

Towards An Ultra-Low-Power Architecture Using Single-Electron Tunneling Transistors

Changyun Zhu*, Zhenyu (Peter) Gu†, Li Shang*, Robert P. Dick†, and Robert G. Knobel‡

*ECE Department
Queen's University
Kingston, ON K7L 3N6, Canada
4cz1@qlink.queensu.ca, li.shang@queensu.ca

†EECS Department
Northwestern University
Evanston, IL 60208, U.S.A.
{zgu646, dickrp}@eecs.northwestern.edu

‡Physics Department
Queen's University
Kingston, ON K7L 3N6, Canada
knobel@physics.queensu.ca

ABSTRACT

Minimizing power consumption is vitally important in embedded system design; power consumption determines battery lifespan. Ultra-low-power designs may even permit embedded systems to operate without batteries, e.g., by scavenging energy from the environment. Moreover, managing power dissipation is now a key factor in integrated circuit packaging and cooling. As a result, embedded system price, size, weight, and reliability are all strongly dependent on power dissipation.

Recent developments in nanoscale devices open new alternatives for low-power embedded system design. Among these, single-electron tunneling transistors (SETs) hold the promise of achieving the lowest power consumption. However, SETs impose unique design constraints that strongly influence architectural and circuit-level decisions. Unfortunately, most analysis of SETs has focused on single devices instead of architectures, making it difficult to determine whether they are appropriate for low-power embedded systems.

This article presents possible uses of SETs in high-performance and battery-powered embedded system design. The resulting fault-tolerant, hybrid SET/CMOS, reconfigurable architecture can be tailored to specific requirements and allows trade-offs among power consumption, performance, operation temperature, fabrication cost, and reliability. This work is a first step in evaluating the system-level potential of reducing power consumption by using SETs.

Categories and Subject Descriptors: B.7.1 [Integrated Circuits]: Types and Design Styles—Advanced technologies; B.8.2 [Performance and Reliability]: Performance Analysis and Design Aids.

General Terms: Design, performance, reliability.

Keywords: Single electron tunneling transistor (SET), low-power, nanoelectronics, reconfigurable architecture.

I. INTRODUCTION AND MOTIVATION

Power related energy and thermal issues are now central in electronic system design. In high-performance applications, temperature impacts integration density, performance, reliability, power consumption, and cost. For battery-powered embedded systems, energy consumption directly determines the system life time. Power consumption crises were historically solved by moving to new technologies that decreased energy per operation, allowing increases in density and eventually performance. Power and thermal concerns were some of the main motivations for replacing vacuum tubes with semiconductor devices in the 1960s and replacing bipolar junction transistors (BJTs) with CMOS in the 1990s. CMOS is the mainstream fabrication technology used today. As integrated circuit (IC) integration further increases, it will reach fabrication, power consumption, and thermal limits; it may soon be time for another transition to a dramatically different technology.

This work is supported in part by the NSERC Discovery Grants #388694-01 and #298185-04, and in part by the NSF under award CNS-0347941.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 2007, June 4-8, 2007, San Diego, California, USA.
Copyright 2007 ACM 978-1-59593-627-1/07/0006 ... \$5.00.

Device researchers have seen the coming challenges for CMOS devices and evaluated alternative technologies such as carbon nanotube transistors [1], nanowires [2], and single-electron tunneling transistors (SETs) [3]. As projected by International Technology Roadmap for Semiconductors, SETs can potentially achieve the lowest projected energy per switching event of any known computation technology (10^{-18} J). However, their use poses unique challenges in architectural design, circuit design, and fabrication. For instance, SETs are susceptible to reliability problems resulting from $1/f$ noise caused by random background offset charge effects. They have cyclic I-V curves (see Figure 2) that can complicate design but permit highly-efficient implementation of some useful logic functions that have proven inefficient with CMOS. The feature sizes (a few nanometers) required for room-temperature operation make fabrication challenging.

Extensive research has been conducted on fabrication, design, and modeling of SETs. Please refer to the survey by Likharev [3] for more details. Recently, researchers have fabricated room-temperature SETs [4, 5]. This work provides a promising start for SET circuit design. Various SET-based circuit applications, such as logic [6]–[10] and memory [11] demonstrate orders of magnitude improvement in energy efficiency compared to CMOS. Research into SET modeling and simulation has also been an active area. Monte Carlo simulation has been widely used to model SETs. SIMON [12] and MOSES [13] are the two most popular SET simulators which, nevertheless, are not suitable for circuit analysis due to large runtimes when characterizing systems containing more than a few SETs. Uchida et al. proposed an analytical SET model and incorporated it into SPICE [14]. Recently, Inokawa et al. extended this model to a more general form to include asymmetric SETs [15]. Mahapatra et al. proposed a simulation framework for hybrid SET/CMOS circuit design and analysis [16]. These models all match Monte Carlo simulation well.

In this paper, we explore the potential use of SETs in low-power embedded system design. Our work starts from design space characterization of SET-based architectures. We evaluate the impacts of using SETs upon architectural, circuit-level, and device-level design, considering metrics such as energy efficiency, performance, reliability, maximum operating temperature, and ease of fabrication.

Based on our evaluation of the architectural and circuit-level features that can most effectively exploit the strengths of SETs while working within the constraints their use imposes, we propose a fault-tolerant, hybrid SET/CMOS, reconfigurable architecture that we call IceFlex. IceFlex is regular and cell-based, easing nanoscale design and fabrication. It is reconfigurable and modular, permitting compensation for fabrication defects. In addition to compensating for the weaknesses of SETs, IceFlex exploits their strengths, e.g., extremely low power consumption and support for multi-gate devices.

We tailor IceFlex to both high-performance and battery-powered embedded systems and characterize its performance and power consumption when used to implement a number of instruction processors and application-specific cores. Compared to CMOS-based designs, IceFlex improves energy efficiency by $235\times$ on average for high-performance applications and $201\times$ on average for battery-powered applications, while maintaining good performance.

II. SET MODELING

The operation of a single-electron tunneling device is governed by the Coulomb charging effect. As shown in Figure 1, a single-




Fig. 1. SET structure and schematic.




Fig. 2. SET Coulomb oscillation ($C_g = 3.0$ aF, $C_s = C_d = 1.0$ aF, $V_{ds} = 26.7$ mV, and $R_s = R_d = 10$ M Ω).

electron tunneling device consists of a nanometer-scale conductive island embedded in an insulating material. Electrons travel between the island and source (S) or drain (D) through tunnel junctions. When an electron tunnels into the island, the overall electrostatic potential of the island increases by e^2/C_Σ , where e is the elementary charge and C_Σ is the island capacitance. For devices with nanometer-scale islands, the capacitance C_Σ is small. As a result, the electrostatic force due to electron charging is significant and dominates the effect of thermal energy, particularly at low temperatures.

Changes to SET island potential results in an energy gap at Fermi energy, preventing further electron tunneling. This phenomenon is called Coulomb blockade. It prevents current from flowing between source and drain ($I_{ds} = 0$), i.e., the SET is turned off. The Coulomb blockade effect can be overcome by changing the voltage of a conductor capacitively coupled to the island, thereby turning tunneling on and off. As shown in Figure 2, discrete electron charging results in a periodic I-V transfer curve. The drain current changes periodically as a function of the gate voltage, and has peaks and valleys with period of e/C_g . These periodic changes are called Coulomb oscillations.

II.A. SET Modeling

Circuit and architecture design involves extensive large-scale circuit simulation. Despite their accuracy, Monte Carlo methods are not suitable for large-scale circuit analysis due to their high time complexity. In this work, SET circuit modeling and analysis build upon the analytical model developed by Inokawa et al. [15]. This compact model can be incorporated into SPICE. Combined with MOS transistor models, it provides an efficient and accurate simulation solution. Inokawa's model ignores random background offset charge effects. In addition, it does not support multi-gate devices, in which multiple gates are capacitively coupled with a SET. In this work, we incorporate these two effects into Inokawa's model, and use it for hybrid SET/CMOS circuit and architecture design. The I-V characteristics of a SET with island charge equal to n or $n + 1$ electrons follow:

$$I_{DS} = \frac{e}{4R_T C_\Sigma} \frac{(1-r^2)(\tilde{V}_{GS}^2 - \tilde{V}_{DS}^2) \sinh(\tilde{V}_{DS}/\tilde{T})}{(\tilde{V}_{GS} + r\tilde{V}_{DS}) \sinh(\tilde{V}_{GS}/\tilde{T}) - (\tilde{V}_{DS} + r\tilde{V}_{GS}) \sinh(\tilde{V}_{DS}/\tilde{T})} \quad (1)$$

where

$$\tilde{V}_{GS} = \frac{2 \sum C_{G_i} V_{GS_i}}{e} - \frac{(\sum C_{G_i} + C_s - C_D) V_{DS}}{e} - 1 - 2n + \zeta \quad (2)$$

$$\tilde{V}_{DS} = \frac{C_\Sigma V_{DS}}{e}, \quad \tilde{T} = \frac{2k_B T C_\Sigma}{e^2} \quad (3)$$

$$r = \frac{R_D - R_S}{R_D + R_S}, \quad R_T = \frac{2}{\frac{1}{R_S} + \frac{1}{R_D}}, \quad C_\Sigma = C_s + C_D + \sum C_{G_i} \quad (4)$$

In this model, $\sum C_{G_i} V_{GS_i}$ models the Coulomb charging effects of the multiple gate terminals. ζ is a real number that characterizes the random background offset charge effect.

III. ICEFLEX: A FAULT-TOLERANT HYBRID SET/CMOS RECONFIGURABLE ARCHITECTURE

This section describes the design and analysis of IceFlex, the proposed fault-tolerant, hybrid SET/CMOS, reconfigurable architecture. The vast majority of devices in IceFlex are SETs, allowing extremely low power consumption. CMOS devices are sparingly used to improve global interconnect driving strength.

Our evaluation of the architectural constraints imposed by SETs leads to four main conclusions. First, flawless fabrication will be challenging, especially for circuits that operate at room temperature. It is important to simplify fabrication and use post-fabrication adaptation to avoid flawed devices. Second, an unpredictable subset of devices will be susceptible to random background offset charge effect noise: SET-based architectures should have the ability to tolerate run-time errors. Third, SETs have poor driving strength; this must be remedied, especially when driving global interconnect. Fourth, SETs have the ability to efficiently implement some functions that are inefficient using BJT or CMOS, e.g., non-linearly-separable functions and threshold logic can be efficiently implemented using multi-gate devices. SET-based architectures should exploit such special properties in order to improve the efficiency of arithmetic and other logic circuits.

III.A. SET Design Space Characterization

In order to characterize the benefits and limitations of SET circuits and architectures, we analyze the tradeoffs among the following metrics: temperature, performance, power consumption, reliability, and fabrication constraints. This study yields the two design configurations, as shown in Table I. One of these targets high performance applications and the other targets battery-powered embedded systems

III.A.1) Temperature: We evaluated IceFlex at seven temperatures (see Table I). IceFlex is a hybrid SET/CMOS design; the temperature range starts at approximately 40 K to permit reliable operation of the CMOS components. 77 K is achieved by liquid nitrogen cooling. 103 K is the average cloud top temperature. 120 K and below are defined to be cryogenic. At 200 K, functional SET devices have been widely demonstrated in the literature. 250 K is a temperature that might be reached using a stacked Peltier heat pump. 300 K is room temperature.

III.A.2) Capacitance: To observe well-defined Coulomb charging effects, electron charging energy must be higher than the energy of thermal fluctuations, i.e., $C_\Sigma(T) \leq e^2/(10k_B T)$, where k_B is Boltzmann's constant and T is the temperature. At room temperature, this constraint requires an island capacitance below 1 aF, making fabrication challenging but possible [5]. In order to operate voltage-state logic, SETs must exhibit voltage gain, which is equal to the gate capacitance divided by the sum of the junction capacitances: $G = C_G/(C_S + C_D)$. Our results indicate that a gain of 1.5 is sufficient for use in digital logic. Targeting battery-powered systems, using $C_\Sigma \leq e^2/(10k_B T)$ and $G = 1.5$, the maximum allowed gate and junction capacitances are derived and shown in Table I (column 2 and 3, respectively). When the island capacitance is 0.5 aF, the island diameter is 3–4 nm [17]. Therefore, room-temperature operation would require an island with a diameter of at most a few nanometers.

The performance of SETs degrades as device capacitance increases. We assume the capacitances at 300 K are the minimal allowed gate and junction capacitances. For high-performance applications, these minimal gate and junction capacitances are used at all the temperature settings and shown in Table I (column 6 and 7, respectively).

III.A.3) Voltage: Consider a SET biased via a second gate, such that a V_{GS} of zero places it in the middle of the positive voltage coefficient (PVC) region in Figure 2. In this case, the maximum range of current values can be traversed by letting V_{GS} (i.e., V_{in}) vary in the range $[-e/(4C_G), e/(4C_G)]$. At all but the lowest temperatures, this range also provides near-optimal sensitivity to V_{GS} . Therefore, we use this range. Once the range of V_{GS} is known, a V_{SS} of $-e/(4C_G)$ and a

TABLE I
DESIGN SPACE CHARACTERIZATION

Temp. (K)	Low power				High performance			
	Capacitance (aF)		Voltage (mV)	Resist. (MΩ)	Capacitance (aF)		Voltage (mV)	Resist. (kΩ)
	C_G	C_D	V_{dd}, V_{in} $e/4C_G$	R_S R_D	C_G	C_D	V_{dd}, V_{in} $e/4C_G$	R_S R_D
40	2.78	0.93	14.36	10	0.37	0.12	107.70	100
77	1.45	0.48	27.65	10	0.37	0.12	107.70	100
103	1.08	0.36	36.99	10	0.37	0.12	107.70	100
120	0.93	0.31	43.09	10	0.37	0.12	107.70	100
200	0.56	0.19	71.82	10	0.37	0.12	107.70	100
250	0.45	0.15	89.77	10	0.37	0.12	107.70	100
300	0.37	0.12	107.70	10	0.37	0.12	107.70	100

V_{DD} of $e/(4C_G)$ naturally follow, shown in Table I (column 4 and 8). Note that a bias voltage applied via a second gate can be used to shift the zero V_{GS} point from the PVC to negative voltage coefficient (NVC) region in Figure 2, permitting NMOS-like or PMOS-like behavior.

III.A.4) Junction Resistance: To observe single-electron charging effects, electrons must be confined in the island, which requires that the junction resistance be much higher than the quantum resistance, i.e., $R_S, R_D \gg h/e^2$, $h/e^2 = 25.8\text{ k}\Omega$, where h is Planck's constant. Therefore, SETs have high resistances and low driving currents. In this work, we pick two resistance settings: 100 KΩ for high-performance applications and 10 MΩ for battery-powered systems, shown in Table I (column 5 and 9).

III.A.5) Reliability Implications: Researchers have pointed out the dangers posed by thermal noise as charging (state change) energy approaches thermal energy. The charging energies of the devices in the proposed architecture are an order of magnitude greater than the thermal energy ($10k_B T$). We explicitly consider the effects of thermal noise; they are reflected in our design decisions and their impacts on power consumption and performance are considered. At charging energies over $10k_B T$, the model we use is accurate to within 4% of the time-dependent master equation [14, 18].

Random background offset charge effects [19, 20] are the main barrier to SET reliability. They are observed as $1/f$ noise on SET gate voltages, with some SETs susceptible and others immune. Currently, the distribution of random background offset charges can only be determined after fabrication [3]. Susceptible SETs may suffer soft errors infrequently, e.g., only once per day. Several recent devices have shown improved immunity to this noise [20], with operation essentially unchanged over several weeks. In this work, we use architectural techniques to reduce the probability of failure using an entirely SET-based design. SETs are used in parallel to exploit the lack of SET-to-SET correlation in random background offset charge effects.

III.B. IceFlex Design

In this section, we present the architecture and circuit design of IceFlex. The microarchitecture of IceFlex is shown in Figure 3. Each cell is a SET logic block (SELB) composed of the following components: (1) multi-gate SET-based reconfigurable look-up tables that can realize arbitrary n -input Boolean functions; (2) a SET-based reconfiguration memory array that caches multiple configuration contexts to support efficient run-time reconfiguration; (3) a SET-based arithmetic unit that allows efficient implementations of non-linearly separable arithmetic operations; (4) multiple reconfigurable interconnect resources, including SET local interconnects, hybrid SET/CMOS global interconnects, SET switch fabric, and SET registers. In addition, IceFlex is also equipped with (5) SET threshold logic-based majority voting logic units (MVL) to compensate for run-time reliability problems.

Next we detail the design of each component of IceFlex and discuss various tradeoffs in circuit-level and architecture-level design.

III.B.1) Multi-Gate SET Reconfigurable Lookup Table Component: Each SELB is equipped with l sets of n -input reconfigurable look-up tables. Each look-up table can realize an arbitrary n -input Boolean function. The basic structure of the look-up table consists of an m -to-1 multi-gate SET multiplexer tree ($m = 2^n$) and an m -bit SET storage cell.

The proposed multi-gate SET multiplexer tree differs from existing CMOS designs in the following way. A CMOS m -to-1 multiplexer tree requires $\lceil \log_2 m \rceil$ stages of transmission gates, plus buffers to meet the




Fig. 3. IceFlex microarchitecture.




Fig. 4. Multi-gate SET multiplexer tree.

required driving strength. SETs may have multiple gate terminals. As described in Equation 2, the overall gate charging effect is a function of $\sum C_{G_i} V_{GS_i}$. Therefore, using SETs, multiple control signals (the select signals of multiplexer) can be integrated into a single SET.

Figure 4 shows the proposed SET multi-gate multiplexer tree design. The basic building block is a q -to-1 multi-gate single-stage multiplexer. Each of the q paths consists of a single multi-gate SET controlled by $\lceil \log_2 q \rceil$ select signals. Figure 4 also shows a design case for $q = 4$. The output SET buffer is used to improve the driving strength.

III.B.2) SET Configuration Memory: In IceFlex, run-time reconfiguration is enabled by SET configuration memory. At run-time, one set of configuration bits are fetched from the memory to program SELB logic and interconnect. Multiple configuration sets may be stored, permitting reconfiguration without off-chip memory access.

Figure 5 shows the circuit structure of the IceFlex configuration storage. Each storage cell contains a dual-island SET [3], i.e., capacitively-coupled primary and secondary SETs. By controlling V_{CG} , electrons can tunnel through the control gate and charge the island of the secondary SET. The charge state of the secondary SET is able to shift the phase of the Coulomb oscillations of the primary gate thereby shifting the I-V curve, i.e., storing a bit of data.

We designed a dual-island based set buffer to hold the current configuration. As shown to the right of Figure 5, this buffer uses two biasing voltages, V_{G_2} and $-V_{G_2}$, and behaves like a complementary SET inverter. During run-time reconfiguration, for each dual-island SET, the corresponding configuration bit stored in the configuration memory updates the island charge of its secondary SET and conductivity of the primary SET, thereby controlling the buffer output.

III.B.3) Efficient SET Implementations of Non-Unate Functions and Implications on Arithmetic: SETs have the ability to support efficient implementation of some critical logic functions that have long frustrated designers using threshold logic, BJT, and CMOS technologies. Most conventional transistors have either non-decreasing or non-increasing I-V curves. As a result, numerous devices are required to implement Boolean functions that are not unate, i.e., linearly separable. However, such functions are widely used, especially in digital arithmetic. The periodic nature of SET I-V curves can be exploited for efficient implementation of non-unate functions such as exclusive-or.

The most efficient CMOS static pass-transistor logic design of a two-input exclusive-or gate in general use requires four transistors [21]. Moreover, it relies on strong input signals because it is not capable of signal restoration. A restoring version would require eight transistors. In contrast, it is possible to implement a two-transistor SET-based exclusive-or gate that is structurally equivalent to a CMOS inverter. However, each SET has two gates, each of which is connected to one of the exclusive-or inputs. This design is capable of signal restoration.




Fig. 5. SET configuration memory.

By appropriately adjusting the gate capacitances, the device can be adjusted such that switching a single gate will result in a 180° phase shift in the I-V curve (see Figure 2).

In SET-based architectures, we propose the use of fast carry chains based on the proposed exclusive-or (sum) computation logic. We have found that this design is 71.8% more energy-efficient and 39.4% faster than a design based on a conventional CMOS-style exclusive-or sum implementation, when both are implemented using SETs. Carry-out logic is equivalent to 2/3 majority vote logic.

III.B.4) Reconfigurable Interconnect Network: IceFlex contains multiple reconfigurable interconnect resources, including SET local interconnects, hybrid SET/CMOS global interconnects, and SET switch fabric. Static power consumption dominates in SET-based interconnect due to the impact of the thermal energy on device conductance. It scales with wireload because high wireload requires low junction resistance. In contrast, SET dynamic power is low. Compared to SETs, CMOS has lower static power consumption, but higher capacitance and dynamic power consumption. Therefore, dynamic power dominates in the hybrid SET/CMOS-based design. Detailed circuit analysis shows that, under the same performance constraint, the SET-based design is more energy efficient for local interconnects and the hybrid design is more energy efficient for long wires. In IceFlex, local interconnects are driven by SET buffers. Three local interconnect lengths are supported: single, double, and hex. The driving strengths of SET buffers are selected to permit constant latency across different interconnect lengths to simplify placement and routing.

We propose a Hybrid SET/CMOS design to optimize global interconnect driving strength and energy efficiency. An interconnect is driven by one SET buffer (SINV1) and two CMOS buffers (CINV1 and CINV2) in series. Global interconnect terminates in a SET buffer (SINV2). Since the voltage range of SET logic is smaller than that of CMOS logic, both MOS transistors are within the switching region with high short-circuit power much of the time. To reduce short-circuit power consumption, CINV1 is designed to satisfy the following two constraints: (1) $V_{tn} + |V_{tp}| > V_{dd} - V_{ss}$ ensures that at least one MOS transistor is off at all times, reducing static power consumption and (2) the output signal range of SINV1 is greater than $V_{tn} + |V_{tp}| - (V_{dd} - V_{ss})$. Therefore, the NMOS (PMOS) transistor of CINV1 is conductive when the output of SINV1 is high (low) to provide enough driving strength to CINV2. CINV1 serves as a signal amplifier and CINV2 provides driving strength. CINV2 cannot drive the input SET logic of a SELB directly. Since the output voltage range of CINV2 is much larger than the period of a SET I-V curve: e/C_G . To solve this problem, we design a special SET inverter, SINV2, for which the gate and island are separated by a large distance in order to reduce the gate capacitance, C_G . Thus, e/C_G can match the output signal range of CMOS inverter CINT2.

Each SELB is equipped with a reconfigurable input switch fabric that connects local and global interconnects. This switch fabric is implemented with a multi-gate SET multiplexor tree.

III.B.5) Design and Modeling of IceFlex Majority Voting Logic: For SETs, noise resulting from random background offset charge effects is the primary reliability concern. We believe it likely that the random background offset problem will ultimately be dealt with by a combination of improved fabrication technology [19, 20], post-fabrication testing [22, 23] to identify and avoid a subset of the affected SETs, and run-time fault-tolerance via conventional structural redundancy [24, 25] or recent advances in probabilistic computation [26].

IceFlex provides for regular structural redundancy supported by SET MVL. SET MVLs are constructed from multi-gate SETs. Each SET

TABLE II
IMPACT OF MAJORITY VOTE LOGIC ON SELB FAULT PROBABILITY

Majority vote inputs	3	5	7
Raw fail probability	6.38×10^{-3}	6.38×10^{-3}	6.38×10^{-3}
Best probability	1.22×10^{-4}	2.57×10^{-6}	5.71×10^{-8}
SET MVL probability	1.22×10^{-4}	2.69×10^{-6}	1.77×10^{-7}

pull-up gate should be placed sufficiently far from the island. This requires the majority of the gates to be high in order to turn on the SET. The converse is true of the pull-down gates. In addition, multiple SETs are used in parallel in order to permit the failure of an MVL SET while still producing correct results.

We now consider the fault model for IceFlex SELBs. Every path from SELB input to output contains 64 SETs. Likharev estimates the long-term density of background offset charge susceptible SETs [3]. We follow his assumptions but correct a typographical error in that article, arriving at one susceptible SET in 10,000. Table II shows SELB failure probabilities using 2/3, 3/5, and 4/7 structural redundancy schemes. For each redundancy setup, we consider the effect of using no MVL (Raw final probability) fault-free MVL (Best probability), and SET MVL. Note that these are the probabilities of an SELB ever failing due to random background offset charge effects, not the probabilities of failure per unit time or per operation. In reality, many affected SELBs would be detected at synthesis time and avoided by synthesis software.

As shown in Table II it is possible for a seven-input SET MVL with redundant SELBs to reduce the failure rate to 1.77×10^{-7} . Given recent trends in noise-resistant SET design and fabrication, it seems likely that a less aggressive fault tolerance configuration will be necessary in the future [19, 20]. However, when we later consider the impact of fault tolerance on energy efficiency and performance, we assume the use of seven-input SET MVLs for every SELB stage.

IV. EXPERIMENTAL RESULTS

In this section, we evaluate the suitability of using SETs in low-power embedded systems. We start from the microarchitecture characterization of IceFlex. IceFlex is then used as a testbed to characterize the benefits and limitations of SETs for both high-performance and battery-powered embedded applications.

IV.A. Characterization of the IceFlex Architecture

We evaluate the performance and power consumption of IceFlex using HSPICE. For SET circuitry, the SPICE model and device parameters are described in Equation 1 and Table I. For CMOS logic and metal wire, we use the 65 nm Berkeley BSIM4 predictive technology model, which models the temperature impact on MOS devices.

Table III summarizes the performance and power characterization of the logic components and interconnect fabric of IceFlex, including multi-gate SET reconfigurable lookup table (LUT), SET register (Register), SET and CMOS 4/7 majority voting logic (MVL), multi-gate exclusive-or (MG), fast carry-out logic (CO), and SET input switch fabric (ISF), SET local interconnect (Single, Double, and Hex), and hybrid SET/CMOS global interconnect (Global). These results indicate that IceFlex has high energy efficiency, good performance, and high flexibility in terms of performance and energy efficiency tradeoff.

For the low-power setting, the power consumptions of SET-based logic components are within the range of nano-Watts. The power consumptions of the SET-based local interconnects are consistently below 1 nW. The hybrid SET/CMOS global interconnect has the highest power consumption. This is a result high global wire capacitance and high CMOS buffer power consumption. All components in the low-power version of IceFlex have latencies in the range of nanoseconds.

SETs have high junction resistance. Using the high-performance setting, SET junction resistances of 100 k Ω are used, reducing SET-based logic components latencies below 100 ps. Even though the resistance scaling results in a 100 \times increase in power, as demonstrated in Section IV-B, the overall energy efficiency of IceFlex is still orders of magnitude higher than that of CMOS-based solutions.

IV.B. Characterization of High-Performance and Battery-Powered Embedded Applications

In this section, we characterize the performance and power consumption of IceFlex when used to implement numerous general-purpose and application-specific processor cores. We evaluate the suitability

TABLE III
CHARACTERIZATION OF ICEFLEX MICROARCHITECTURE

		Low power							High performance							
		40 K	77 K	103 K	120 K	200 K	250 K	300 K	40 K	77 K	103 K	120 K	200 K	250 K	300 K	
Latency (ns)	LUT	9.38	8.25	7.93	7.96	7.59	7.51	7.37	0.10	0.09	0.08	0.07	0.06	0.07	0.07	
	Register	2.09	1.51	1.37	1.31	1.37	1.02	1.19	0.02	0.02	0.02	0.02	0.02	0.02	0.01	
	7-INPUT MVL	0.65	0.65	0.65	0.65	0.65	0.65	0.65	0.01	0.01	0.01	0.01	0.01	0.01	0.01	
	SET-MVL	1.85	1.83	1.83	1.82	1.82	1.82	1.81	0.04	0.03	0.03	0.02	0.02	0.02	0.02	
	Arithmetic	MG	1.31	1.29	1.29	1.29	1.28	1.28	0.02	0.02	0.02	0.02	0.01	0.01	0.01	
	Logic		CO	1.85	1.83	1.83	1.82	1.82	1.81	0.04	0.03	0.03	0.02	0.02	0.02	0.02
	ISF		7.50	6.60	6.34	6.37	6.07	6.01	5.90	0.27	0.10	0.08	0.08	0.06	0.06	0.06
Single, Double, Hex		3.82	3.65	3.58	3.59	3.54	3.56	3.74	0.3	0.3	0.3	0.3	0.2	0.2	0.3	
Global		15.04	27.32	15.50	28.09	17.27	12.61	11.10	0.34	0.36	0.42	0.39	1.50	2.50	1.43	
Power (nW)	LUT	0.08	0.31	0.55	0.75	2.09	3.26	4.70	5.65	20.94	36.44	47.59	96.51	346.28	442.85	
	Register	0.03	0.10	0.17	0.23	0.64	1.01	1.45	4.32	18.88	30.31	38.79	66.73	113.84	135.93	
	7 INPUT-MVL	0.02	0.08	0.14	0.19	0.54	0.84	1.21	2.29	14.55	26.54	34.81	72.94	94.39	113.51	
	SET-MVL	3.45E-3	0.01	0.02	0.03	0.09	0.13	0.19	1.57	4.11	6.61	8.39	17.06	22.21	27.08	
	Arithmetic	MG	2.90E-3	0.01	0.02	0.03	0.07	0.11	0.16	0.70	1.06	1.59	2.15	7.00	11.13	15.52
	Logic		CO	3.45E-3	0.01	0.02	0.03	0.09	0.13	0.19	1.57	4.11	6.61	8.39	17.06	22.21
	ISF		0.31	1.16	2.07	2.82	7.85	12.25	17.65	37.57	193.39	337.83	440.15	948.82	1300.60	1662.85
	Single		0.00	0.01	0.02	0.03	0.09	0.14	0.19	7.24	9.57	11.40	12.56	17.38	20.17	23.24
	Double		0.01	0.03	0.05	0.06	0.18	0.28	0.38	14.47	19.14	22.81	25.11	34.77	40.34	46.48
	Hex		0.01	0.05	0.09	0.13	0.36	0.56	0.76	28.94	38.28	45.62	50.22	69.54	80.68	92.97
	Global		372.92	84.76	184.52	35.78	42.43	66.12	60.22	9253.10	7946.40	6628.50	6746.50	2173.50	1530.00	3351.90

TABLE IV
ICEFLEX PERFORMANCE AND POWER CONSUMPTION AT ROOM TEMPERATURE

Benchmarks	FPGA			IceFlex											
	Xilinx Vertex-II XC2V2000			Non-Redundant Battery Powered			Non-Redundant High Performance			Redundant Battery Powered			Redundant High Performance		
	Freq (MHz)	Power (mW)	Energy req. (J/cycle)	Freq (MHz)	Power (mW)	Energy req. (J/cycle)	Freq (MHz)	Power (mW)	Energy req. (J/cycle)	Freq (MHz)	Power (mW)	Energy req. (J/cycle)	Freq (MHz)	Power (mW)	Energy req. (J/cycle)
ARM7	20.32	424	2.09×10^{-8}	1.00	0.13	1.33×10^{-10}	109.35	12.47	1.14×10^{-10}	0.96	0.92	9.56×10^{-10}	102.72	86.52	8.42×10^{-10}
ASPIDA DLX	97.09	611	6.29×10^{-9}	5.88	0.09	1.51×10^{-11}	645.16	8.37	1.30×10^{-11}	5.67	0.61	1.08×10^{-10}	606.06	57.69	9.52×10^{-11}
Jam RISC	74.05	469	6.33×10^{-9}	6.54	0.06	8.83×10^{-12}	716.85	5.44	7.59×10^{-12}	6.29	0.40	6.36×10^{-11}	673.40	37.72	5.60×10^{-11}
LEON2 SPARC	66.33	886	1.34×10^{-8}	4.52	0.27	5.87×10^{-11}	496.28	25.01	5.04×10^{-11}	4.36	1.85	4.24×10^{-10}	466.20	174.21	3.74×10^{-10}
Microblaze RISC	88.94	460	5.17×10^{-9}	8.40	0.04	4.61×10^{-12}	921.66	3.65	3.96×10^{-12}	8.09	0.26	3.26×10^{-11}	865.80	24.89	2.87×10^{-11}
miniMIPS	67.96	235	3.46×10^{-9}	4.90	0.11	2.33×10^{-11}	537.63	10.78	2.00×10^{-11}	4.72	0.79	1.67×10^{-10}	505.05	74.44	1.47×10^{-10}
MIPS	62.12	449	7.23×10^{-9}	5.35	0.06	1.04×10^{-11}	586.51	5.22	8.89×10^{-12}	5.15	0.38	7.42×10^{-11}	550.96	36.00	6.53×10^{-11}
Plasma	58.21	468	8.04×10^{-9}	4.52	0.08	1.72×10^{-11}	496.28	7.33	1.48×10^{-11}	4.36	0.54	1.25×10^{-10}	466.20	51.29	1.10×10^{-10}
UCore	105.34	613	5.82×10^{-9}	6.54	0.09	1.31×10^{-11}	716.85	8.06	1.12×10^{-11}	6.29	0.59	9.39×10^{-11}	673.40	55.71	8.27×10^{-11}
YACC	55.68	466	8.37×10^{-9}	9.80	0.07	7.38×10^{-12}	1075.27	6.81	6.34×10^{-12}	9.44	0.50	5.30×10^{-11}	1010.10	47.12	4.67×10^{-11}
AES	158.63	387	2.44×10^{-9}	14.71	0.08	5.62×10^{-12}	1612.90	7.79	4.83×10^{-12}	14.16	0.57	4.04×10^{-11}	1515.15	53.86	3.55×10^{-11}
AVR	55.55	105	1.89×10^{-9}	4.90	0.06	1.29×10^{-11}	537.63	5.94	1.11×10^{-11}	4.72	0.44	9.26×10^{-11}	505.05	41.18	8.15×10^{-11}
CORDIC	209.95	205	9.76×10^{-10}	58.82	0.03	4.44×10^{-13}	6451.61	2.46	3.81×10^{-13}	56.65	0.17	3.08×10^{-12}	6060.61	16.46	2.72×10^{-12}
ECC	30.24	105	3.47×10^{-9}	5.88	0.09	1.46×10^{-11}	645.16	8.07	1.25×10^{-11}	5.67	0.57	1.00×10^{-10}	606.06	53.45	8.82×10^{-11}
FPU	21.96	155	7.06×10^{-9}	1.31	0.26	2.00×10^{-10}	143.37	24.60	1.72×10^{-10}	1.26	1.83	1.45×10^{-9}	134.68	172.10	1.28×10^{-9}
RS	383.73	35	9.12×10^{-11}	29.41	0.00	1.05×10^{-13}	3225.81	0.29	9.04×10^{-14}	28.33	0.02	7.39×10^{-13}	3030.30	1.97	6.52×10^{-13}
USB	132.57	305	2.30×10^{-9}	19.61	0.07	3.53×10^{-12}	2150.54	6.52	3.03×10^{-12}	18.88	0.47	2.50×10^{-11}	2020.20	44.48	2.20×10^{-11}
VC	88.19	775	8.79×10^{-9}	11.76	0.29	2.50×10^{-11}	1290.32	27.75	2.15×10^{-11}	11.33	2.04	1.80×10^{-10}	1212.12	192.19	1.59×10^{-10}
Avg. Energy Imp.						201.38x			234.60x			27.05x			30.85x

of IceFlex for use in both high-performance and battery-powered embedded systems.

For high-performance applications, we consider ARM7: a power-efficient RISC CPU; ASPIDA DLX: a synchronous DLX Core; Jam RISC: a five-stage pipelined RISC CPU; LEON2 SPARC: a SPARC V8 compatible processor; Microblaze: a RISC CPU; miniMIPS, MIPS, Plasma, UCore, and YACC: five MIPS clones with varying features.

For battery-powered applications, we consider AES: a Rijndael IP Core; AVR: a microcontroller compatible with the ATMega103; CORDIC core: a coordinate rotation computer; ECC: an error correction code core; FPU: an IEEE 754 32-bit floating point unit; RS: a Reed Solomon encoder; USB: USB 2.0 function; and VC: a video compression core.

The Xilinx Virtex-II XC2V2000 FPGA is used for comparison. Each application is synthesized with Xilinx ISE to determine the number of required LUTs, maximum frequency, and power consumption, using a switching probability of 10% [27]. This synthesis flow permitted the system-level evaluation of all microarchitectural components in Table III except for multi-gate exclusive-or and fast carry-out logic. We used FPGA synthesis software to estimate the number of IceFlex SELBs required. 16-entry Virtex-II LUTs were used due to their functional (but not structural) similarity to IceFlex SELBs. For each design, the maximum frequency for IceFlex was determined by multiplying the

number of combinational SELBs along the longest combinational path by the delay of an IceFlex SELB plus the delay of a local interconnect. The Xilinx ISE synthesis software did not use global interconnects for any of the synthesized processors. IceFlex power consumption was computed by taking the sum of the power consumptions of all components at the maximum operating frequency.

Table IV shows the operating frequencies, power consumptions, and energy efficiency in Joules per clock cycle of XC2V2000 and various versions of IceFlex for each benchmark application. This table characterizes IceFlex for both high-performance and battery-powered embedded applications. As described in Section III-A.5, recent progress in fabrication is reducing the severity of the background charge effects. Therefore, we show the characteristics of both the spatially-redundant and non-spatially-redundant versions of IceFlex in Table IV.

IV.B.1) Power-Efficient High-Performance Computing: We can draw the following general conclusions from Table IV. For a wide range of processor cores, the IceFlex architecture is capable of achieving energy efficiencies two orders of magnitude better than CMOS-based FPGAs. Peak frequencies ranging from 100 MHz to 1,100 MHz are maintained for all processors.

One might expect the high-performance version of IceFlex to consistently achieve higher frequency but require more Joules per clock cycle than the low-power version of IceFlex. However, it typically

requires slightly fewer Joules per clock cycle, as well. Joules per clock cycle are computed at the maximum clock frequency of each processor. This has the effect of reducing the impact of static power consumption for processors with higher peak frequencies. If one must operate at a low frequency, the power consumption of the low-power version of IceFlex will generally be lower than that of the high-performance version. However, the reported values of Joules per clock cycle at maximum frequency have interesting implications. Although SETs have extremely low power consumption, at room-temperature a large percentage of this power consumption can be attributed to static power (see Figure 2). Therefore, for SET-based architectures that are operated at room temperature, it will generally be more energy efficient to operate at high frequency and frequently enter a power-gated sleep mode than to continuously operate at a low frequency.

In high-performance applications for which parallel computation is appropriate, improved energy efficiency can be traded for improved performance with the same energy budget. For example, given a power budget of 1 W, one could use one LEON2 SPARCs implemented with FPGAs or 40 LEON2 SPARCs implemented with the high-performance variant of IceFlex. This implies an overall performance $301\times$ higher than that of the FPGA version. Taken to its logical extreme, given a power budget of 100 W and one instruction per cycle, one could execute at 1.6 Terra IPS; a 1 kW power budget would permit 16 Terra IPS. These numbers are intended to give the reader some indication of the potential to improve performance given a power budget. In practice some of this performance will be lost due to parallelization inefficiency and off-chip communication latency. A similar comparison can be used for the MIPS processor, for which IceFlex permits a $813\times$ improvement in energy efficiency compared with an FPGA implementation.

IV.B.2) Ultra-Low-Power Embedded systems: From the data in Table IV, we can conclude that the non-redundant, room temperature, low-power version of IceFlex is suitable for use in ultra-low-power applications such as sensor network nodes. In the following analysis, we shall focus on the AVR core, which is representative of the most commonly-used architecture for sensor network nodes. The power consumption of IceFlex is low enough to permit an AVR processor to operate at 4 MHz for 9.7 years using the energy of a single AA battery, i.e., to the shelf life of the battery.

If we assume an energy scavenging volume of 5 cm^3 and use Roundy's power densities of $4\mu\text{W}/\text{cm}^3$ for indoor solar energy, $200\mu\text{W}/\text{cm}^3$ for vibrations, $10\mu\text{W}/\text{cm}^3$ for daily temperature variation, and $0.003\mu\text{W}/\text{cm}^3$ for acoustic noise at 75 dB [28], we find that one sensor network node is capable of scavenging enough energy for an IceFlex AVR processor running at the maximum clock frequency from vibrations or daily temperature variation, at 1.6 MHz from indoor solar energy, and at 1 kHz from 75 dB acoustic noise.

V. CONCLUSIONS

In this article, we have analyzed the impact of using SETs in architecture and circuit design; proposed IceFlex, a fault-tolerant, hybrid SET/CMOS, reconfigurable architecture for use in high-performance and battery-powered embedded systems; and evaluated the energy efficiency, power consumption, and performance of IceFlex in these applications. Although using SETs for computation poses many design challenges, many of these challenges can be solved with the proposed architecture and circuit design techniques. SETs have some unique properties that permit significant improvements in energy efficiency compared with BJT, CMOS based design. In summary, we find that a hybrid SETs/CMOS architecture has the potential to improve energy efficiency in high-performance applications by $235\times$ compared with today's CMOS while permitting operating frequencies that are as high, or higher. This improved energy efficiency can be traded for performance when operating within a power dissipation budget. In battery-powered embedded systems, such as sensor network nodes, SETs have the potential to increase energy efficiency by $201\times$, thereby permitting corresponding increased in battery lifespan or permitting operation on scavenged energy. Although they hold great promise, the practical use of SETs will require additional research into fault tolerance techniques, processing technologies, and novel circuit designs. It is our hope that this article provide a starting point for additional research in this area and reveals the potential advantages of SET-based architectures.

- [1] M. S. Dresselhaus, G. Dresselhaus, and P. Avouris, *Carbon Nanotubes*. Springer-Verlag, Germany, Feb. 2001.
- [2] Y. Huang, et al., "Logic gates and computation from assembled nanowire building blocks," *Nature*, vol. 294, no. 5545, pp. 1313–1317, Nov. 2001.
- [3] K. K. Likharev, "Single-electron devices and their applications," *Proc. IEEE*, vol. 87, no. 4, pp. 606–632, Apr. 1999.
- [4] E. S. Soldatov, et al., "Room temperature molecular single-electron transistor," *JETP Ltrs.*, vol. 64, pp. 556–558, Oct. 1996.
- [5] J.-I. Shirakashi, et al., "Single-electron charging effects in Nb/Nb oxide-based single-electron transistors at room temperature," *Applied Physics Ltrs.*, vol. 72, no. 15, pp. 1893–1895, Apr. 1998.
- [6] J. R. Tucker, "Complementary digital logic based on the Coulomb blockade," *J. Applied Physics*, vol. 72, no. 99, pp. 4399–4413, 1992.
- [7] R. H. Chen, A. N. Korotkov, and K. K. Likharev, "Single-electron transistor logic," *Applied Physics Ltrs.*, vol. 68, no. 14, pp. 1954–1956, Apr. 1996.
- [8] Y.-K. Cho and Y.-H. Jeong, "Single-electron pass-transistor logic with multiple tunnel junctions and its hybrid circuit with MOSFETs," *ETRI J.*, vol. 26, no. 6, pp. 669–672, Dec. 2004.
- [9] S. Mahapatra, A. M. Ionescu, and K. Banerjee, "A quasi-analytical SET model for few electron circuit simulation," *IEEE Electron Device Ltrs.*, vol. 23, no. 6, pp. 366–368, June 2002.
- [10] K. Uchida, et al., "Programmable single-electron transistor logic for future low-power intelligent LSI: proposal and room-temperature operation," *IEEE Trans. Electron Devices*, vol. 50, no. 7, pp. 1623–1630, July 2003.
- [11] K. K. Yadavalli, et al., "Single electron memory devices: toward background charge insensitive operation," *J. Vacuum Science Technology B Microelectronics and Nanometer Structures*, vol. 21, pp. 2860–2864, 2003.
- [12] C. Wasshuber, H. Kosina, and S. Selberherr, "A single-electron device and circuit simulator," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 16, pp. 937–944, Sept. 1997.
- [13] R. H. Chen, "MOSES: a general Monte Carlo simulator for single-electron circuits," *Meeting Abstracts, The Electrochemical Society*, vol. 96, no. 2, p. 576, Oct. 1996.
- [14] K. Uchida, et al., "Analytical single-electron transistor (SET) model for design and analysis of realistic set circuits," *Japanese. J. Applied Physics*, vol. 39, pp. 2321–2324, Apr. 2000.
- [15] H. Inokawa and Y. Takahashi, "A compact analytical model for asymmetric single-electron tunneling transistors," *IEEE Trans. Electron Devices*, vol. 50, no. 2, pp. 455–461, Feb. 2003.
- [16] S. Mahapatra, et al., "A CAD framework for co-design and analysis of CMOS-SET hybrid integrated circuits," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2003, pp. 497–502.
- [17] A. M. Ionescu, S. Mahapatra, and V. Pott, "Hybrid SETMOS architecture with coulomb blockade oscillations and high current drive," *IEEE Electron Device Ltrs.*, vol. 25, no. 6, pp. 411–413, June 2004.
- [18] M. Kirihara, K. Nakazato, and M. Wagner, "Hybrid circuit simulator including a model for single electron tunneling devices," *Japanese J. of Applied Physics*, vol. 38, no. 4A, Apr. 1999.
- [19] V. A. Krupenin, et al., "Aluminum single electron transistors with islands isolated from a substrate," *J. of Low Temperature Physics*, vol. 118, no. 5/6, Dec. 1999.
- [20] N. M. Zimmerman, et al., "Excellent charge offset stability in Si-based SET transistors," in *Proc. Precision Electromagnetic Measurements*, Nov. 2002, pp. 124–125.
- [21] J. M. Rabaey, *Digital Integrated Circuits*. Prentice-Hall, NJ, 1998.
- [22] K. R. Brown, L. Sun, and B. E. Kane, "Electric-field-dependent spectroscopy of charge motion using a single-electron transistor," *Applied Physics Ltrs.*, vol. 88, 2006 May.
- [23] S. C. Goldstein and M. Budiuh, "Nanofabrics: spatial computing using molecular electronics," in *Proc. Int. Symp. Computer Architecture*, June 2001, pp. 178–189.
- [24] D. Bhaduri and S. K. Shukla, "NANOPRISM: a tool for evaluating granularity vs. reliability trade-offs in nano architectures," in *Proc. Great Lakes Symp. VLSI*, Mar. 2004, pp. 109–112.
- [25] A. DeHon, "Array-based architecture for FET-based nanoscale electronics," *IEEE Trans. Nanotechnology*, vol. 2, no. 1, Mar. 2003.
- [26] R. I. Bahar, J. Mundy, and J. Chen, "A probabilistic-based design methodology for nanoscale computation," in *Proc. Int. Conf. Computer-Aided Design*, Nov. 2003, pp. 480–486.
- [27] "Xilinx XPower," <http://www.xilinx.com>.
- [28] S. Roundy, P. K. Wright, and J. Rabaey, "A study of low level vibrations as a power source for wireless sensor nodes," *Computer Communications*, vol. 26, Oct. 2003.