

# Charge-Recovery Computing on Silicon

Suhwan Kim, *Member, IEEE*, Conrad H. Ziesler, *Member, IEEE*, and  
Marios C. Papaefthymiou, *Senior Member, IEEE*

**Abstract**—Three decades ago, theoretical physicists suggested that the controlled recovery of charges could result in electronic circuitry whose power dissipation approaches thermodynamic limits, growing at a significantly slower pace than the  $fCV^2$  rate for CMOS switching power. Early engineering research in this field, which became generally known as adiabatic computing, focused on the asymptotic energetics of computation, exploring VLSI designs that use reversible logic and adiabatic switching to preserve information and achieve nearly zero power dissipation as operating frequencies approach zero. Recent advances in CMOS VLSI design have taken us to real working chips that rely on controlled charge recovery to operate at substantially lower power dissipation levels than their conventional counterparts. Although their origins can be traced back to the early adiabatic circuits, these charge-recovering systems approach energy recycling from a more practical angle, shedding reversibility to achieve operating frequencies in the hundreds of MHz with relatively low overhead. Among other charge-recovery designs, researchers have demonstrated micro-controllers, standard-cell ASICs, SRAMs, LCD panel drivers, I/O drivers, and multi-GHz clock networks. In this paper, we present an overview of the field and focus on two chip designs that highlight some of the promising charge recovering techniques in practice.

**Index Terms**—Energy-recovering circuits, adiabatic computing, reversible logic, resonant systems, energy efficient computing, voltage scaling.

## 1 INTRODUCTION

ENERGY efficiency has become a major design concern in high-performance and mobile computer systems. Excessive power dissipation requires increasingly large, heavy, expensive, and noisy cooling machinery including special packages, heat sinks, heat pipes, and fans. Excessive energy consumption on mobile computer systems results in increasingly large, heavy, and expensive batteries, power-conversion circuits, or fuel-cells, which themselves may introduce further heat removal issues. Several effective power management design techniques have been developed over the past few years, including lowering the supply voltage. As process scaling continues below 90nm, however, it becomes more difficult to scale the supply voltage for several reasons.

To maintain high transistor drive current and thus achieve performance improvements, transistor thresholds must be scaled along with the supply voltage. However, threshold voltage scaling results in a substantial increase in subthreshold leakage current [1]. Furthermore, the uncertainty related to variations in process, voltage, and temperature has the effect of reducing the available range over which supply voltage may be scaled. As a result, there is a demand on novel circuits whose power saving mechanism is not heavily dependent on further supply voltage scaling.

Long before power consumption became a high-priority objective in computer system design, theoretical physicists had been exploring the fundamental connections of computation and power dissipation. The somewhat astonishing result of these early investigations into the energetics of computation is that the minimum energy requirements of a computation are proportional to the number of information bits destroyed during its course. Thus, if a computation could be somehow implemented without loss of information, its energy requirements could potentially be reduced to zero. Bennett and Landauer at IBM were able to show, in theory, that, by performing computations in a *reversible* manner, no information is destroyed and thus potentially zero energy would be needed [2]. Further work demonstrated concrete transformations that can map ordinary computations into reversible computations.

The idea of a reversible computation is straightforward. A system is reversible if no information about its state is lost at any time during its transformation. For example, if  $2 + 2$  is computed, a reversible computing system would pass on 4 as the answer, exactly as a standard one would. Furthermore, a reversible computing system would also save at least one of the operands, so it could reverse the computation. If it didn't, then it would not necessarily know if the two operands were 1 and 3, 2 and 2, or 0 and 4, resulting in loss of information and unrecoverable energy costs.

While reversibility is required for zero energy consumption, it is by no means sufficient. Since the mere transfer of a charge across a voltage difference is the result of energy exchange, some circuit embodiment is still needed in addition to reversibility to actually compute with zero dissipation. The set of circuit design techniques targeted at the implementation of computations with minimal (asymptotically zero) power consumption during charge transfer is

- S. Kim is with the Inter-University Semiconductor Research Center, Seoul National University, Seoul 151-744, Korea. E-mail: suhwan@ee.snu.ac.kr.
- C.H. Ziesler is with the Multigig, Inc., Scotts Valley, CA 95066. E-mail: conrad.ziesler@multigig.com.
- M.C. Papaefthymiou is with the Advanced Computer Architecture Laboratory, University of Michigan, Ann Arbor, MI 48109. E-mail: marios@eecs.umich.edu.

Manuscript received 4 Feb. 2004; revised 5 Oct. 2004; accepted 19 Nov. 2004; published online 15 Apr. 2005.

For information on obtaining reprints of this article, please send e-mail to: [tc@computer.org](mailto:tc@computer.org), and reference IEEECS Log Number TCSI-0040-0204.




Fig. 1. Charging an RC tree with a switch when  $V_{dd} = 2V_A$ .

generally known as *adiabatic switching* or *adiabatic charging*. The use of the word adiabatic (Greek for “impassable”) is suggestive of the thermodynamic principle of state change with no loss of gain or heat.

The principle of adiabatic switching can be best explained by contrasting it with conventional dissipative switching. Fig. 1 shows how energy is dissipated during a switching transition in static CMOS circuits. The transition of a circuit node from LOW to HIGH can be modeled as charging an RC tree through a switch, where  $C$  is the capacitance of the node and  $R$  is the resistance of the switch and interconnect. When the switch is closed, a high voltage ( $2V_A$ ) is applied across  $R$  and current starts flowing suddenly through  $R$ . After a short period of time,  $C$  is charged to a constant supply voltage  $V_{dd}$ . The energy taken from the power supply is  $CV_{dd}^2$ , but only half of that,  $1/2CV_{dd}^2$ , is stored in  $C$ . The other half is dissipated in  $R$ .

Now, consider the circuit and current waveform shown in Fig. 2. Notice that, in contrast to conventional charging, the transition has been slowed down by using a time-varying voltage source instead of a fixed supply. By spreading out the charge transfer more evenly over the entire time available, peak current is greatly reduced. Consequently, the overall energy dissipated in the transition has been reduced to being proportional to

$$\frac{RC}{T_s} CV_{dd}^2,$$

where  $R$  is the effective resistance of the driver device,  $C$  is the capacitance to be switched,  $T_s$  is the time over which the switching occurs, and  $V_{dd}$  is the voltage to be switched across. The constant of proportionality is related to the exact shape of the time-varying voltage source waveform and can be calculated by direct integration.

Ideally then, by increasing the time  $T_s$  over which computation is performed and by using reversible logic to avoid the destruction of information, it should be possible to create a circuit which computes with vanishingly low energy dissipation as the time allowed for that computation




Fig. 2. Adiabatic charging of an RC tree when  $V_{PC} = V_A[\sin(\omega t + 3\pi/2) + 1]$ .

extends indefinitely. Known in the field as “asymptotically zero energy consumption,” this possibility might have remained a mere theoretical curiosity had not a dedicated community of researchers worked on creating first theoretical and later practical circuit implementations of logic and state elements. These circuit implementations applied some of the principles of reversible computing and adiabatic charging to achieve low, but nonzero, dissipation for computations performed over fixed amounts of time. Over time, the strict requirements of reversibility were dropped, giving way to engineering compromises that have led to practical systems. Because some of the energy in these circuits (in the form of charge stored on capacitances) was being recovered instead of dissipated, the terms *charge recovery* or *energy recycling* began to be used to describe these circuits. Broadly speaking, the term charge recovery is nowadays being used to describe systems that reclaim some of the energy that is stored in their capacitors during a computation and reuse it on subsequent computations. These systems are not necessarily reversible.

This paper highlights some of the more practical charge-recovering circuits and prototype chips for which results have been published. It then focuses on two such chips that have been developed in the Advanced Computer Architecture Laboratory at the University of Michigan. The first chip is a dynamic multiplier designed in a source-coupled adiabatic dynamic logic [3], [4]. The second chip is a resonant-clocked ASIC for the Discrete Wavelet Transform that has been synthesized using industry-standard tools [5]. Neither of the two chips performs reversible computations. Both chips operate with a true single-phase clocking scheme, relying on on-chip circuitry to generate a single power-clock waveform of sinusoidal shape. With measured power savings at clock rates in the 100-300MHz range, these prototypes provide tangible evidence that charge recovery approaches can yield highly energy-efficient designs in practice. They also provide experimental evidence revealing

another important and largely unexplored advantage of charge-recovery circuitry, namely, their low electromagnetic interference.

The remainder of this paper has four sections. Section 2 discusses a number of key challenges in the design of charge-recovery VLSI systems. Section 3 describes some of the pioneering work in the field that resulted in charge-recovery circuit topologies with multiple charge-recovering clock phases. The two single-phase designs from Michigan are described in Section 4. Section 5 summarizes the paper and concludes with directions for further research.

## 2 KEY CHALLENGES

There are two central challenges to the design of efficient energy recovering VLSI systems. First, the circuit implementations of the time-varying power sources should be highly efficient. These so-called *power-clock generators* are typically implemented as *LC tanks*, with *L* and *C* provided in part by design and by the intrinsic characteristics of the circuit. A key requirement for power-clock generators is the ability to transfer energy bidirectionally to and from the energy tank and the power-clock node without dissipating much of that energy in the process. This bidirectional charge transfer can be accomplished with high efficiency through the resonant power-clock waveform that is generated when the *LC* tank oscillates either freely or under a periodic signal. Since many VLSI systems have high throughput requirements, power-clock generators should be capable of maintaining high-frequency power-clock waveforms, balancing speed with the need to keep the transitions of the power-clock gradual so that as much energy as possible is recovered. Depending on the granularity of charge recovery, power-clock signals may need to be routed to nearly every logic gate in the system. To ensure correct operation and efficient charge recovery, precise timing relationships must be maintained on these global high-speed signals. This picture is further complicated if multiple power-clock phases are used, especially if the logic creates data-dependent variations in the current of the power-clock signal.

The second main challenge is that computations should be implemented by low-overhead circuit structures that use standard MOSFET devices. These circuit topologies must allow for some of the energy stored within them to be recovered, typically using a shared time-varying power source or sources. To provide maximum efficiency of recovery, they should be designed so that they present as balanced a load to the power-clock generator as possible. Furthermore, to provide further power savings to evolutionary advances, rather than competing with them, they should be compatible with conventional power management approaches.

Despite these challenges, the untapped potential for very low dissipation has retained interest in the field, leading to important advances over time. Early work pursued reversible designs such as the circuit families SCRL (Split-level Charge-Recovery Logic) [6] and RERL (Reversible Energy

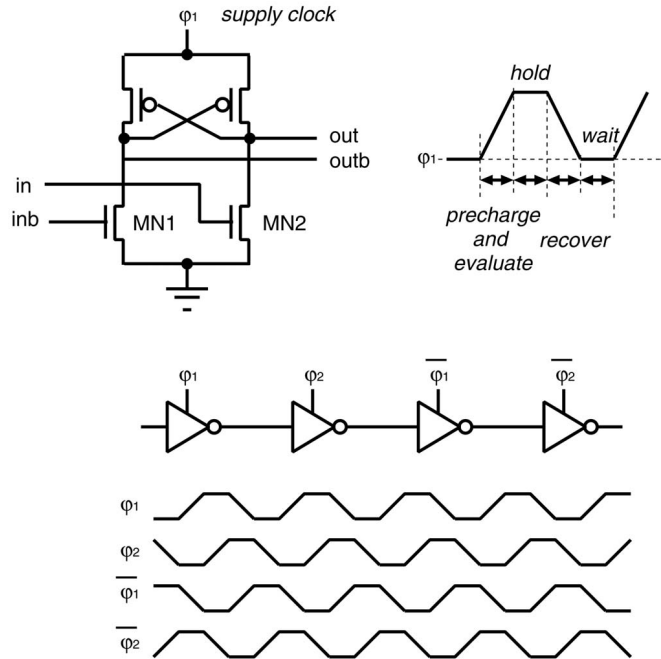


Fig. 3. Basic structure of ECRL circuit.

Recovery Logic) [7] and the Pendulum reversible architecture [8]. The main idea in these investigations was to create a “backward” version of each original “forward” circuit  $F$  to compute the inverse function  $F^{-1}$ . After the output node of each stage in the original circuit is gradually driven to a high or low voltage, it is handed off to the reversing circuitry, which gradually restores its voltage to its original value. Thus, in the forward circuit, charge moves toward the end of the pipeline, while, in the reverse circuit, charge is recycled back to the beginning. This early work represents important pioneering advances in the field and, at the same time, exposes the two main challenges in charge-recovery design.

## 3 ENERGY RECOVERY WITH MULTIPLE-PHASE POWER-CLOCK

With an eye on reducing complexity, several researchers turned their attention to nonreversible designs with non-zero asymptotic dissipation [9]. Early work in this area focused primarily on adiabatic switching and dynamic topologies, resulting in a variety of high-performance circuit families, such as 2N-2N2D, 2N-2P, 2N-2N2P, ECRL, PAL, and CAL, that operate with multiple power-clock phases [10], [11], [12], [13], [14]. Some of these families have been demonstrated to work on CMOS chips fabricated in standard CMOS processes.

A representative member of these multiple-phase charge-recovering families is the Efficient Charge Recovery Logic (ECRL) from [12], shown in Fig. 3. ECRL circuits require four power-clock phases for correct operation. In contrast to previous dynamic energy-recovering families in which energy is delivered during precharging and recovered during evaluation [10], [11], ECRL performs precharge




Fig. 4. Energy-recovery (E-R) latch and a sketch of its signal timing when  $V_{in}$  is high.

and evaluation simultaneously. ECRL was used to build nontrivial functions such as a 16-bit carry-lookahead adder in an energy-recovery circuit style, but no silicon results have been reported to date.

Pass-Transistor Adiabatic Logic (PAL) and Clocked CMOS Adiabatic Logic (CAL) are variants of ECRL [13], [14]. PAL is a dual-rail logic that operates with two power-clock phases and uses pass-transistor NMOS functional blocks. A 1,600-stage PAL shift register fabricated in  $1.2\ \mu\text{m}$  CMOS technology has been experimentally verified at 10 MHz. CAL is a dual-rail logic that operates with a single-phase power-clock and two auxiliary clock waveforms. A chain of 736 CAL inverters has been fabricated in  $1.2\ \mu\text{m}$  CMOS technology, achieving a maximum operating frequency of 50 MHz when driven by a sinusoidal power-clock that was supplied from external functional generators.

The architecture and circuit implementation of two energy-recovering microprocessors, AC-1 and MD1, have been described in [15], [16]. These microprocessors apply adiabatic charging on two different generations of clock-powered logic and have been fabricated in a  $0.5\ \mu\text{m}$  n-well and a  $0.6\ \mu\text{m}$  n-well CMOS chip, respectively. Measurements of the AC-1 chip show correct operation between 35.5 MHz and 58.8 MHz. Using an external "blip circuit" to generate two-phase power-clocks, the MD1 chip achieved correct operation at 8.5 and 15.8 MHz.

Fig. 4 shows the circuit diagram of the Energy-Recovery (E-R) latch, one of the key components used in both AC-1 and MD1. E-R latches the input data and transfers charge from a clock line to a load capacitance and back again. Charge can be transferred adiabatically and conditionally to a load capacitance by means of a charge controlled switch that is implemented as a bootstrapped nFET. The rationale for this bootstrapped circuit topology is to minimize the dissipation in the switch so that as much as possible of the energy supplied to the load capacitance can be recovered.

For the correct operation of the E-R latch, nonoverlapping two-phase power-clocks ( $\Phi_I$  and  $\Phi_D$ ) are required.

#### 4 TRUE SINGLE-PHASE ENERGY RECOVERY

This section describes two prototype chips developed at Michigan that demonstrate different aspects in the application of charge recovery. The first chip is a full-custom multiplier implemented in a dynamic source-coupled adiabatic logic family. The second chip is a fully automated ASIC with a resonant clock distribution network. This ASIC uses an energy recovering flip-flop that has been specifically designed to operate efficiently with a sinusoidal clock waveform.

Both chips rely on a true single-phase power-clock waveform to provide power and synchronization. Single-phase energy recovery enjoys several distinct advantages over multiphase approaches. First, single-phase clocking imposes less stringent requirements on clock distribution than multiple-phase clock schemes. Most notably, the need to control the relative skew of multiple clock phases is completely eliminated. Second, power-clock generation is more efficient with a single phase than with multiple phases. Specifically, data dependencies may cause significant variations in power-clock currents and can be detrimental to the efficiency of the power-clock generator. In single-phase systems, tolerating such variations through feedback control in the power-clock generator presents fewer challenges than in multiphase systems.

None of the two chips performs any reversible computation in its logic blocks. They both rely on some notion of "balanced operation," however, to allow for the efficient recycling of charge. In particular, the multiplier uses dual-rail gates, with charge recovery taking place all the way to the gate fanouts. At each gate, the loading of the two fanout nets is essentially the same. Thus, at each cycle of the computation, the same amount of charge is stored on one or the other rail, depending on the output of the function, and the overall load of the power-clock generator remains relatively stable. In the resonant-clocked ASIC, the balanced operation is performed by the clock distribution network. In some sense, this network performs a "reversible computation" since the state of the clock toggles between 0 and 1 every clock cycle.

##### 4.1 Dynamic Multiplier

Using the source-coupled adiabatic logic family SCAL-D [17] and a  $0.5\ \mu\text{m}$  single-poly triple-metal n-well CMOS process supplied through MOSIS, we recently designed an 8-bit charge-recovery multiplier. In HSPICE simulations with parasitics extracted from layout, the multiplier achieved operating speeds up to 200MHz, and test results (limited by our experimental setup) confirmed its correct operation for clock rates up to 140MHz. Our chip was first presented in [3], while a more detailed description and comparative evaluation followed in [4]. In this section, we summarize and highlight key aspects of the design.




Fig. 5. Microphotograph of the true single-phase multiplier chip.

A die photo and simplified floorplan of our chip are shown in Figs. 5 and 6, respectively. Each die includes two 8-bit charge-recovery multipliers with charge-recovery built-in logic block observer (BILBO) circuitry, a power-clock generator, adiabatic-to-digital converters (denoted by A/D), and conventional I/O pads. Total chip area is  $4.83\text{mm}^2$  ( $= 2.60\text{mm} \times 1.84\text{mm}$ ).

Fig. 7 shows a typical pipeline of gates from the logic family SCAL-D that was used in the multiplier chip. This particular pipeline implements a 2-bit shift-register. Each pipeline stage consists of one PMOS and one NMOS SCAL-D gate, each computing on the rising and falling edges of the single-phase power-clock signal, respectively. Notice that this circuit is micropipelined, with fine-grain charge recovery taking place all the way from the fanouts of the individual gates. The level-shifter blocks LSP and LSN provide appropriate biases for the correct operation of the PMOS and NMOS gates.

Fig. 8 shows the power-clock generator circuit that was used with the multiplier chip. Except for the inductor, all clock generator circuitry was integrated on the chip. The power-clock generator relies on Class-E amplifier principles

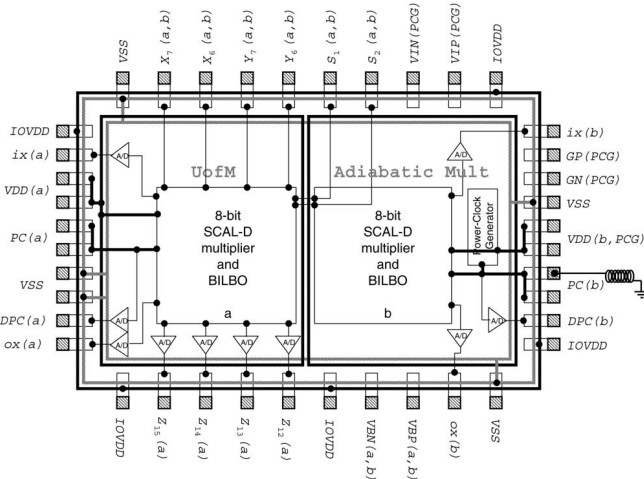


Fig. 6. Simplified floorplan of the chip in Fig. 5.




Fig. 7. Shift-register in adiabatic logic family SCAL-D.

to achieve high energy efficiency. Specifically, switches  $S_1$  and  $S_2$  replenish the energy of the power-clock waveform every cycle, conducting current when the voltage difference across their terminals is minimal. The timing of the two switches is controlled by a pulse alternator circuit that receives a pulse train from the ring oscillator at a rate equal to the target operating frequency and generates two pulse trains of the same rate and 180 degrees out of phase.

Fig. 9 demonstrates the correct operation of the SCAL-D chip at an operating frequency of 130MHz. This figure shows the power-clock (channel 1), a BILBO control signal (channel 2), and the output sequences of BILBO-1 and BILBO-2 (channels 3 and 4).

Fig. 10 gives the energy consumption of the 8-bit SCAL-D multiplier and associated BIST logic for various PMOS and NMOS biasing voltages. Measurements in the 40-130MHz frequency range were obtained using an external source of power-clock, with both the amplitude of the sinusoidal power-clock  $V_{PC}$  and the constant supply voltage  $V_{dd}$  set to 3.0V. At 140MHz, the measurement was obtained for the combined clock generator and multiplier system also at 3.0V.

To evaluate the relative efficiency of our charge recovery multiplier, we synthesized and simulated conventional static CMOS multipliers with two, four, and eight pipeline stages. At equal throughputs, our HSPICE simulations of voltage-scaled designs predicted energy savings up to a factor of 4, as shown in Fig. 11. In all cases, the latency of our charge recovery multiplier was 15 cycles.

### 4.2 Resonant-Clocked ASIC

We recently designed and tested a synthesized ASIC that performs a 7-bit Discrete Wavelet Transform. The chip was fabricated in a  $0.25\mu\text{m}$  bulk-CMOS process through MOSIS. Comprising close to 4,000 gates, our ASIC is clocked by a




Fig. 8. Power-clock generator for multiplier chip.




Fig. 9. Measured waveforms of 8-bit SCAL-D multiplier with associated BIST logic in self-test mode at 130MHz.

resonant charge-recovering waveform of sinusoidal shape. Fig. 12 shows a microphotograph of our resonant-clocked chip [5]. The lower left corner of the die contains our experimental energy-recovering design that consists of an ASIC core, an on-chip resonant clock generator, and some testing logic. The energy recovering flip-flops are driven by a resonant waveform generated using an off-chip surface-mount inductor and the on-chip power-clock generator. This chip was fabricated in a  $0.25\mu\text{m}$  standard bulk-CMOS process.

A schematic of the energy-recovering flip-flop used in our ASIC is shown in Fig. 13. This flip-flop consists of a charge recovering dynamic buffer driving a pair of cross-coupled NOR gates as the static latch element. Our flip-flop latches on rising pulses of power-clock. The input needs to be stable by the time power-clock is roughly half way to its peak and should be held stable until power-clock is at the




Fig. 10. Measured energy consumption per cycle for various PMOS/NMOS biasing voltages.

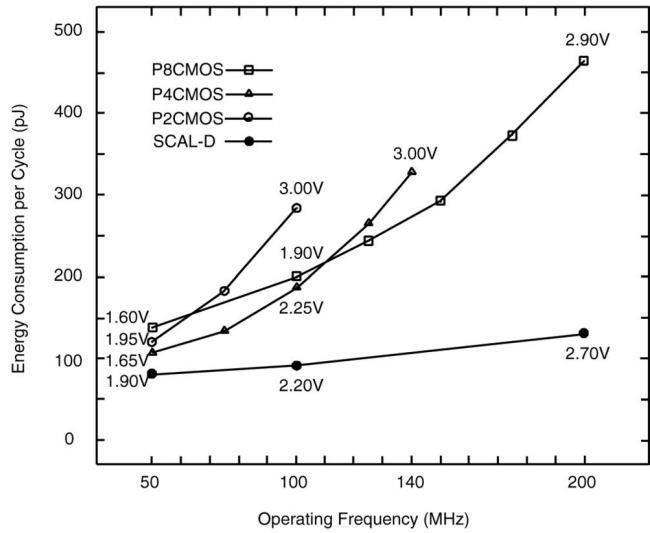


Fig. 11. Energy consumption per cycle for charge recovery multiplier and conventional pipelined counterparts in static CMOS.




Fig. 12. Microphotograph of the resonant clocked ASIC chip.

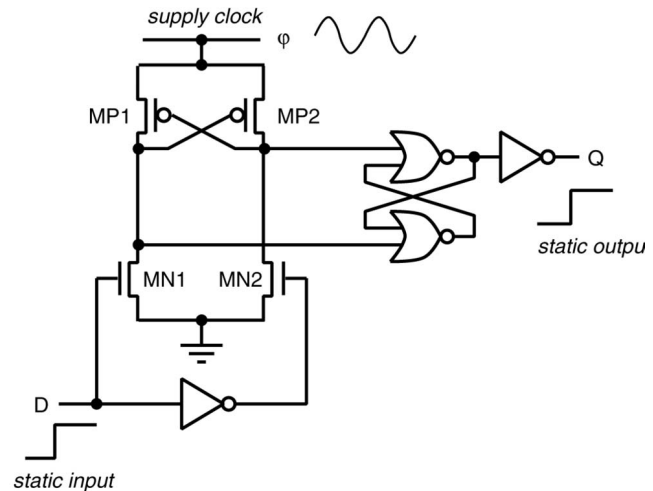


Fig. 13. Energy recovering sinusoidal flip-flop used in the resonant clocked ASIC chip.

peak. The flip-flop draws more current from the power-clock when active (i.e., the data is changing), thus changing the effective load seen by the power-clock generator.

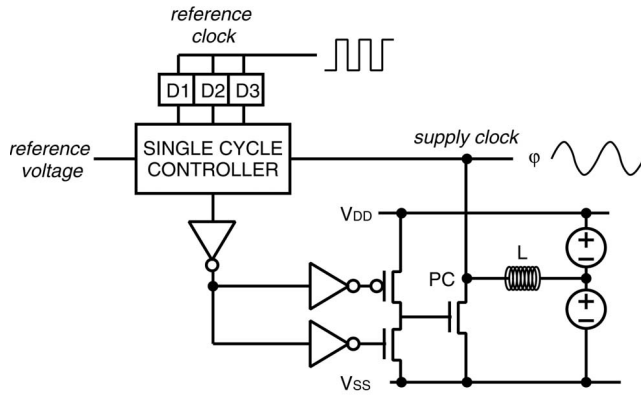


Fig. 14. Clock generator used in ASIC chip.




Fig. 15. Measured power consumption resonant-clocked ASIC.

Our chip includes a single-cycle feedback control resonant power-clock generator, shown in Fig. 14, that is capable of reacting to changes in its load. The amplitude of the power-clock signal is sampled and compared against a reference level. The result of this comparison is used to decide, on a cycle-by-cycle basis, whether or not to turn on the main NMOS power-switch to pump more energy into the power-clock. This control is critical for achieving ultra-low dissipation when the ASIC is idling.

Fig. 15 shows the measured energy dissipation of the clock network in our resonant-clocked ASIC chip at several frequencies between 100MHz and 300MHz. At each frequency point, the voltage was scaled down to the minimum required for correct operation. The inductor and DC supplies were connected externally. For reference, we plot a quadratic curve fit to the function  $fCV_{dd}^2$  evaluated at each of the voltage, frequency pairs. This curve represents the dissipation required to drive the same clock capacitance if charge recovery techniques were not used. At  $f = 300\text{MHz}$ , the clock was overdriven using an inductance value larger than  $1/C(2\pi\omega)^2$ , resulting in suboptimal power dissipation at that frequency. At 205MHz, the measured clock power dissipation was 4.5mW, about 5 times less than required to drive the same clock capacitance with conventional means. These dramatic power savings are due to operation near the resonance of the inductor in conjunction with the clock-capacitance.




Fig. 16. Measured power-clock spectrum, 200MHz.

In addition to reduced power dissipation, charge recovery circuitry has the potential to operate with substantially reduced electromagnetic interference. To provide empirical evidence in support of this largely unexplored fact, we analyzed the spectrum of the measured power-clock waveform when resonating at 200MHz. The spectrum obtained is shown in Fig. 16, zoomed in on the region of interest from 0 to 1GHz. This data was obtained by recording 100,000 voltage samples at 100ps/sample at the off-chip inductor terminal. Assuming linear characteristics from the parasitic elements between the inductor terminal and the on-chip clock network, this data should be proportional to the actual clock signal on-chip. The graph shows the presence of substantially attenuated odd and even harmonics. Specifically, the first three harmonics are 22dB, 36dB, and 43dB below the fundamental, respectively. In contrast, the first harmonic of a square waveform at 600MHz is about 12dB below the fundamental. The spur at roughly 10MHz can be attributed to a periodicity in the datapath self-test activity as it corresponds roughly to the spectrum of one of the self-test signature outputs. An alternate hypothesis is that it results from some coupling with one of the I/O pads slewing.

## 5 CONCLUSION

The field of charge recovery has come a long way from its origins in the physics of computation. Originally merely a theoretical curiosity, it is approaching maturity, with several researchers independently producing functional prototypes in conventional silicon CMOS fabrication processes. In contrast to other low-power design techniques, which try to reduce power by performing less wasted computation, slowing down the computation, or lowering supply voltage, charge recovery fundamentally changes the shape of the dissipation versus delay and area trade-offs, allowing for sub- $fCV_{dd}^2$  dissipation with the help of novel circuits implemented in standard CMOS processes.

Although this paper has focused primarily on charge recovery circuitry for datapaths and finite-state machines, energy-recycling design technologies have shown great promise in reducing the power consumption of other key components in digital computing systems. Rotary traveling-wave oscillators have been demonstrated in silicon, enabling the recovery of charges from clock distribution lines operating at tens of GHz [18]. In this approach, energy is recycled using the distributed inductance of the clock network, which behaves as a transmission line. Charge recovery has also been applied to the design of low-power memory arrays, including SRAMs and register files [19], [20], [21], [22], [23]. Furthermore, charge recovery has been applied to the design of low-power drivers for LCD displays [24].


The future of charge recovery appears more promising than ever. With the advent of design automation tools and the increasing familiarization of designers with charge recovery design technologies, we can expect commercial chips in certain power or energy-sensitive areas to adopt some form of energy recycling in pursuit of higher efficiency.

## ACKNOWLEDGMENTS

Suhwan Kim would like to express thanks for the financial support from the Nano-Systems Institute (NSI-NCRC) program sponsored by the Korea Science and Engineering Foundation (KOSEF). Conrad H. Ziesler and Marios C. Papaefthymiou would like to acknowledge financial support from the US Army Research Office under Grant Nos. DAAD19-99-1-0304 and DAAD19-03-1-0122.

## REFERENCES

- [1] K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand, "Leakage Current Mechanism and Leakage Reduction Techniques in Deep-Submicrometer CMOS Circuits," *Proc. IEEE*, vol. 91, no. 2, pp. 305-327, Feb. 2003.
- [2] C.H. Bennett and R. Landauer, "The Fundamental Physical Limits of Computation," *Scientific Am.*, vol. 253, no. 1, pp. 38-46, July 1985.
- [3] S. Kim, C. Ziesler, and M.C. Papaefthymiou, "A True Single-Phase 8-Bit Adiabatic Multiplier," *Proc. 38th ACM/IEEE Design Automation Conf.*, pp. 758-763, June 2001.
- [4] S. Kim, C.H. Ziesler, and M.C. Papaefthymiou, "A True Single-Phase Energy-Recovery Multiplier," *IEEE Trans. VLSI Systems*, vol. 11, no. 2, pp. 194-207, Apr. 2003.
- [5] C. Ziesler, J. Kim, V. Sathe, and M.C. Papaefthymiou, "A 225MHz Resonant Clocked ASIC Chip," *Proc. Int'l Symp. Low-Power Electronics and Design*, pp. 48-53, Aug. 2003.
- [6] S.G. Younis, "Asymptotically Zero Energy Computing Using Split-Level Charge Recovery Logic," PhD thesis, Massachusetts Inst. of Technology, 1994.
- [7] J. Lim, D. Kim, and S. Chae, "A 16-Bit Carry-Lookahead Adder Using Reversible Energy Recovery Logic for Ultra-Low-Energy Systems," *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 898-903, June 1999.
- [8] C. Vieru, "Pendulum: A Reversible Computer Architecture," SM thesis, Massachusetts Inst. of Technology, 1995.
- [9] W.C. Athas, L.J. Svensson, J.G. Koller, N. Tzartzanis, and Y. Chou, "Low-Power Digital Systems Based on Adiabatic-Switching Principles," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 2, no. 4, pp. 398-406, Dec. 1994.
- [10] A. Kramer, J.S. Denker, S.C. Avery, A.G. Dickinson, and T.R. Wik, "Adiabatic Computing with the 2N-2N2D Logic Family," *Digest Technical Papers IEEE Symp. VLSI Circuits*, pp. 25-26, Apr. 1994.
- [11] A. Kramer, J.S. Denker, B. Flower, and J. Moroney, "2nd Order Adiabatic Computation with 2N-2P and 2N-2N2P Logic Circuits," *Proc. Int'l Symp. Low Power Design*, pp. 191-196, 1995.
- [12] Y. Moon and D. Jeong, "An Efficient Charge Recovery Logic Circuit," *IEEE J. Solid-State Circuits*, vol. 31, no. 4, pp. 514-522, Apr. 1996.
- [13] V.G. Oklobdzija and D. Maksimovic, "Pass-Transistor Adiabatic Logic Using Single Power-Clock Supply," *IEEE Trans. Circuits and Systems-II: Analog and Digital Signal Processing*, vol. 44, no. 10, pp. 842-846, Oct. 1997.
- [14] D. Maksimovic, V.G. Oklobdzija, B. Nikolic, and K.W. Current, "Clocked CMOS Adiabatic Logic with Integrated Single-Phase Power-Clock Supply," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 8, no. 4, pp. 460-463, Aug. 2000.
- [15] W.C. Athas, N. Tzartzanis, L.J. Svensson, and L. Peterson, "A Low-Power Microprocessor Based on Resonant Energy," *IEEE J. Solid-State Circuits*, vol. 32, no. 11, pp. 1693-1701, Nov. 1997.
- [16] W. Athas, N. Tzartzanis, W. Mao, L. Peterson, R. Lal, K. Chong, J.-S. Moon, L. Svensson, and M. Bolotski, "The Design and Implementation of a Low-Power Clock-Powered Microprocessor," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1561-1570, Nov. 2000.
- [17] S. Kim and M.C. Papaefthymiou, "Single-Phase Source-Coupled Adiabatic Logic," *Proc. Int'l Symp. Low-Power Electronics and Design*, pp. 97-99, Aug. 1999.
- [18] J. Wood, T. Edwards, and S. Lipa, "Rotary Travelling-Wave Oscillator Arrays: A New Clock Technology," *IEEE J. Solid-State Circuits*, vol. 36, no. 11, pp. 1654-1665, Nov. 2001.
- [19] Y. Moon and D.K. Jeong, "A 32x32-b Adiabatic Register File with Supply Clock Generator," *IEEE J. Solid-State Circuits*, vol. 33, no. 5, pp. 696-701, May 1998.
- [20] D. Somasekhar, Y. Ye, and K. Roy, "An Energy Recovering Static RAM Memory Core," *Proc. Int'l Symp. Low Power Design*, pp. 9-11, Oct. 1995.
- [21] N. Tzartzanis, W.C. Athas, and L. Svensson, "A Low-Power SRAM with Resonantly Powered Data, Address, Word, and Bit Lines," *Proc. European Solid-State Circuits Conf.*, pp. 336-339, Sept. 2000.
- [22] J. Kim, C.H. Ziesler, and M.C. Papaefthymiou, "Energy Recovering Static Memory," *Proc. Int'l Symp. Low-Power Electronics and Design*, pp. 92-97, Aug. 2002.
- [23] J. Kim and M.C. Papaefthymiou, "Constant-Load Energy Recovery Memory for Efficient High-Speed Operation," *Proc. Int'l Symp. Low Power Electronics and Design*, pp. 240-243, Aug. 2004.
- [24] J. Ammer, M. Bolotski, P. Alvelda, and T. Knight, "A 160x120 Pixel Liquid-Crystal-on-Silicon Microdisplay with an Adiabatic DAC," *Proc. IEEE Int'l Solid-State Circuits Conf.*, pp. 212-213, Nov. 1999.



**Suhwan Kim** (M'01) received the BS and MS degrees in electrical engineering and computer science from Korea University, Seoul, Korea, in 1990 and 1992, respectively, and the PhD degree in electrical engineering and computer science from the University of Michigan, Ann Arbor, in 2001. From 1993 to 1999, he was with LG Electronics, Seoul, Korea, where he designed several multimedia system-on-a-chip (SoC), including an MPEG2 CODEC for audio, video, and system. From 2001 to 2004, he was a research staff member at the IBM T.J. Watson Research Center, Yorktown Heights, New York. In 2004, he joined Seoul National University, Seoul, Korea, where he is currently an assistant professor of electrical engineering. His research interests encompass high-performance and low-power Analog and Mixed Signal (AMS) circuits and technology and low-power design methodologies for high-performance VLSI signal processing. He received the 1991 Best Student Paper Award of the IEEE Korea Section and the First Prize (Operational Category) in the VLSI Design Contest of the 38th ACM/IEEE Design Automation Conference. He served as the Technical Program cochair for the IEEE International SOC Conference. He has also participated several times on the Technical Program Committee of the IEEE International SOC Conference and the International Symposium on Low-Power Electronics and Design. He is a member of the IEEE.





**Conrad H. Ziesler** (M'05) received the BS degree in electrical engineering in 1999 from the California Institute of Technology, Pasadena, and the MS and PhD degrees in electrical engineering and computer science in 2002 and 2004 from the University of Michigan, Ann Arbor. He received the First Prize in the VLSI Design Contest of the 38th ACM/IEEE Design Automation Conference, and an NFS Graduate Research Honorable Mention. He is currently a

member of the technical staff at Multigig, Inc. His research at the Advanced Computer Architecture Laboratory in Ann Arbor included work on novel high-speed, ultra-low power adiabatic logic families. Previously, he worked in the microdevices laboratory at NASA's Jet Propulsion Laboratory, focusing on miniaturizing and integrating electron-beam devices with radiation sensors. He is a member of the IEEE.



**Marios C. Papaefthymiou** (M'93, SM'02) received the BS degree in electrical engineering from the California Institute of Technology in 1988 and the SM and PhD degrees in electrical engineering and computer science from the Massachusetts Institute of Technology in 1990 and 1993, respectively. After a three-year term as an assistant professor at Yale University, he joined the University of Michigan, where he is currently an associate professor of electrical engineering and computer science and director of the Advanced Computer Architecture Laboratory. His research interests encompass algorithmic, architectural, and circuit issues in the design of VLSI systems with a primary focus on energy and timing optimization. He is also active in the field of parallel and distributed computing. Together with his students, he received a Best Paper Award from the 32nd ACM/IEEE Design Automation Conference and the First Prize (Operational Category) in the VLSI Design Contest of the 38th ACM/IEEE Design Automation Conference. He has served multiple terms as an associate editor for the *IEEE Transactions on the Computer-Aided Design of Integrated Circuits*, the *IEEE Transactions on Computers*, and the *IEEE Transactions on VLSI Systems*. He has served as the general chair and as the Technical Program chair for the ACM/IEEE International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems. He is a senior member of the IEEE.

▷ For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).