frequency of a microstructure is given by $\omega_r = \sqrt{k/m}$, where k is the stiffness and m is the proof mass. Thus, high resonance frequency can be achieved either by reducing the device size (smaller m) or increasing the stiffness (larger k). Smaller size leads to smaller $C(x)$ and smaller electrostatic drive force, while stiffer structures will require larger drive force or larger chip area. Careful design and simulation is needed to optimize the tradeoffs. In order to maximize the capacitance changes in a given area, a structural design such as shown in inset B of Figure 3 will be used, where the arm plate, target plate and drive plate are made of beam networks. The beam width of the arm plate is smaller than that of the target and drive plates. Thus, the beams of the arm plate will be completely overlapped with those of the target plate, resulting in a capacitance plateau which in turn can generate a trapezoidal voltage waveform.

For a rough estimation, if we consider a 10 MHz resonator with a Q of 10,000, vibration amplitude of about $3 \mu\text{m}$ can be achieved in a simple design with 5V drive voltage and an area of about $100\mu\text{m}$ by $50\mu\text{m}$. The corresponding maximum $C(x)$ is about 10 fF. We will also explore alternative designs offering higher capacitance-per-area. We will investigate methods for scalably synchronizing arrays of resonators, including techniques such as mechanical coupling, as well as phase-locked loop feedback circuits to tune the resonant frequencies electrostatically.

Finally, we know that the adiabatic approach to low power in FETs has long-term viability. Figure 4 below shows the result of an analysis we have done indicating that FET-based adiabatics remains viable at the nanoscale, and can achieve arbitrarily small, sub- kT dissipation as the Q of the resonator is increased [12].

9. Approach. As a proof of concept, we propose to carry out the detailed design optimization, analysis, and simulation of a prototype digital system using our MEMS-powered adiabatics approach. Below is an outline of steps to be taken in our present research plan.

- 1. Choose a target application.** Candidates currently being considered include a simple DSP processor core that would be suited for low-power mobile and space applications. A wide range of total application power constraints will also be considered ($1 \mu\text{W}$ to 100 W), in order-of-magnitude increments.
- 2. Choose a process technology model.** We will model presently commercially available VLSI/MEMS process technologies, as well as projected technology nodes located 5, 10, and 15 years out.
- 3. Design the high-level digital architecture.** This includes breaking down the core DSP design into functional units of roughly similar size whose internal design will be mostly adiabatic.

4. **Global system optimization.** Analytical methods and numerical tools that we have already developed and used extensively in previous studies will be used to simultaneously optimize the following parameters so as to attain maximum estimated system-level cost-efficiency (ops/dollar, given device lifetime) for our design, for each technology generation studied:

- Degree of reversibility in the digital logic (reversible ops/irreversible op). This determines the circuit-complexity overhead imposed by the adiabaticity constraint, according to known relations.
- Power supply voltage.
- Clock frequency of the adiabatic logic.
- Geometric & architectural parameters of MEMS resonant elements (*e.g.* area- and gates-per-resonator).

The optimization analysis will also yield accurate estimates of the performance (ops/sec) and cost-performance (performance/dollar) for correspondingly optimized adiabatic and non-adiabatic designs for each power level and technology generation studied.

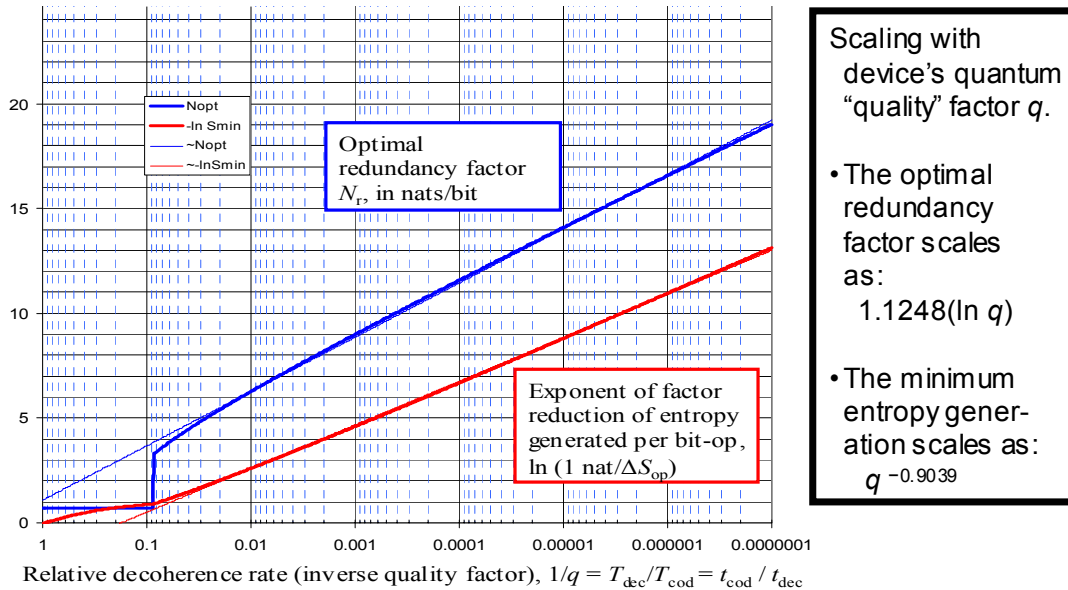


Figure 4. Adiabatic FETs can beat the kT limit. Results from a recent analysis [12] showing that even when accounting for thermally-activated leakage, adiabatic logic based on field-effect devices can still scale to levels of energy dissipation per operation that are arbitrarily small, limited only by the Q of the coupled logic+resonator system (related to the local quantum decoherence rate). In order to suppress thermally-activated leakage, it turns out that the device must (perhaps surprisingly) actually be made somewhat *larger* (not smaller), at least in terms of physical information content, as the required energy dissipation decreases. Fortunately, the size increases only by a logarithmic factor (related to the exponential falloff of leakage currents with voltage) and ends up being overwhelmed by the benefit from the reduced leakage. The significance of this result is that adiabatic nanoscale CMOS (*e.g.* perhaps based on semiconducting carbon nanotubes or crystalline nanowires), coupled with motion-quantized NEMS resonators, can potentially be a viable method to achieve cost-effective computation with significantly less than $k_B T$ energy dissipation per logic operation. There is no definite upper limit to the Q 's that might be achieved in well-engineered quantum systems, and so there is still no known fundamental lower limit on energy dissipation per operation. A future of unbounded power-performance improvements is thus, in principle, theoretically possible, as far as we know for certain.

5. **Design scenario selection.** From the results of (4), one or more specific scenarios (each specified by a total power constraint, a target technology generation, and a choice of figure of merit) will be selected for further study. The scenario(s) selected will be those that most realistically illustrate the commercial viability of our approach. Based on past analyses such as [3], we are presently highly confident that this power/technology regime will indeed include scenarios in which our approach will be highly advantageous, in terms of one or more commercially relevant figures of merit.

6. **Detailed design.** For each of the selected scenario(s) (in order of decreasing relevance), we will flesh out our architecture into a detailed design. This includes several subtasks.

- Standard VLSI tools available at ECE's facility, such as Cadence, will be used, augmented by our own custom adiabatic design tools, to flesh out the CMOS adiabatic logic circuits and simulate them in SPICE (at first with idealized power/clock waveforms) to validate the expected level of dissipation for the partially-adiabatic logic.
- Meanwhile, available MEMS design and simulation tools including Coventorware and NODAS will be used to create custom resonator designs that produce the desired quality and waveform shape when coupled to a lumped RC model of the expected CMOS load.

c. Finally, the designs from (a) and (b) will be integrated and simulated together to verify the expected low power dissipation of the overall design.

7. **Prototype proposal.** Assuming that our detailed simulation study confirms that (as we expect) our approach will be cost-effective for a near-term low-power product, we will proceed to write a proposal to SRC or other appropriate organization to obtain funding to fabricate a prototype of our design in a commercially available process. We are consulting with Dr. Zhiyu Pan, our collaborator at Robert Bosch Corp., a company that offers a CMOS-MEMS integrated fabrication process, which will provide process information and cost estimates for research prototypes and an eventual commercial rollout of this product concept.

10. Research output. The output of our work would include the following:

- Detailed scenarios including both near-term low-power applications and long-term high-performance applications in which it can be shown that adiabatics using our approach is commercially viable.
- Examples of optimized adiabatic designs (using the new design styles we have developed here at UF) for important, widely-used digital elements such as DSP functional units.
- A detailed design of a MEMS resonator suitable for driving adiabatic circuits.
- Projections on scaling this approach to the nanoscale, at which quantum effects become more important.

11. Executive Summary. Continuing to improve computing power-performance above and beyond the roadmap is a major challenge. In any device technology, the quantum-physical information per encoded bit gives a lower bound on the free-energy loss when that bit is erased. This physical information can never be less than $k_B \ln 2$, and may level off at even a much higher value due to practical concerns with leakage and thermal noise, which limit e.g. voltage scaling. It is today increasingly widely accepted that adiabatic techniques using reversible logic offer the *only* potential way to continue improving performance-per-power once physical information per bit levels out. All previous approaches to adiabatic logic have been impractical, since a high- Q clock/power technology was lacking, and a proper system design optimization was not done. We propose a novel approach of using shape-tailored high- Q MEMS resonators to yield a near-ideal voltage waveform for driving adiabatic CMOS. We will do analysis, detailed design and simulation of a sample DSP implementation which will prove that our approach can be commercially viable within and beyond the timeframe of the roadmap. Our team is highly interdisciplinary and very highly qualified to carry out this work. Our work will lead to a subsequent product-development proposal, if near-term feasibility is indicated.

-
- 1 Michael P. Frank, Carlin Vieri, M. Josephine Ammer, Nicole Love, Norman H. Margolus, and Thomas F. Knight, Jr., "A Scalable Reversible Computer in Silicon," in Calude, Casti, and Dineen, eds., *Unconventional Models of Computation*, Springer, 1998, pp. 183-200.
 - 2 Michael P. Frank, *Reversibility for Efficient Computing*, Ph.D. Thesis, MIT EECS dept./AI lab, June 1999. <http://www.cise.ufl.edu/~mpf/manuscript>.
 - 3 Michael P. Frank, "Nanocomputer Systems Engineering," proceedings of the 2003 Nanotechnology Conference and Trade Show, held Feb. 23-27, 2003, San Francisco, CA. <http://www.cise.ufl.edu/research/revcomp/theory/NanoTech2003/Frank-NanoTech2003.doc>.
 - 4 Boyd G. Watkins, "A low-power multiphase circuit technique," *IEEE JSSC*, pp. 213-220, Dec. 1967.
 - 5 Edward F. Fredkin and Tommaso Toffoli, "Design principles for achieving high-performance submicron digital technologies," Proposal to DARPA, MIT Information Mechanics group, Nov. 1978.
 - 6 Charles L. Seitz, Alexander H. Frey, Sven Mattison, Steve D. Rabin, Don A. Speck, and Jan L. A. van de Snepscheut. "Hot-clock nMOS," in Henry Fuchs, ed., *1985 Chapel Hill Conference on Very Large Scale Integration*, pp. 1-17, Computer Science Press, 1985.
 - 7 J. G. Koller and William C. Athas, "Adiabatic switching, low energy computing, and the physics of storing and erasing information," in *PhysComp '92: Proceedings of the Workshop on Physics and Computation, Oct. 2-4, 1992, Dallas*, pp. 267-270. IEEE Computer Society Press, 1992.
 - 8 Saed G. Younis, *Asymptotically Zero Energy Computing Using Split-Level Charge Recovery Logic*, Ph.D. thesis, MIT EECS dept./AI lab, 1994.
 - 9 Paul M. Solomon and David J. Frank, "Power measurements of adiabatic circuits by thermoelectric technique," in *Symposium on Low Power Electronics*, pp. 18-19, 1995.
 - 10 Jing Wang, Zeying Ren, C.T.-C. Nguyen, "Self-aligned 1.14-GHz vibrating radial-mode disk resonators," 12th International Conference on Solid-State Sensors, Actuators and Microsystems (TRANSDUCERS'03), Boston, MA, June 9-12, 2003, pp. 947-950.
 - 11 Kun Wang, Ark-Chew Wong, C.T.-C. Nguyen, "VHF free-free beam high- Q micromechanical resonators," *Journal of Microelectromechanical Systems*, 9(3):347-360, Sep. 2000.
 - 12 Michael P. Frank, "Scaling of Energy Efficiency with Decoherence Rate in Closed, Self-Timed Reversible Computing," UF Reversible Computing Project Memo #M18, Nov. 2002, <http://www.cise.ufl.edu/research/revcomp/memos/Memo18-Timing.doc>.